



ORACLE®
Cloud Infrastructure

User Guide

11/26/2019

Copyright © 2016, 2019, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement

between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

CONTENTS

CHAPTER 1 About Oracle Cloud Infrastructure	19
Prefer Online Help?	21
Need API Documentation?	21
CHAPTER 2 Oracle Cloud Infrastructure's Free Tier	22
Free Trial	22
Always Free Resources	23
Upgrading to a Paid Account	25
Additional Information	25
Details of the Always Free Resources	25
Frequently Asked Questions: Oracle Cloud Infrastructure Free Tier	30
CHAPTER 3 Oracle Cloud Infrastructure for Government	33
For All Government Cloud Customers	33
Government Cloud with FedRAMP Authorization	43
Federal Government Cloud with Impact Level 5 Authorization	50
CHAPTER 4 Service Essentials	58
Security Credentials	59
Regions and Availability Domains	62
IP Address Ranges	69
Resource Identifiers	75
Resource Tags	79
Resource Monitoring	84

Table of Contents

Service Limits	104
Viewing All Resources in a Compartment	127
Compartment Quotas	132
Work Requests	151
Console Announcements	155
Prerequisites for Oracle Platform Services on Oracle Cloud Infrastructure	163
Billing and Payment Tools Overview	175
Console Cookies and Local Storage	190
My Services Use Cases	192
CHAPTER 5 Archive Storage	245
Overview of Archive Storage	245
CHAPTER 6 Audit	251
Overview of Audit	251
Contents of an Audit Log Event	253
Viewing Audit Log Events	262
Setting Audit Log Retention Period	267
Bulk Export of Audit Log Events	269
CHAPTER 7 Block Volume	273
Overview of Block Volume	273
iSCSI Commands and Information	282
Volume Groups	284
Creating a Volume	296
Attaching a Volume	299
Connecting to a Volume	305
Listing Volumes	314
Listing Volume Attachments	316
Listing Boot Volume Attachments	317
Renaming a Volume	318
Resizing a Volume	319
Overview of Block Volume Backups	329

Table of Contents

Backing Up a Volume	338
Policy-Based Backups	340
Copying a Volume Backup Between Regions	352
Restoring a Backup to a New Volume	356
Cloning a Volume	359
Disconnecting From a Volume	363
Detaching a Volume	364
Deleting a Volume	366
Move Block Volume Resources Between Compartments	367
Block Volume Performance	372
Block Volume Metrics	398
CHAPTER 8 Compute	402
Overview of the Compute Service	402
Best Practices for Your Compute Instance	407
Protecting Data on NVMe Devices	420
Boot Volumes	431
Oracle-Provided Images	470
Compute Shapes	505
Installing and Running Oracle Ksplice	513
Managing Custom Images	514
Image Import/Export	525
Bring Your Own Image (BYOI)	533
OS Kernel Updates	554
Managing Key Pairs on Linux Instances	557
Creating an Instance	560
Managing Compute Instances	579
Autoscaling	599
Managing Cluster Networks	609
Dedicated Virtual Machine Hosts	615
Connecting to an Instance	623
Instance Console Connections	628

Table of Contents

Adding Users on an Instance	638
Displaying the Console for an Instance	641
Getting Instance Metadata	642
Updating Instance Metadata	645
Renaming an Instance	648
Moving a Compute Instance to a New Host	649
Migrating an Instance from a Local to Remote Boot Volume	654
Moving Compute Resources to a Different Compartment	658
Stopping and Starting an Instance	663
Terminating an Instance	667
Enabling Monitoring for Compute Instances	669
Compute Metrics	685
Compute Performance	701
Compute Health Monitoring for Bare Metal Instances	704
Microsoft Licensing on Oracle Cloud Infrastructure	707
Updating the Linux iSCSI Service to Restart Automatically	721
CHAPTER 9 Container Engine for Kubernetes	727
Overview of Container Engine for Kubernetes	727
Preparing for Container Engine for Kubernetes	730
About Kubernetes Clusters and Nodes	782
Creating a Kubernetes Cluster	784
Downloading a kubeconfig File to Enable Cluster Access	794
Modifying a Kubernetes Cluster	801
Deleting a Kubernetes Cluster	803
Monitoring Clusters	804
Monitoring Operations of Container Engine for Kubernetes	807
Accessing a Cluster Using kubectl	808
Starting the Kubernetes Dashboard	809
Deploying a Sample Nginx App on a Cluster Using kubectl	812
Pulling Images from Registry during Deployment	813
Encrypting Kubernetes Secrets At Rest in Etcd	816

Table of Contents

Connecting to Worker Nodes Using SSH	820
About Access Control and Container Engine for Kubernetes	823
Kubernetes Versions and Container Engine for Kubernetes	828
'Upgrading' the image running on worker nodes by creating a new node pool	834
Creating Load Balancers to Distribute Traffic Between Cluster Nodes	835
Creating a Persistent Volume Claim	845
Adding OCI Service Broker for Kubernetes to Clusters	849
Example: Setting Up an Ingress Controller on a Cluster	851
CHAPTER 10 Data Transfer	861
Overview of Data Transfer Service	861
Disk Data Transfer	863
Appliance Data Transfer	920
Troubleshooting	990
CHAPTER 11 Database	995
Overview of the Database Service	995
Overview of Autonomous Database	1001
Exadata Cloud at Customer	1093
Exadata DB Systems	1239
Bare Metal and Virtual Machine DB Systems	1392
Database Metrics	1692
Using Performance Hub to Analyze Database Performance in Oracle Cloud Infrastructure	1699
Migrating Databases to the Cloud	1705
Troubleshooting	1769
CHAPTER 12 DNS and Traffic Management	1798
Overview of the DNS Service	1798
Overview of the Traffic Management Steering Policies Service	1840
CHAPTER 13 Email Delivery	1891
Overview of the Email Delivery Service	1891

Table of Contents

Getting Started with Email Delivery	1899
Generate SMTP Credentials for a User	1911
Managing Approved Senders	1912
Configure SPF	1915
Configure SMTP Connection	1917
Managing the Suppression List	1918
Email Delivery Metrics	1920
Integrating Oracle Application Express with Email Delivery	1923
Integrating Postfix with Email Delivery	1925
Integrating Oracle Enterprise Manager with Email Delivery	1928
Integrating Mailx with Email Delivery	1930
Integrating Swaks with Email Delivery	1931
Integrating JavaMail with Email Delivery	1934
Integrating Sendmail with Email Delivery	1939
Integrating PeopleSoft with Email Delivery	1943
Integrating Python with Email Delivery	1948
Troubleshooting Email Delivery	1951
Deliverability Best Practices	1957
CHAPTER 14 Events	1963
Overview of Events	1963
Getting Started with Events	1968
Matching Events with Filters	1985
Events and IAM Policies	1992
Managing Rules for Events	1994
Contents of an Event Message	2019
Services that Produce Events	2022
Events Metrics	2100
CHAPTER 15 File Storage	2106
Overview of File Storage	2106
About Security	2114

Table of Contents

Creating File Systems	2137
Mounting File Systems	2151
Managing File Systems	2183
Managing Mount Targets	2202
Managing Snapshots	2239
Paths in File Systems	2246
Troubleshooting Your File System	2248
CHAPTER 16 Functions	2273
Overview of Functions	2273
Oracle Functions Concepts	2276
How Oracle Functions Works	2278
Preparing for Oracle Functions	2282
Creating, Deploying, and Invoking a Helloworld Function	2316
Creating Applications	2320
Creating and Deploying Functions	2322
Creating Functions from Existing Docker Images	2325
Viewing Functions and Applications	2329
Invoking Functions	2332
Storing and Viewing Function Logs	2339
Updating Functions	2342
Deleting Applications and Functions	2346
Passing Custom Configuration Parameters to Functions	2350
Accessing File Systems from Running Functions	2354
Accessing Other Oracle Cloud Infrastructure Resources from Running Functions	2355
Permissions Granted to Containers Running Functions	2363
Invoking Oracle Functions from Other Oracle Cloud Infrastructure Services	2364
Changing Oracle Functions Default Behavior	2364
Differences between Oracle Functions and Fn Project	2366
Troubleshooting Oracle Functions	2368
Function Metrics	2379

Table of Contents

CHAPTER 17 Health Checks	2387
Overview of the Health Checks Service	2387
Getting Started With the Health Checks API	2390
Managing Health Checks	2398
Health Checks Metrics	2404
CHAPTER 18 IAM	2411
Overview of Oracle Cloud Infrastructure Identity and Access Management	2411
Getting Started with Policies	2428
How Policies Work	2431
Common Policies	2441
Advanced Policy Features	2465
Policy Syntax	2471
Policy Reference	2477
User Credentials	2674
Federating with Identity Providers	2677
User Provisioning for Federated Users	2760
Managing User Capabilities for Federated Users	2766
Calling Services from an Instance	2771
Managing Users	2777
Managing Groups	2787
Managing Dynamic Groups	2793
Managing Compartments	2803
Managing Regions	2822
Managing Platform Services Regions	2826
Managing the Tenancy	2829
Managing Policies	2831
Managing User Credentials	2838
Managing Authentication Settings	2851
Managing Multi-Factor Authentication	2855

Table of Contents

CHAPTER 19 Key Management	2866
Overview of Key Management	2866
Managing Keys	2870
Managing Vaults	2891
Using Keys	2900
Key Management Metrics	2904
CHAPTER 20 Load Balancing	2909
Overview of Load Balancing	2909
How Load Balancing Policies Work	2923
Connection Management	2926
HTTP "X-" Headers	2928
Session Persistence	2930
Managing a Load Balancer	2936
Managing Backend Sets	2959
Managing Backend Servers	2968
Managing Load Balancer Listeners	2978
Managing Request Routing	2983
Managing Rule Sets	2998
Managing SSL Certificates	3017
Editing Health Check Policies	3029
Viewing the State of a Work Request	3036
Load Balancing Metrics	3038
CHAPTER 21 Marketplace	3053
Overview of Marketplace	3053
Working with Listings	3054
Viewing Terms of Use Agreements for Deployed Applications	3061
CHAPTER 22 Monitoring	3064
Monitoring Overview	3064
Viewing Default Metric Charts	3084
Building Metric Queries	3130

Table of Contents

Publishing Custom Metrics	3160
Managing Alarms	3164
Best Practices for your Alarms	3198
Monitoring Query Language (MQL) Reference	3201
CHAPTER 23 Networking	3210
Overview of Networking	3210
Scenario A: Public Subnet	3224
Scenario B: Private Subnet with a VPN	3234
Scenario C: Public and Private Subnets with a VPN	3250
Transit Routing: Access to Multiple VCNs in the Same Region	3270
Transit Routing: Private Access to Oracle Services	3315
FastConnect with Multiple DRGs and VCNs	3342
VCNs and Subnets	3349
Ways to Secure Your Network	3362
Access Control	3365
Security Rules	3369
Network Security Groups	3386
Security Lists	3403
Virtual Network Interface Cards (VNICs)	3413
VNIC Metrics	3430
Private IP Addresses	3439
Public IP Addresses	3458
IPv6 Addresses	3473
DNS in Your Virtual Cloud Network	3512
DHCP Options	3522
Route Tables	3531
Dynamic Routing Gateways (DRGs)	3545
Routing Details for Connections to Your On-Premises Network	3555
VPN Connect	3560
VPN Connect Overview	3561
VPN Connect Quickstart	3571

Table of Contents

Supported IPSec Parameters	3579
Supported Encryption Domain or Proxy ID	3584
Setting Up VPN Connect	3587
CPE Configuration	3620
Working with VPN Connect	3912
VPN Connect FAQ	3925
Using the API for VPN Connect	3926
VPN Connect Metrics	3929
VPN Connect Troubleshooting	3933
FastConnect	3943
FastConnect Overview	3944
FastConnect Requirements	3959
FastConnect Redundancy Best Practices	3967
FastConnect: With an Oracle Provider	3977
FastConnect: With a Third-Party Provider	3991
FastConnect: Colocation with Oracle	4013
FastConnect Public Peering Advertised Routes	4034
FastConnect Metrics	4061
FastConnect Troubleshooting	4068
Internet Gateway	4073
NAT Gateway	4082
Access to Oracle Services: Service Gateway	4092
Access to Other VCNs: Peering	4111
Local VCN Peering (Within Region)	4116
Remote VCN Peering (Across Regions)	4140
Access to Oracle Cloud Infrastructure Classic	4159
Connection Over Oracle Network	4159
Connection Over IPSec VPN	4168
Access to Microsoft Azure	4177
Access to Other Clouds with Libreswan	4195
Network Performance	4212
Troubleshooting	4214

Table of Contents

CHAPTER 24 Notifications	4228
Notifications Overview	4228
Managing Topics and Subscriptions	4235
Publishing Messages	4246
Notifications Metrics	4248
CHAPTER 25 Object Storage	4253
Overview of Object Storage	4253
Understanding Object Storage Namespaces	4259
Managing Buckets	4261
Managing Objects	4296
Copying Objects	4317
Using Pre-Authenticated Requests	4324
Using Multipart Uploads	4333
Using Object Lifecycle Management	4340
Object Storage Metrics	4352
Hadoop Support	4361
Designating Compartments for the Amazon S3 Compatibility and Swift APIs	4362
Amazon S3 Compatibility API	4366
CHAPTER 26 Registry	4373
Overview of Registry	4373
Preparing for Registry	4376
About Images	4378
About Repositories	4378
Creating a Repository	4379
Pushing Images Using the Docker CLI	4381
Pulling Images Using the Docker CLI	4386
Pulling Images from Registry during Kubernetes Deployment	4388
Viewing Images and Image Details	4391
Deleting an Image	4392
Retaining and Deleting Images Using Retention Policies	4393

Table of Contents

Deleting a Repository	4400
Getting an Auth Token	4400
Policies to Control Repository Access	4401
CHAPTER 27 Overview of Resource Manager	4405
Key Concepts	4405
Generalized Workflow	4407
Ways to Access Resource Manager	4409
Authentication and Authorization	4409
Limits on Resource Manager Resources	4410
Managing Stacks and Jobs	4411
Sample: Creating a Compute Instance Using Resource Manager	4439
Using Remote Exec	4461
Terraform Configurations for Resource Manager	4465
CHAPTER 28 Search	4473
Overview of Search	4473
Search Language Syntax	4478
Sample Queries	4486
Querying Resources	4491
Troubleshooting Search	4497
CHAPTER 29 Security Guide and Announcements	4499
Oracle Cloud Infrastructure Security Guide	4499
Oracle Cloud Security Response to Intel L1TF Vulnerabilities	4614
Oracle Cloud Security Response to Intel Microarchitectural Data Sampling (MDS) Vulnerabilities	4624
CHAPTER 30 Storage Gateway	4635
Overview of Storage Gateway	4635
Features of Storage Gateway	4643
Getting Started With Storage Gateway	4647
Configuring the Cache for File Systems	4648

Table of Contents

Understanding Storage Gateway Performance	4659
Interacting With Object Storage	4663
Installing Storage Gateway	4667
Logging In to the Storage Gateway Management Console	4680
Creating Your First File System	4682
Managing File Systems	4688
Managing Storage Gateway	4704
Using Storage Gateway File Management Operations	4708
Monitoring Storage Gateway	4715
Using Storage Gateway Cloud Sync	4720
Best Practices for Using Storage Gateway	4728
Troubleshooting Storage Gateway	4729
Upgrading Storage Gateway	4733
Uninstalling Storage Gateway	4741
Getting Help with Storage Gateway	4742
CHAPTER 31 Streaming Service Overview	4744
Streaming Usage Scenarios	4744
Streaming Concepts	4745
How Streaming Works	4746
Stream Archiving	4747
Limits on Streaming Resources	4747
Resource Identifiers	4748
Ways to Access Oracle Cloud Infrastructure	4749
Using Streaming	4749
Authentication and Authorization	4749
Tagging Resources	4750
Moving Resources to a Different Compartment	4750
Managing Streams	4750
Publishing Messages	4754
Consuming Messages	4756
Using the Streaming SDK	4758

Table of Contents

Streaming Metrics	4767
CHAPTER 32 Tagging	4776
Tagging Overview	4776
Managing Tags and Tag Namespaces	4781
Using Cost-Tracking Tags	4799
Using Predefined Values	4802
Using Tag Variables	4806
Managing Tag Defaults	4807
CHAPTER 33 Web Application Firewall	4817
Overview of the Web Application Firewall Service	4817
Getting Started with WAF	4821
Managing WAF Policies	4830
Origin Management	4836
Bot Management	4840
WAF Protection Rules	4848
Access Control	4883
Caching Rules	4891
Threat Intelligence	4895
Settings	4899
Logs	4902
CHAPTER 34 Developer Tools	4912
Software Development Kits and Command Line Interface	4912
Other Tools and Plug-ins	4984
Properties	4994
Tools Configuration	5068
REST APIs	5078
GLOSSARY	5179
RELEASE NOTES	5202

CHAPTER 1 About Oracle Cloud Infrastructure

Oracle Cloud Infrastructure provides bare metal cloud infrastructure that lets you create networking, compute, and storage resources for your enterprise workloads.

If you're new to Oracle Cloud Infrastructure and would like to learn some key concepts and take a quick tutorial, see the *Oracle Cloud Infrastructure Getting Started Guide*.

If you're ready to create cloud resources such as users, access controls, cloud networks, instances, and storage volumes, this guide is right for you. It provides the following information about using Oracle Cloud Infrastructure:

Service	What's Covered	Chapter
Archive Storage	Preserving cold data.	Archive Storage
Audit	Logging activity in your cloud.	Audit
Block Volume	Adding storage capacity to instances.	Block Volume
Compute	Launching compute instances and connecting to them by using an SSH key pair.	Compute
Container Engine for Kubernetes	Defining and creating Kubernetes clusters to enable the deployment, scaling, and management of containerized applications.	Container Engine for Kubernetes
Data Transfer	Migrating large volumes of data.	Data Transfer
Database	Creating and managing database systems and Oracle Databases.	Database

CHAPTER 1 About Oracle Cloud Infrastructure

Service	What's Covered	Chapter
Edge Services	Encompasses several services that allow you to manage, secure, and maintain your domains and endpoints.	Edge Services
Email Delivery	Sending large volume email.	Email Delivery
Events	Creating automation in your tenancy.	Events
File Storage	Managing shared file systems, mount targets, and snapshots.	File Storage
Functions	Building and deploying applications and functions.	Functions
IAM	Setting up administrators, users, and groups and specifying their permissions to access to cloud resources.	IAM
Key Management	Creating and managing encryption keys and key vaults to control the encryption of your data.	Key Management
Load Balancing	Setting up load balancers, listeners, backend sets, certificate bundles, and managing health check policies.	Load Balancing
Monitoring	Querying metrics and managing alarms to monitor the health, capacity, and performance of your cloud resources.	Monitoring
Networking	Setting up cloud networks, subnets, gateways, route tables, and security lists.	Networking

Service	What's Covered	Chapter
Notifications	Setting up topics and subscriptions, and publishing messages.	Notifications
Object Storage	Creating and managing buckets to store objects, and uploading and accessing data files.	Object Storage
Registry	Storing, sharing, and managing development artifacts like Docker images in an Oracle-managed registry.	Registry
Search	Searching for Oracle Cloud Infrastructure resources using free text search or advanced queries.	Search
Tagging	Adding metadata tags to your resources.	Tagging

For a description of the terminology used throughout this guide, see the [GLOSSARY](#).

Prefer Online Help?

The information in this guide and the *Getting Started Guide* is also available in the online help at <https://docs.cloud.oracle.com/iaas/Content/home.htm>.

Need API Documentation?

For general information, see [REST APIs](#). For links to the detailed service API documentation, see the online help at <https://docs.cloud.oracle.com/iaas/Content/home.htm>.

CHAPTER 2 Oracle Cloud Infrastructure's Free Tier

Oracle Cloud Infrastructure's Free Tier is composed of a free promotional trial that allows you to explore a wide range of Oracle Cloud Infrastructure products, and a set of Always Free offers that never expire.

Free Trial

The Free Trial provides you with \$300 of cloud credits that are valid for up to 30 days. You may spend these credits on any eligible Oracle Cloud Infrastructure service.

Getting Started

[Start for Free](#)

For more information, and to see a complete list of services available to you during the trial, visit the [Free Trial website](#).



Tip

During sign up, choose the home region carefully. Most Always Free resources can be provisioned only in your home region.

For security purposes, most users will need a mobile phone number and a credit card to create an account. Your credit card will not be charged unless you upgrade your account.

What Happens When Your Trial Period Ends

After your trial ends, your account remains active. There is no interruption to the availability of the [Always Free Resources](#) you have provisioned. You can terminate and re-provision Always Free resources as needed.

Paid resources that were provisioned with your credits during your free trial are reclaimed by Oracle unless you upgrade your account.

Pay as You Go accounts are available with no commitment, or contact an [Oracle sales representative](#) in your location to learn about monthly and annual flex accounts that offers discounted pricing. See the [pricing details](#) learn more.

Always Free Resources

All Oracle Cloud Infrastructure accounts (whether free or paid) have a set of resources that are free of charge for the life of the account. These resources display the **Always Free** label in the Console.

Using the Always Free resources, you can provision a [virtual machine \(VM\) instance](#), an [Oracle Autonomous Database](#), and the networking, load balancing, and storage resources needed to support the applications that you want build. With these resources, you can do things like run small-scale applications or perform proof-of-concept testing.

The following list summarizes the Oracle Cloud Always Free-eligible resources that you can provision in your tenancy:

- **Compute** (up to two instances)
- **Autonomous Database** (up to two database instances)
- **Load Balancing** (one load balancer)
- **Block Volume** (up to 100 GB total storage)
- **Object Storage** (up to 20 GiB)

For detailed information about the Always Free resources, see [Details of the Always Free Resources](#).

You can find your tenancy's limits for Always Free resources in the Console. To check these limits: Open the navigation menu. Under **Governance and Administration**, go to **Limits, Quotas and Usage**.

Quickly Spinning Up Your Environment Using Terraform

Oracle offers a Terraform configuration file that lets you automatically create the full set of Always Free resources in a few minutes. You don't need to have experience with Terraform to use this file. To set up your Oracle Cloud environment, you log in to your account and then upload the file in the Console using the [Resource Manager](#) service. Resource Manager reads the file and provisions the Always Free resources for you with the settings and configuration you need to start creating applications in the cloud.

Note that Terraform refers to the set of resources being provisioned as a "stack". For a general introduction to Terraform and the "infrastructure-as-code" model, see [Terraform: Write, Plan, and Create Infrastructure as Code](#).

To provision your Always Free using Terraform and Resource Manager

1. Download the Always Free Terraform configuration file (a .zip file) to your computer. To do this, navigate to <https://github.com/oracle/oci-quickstart-cloudnative> and follow the instructions in the `readme.md` file included in this repository.
2. Log into your Oracle Cloud Infrastructure account.
3. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
4. Click the **Create Stack** button to open the **Create Stack** dialog.
5. Add your downloaded configuration (.zip) file, either by dragging and dropping it onto the dialog's control, or by clicking **Browse** and navigating to the file location.
6. Optionally, provide a name for the new stack. If you don't provide a name, a default name is provided on the server.
7. Optionally, provide a description for the stack.

8. Optionally, select a different compartment from your current compartment in which to create the stack. To do so, select a compartment from the **Create In Compartment** drop-down.
9. Click **Next** to proceed to the **Configure Variables** panel.
10. The variables displayed in the **Configure Variables** panel are auto-populated from the Terraform file that you uploaded. You don't need to change these variables if you are provisioning your Always Free resources using the Terraform file provided by Oracle.
11. Click **Next** to proceed to the **Review** panel.
12. Verify your stack configuration, then click **Create** to create your stack.

Your set of Always Free resources should take no more than a few minutes to provision.

Upgrading to a Paid Account

You can upgrade to a paid account at any time through the Oracle Cloud Infrastructure Console. To do so, click the **Upgrade** link in the banner at the top of the Console web page. If you don't see an **Upgrade** link on the page you are viewing, you can click the **Oracle Cloud** logo at the top of the Console and then look for the Upgrade link in the sidebar on the right side of the Console home page.

Additional Information

See [Frequently Asked Questions: Oracle Cloud Infrastructure Free Tier](#) for answers to your questions about Free Tier accounts and resources.

Details of the Always Free Resources

This topic provides reference information on Oracle Cloud Infrastructure's Always Free resources.

Compute

All tenancies get two Always Free [Compute virtual machine \(VM\) instances](#).

Details of the Always Free Compute instance

- **Shape:** VM.Standard.E2.1.Micro
- **Processor:** 1/8th of an OCPU with the ability to use additional CPU resources
- **Memory:** 1 GB
- **Networking:** Includes one VNIC with one public IP address and up to 480 Mbps network bandwidth
- **Operating System:** Your choice of one of the following Always Free-eligible operating systems:
 - Oracle Linux
 - Canonical Ubuntu Linux
 - CentOS Linux



Tip

The Linux operating systems labeled "Always Free Eligible" in the Console are compatible with Always Free Compute instances and incur no licensing fees. These operating systems are also compatible with paid resources and are available to users of paid accounts. To provision a Compute instance with an operating system that is not Always Free-eligible, you must have a paid account or a Free Trial account with available credits.

See [Oracle-Provided Images](#) for more information about the available operating systems. For steps to create an Always Free-eligible Compute instance, see "Tutorial - Launching Your First Linux Instance" in the *Oracle Cloud Infrastructure Getting Started Guide*.

Database

All tenancies get two Always Free Oracle [Autonomous Databases](#). The Autonomous Databases use [serverless deployment](#) (meaning Oracle handles the database infrastructure provisioning and maintenance). For current regional availability, see [Always Free Availability](#).

Details of the Always Free Oracle Autonomous Database instance

- **Processor:** 1 Oracle CPU processor (cannot be scaled)
- **Memory:** 8 GB RAM
- **Database Storage:** 20 GB storage (cannot be scaled)
- **Workload Type:** Your choice of either the [transaction processing](#) or [data warehouse](#) workload type
- **Maximum Simultaneous Database Sessions:** 20



Tip

Always Free Autonomous Databases can be upgraded to paid instances after provisioning if you need features like storage or CPU scaling.

See [To create an Always Free Autonomous Database](#) for steps to create an Always Free Autonomous Database.

Load Balancing

All tenancies get one Always Free 10 Mbps [load balancer](#).

Details of the Always Free load balancer

- **Shape:** Micro (10 Mbps)
- **Listeners:** 10
- **Virtual Hostnames:** 10
- **Backend Sets:** 10
- **Backend Servers:** 128

For information about provisioning an Always Free load balancer, see [Getting Started with Load Balancing](#).

Block Volume

All tenancies receive a total of 100 GB of Always Free [Block Volume](#) storage, and five volume backups. These amounts apply to both boot volumes and block volumes combined. When you provision a Compute instance, the instance automatically receives a 50 GB [boot volume](#) for storage. You can also create and attach block volumes to expand the storage capacity of a Compute instance. For more information, see [Creating a Volume](#) and [Attaching a Volume](#).

Details of the Always Free Block Volume resources

- 100 GB total of combined boot volume and block volume Always Free Block Volume storage.
- Five total volume backups (boot volume and block volume combined).

When you create a Compute instance, the default [boot volume](#) size for the instance is 50 GB, which counts towards your allotment of 100 GB. You can customize the instance's boot volume size up to 100 GB; however, this will use up your full allotment of storage for Always Free Block Volume resources. Also, because the minimum boot volume size allowed for Compute instances is 50 GB, launching two instances will use all your Always Free Block Volume resources. Alternatively, you can launch one instance with the default boot volume size of 50

GB, and then create and attach a 50 GB block volume to expand the storage capacity of the instance. For more information, see [Creating a Volume](#) and [Attaching a Volume](#). Although it is possible to mix paid and Always Free resources, Oracle does not recommend this. If you have used up your allotment of Always Free Block Volume resources, you can free up block storage resources by terminating an Always Free instance and deleting the boot volume, or terminating an Always Free block volume.

You can have a maximum of five Always Free volume backups at any time. This applies to both boot volume and block volume backups. For example, you could have three boot volume backups for your Always Free instance and two block volume backups for your Always Free block volumes. In this example, if you try to create new backups, the operation will fail with an error until you delete existing Always Free volume backups. For more information about volume backups, see [Overview of Block Volume Backups](#) and [Overview of Boot Volume Backups](#).

Object Storage

All tenancies get a total of 20 GiB (gibibytes) of Always Free Object Storage.

Details of the Always Free Object Storage resources

If you have a free account (including trial accounts), Always Free Object Storage includes the following:

- 20 GiB of combined [Object Storage](#) and [Archive Storage](#)
- 50,000 [Object Storage API](#) requests per month

If you have a paid account, Always Free Object Storage includes the following:

- 10 GiB of [Object Storage](#)
- 10 GiB of [Archive Storage](#)
- 50,000 [Object Storage API](#) requests per month



Important

If you are participating in an Oracle Cloud Free Trial, you can store unlimited data and can use 20 GiB for free (your usage of the first 20 GiB incurs no deduction of your initial \$300 trial credit balance). Upgrade to a paid account to continue access to unlimited storage. If you do not upgrade before your trial ends, your free account will be limited to 20 GiB of combined Object Storage and Archive Storage. If you are using more than the 20 GiB limit when your Free Trial ends, all of your objects will be deleted. You can then upload objects until you reach your Always Free usage limits.

See [Putting Data into Object Storage](#) for instructions on using your Always Free Object Storage resources.

Frequently Asked Questions: Oracle Cloud Infrastructure Free Tier

I just signed up and I cannot access specific services. What can I do?

Registering your account with all services and regions can take a few minutes. Check again after a few minutes have passed.

How do I change which resources I want to designate as Always Free?

In short, you cannot. Eligible resources are designated Always Free when they are created. After you provision an Always Free resource, the Always Free status is not transferable to another existing resource. However, you can delete an existing Always Free resource in order to create a new Always Free resource in its place.

What happens when my Free Trial expires or my credits are used up?

When you've reached the end of your 30 day trial, or used all of your Free Trial credits (whichever comes first), you will no longer be able to create new paid resources. However, your account will remain active. Your existing resources will continue to run for a few days, allowing you to upgrade your account and keep your resources before they're reclaimed by Oracle. (Note that reclaimed resources cannot be recovered—they are permanently deleted.)

Resources identified as Always Free will not be reclaimed. After your Free Trial expires, you'll continue to be able to use and manage your existing Always Free resources, and create new Always Free resources according to tenancy limits.

If I upgrade, do I keep my Free Trial credit balance?

Yes, if you upgrade during the Free Trial period, you will not be billed until your remaining credit balance is exhausted. You will be notified by email when billing begins.

After I upgrade my account, can I downgrade?

There is no option to downgrade your account. However, with a paid account, you'll continue to have access to Always Free resources, and you'll only pay for the standard resources you use. No minimums and no prepayment are required for your paid account.

My resources no longer appear. How can I restore them?

If you have a Free Tier account and your resources no longer appear, it is likely that your Free Trial has expired and your paid resources have been reclaimed (terminated). You can verify this if this is the case by doing the following:

1. Log in to the Console
2. Check for a banner at the top of the Console with the following text: "You are using a Free Tier account. To access all services and resources, upgrade to a paid account."

If you see this message, your resources have been reclaimed and cannot be restored.

Is it possible to extend my Free Trial?

If you need additional credits or time, you can schedule a call with an Oracle sales representative using the Upgrade page in the Console. Sales representatives have the authority to extend trials or issue additional credits if appropriate.

If you don't see an **Upgrade** link on the Console page you are viewing, you can click the **Oracle Cloud** logo at the top of the Console and then look for the Upgrade link in the sidebar on the right side of the page.

Is my Free Tier account eligible for support?

Community support through our [forums](#) is available to all customers. Customers using only Always Free resources are not eligible for Oracle Support. Limited support is available to Free Tier accounts with Free Trial credits. After you use all of your credits or after your trial period ends (whichever comes first), you must upgrade to a paid account to access Oracle Support. If you choose not to upgrade and continue to use Always Free Services, you will not be eligible to raise a service request in My Oracle Support. See [Getting Help and Contacting Support](#).

CHAPTER 3 Oracle Cloud Infrastructure for Government

Oracle Cloud Infrastructure for government provides cloud services for two levels of government operators:

- [Oracle Cloud Infrastructure Government Cloud with FedRAMP authorization](#)
- [Oracle Cloud Infrastructure Federal Government Cloud with Impact Level 5 \(IL5\) authorization](#)

For All Government Cloud Customers

This topic contains information common to both the [Government Cloud with FedRAMP authorization](#) and to the [Federal Government Cloud with IL5 authorization](#).

Shared Responsibilities

Oracle Cloud Infrastructure for government offers best-in-class security technology and operational processes to secure its enterprise cloud services. However, for you to securely run your workloads, you must be aware of your security and compliance responsibilities. By design, Oracle provides security of cloud infrastructure and operations (cloud operator access controls, infrastructure security patching, and so on), and you are responsible for securely configuring your cloud resources. Security in the cloud is a shared responsibility between you and Oracle.

For more information about shared responsibilities in the Oracle Cloud, see the following white papers:

- [Making Sense of the Shared Responsibility Model](#)
- [Oracle Cloud Infrastructure Security](#)

Setting Up an Identity Provider for Your Tenancy

As a Government Cloud customer, you must bring your own identity provider that meets your agency's compliance requirements and supports common access card/personal identity verification card (CAC/PIV) authentication. You can federate Oracle Cloud Infrastructure with SAML 2.0 compliant identity providers that also support CAC/PIV authentication. For instructions on setting up a federation, see [Federating with Identity Providers](#).

Remove the Oracle Cloud Infrastructure Default Administrator User and Any Other Non-Federated Users

When your organization signs up for an Oracle account and Identity Domain, Oracle sets up a *default administrator* for the account. This person will be the first IAM user for your company and will have full administrator access to your tenancy. This user can set up your federation.

After you have successfully set up the federation with your chosen identity provider, you can delete the default administrator user and any other IAM service local users you might have added to assist with setting up your tenancy. Deleting the local, non-federated users ensures that only users in your chosen identity provider can access Oracle Cloud Infrastructure.

To delete the default administrator:

1. Sign in to the Console through your identity provider.

More details

- a. Open a supported browser and go to the Government Cloud Console URL.
 - b. Enter your **Cloud Tenant** and click **Continue**.
 - c. On the **Single Sign-On** pane, select your identity provider and click **Continue**. You will be redirected to your identity provider to sign in.
 - d. Enter your user name and password.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. The list of users is displayed.

3. On the **User Type** filter, select only **Local Users**.
4. For each local user, go to the the Actions icon (three dots) and click **Delete**.

Using a Common Access Card/Personal Identity Verification Card to Sign in to the Console

After you set up CAC/PIV authentication with your identity provider and successfully federate with Oracle Cloud Infrastructure, you can use your CAC/PIV credentials to sign in to the Oracle Cloud Infrastructure Console. See your identity provider's documentation for the specific details for your implementation.

In general, the sign in steps are:

1. Insert your CAC/PIV card into your card reader.
2. Navigate to the Oracle Cloud Infrastructure [Console sign in page](#).
3. If prompted, enter your **Cloud Tenant** name and click **Continue**.
4. Select the Single Sign-On provider and click **Continue**.
5. On your identity provider's sign on page, select the appropriate card, for example, PIV Card.
6. If presented with a certificate picker, choose the appropriate certificate or other attributes set up by your organization.
7. When prompted, enter the PIN.

IPv6 Support for Virtual Cloud Networks

US Government Cloud customers have the option to enable IPv6 addressing for their VCNs. For more information, see [IPv6 Addresses](#).

Setting Up Secure Access for Compute Hosts

You can set up CAC/PIV authentication using third-party tools to enable multi-factor authentication for securely connecting to your compute hosts. Example tools include PuTTY-

CAC for Windows and Open SC for macOS. For more information see the U.S. Government website, [PIV Usage Guidelines](#).

Enabling FIPS Mode for Your Operating System

Government Cloud customers are responsible for enabling FIPS mode for the operating systems on their Compute hosts. To make your operating system compliant with Federal Information Processing Standard (FIPS) Publication 140-2, follow the guidelines for your operating system:

Oracle Linux

Follow the guidance provided at [Enabling FIPS Mode on Oracle Linux](#).

Ubuntu

Follow the guidance provided at [Ubuntu Security Certifications](#).

Windows Server 2008 and 2012

Follow the guidance provided at [Data Encryption for Web console and Reporting server Connections](#).

Windows Server 2016

First, follow the guidance provided at [How to Use FIPS Compliant Algorithms](#).

Next, go to the Microsoft document, [FIPS 140 Validation](#) and navigate to the topic [Information for System Integrators](#). Follow the instructions under "Step 2 – Setting FIPS Local/Group Security Policy Flag" to complete the FIPS enablement.

CentOS

The following guidance is for enabling FIPS on CentOS 7.5. These procedures are valid for both VM and bare metal instances, and only in NATIVE mode. These procedures can be modified for both Emulated and PV modes as needed. Note that this procedure provides an

instance that contains the exact FIPS cryptographic modules EXCEPT kernel. However, the kernel module is the same major/minor version but is accelerated in revision, so can be considered compliant under most FIPS compliant models.

After you complete this procedure, Oracle strongly recommends that you do NOT run system-wide yum updates. The system-wide update will remove the FIPS modules contained herein.

Verify that the version of the kernel, FIPS modules, and FIPS software are at the minimum version:

1. Validate the current version of the kernel package meets the requirement:
 - a. Current version: `kernel-3.10.0-693.el7`
 - b. Execute `rpm -qa | grep kernel-3`
2. Execute the following and validate the major or minor version is the same as the requirements.
 - a. Run

```
yum list <package_name>
```

- b. Verify that the major/minor version matches the required ones.

Required packages and versions are:

- `fipscheck - fipscheck-1.4.1-6.el7`
- `hmacalc - hmacalc-0.9.13-4.el7`
- `dracut-fips - dracut-fips-033-502.el7`
- `dracut-fips-aesni - dracut-fips-aesni-033-502.el7`

- c. For each version of package that is not installed, run

```
yum install <package_name>
```

3. Download and install the following packages:
 - a. Packages already installed as part of the image:
 - i. Create a directory called `preinstall`.
 - ii. Download the following packages into this directory:

```
openssl, openssl-libs – 1.0.2k-8.el7  
nss, nss-tools, nss-sysinit – 3.28.4-15.el7_4  
nss-util – 3.28.4-3.el7  
nss-softokn, nss-softokn-freebl – 3.28.3-8.el7_4  
openssh, openssh-clients, openssh-server – 7.4p1-11.el7
```

- iii. In the preinstall directory, run

```
yum - -nogpgcheck downgrade *.rpm
```

- b. Packages to be added to the image:

- i. Create a directory called `newpackages`.
- ii. Download the following packages into this directory:

```
libreswan – 3.20-3.el7  
libgcrypt – 1.5.3-14.el7  
gnutls – 3.3.26-9.el7  
gmp – 6.0.0-15.el7  
nettle – 2.7.1-8.el7
```

- iii. In the `newpackages` directory, run

```
yum - -nogpgcheck localinstall *.rpm
```

The URLs for the packages used for this installation are:

Preinstall:

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-3.28.4-15.el7_4.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-util-3.28.4-3.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-tools-3.28.4-15.el7_4.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-sysinit-3.28.4-15.el7_4.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-softokn-freebl-3.28.3-8.el7_4.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nss-softokn-3.28.3-8.el7_4.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/openssl-1.0.2k-8.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/openssl-libs-1.0.2k-8.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/openssh-7.4p1-11.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/openssh-clients-7.4p1-11.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/openssh-server-7.4p1-11.el7.x86_64.rpm

Newpackages:

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/libreswan-3.20-3.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/libgcrypt-1.5.3-14.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/gnutls-3.3.26-9.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/gmp-6.0.0-15.el7.x86_64.rpm

http://linuxsoft.cern.ch/cern/centos/7/updates/x86_64/Packages/nettle-2.7.1-8.el7.x86_64.rpm

Kernel FIPS module and initramfs validation installation.

Perform this procedure as root:

1. Regenerate dracut:

```
dracut -f -v
```

2. Add the `fips` argument to the end of the default kernel boot command line:

- a. Edit `/etc/default/grub`
- b. At the end of the line starting with "GRUB_CMDLINE_LINUX", add

```
fips=1
```

inside the double quotes of the command.

- c. Save the result.

3. Generate a new `grub.cfg`:

```
grub2-mkconfig -o /etc/grub2-efi.cfg
```

Configure SSH to limit the encryption algorithms.

1. Sudo to root.
2. Edit `/etc/ssh/sshd_config`.
3. Add the following lines to the bottom of the file:

```
Protocol 2  
Ciphers aes128-ctr,aes192-ctr,aes256-ctr,aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc  
Macs hmac-sha1
```

4. Reboot the instance.
5. After instance has rebooted, validate that FIPS mode has been enabled in the kernel:
 - a. Sudo to root.
 - b. Run the following command:

```
cat /proc/sys/crypto/fips-enabled
```

The result should be '1'.

CHAPTER 3 Oracle Cloud Infrastructure for Government

To further secure CentOS7/RHEL 7.x systems as required by individual agency guidance, follow the checklist contained in the OpenSCAP guide. This guide can be found here:

<https://static.open-scap.org/ssg-guides/ssg-centos7-guide-index.html>

The STIG for evaluating compliance under multiple profiles can be found here:

<https://iase.disa.mil/stigs/os/unix-linux/Pages/index.aspx> . Use the Red Hat Linux 7.x STIG for CentOS 7.5 releases.

Required VPN Connect Parameters for Government Cloud

If you use [VPN Connect](#) with the Government Cloud, you must configure the IPsec connection with the following FIPS-compliant IPsec parameters.

For some parameters, Oracle supports multiple values, and the recommended one is highlighted in *red italics*.

Oracle supports the following parameters for IKEv1 or IKEv2. Check the documentation for your particular CPE to confirm which parameters the CPE supports for IKEv1 or IKEv2.

Phase 1 (ISAKMP)

Parameter	Options
ISAKMP protocol	Version 1
Exchange type	Main mode
Authentication method	Pre-shared keys
Encryption algorithm	<i>AES-256-cbc</i> AES-192-cbc AES-128-cbc

Parameter	Options
Authentication algorithm	<i>SHA-2 384</i> SHA-2 256 SHA-1 (also called SHA or SHA1-96)
Diffie-Hellman group	group 14 (MODP 2048) group 19 (ECP 256) <i>group 20 (ECP 384) *</i>
IKE session key lifetime	28800 seconds (8 hours)
* Group 20 will be supported in all Oracle Cloud Infrastructure regions very soon.	

Phase 2 (IPSec)

Parameter	Options
IPSec protocol	ESP, tunnel mode
Encryption algorithm	<i>AES-256-gcm</i> AES-192-gcm AES-128-gcm AES-256-cbc AES-192-cbc AES-128-cbc
Authentication algorithm	If using GCM (Galois/Counter Mode), no authentication algorithm is required because authentication is included with GCM encryption. If not using GCM, use HMAC-SHA-256-128.

Parameter	Options
IPSec session key lifetime	3600 seconds (1 hour)
Perfect Forward Secrecy (PFS)	enabled, group 14

Oracle's BGP ASN

This section is for network engineers who configure an edge device for FastConnect or VPN Connect.

Oracle's BGP ASN for the Government Cloud depends on the authorization level:

- Government Cloud: 6142
- Federal Government Cloud (Impact Level 5 authorization): 20054

Government Cloud with FedRAMP Authorization

This topic contains information specific to Oracle Cloud Infrastructure Government Cloud with FedRAMP authorization.

Authorizations

Oracle Cloud Infrastructure Government Cloud has obtained the following authorizations:

- FedRAMP Moderate
- DISA Impact Level 2

For information about the Federal Government Cloud, see [Federal Government Cloud with Impact Level 5 Authorization](#).

Regions

The region names and identifiers for the Government Cloud with FedRAMP authorization are shown in the following table:

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
US Gov East (Ashburn)	us-langley-1	Ashburn, VA	LFI	OC2	1
US Gov West (Phoenix)	us-luke-1	Phoenix, AZ	LUF	OC2	1

After your tenancy is created in one of these regions, you can subscribe to the other region. Tenancies in the FedRAMP-authorized regions cannot subscribe to the commercial regions, or to the [Federal Government Cloud regions](#). For information about subscribing to a region, see [Managing Regions](#).

Console Sign-in URLs

To sign in to the FedRAMP-authorized Government Cloud, enter one of the following URLs in a [supported browser](#):

- <https://console.us-langley-1.oraclegovcloud.com/>
- <https://console.us-luke-1.oraclegovcloud.com/>



Note

When you're logged in to the Console for one of the Government Cloud regions, the browser times out after 15 minutes of inactivity, and you need to sign in again to use the Console.

Government Cloud API Reference and Endpoints

Oracle Cloud Infrastructure Government Cloud has these APIs and corresponding regional endpoints:

Core Services (covering Networking, Compute, and Block Volume)

The Networking, Compute, and Block Volume services are accessible with the following API:

Core Services API

[API reference](#)

- <https://iaas.us-langley-1.oraclegovcloud.com>
- <https://iaas.us-luke-1.oraclegovcloud.com>

Database API

[API reference](#)

- <https://database.us-langley-1.oraclegovcloud.com>
- <https://database.us-luke-1.oraclegovcloud.com>

You can track the progress of long-running Database operations with the [Work Requests](#) API.

IAM API

[API reference](#)

- <https://identity.us-langley-1.oraclegovcloud.com>
- <https://identity.us-luke-1.oraclegovcloud.com>



Note

Use the Endpoint of Your Home Region for All IAM API Calls

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region. Your home region is where your IAM resources are defined. When you subscribe to a new region, your IAM resources are replicated in the new region, however, the master definitions reside in your home region and can only be changed there. Make all IAM API calls against your home region endpoint. The changes automatically replicate to all regions. If you try to make an IAM API call against a region that is not your home region, you will receive an error.

Key Management API

[API reference](#)

- <https://kms.us-langley-1.oraclegovcloud.com>
- <https://kms.us-luke-1.oraclegovcloud.com>

In addition to these endpoints, each vault has a unique endpoint for create, update, and list operations for keys. This endpoint is referred to as the control plane URL or management endpoint. Each vault also has a unique endpoint for cryptographic operations. This endpoint is known as the data plane URL or the cryptographic endpoint.

Object Storage and Archive Storage APIs

Both Object Storage and Archive Storage are accessible with the following APIs:

Object Storage API

[API reference](#)

- <https://objectstorage.us-langley-1.oraclegovcloud.com>
- <https://objectstorage.us-luke-1.oraclegovcloud.com>

Amazon S3 Compatibility API

[API reference](#)

- https://<object_storage_namespace>.compat.objectstorage.us-langley-1.oraclegovcloud.com
- https://<object_storage_namespace>.compat.objectstorage.us-luke-1.oraclegovcloud.com



Tip

See [Understanding Object Storage Namespaces](#) for information regarding how to find your Object Storage namespace.

Swift API (for use with Oracle RMAN)

- <https://swiftobjectstorage.us-langley-1.oraclegovcloud.com>
- <https://swiftobjectstorage.us-luke-1.oraclegovcloud.com>

Work Requests API (for Compute and Database work requests)

[API reference](#)

- <https://iaas.us-langley-1.oraclegovcloud.com>
- <https://iaas.us-luke-1.oraclegovcloud.com>

Services Not Supported in Oracle Cloud Infrastructure Government Cloud

The following services are currently not available for tenancies in the Government Cloud:

Core Infrastructure services and features not available:

- Compute service features:
 - Autoscaling
- Data Transfer service
- File Storage service

Database services not available:

- Autonomous Data Warehouse
- Autonomous Transaction Processing
- Data Safe

Data and AI services not available:

- Digital Assistant

Solutions and Platform services not available:

- Analytics Cloud
- Analytics for Applications
- Container Engine for Kubernetes
- Content and Experience
- DNS Zone Management
- Email Delivery

CHAPTER 3 Oracle Cloud Infrastructure for Government

- Events
- Functions
- Health Checks
- Integration
- Marketplace
- Monitoring
- Notifications
- Registry
- Resource Manager
- Streaming
- Traffic Management Steering Policies

Governance and Administration features not supported

- Auto-federation with Oracle Identity Cloud Service
- WAF service

Infrastructure Tools

- Oracle Cloud Infrastructure Terraform Provider

Integration with Oracle SaaS and PaaS services, including those listed here: [Getting Started with Oracle Platform Services](#)

Additional Information for Government Cloud Customers

- [Shared Responsibilities](#)
- [Setting Up an Identity Provider for Your Tenancy](#)
- [Using a Common Access Card/Personal Identity Verification Card to Sign in to the Console](#)

- [IPv6 Support for Virtual Cloud Networks](#)
- [Setting Up Secure Access for Compute Hosts](#)
- [Enabling FIPS Mode for Your Operating System](#)
- [Required VPN Connect Parameters for Government Cloud](#)
- [Oracle's BGP ASN](#)

Federal Government Cloud with Impact Level 5 Authorization

This topic contains information specific to Oracle Cloud Infrastructure Federal Government Cloud.

Compliance with Defense Cloud Security Requirements

Oracle Cloud Infrastructure Federal Government Cloud supports applications that require Impact Level 5 (IL5) data, as defined in the Department of Defense [Cloud Computing Security Requirements Guide](#) (SRG).

Federal Government Cloud Regions

The region names and identifiers for the Oracle Cloud Infrastructure Federal Government Cloud regions are shown in the following table:

Region Name	Region Identifier	Region Key	Realm Key	Availability Domains
US DoD East (Ashburn)	us-gov-ashburn-1	ric	OC3	1
US DoD North (Chicago)	us-gov-chicago-1	pia	OC3	1
US DoD West (Phoenix)	us-gov-phoenix-1	tus	OC3	1

After your tenancy is created in one of the Federal Government Cloud regions, you can subscribe to the other regions in the Federal Government Cloud. Federal Government Cloud tenancies cannot subscribe to any Oracle Cloud Infrastructure regions not belonging to the OC3 realm. For information about subscribing to a region, see [Managing Regions](#).

Federal Government Cloud Console Sign-in URLs

To sign in to the Oracle Cloud Infrastructure Federal Government Cloud, enter one of the following URLs in a [supported browser](#):

- <https://console.us-gov-ashburn-1.oraclegovcloud.com/>
- <https://console.us-gov-chicago-1.oraclegovcloud.com/>
- <https://console.us-gov-phoenix-1.oraclegovcloud.com/>



Note

When you're logged in to the Console for one of the Government Cloud regions, the browser times out after 15 minutes of inactivity, and you need to sign in again to use the Console.

Government Cloud API Reference and Endpoints

Oracle Cloud Infrastructure Government Cloud has these APIs and corresponding regional endpoints:

Core Services (covering Networking, Compute, and Block Volume)

The Networking, Compute, and Block Volume services are accessible with the following API:

Core Services API

[API reference](#)

- <https://iaas.us-gov-ashburn-1.oraclegovcloud.com>
- <https://iaas.us-gov-chicago-1.oraclegovcloud.com>
- <https://iaas.us-gov-phoenix-1.oraclegovcloud.com>

Database API

[API reference](#)

- <https://database.us-gov-ashburn-1.oraclegovcloud.com>
- <https://database.us-gov-chicago-1.oraclegovcloud.com>
- <https://database.us-gov-phoenix-1.oraclegovcloud.com>

You can track the progress of long-running Database operations with the [Work Requests](#) API.

IAM API

[API reference](#)

- <https://identity.us-gov-ashburn-1.oraclegovcloud.com>
- <https://identity.us-gov-chicago-1.oraclegovcloud.com>
- <https://identity.us-gov-phoenix-1.oraclegovcloud.com>



Note

Use the Endpoint of Your Home Region for All IAM API Calls

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region. Your home region is where your IAM resources are defined. When you subscribe to a new region, your IAM resources are replicated in the new region, however, the master definitions reside in your home region and can only be changed there. Make all IAM API calls against your home region endpoint. The changes automatically replicate to all regions. If you try to make an IAM API call against a region that is not your home region, you will receive an error.

Key Management API

[API reference](#)

- <https://kms.us-gov-ashburn-1.oraclegovcloud.com>
- <https://kms.us-gov-chicago-1.oraclegovcloud.com>
- <https://kms.us-gov-phoenix-1.oraclegovcloud.com>

In addition to these endpoints, each vault has a unique endpoint for create, update, and list operations for keys. This endpoint is referred to as the control plane URL or management endpoint. Each vault also has a unique endpoint for cryptographic operations. This endpoint is known as the data plane URL or the cryptographic endpoint.

Object Storage and Archive Storage APIs

Both Object Storage and Archive Storage are accessible with the following APIs:

Object Storage API

[API reference](#)

- <https://objectstorage.us-gov-ashburn-1.oraclegovcloud.com>
- <https://objectstorage.us-gov-chicago-1.oraclegovcloud.com>
- <https://objectstorage.us-gov-phoenix-1.oraclegovcloud.com>

Amazon S3 Compatibility API

[API reference](#)

- https://<object_storage_namespace>.compat.objectstorage.us-gov-ashburn-1.oraclegovcloud.com
- https://<object_storage_namespace>.compat.objectstorage.us-gov-chicago-1.oraclegovcloud.com
- https://<object_storage_namespace>.compat.objectstorage.us-gov-phoenix-1.oraclegovcloud.com



Tip

See [Understanding Object Storage Namespaces](#) for information regarding how to find your Object Storage namespace.

Swift API (for use with Oracle RMAN)

- <https://swiftobjectstorage.us-gov-ashburn-1.oraclegovcloud.com>
- <https://swiftobjectstorage.us-gov-chicago-1.oraclegovcloud.com>
- <https://swiftobjectstorage.us-gov-phoenix-1.oraclegovcloud.com>

Work Requests API (for Compute and Database work requests)

[API reference](#)

- <https://iaas.us-gov-ashburn-1.oraclegovcloud.com>
- <https://iaas.us-gov-chicago-1.oraclegovcloud.com>
- <https://iaas.us-gov-phoenix-1.oraclegovcloud.com>

Services Not Supported in Oracle Cloud Infrastructure Federal Government Cloud

Currently, the following services are not available for tenancies in the Federal Government Cloud:

Core Infrastructure services and features not available:

- Compute service features:
 - Autoscaling
- FastConnect
- Data Transfer service
- File Storage service

Database services not available:

- Autonomous Data Warehouse
- Autonomous Transaction Processing
- Data Safe

Data and AI services not available:

- Digital Assistant

Solutions and Platform services not available:

CHAPTER 3 Oracle Cloud Infrastructure for Government

- Analytics Cloud
- Analytics for Applications
- Container Engine for Kubernetes
- Content and Experience
- DNS Zone Management
- Email Delivery
- Events
- Functions
- Health Checks
- Integration
- Marketplace
- Monitoring
- Notifications
- Registry
- Resource Manager
- Streaming
- Traffic Management Steering Policies

Governance and Administration features not supported

- Auto-federation with Oracle Identity Cloud Service
- WAF service

Infrastructure Tools

- Oracle Cloud Infrastructure Terraform Provider

Integration with Oracle SaaS and PaaS services, including those listed here: [Getting Started with Oracle Platform Services](#).

Additional Information for Federal Government Cloud Customers

- [Shared Responsibilities](#)
- [Setting Up an Identity Provider for Your Tenancy](#)
- [Using a Common Access Card/Personal Identity Verification Card to Sign in to the Console](#)
- [IPv6 Support for Virtual Cloud Networks](#)
- [Setting Up Secure Access for Compute Hosts](#)
- [Enabling FIPS Mode for Your Operating System](#)
- [Required VPN Connect Parameters for Government Cloud](#)
- [Oracle's BGP ASN](#)

CHAPTER 4 Service Essentials

The following topics provide essential information that applies across Oracle Cloud Infrastructure.

[Security Credentials](#)

The types of credentials you'll use when working with Oracle Cloud Infrastructure.

[Regions and Availability Domains](#)

An introduction to the concepts of regions and availability domains.

[Resource Identifiers](#)

A description of the different ways your Oracle Cloud Infrastructure resources are identified.

[Resource Monitoring](#)

Information about how to monitor your resources.

[Resource Tags](#)

Information about Oracle Cloud Infrastructure tags and how to apply them to your resources.

[Compartment Quotas](#)

Information about how to control resource consumption within compartments using quotas.

[Compartment Explorer](#)

View all resources in a selected compartment, across regions.

Service Limits

A list of the default limits applied to your cloud resources and how to request an increase.

Console Announcements

Information about the announcements that occasionally appear in the Oracle Cloud Infrastructure Console

Prerequisites for Oracle Platform Services on Oracle Cloud Infrastructure

Instructions for setting up the resources required when running an Oracle Platform Service on Oracle Cloud Infrastructure.

Billing and Payment Tools Overview

Information about billing and payment tools that you can use to analyze your service usage and manage your costs.

My Services Use Cases

Use cases for the [Oracle Cloud My Services API](#), to help you interact programmatically with My Services.

Security Credentials

This section describes the types of credentials you'll use when working with Oracle Cloud Infrastructure.

Console Password

- **What it's for:** Using the Console.
- **Format:** Typical password text string.
- **How to get one:** An administrator will provide you with a one-time password.
- **How to use it:** Sign in to the Console the first time with the one-time password, and then change it when prompted. Requirements for the password are displayed there. The one-time password expires in seven days. If you want to change the password later, see [To change your Console password](#). Also, you or an administrator can reset the password in the Console or with the API (see [To create or reset another user's Console password](#)). Resetting the password creates a new one-time password that you'll be prompted to change the next time you sign in to the Console. If you're blocked from signing in to the Console because you've tried 10 times in a row unsuccessfully, contact your administrator.
- **Note for Federated Users:** Federated users do not use a Console password. Instead, they sign in to the Console through their identity provider.

API Signing Key

- **What it's for:** Using the API (see [Software Development Kits and Command Line Interface](#) and [Request Signatures](#)).
- **Format:** RSA key pair in PEM format (minimum 2048 bits required).
- **How to get one:** See [Required Keys and OCIDs](#).
- **How to use it:** In the Console, copy and paste the contents of the PEM public key file from the key pair (see [How to Upload the Public Key](#)). Then use the private key with the SDK or with your own client to sign your API requests. Note that after you've uploaded your first API key in the Console, you can use the API to upload any additional ones you want to use. If you provide the wrong kind of key (for example, your instance SSH key, or a key that isn't at least 2048 bits), you'll get an `InvalidKey` error.
- **Example:** The PEM public key looks something like this:

```
-----BEGIN PUBLIC KEY-----
```

```
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAOtFqF...  
...  
-----END PUBLIC KEY-----
```

Instance SSH Key

- **What it's for:** Accessing a compute instance.
- **Format:** For [Oracle-provided images](#), these SSH key types are supported: RSA, DSA, DSS, ECDSA, and Ed25519. If you [bring your own image](#), you're responsible for managing the SSH key types that are supported.
For RSA, DSS, and DSA keys, a minimum of 2048 bits is recommended. For ECDSA keys, a minimum of 128 bits is recommended.
- **How to get one:** See [Creating a Key Pair](#).
- **How to use it:** When you launch an instance, provide the public key from the key pair.
- **Example:** An RSA public key looks something like this:

```
ssh-rsa AAAAB3BzaC1yc2EAAAADAQABAAQD9BRwrUiLDki6P0+jZhwsjS2muM...  
... jane.smith@example.com
```

Auth Token

- **What it's for:** Authenticating with third-party APIs that do not support Oracle Cloud Infrastructure's signature-based authentication. For example, use an auth token as your password with Swift clients.
- **Format:** Typical password text string.
- **How to get one:** See [Working with Auth Tokens](#).
- **How to use it:** Usage depends on the service you are authenticating with. Typically, you authenticate with third-party APIs by providing your Oracle Cloud Infrastructure Console login, your auth token provided by Oracle, and your organization's Oracle tenant name.

Regions and Availability Domains

This topic describes the physical and logical organization of Oracle Cloud Infrastructure resources.

About Regions and Availability Domains

Oracle Cloud Infrastructure is hosted in regions and availability domains. A region is a localized geographic area, and an availability domain is one or more data centers located within a region. A region is composed of one or more availability domains. Most Oracle Cloud Infrastructure resources are either region-specific, such as a virtual cloud network, or availability domain-specific, such as a compute instance. Traffic between availability domains and between regions is encrypted.

Availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

The availability domains within the same region are connected to each other by a low latency, high bandwidth network, which makes it possible for you to provide high-availability connectivity to the internet and on-premises, and to build replicated systems in multiple availability domains for both high-availability and disaster recovery.

Oracle is adding multiple cloud regions around the world to provide local access to cloud resources for our customers. To accomplish this quickly, we've chosen to launch regions in new geographies with one availability domain.

As regions require expansion, we have the option to add capacity to existing availability domains, to add additional availability domains to an existing region, or to build a new region. The expansion approach in a particular scenario is based on customer requirements as well as considerations of regional demand patterns and resource availability.

For any region with one availability domain, a second availability domain or region in the same country or geo-political area will be made available within a year to enable further options for disaster recovery that support customer requirements for data residency where they exist.

CHAPTER 4 Service Essentials

Regions are independent of other regions and can be separated by vast distances—across countries or even continents. Generally, you would deploy an application in the region where it is most heavily used, because using nearby resources is faster than using distant resources. However, you can also deploy applications in different regions for these reasons:

- To mitigate the risk of region-wide events such as large weather systems or earthquakes.
- To meet varying requirements for legal jurisdictions, tax domains, and other business or social criteria.

Regions are grouped into realms. Your tenancy exists in a single realm and can access all regions that belong to that realm. You can't access regions that are not in your realm. Currently, Oracle Cloud Infrastructure has three realms: the commercial realm, and two realms for Government Cloud: [FedRAMP authorized](#) and [IL5 authorized](#).

The following table lists the regions in the Oracle Cloud Infrastructure commercial realm:

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
Australia East (Sydney)	ap-sydney-1	Sydney, Australia	SYD	OC1	1
Brazil East (Sao Paulo)	sa-saopaulo-1	Sao Paulo, Brazil	GRU	OC1	1
Canada Southeast (Toronto)	ca-toronto-1	Toronto, Canada	YYZ	OC1	1
Germany Central (Frankfurt)	eu-frankfurt-1	Frankfurt, Germany	FRA	OC1	3
India West (Mumbai)	ap-mumbai-1	Mumbai, India	BOM	OC1	1
Japan East (Tokyo)	ap-tokyo-1	Tokyo, Japan	NRT	OC1	1

CHAPTER 4 Service Essentials

Region Name	Region Identifier	Region Location	Region Key	Realm Key	Availability Domains
South Korea Central (Seoul)	ap-seoul-1	Seoul, South Korea	ICN	OC1	1
Switzerland North (Zurich)	eu-zurich-1	Zurich, Switzerland	ZRH	OC1	1
UK South (London)	uk-london-1	London, United Kingdom	LHR	OC1	3
US East (Ashburn)	us-ashburn-1	Ashburn, VA	IAD	OC1	3
US West (Phoenix)	us-phoenix-1	Phoenix, AZ	PHX	OC1	3

To subscribe to a region, see [Managing Regions](#).

For a list of the Oracle Government Cloud regions, see [Government Cloud with FedRAMP Authorization](#) and [Federal Government Cloud with Impact Level 5 Authorization](#).



Note

Your Tenancy's Availability Domain Names

Oracle Cloud Infrastructure randomizes the availability domains by tenancy to help balance capacity in the data centers. For example, the availability domain labeled PHX-AD-1 for tenancyA may be a different data center than the one labeled PHX-AD-1 for tenancyB. To keep track of which availability domain corresponds to which data center for each tenancy, Oracle Cloud Infrastructure uses tenancy-specific prefixes for the availability domain names. For example: the availability domains for your tenancy are something like *Uocm*:PHX-AD-1, *Uocm*:PHX-AD-2, and so on.

To get the specific names of your tenancy's availability domains, use the [ListAvailabilityDomains](#) operation, which is available in the IAM API. You can also see the names when you use the Console to [launch an instance](#) and choose which availability domain to launch the instance into.

Fault Domains

A fault domain is a grouping of hardware and infrastructure within an availability domain. Each availability domain contains three fault domains. Fault domains let you distribute your instances so that they are not on the same physical hardware within a single availability domain. A hardware failure or Compute hardware maintenance event that affects one fault domain does not affect instances in other fault domains.

To control the placement of your compute instances, bare metal DB system instances, or virtual machine DB system instances, you can optionally specify the fault domain for a new instance at launch time. If you do not specify the fault domain, the system selects one for

you. To change the fault domain for an instance, terminate it and launch a new instance in the preferred fault domain.

Use fault domains to do the following things:

- Protect against unexpected hardware failures.
- Protect against planned outages due to Compute hardware maintenance.

For more information:

- For recommendations on using fault domains when provisioning application and database servers, see [Fault Domains](#) in [Best Practices for Your Compute Instance](#).
- For more information about using fault domains when provisioning Oracle bare metal and virtual machine DB systems, see [Fault Domain Considerations for 2-node Virtual Machine DB Systems](#) and [Availability Domain and Fault Domain Considerations for Data Guard](#).

Service Availability Across Regions

All Oracle Cloud Infrastructure regions offer core infrastructure services, including the following:

- Compute: Compute (Intel based bare metal & VM, DenseIO & Standard), Container Engine for Kubernetes, Registry
- Storage: Block Volume, File Storage, Object Storage, Archive Storage
- Networking: Virtual Cloud Network, Load Balancing, FastConnect (specific partners as available and requested)
- Database: Database, Exadata Cloud Service, Autonomous Data Warehouse, Autonomous Transaction Processing
- Edge: DNS
- Platform: Identity and Access Management, Tagging, Audit, Work Requests

Generally available cloud services beyond those in the preceding list are made available based on regional customer demand. Any service can be made available within a maximum of

three months, with many services deploying more quickly. New cloud services are made available in regions as quickly as possible based on a variety of considerations including regional customer demand, ability to achieve regulatory compliance where applicable, resource availability, and other factors. Because of our low latency interconnect backbone, customers can use cloud services in other geographic regions with effective results when they are not available in their home region, provided that data residency requirements do not prevent them from doing so. We regularly work with customers to help ensure effective access to required services.

Resource Availability

The following sections list the resource types based on their availability: global across regions, within a single region, or within a single availability domain.



Tip

In general: IAM resources are global. DB Systems, instances, and volumes are specific to an availability domain. Everything else is regional. Exception: Subnets were originally designed to be specific to an availability domain. Now, you can create [regional subnets](#), which is what Oracle recommends.

Global Resources

- API signing keys
- compartments
- dynamic groups
- federation resources
- groups
- policies

- tag namespaces
- tag keys
- users

Regional Resources

- alarms
- applications
- buckets: Although buckets are regional resources, they can be accessed from any location if you use the correct region-specific Object Storage URL for the API calls.
- clusters
- cloudevents-rules
- customer-premises equipment (CPE)
- DHCP options sets
- dynamic routing gateways (DRGs)
- encryption keys
- functions
- images
- internet gateways
- jobs
- key vaults
- load balancers
- local peering gateways (LPGs)
- metrics
- NAT gateways
- network security groups
- node pools

- repositories
- reserved public IPs
- route tables
- security lists
- service gateways
- stacks
- subnets: When you create a subnet, you choose whether it's [regional or specific to an availability domain](#). Oracle recommends using regional subnets.
- subscriptions
- topics
- virtual cloud networks (VCNs)
- volume backups: They can be restored as new volumes to any availability domain within the same region in which they are stored.

Availability Domain-Specific Resources

- DB Systems
- ephemeral public IPs
- instances: They can be attached only to volumes in the same availability domain.
- subnets: When you create a subnet, you choose whether it is [regional or specific to an availability domain](#). Oracle recommends using regional subnets.
- volumes: They can be attached only to an instance in the same availability domain.

IP Address Ranges

This topic provides information about public IP address ranges for services that are deployed in Oracle Cloud Infrastructure. Allow traffic to these CIDR blocks to ensure access to the services.

Endpoints for [Oracle YUM repos](#) and the [Oracle Container Registry](#) are listed on this page. You can use DNS lookup to determine the public IP address for each endpoint.

Public IP Addresses for VCNs and the Oracle Services Network

Public IP address ranges for VCNs and the Oracle Services Network are published to a JSON file which you can download and view manually or consume programmatically.

The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. A [service gateway](#) offers private access to the Oracle Services Network from workloads in your VCN and your on-premises network. The published addresses correspond to the [service CIDR label](#) called **All <region> Services in Oracle Services Network**. For a list of the services available with a service gateway, see [Service Gateway: Supported Cloud Services in Oracle Services Network](#).

Downloading the JSON File

[Use this link to download the current list of public IP ranges.](#)

You can poll the published file to check for new IP address ranges as frequently as every 24 hours. We recommend that you poll the published file at least weekly.

JSON File Contents and Syntax

IP addresses are published in the `public_ip_ranges.json` file with the fields in the following table.

Example of the `public_ip_ranges.json` file

```
{
  "last_updated_timestamp": "2019-11-18T19:55:47.204985",
  "regions": [
    {
      "region": "us-phoenix-1",
      "cidrs": [
        {
```

CHAPTER 4 Service Essentials

```
        "cidr": "129.146.0.0/21",
        "tags": [
            "OCI"
        ]
    },
    {
        "cidr": "134.70.8.0/21",
        "tags": [
            "OSN",
            "OBJECT_STORAGE"
        ]
    },
]
}
{
    "region": "us-ashburn-1",
    "cidrs": [
        {
            "cidr": "129.213.8.0/21",
            "tags": [
                "OCI"
            ]
        },
        {
            "cidr": "134.70.24.0/21",
            "tags": [
                "OSN",
                "OBJECT_STORAGE"
            ]
        }
    ]
}
]
```

CHAPTER 4 Service Essentials

Field Name	Definition	Type	Example
last_updated_timestamp	File creation time in ISO 8601 format. Expressed as <i><date>T<time></i>	string	"last_updated_timestamp": "2019-11-18T19:55:47.204985"
regions	IP CIDR ranges grouped by region.	array	See preceding Example of the public_ip_ranges.json file
region	The region of the IP CIDR ranges. Valid values: Any region in the Oracle Cloud Infrastructure commercial realm. For a complete list of regions, see Regions and Availability Domains .	string	"region": "us-phoenix-1"
cidrs	A group of IP address CIDR ranges.	array	See preceding Example of the public_ip_ranges.json file

Field Name	Definition	Type	Example
cidr	One or more IPv4 IP addresses expressed in CIDR notation.	string	"cidr": "147.154.0.0/18"
tags	<p>The services associated with the IP address CIDR range.</p> <p>Valid values:</p> <ul style="list-style-type: none"> OCI: The VCN CIDR blocks. OSN: The CIDR block ranges for the Oracle Services Network. OBJECT_STORAGE: The CIDR block ranges used by the Object Storage service. For more information, see Overview of Object Storage. 	array of string values	"tags": ["OCI"]

Filtering the JSON file contents

After you download the JSON file, you can use a command line tool such as `jq` to filter the contents.

[Download jq](#)

CHAPTER 4 Service Essentials

Here are some examples of how you can use the tool to find and filter the information you need:

Find the creation date of the JSON file:

```
jq .last_updated_timestamp < public_ip_ranges.json
```

Get all IPv4 addresses for a specific region:

```
jq -r '.regions[] | select (.region=="us-phoenix-1") | .cidrs[] | select (.cidr | contains(".")) | .cidr' < public_ip_ranges.json
```

Public IP Addresses for the Oracle YUM Repos

The Oracle YUM repos have the following regional public endpoints.

Region	YUM Server Endpoint
Australia East (Sydney)	https://yum-ap-sydney-1.oracle.com
Canada Southeast (Toronto)	https://yum-ca-toronto-1.oracle.com
Germany Central (Frankfurt)	https://yum-eu-frankfurt-1.oracle.com
India West (Mumbai)	https://yum-ap-mumbai-1.oracle.com
Japan East (Tokyo)	https://yum-ap-tokyo-1.oracle.com
South Korea Central (Seoul)	https://yum-ap-seoul-1.oracle.com
UK South (London)	https://yum-uk-london-1.oracle.com
US East (Ashburn)	https://yum-us-ashburn-1.oracle.com
US West (Phoenix)	https://yum-us-phoenix-1.oracle.com

You can use DNS lookup to determine the public IP address for each endpoint.

Resource Identifiers

This chapter describes the different ways your Oracle Cloud Infrastructure resources are identified.

Oracle Cloud IDs (OCIDs)

Most types of Oracle Cloud Infrastructure resources have an Oracle-assigned unique ID called an *Oracle Cloud Identifier* (OCID). It's included as part of the resource's information in both the Console and API.



Important

To use the API, you need the OCID for your tenancy. For information about where to find it, see the next section.

OCIDs use this syntax:

```
ocid1.<RESOURCE TYPE>.<REALM>.[REGION][.FUTURE USE].<UNIQUE ID>
```

- **ocid1:** The literal string indicating the version of the OCID.
- **resource type:** The type of resource (for example, *instance*, *volume*, *vcn*, *subnet*, *user*, *group*, and so on).
- **realm:** The realm the resource is in. A *realm* is a set of regions that share entities. Possible values are `oc1` for the [commercial realm](#), `oc2` for the [Government Cloud realm](#), or `oc3` for the [Federal Government Cloud realm](#). The regions in the commercial realm (OC1) belong to the domain `oraclecloud.com`. The regions in the Government Cloud (OC2) belong to the domain `oraclegovcloud.com`.
- **region:** The region the resource is in (for example, `phx`, `iad`, `eu-frankfurt-1`). With the introduction of the Frankfurt region, the format switched from a three-character code to a longer string. This part is present in the OCID only for regional resources or those specific to a single availability domain. If the region is not applicable to the resource, this part might be blank (see the example tenancy ID below).

CHAPTER 4 Service Essentials

- **future use:** Reserved for future use. Currently blank.
- **unique ID:** The unique portion of the ID. The format may vary depending on the type of resource or service.

Example OCIDs

Tenancy:

```
ocid1.tenancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdsq
```

Instance:

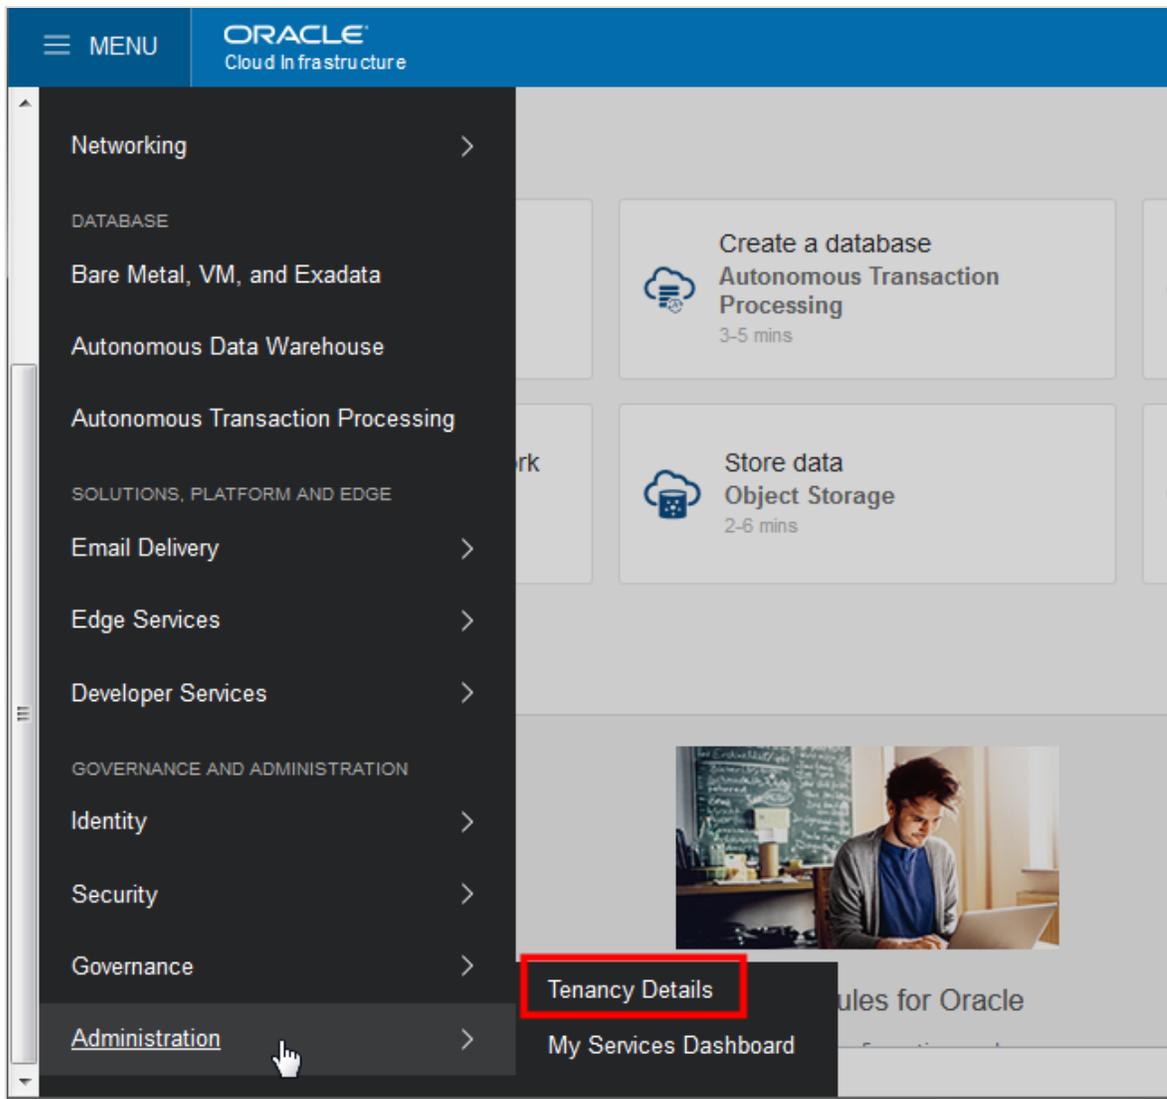
```
ocid1.instance.oc1.phx.abuw41jr1sfiqw6vzzxb43vyypt4pkodawglp3wqxjqofakrwvou52gb6s5a
```

Where to Find Your Tenancy's OCID

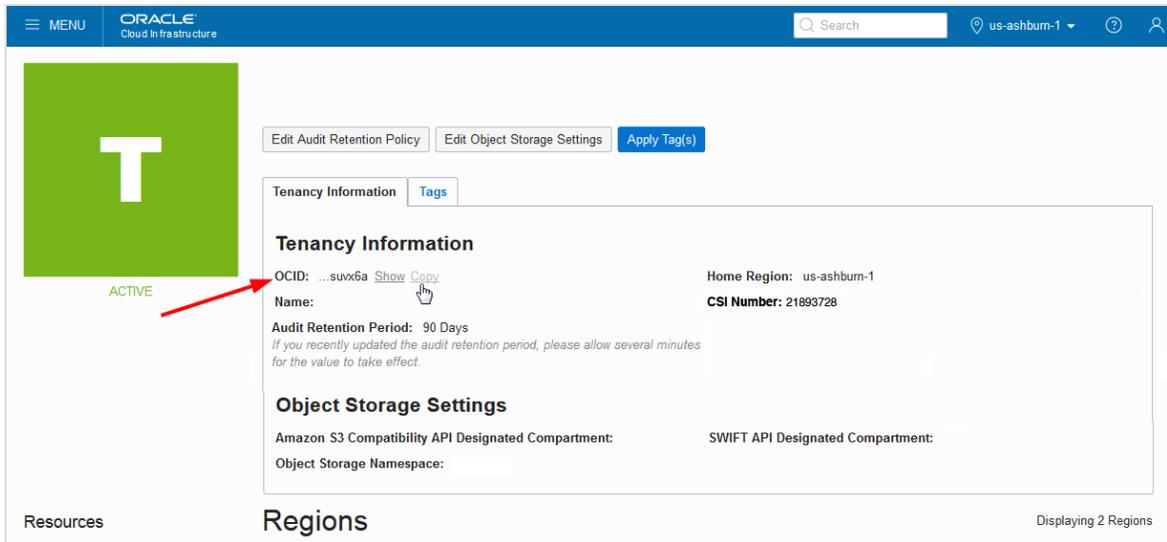
If you use the Oracle Cloud Infrastructure API, you need your tenancy's OCID in order to sign the API requests. You also use the tenancy ID in some of the IAM API operations.

Get the tenancy OCID from the Oracle Cloud Infrastructure Console on the **Tenancy Details** page:

1. Open the navigation menu, under Governance and Administration, go to **Administration** and click **Tenancy Details**.



2. The tenancy OCID is shown under **Tenancy Information**. Click **Copy** to copy it to your clipboard.



The tenancy OCID looks something like this (notice the word "tenancy" in it):

```
ocid1.tenancy.oc1..<unique_ID>
```

Name and Description (IAM Only)

The IAM service requires you to assign a unique, unchangeable *name* to each of your IAM resources (users, groups, dynamic groups, federations, and policies). The name must be unique within the scope of the type of resource (for example, you can only have one user with the name BobSmith). Notice that this requirement is specific to IAM and not the other services.

The name you assign to a user at creation is their login for the Console.

You can use these names instead of the OCID when writing a policy (for example, `Allow group <GROUP NAME> to manage all-resources in compartment <COMPARTMENT NAME>`).

In addition to the name, you must also assign a *description* to each of your IAM resources (although it can be an empty string). It can be a friendly description or other information that helps you easily identify the resource. The description does not have to be unique, and you can change it whenever you like. For example, you might want to use the description to store

the user's email address if you're not already using the email address as the user's unique name.

Display Name

For most of the Oracle Cloud Infrastructure resources you create (other than those in IAM), you can optionally assign a *display name*. It can be a friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. The Console shows the resource's display name along with its OCID.

Resource Tags

When you have many resources (for example, instances, VCNs, load balancers, and block volumes) across multiple compartments in your tenancy, it can become difficult to track resources used for specific purposes, or to aggregate them, report on them, or take bulk actions on them. *Tagging* allows you to define keys and values and associate them with resources. You can then use the tags to help you organize and list resources based on your business needs.

There are two types of tags:

[Defined tags](#) are set up in your tenancy by an administrator. Only users granted permission to work with the defined tags can apply them to resources.

[Free-form tags](#) can be applied by any user with permissions on the resource.

For more detailed information about tags and their features, see [Tagging Overview](#).



Tip

Watch a video to introduce you to the concepts and features of tagging: [Introduction to Tagging](#).

Working with Resource Tags



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

To add a defined tag to an existing resource

To apply a defined tag, you must have permission to use the namespace.

1. Open the Console, go to the details page of the resource you want to tag.
For example, to tag a compute instance: Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. A list of the instances in your current compartment is displayed. Find the instance that you want to tag, and click its name to view its details page.
2. Click **Apply Tags**.
3. In the **Apply Tags to the Resource** dialog:
 - a. Select the **Tag Namespace**.
 - b. Select the **Tag Key**.
 - c. In **Value**, either enter a value or select one from the list.
 - d. To apply another tag, click **+ Additional Tag**.
 - e. When finished adding tags, click **Apply Tag(s)**.

To add a free-form tag to an existing resource

1. Open the Console, go to the details page of the resource you want to tag.

For example, to tag a compute instance: Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. A list of the instances in your current compartment is displayed. Find the instance that you want to tag, and click its name to view its details page.

2. Click **Apply Tags**.
3. In the **Apply Tags to the Resource** dialog:
 - a. Select **None (apply a free-form tag)**.
 - b. Enter the **Tag Key**.
 - c. Enter a **Value**.
 - d. To apply another tag, click **+ Additional Tag**.
 - e. When finished adding tags, click **Apply Tag(s)**.

To add a tag during resource creation

You can apply tags during resource creation. The location of the **Apply Tag(s)** option in the dialog varies by resource. The general steps are:

1. In the resource Create dialog, click **Apply Tags**.
On some resources, you have to click **Show Advanced Options** to apply a tag.
2. In the **Apply Tags to the Resource** dialog:
 - a. Select the **Tag Namespace**, or select **None** to apply a free-form tag.
 - b. Select or enter the **Tag Key**.
 - c. In **Value**, either enter a value or select one from the list.
 - d. To apply another tag, click **+ Additional Tag**.
 - e. Click **Apply Tag(s)**.

To filter a list of resources by a tag

Open the Console, click the service name and then click the resource you want to view. The

left side of the page shows all the filters currently applied to the list.

For example, to view compute instances: Click **Compute** and then click **Instances**, to see the list of instances in your current compartment.

To filter a list of resources by a defined tag

1. Next to **Tag Filters**, click **add**.
2. In the **Apply a Tag Filter** dialog, enter the following:
 - a. **Namespace:** Select the tag namespace.
 - b. **Key:** Select a specific key.
 - c. **Value:** Select from the following:
 - **Match Any Value** - returns all resources tagged with the selected namespace and key, regardless of the tag value.
 - **Match Any of the Following** - returns resources with the tag value you enter in the text box. Enter a single value in the text box. To specify multiple values for the same namespace and key, click **+** to display another text box. Enter one value per text box.
 - d. Click **Apply Filter**.

To filter a list of resources by a free-form tag

1. Next to **Tag Filters**, click **add**.
2. In the **Apply a Tag Filter** dialog, enter the following:
 - a. **Key:** Enter the tag key.
 - b. **Value:** Select from the following:
 - **Match Any Value** - returns all resources tagged with the selected free-form tag key, regardless of the tag value.

- **Match Any of the Following** - returns resources with the tag value you enter in the text box. Enter a single value in the text box. To specify multiple values for the same key, click **+** to display another text box. Enter one value per text box.
- c. Click **Apply Filter**.

To update a tag applied to a resource

1. Open the Console, click the service name and then click the resource you want to view. For example, to view compute instances: Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. A list of the instances in your current compartment is displayed. Find the instance that you want to update, and click its name to view its details page.
2. Click **Tags**.
The list of tags applied to the resource is displayed.
3. Find the tag you want to update and click the pencil icon next to it.
4. Enter or select a new value.
5. Click **Save**.

To remove a tag from a resource

1. Open the Console, click the service name and then click the resource you want to view. For example, to view a compute instance: Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. A list of the instances in your current compartment is displayed. Find the instance that you want to remove the tag from, and click its name to view its details page.
2. Click **Tags**.
The list of tags applied to the resource is displayed.
3. Find the tag you want to remove and click the pencil icon next to it.

4. Click **Remove Tag**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To apply a tag to a resource using the API, use the appropriate resource's `create` or `update` operation.

Resource Monitoring

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources when needed using queries or on a passive basis using alarms. Queries and alarms rely on metrics emitted by your resource to the Monitoring service.

Prerequisites

- **IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- **Metrics exist in Monitoring:** The resources that you want to monitor must emit metrics to the Monitoring service.

- Compute instances: To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).

Working with Resource Monitoring

Not all resources support monitoring. See [Supported Services](#) for the list of resources that support the Monitoring service, which is required for queries and alarms used in monitoring.

The Monitoring service works with the Notifications service to notify you when metrics breach. For more information about these services, see [Monitoring Overview](#) and [Notifications Overview](#).

To view default metric charts for a resource

On the page for the resource of interest, under **Resources**, click **Metrics**.

For example, to view metric data for a Compute instance:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance you're interested in.
3. On the instance detail page, under **Resources**, click **Metrics**.
A chart is shown for each metric. For a list of metrics related to Compute instances, see [Compute Instance Metrics](#).

The Console displays the last hour of metric data for the selected resource. A chart is shown for each metric emitted by the selected resource.

For a list of metrics emitted by your resource, see [Supported Services](#).

To view default metric charts for a set of resources

Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.

The **Service Metrics** page displays the default charts for all resources in the first accessible **Compartment** and **Metric Namespace**. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

Don't see all expected resources or metrics?

- Try a [different time range](#).
- Make sure the correct **Compartment** is selected.
On the **Service Metrics** page, metric namespaces are shown only when associated resources exist in the selected compartment. For example, the `oci_autonomous_database` namespace is shown only when Autonomous Databases exist in the selected compartment.
- Confirm that the missing resources are emitting metrics. See [Enabling Monitoring for Compute Instances](#).
- Review limits information. Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

To create a query

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
The **Metrics Explorer** page displays an empty chart with fields to build a query.
2. Fill in the fields for a new query.
 - **Compartment**: The compartment containing the resources that you want to monitor. By default, the first accessible compartment is selected.
 - **Metric Namespace**: The service or application emitting metrics for the resources that you want to monitor.

- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Metric Name**: The name of the metric. Only one metric can be specified. Metric selections depend on the selected compartment and metric namespace.
Example: **CpuUtilization**
- **Interval**: The aggregation window.

Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour

**Note**

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

Maximum time range returned for a query

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

Interval	Default resolution (metric queries)	Maximum time range returned
1h	1 hour	90 days
5m	5 minutes	30 days
1m	1 minute	7 days



See examples of returned data

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No resolution or end time is specified, so the resolution defaults to the interval value of 1m, and the end time defaults to the current time (2019-01-08T10:00:00.789Z). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (2019-01-06T10:00:00.789Z), the earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

- **Statistic:** The aggregation function.

Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
 - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Aggregate Metric Streams:** Aggregates all results to plot a single aggregated average for all metric streams. This average is plotted as a single line on the metric chart. This operation is helpful when you want to plot a metric as one line for all resources.

3. Click **Update Chart**.

The chart shows the results of your new query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power). Units correspond to the selected metric and do not change by statistic.

Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

4. To customize the y-axis label or range, type the label you want into **Y-Axis Label** or type the minimum and maximum values you want into **Y-Axis Min Value** and **Y-Axis Max Value**.

Only numeric characters are allowed for custom ranges. Custom labels and ranges are not persisted in shared queries (MQL).

5. To view the query as a Monitoring Query Language (MQL) expression, click **Advanced Mode**.

Advanced Mode is located on the right, under the chart.

Use Advanced Mode to edit your query using MQL syntax to [aggregate results by group](#). The MQL syntax also supports additional parameter values. For more information about query parameters in Basic Mode and Advanced Mode, see [Monitoring Query Language \(MQL\) Reference](#).

6. To create another query, click **Add Query** below the chart.

To create an alarm

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click **Create alarm**.



Note

You can also create an alarm from a predefined query on the **Service Metrics** page. Expand **Options** and click **Create an Alarm on this Query**. For more information about service metrics, see [Viewing Default Metric Charts](#).

3. On the **Create Alarm** page, under **Define alarm**, fill in or update the alarm settings:



Note

To toggle between Basic Mode and Advanced Mode, click **Switch to Advanced Mode** or **Switch to Basic Mode** (to the right of **Define Alarm**).

Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description:** The metric to evaluate for the alarm condition.
 - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
 - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.

- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Metric Name**: The name of the metric. Only one metric can be specified. Example: **CpuUtilization**
- **Interval**: The aggregation window, or the frequency at which data points are aggregated.

Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic**: The aggregation function.

Statistic values

- **COUNT**- The number of observations received in the specified time

period.

- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
 - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

Operator values

- **greater than**
 - **greater than or equal to**
 - **equal to**
 - **less than**
 - **less than or equal to**
 - **between** (inclusive of specified values)
 - **outside** (inclusive of specified values)
 - **absent**
- **Value:** The value to use for the condition threshold.
 - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
 - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

4. Set up notifications: Under **Notifications**, fill in the fields.

. Destinations:

- **Destination Service:** The provider of the destination to use for notifications.
Available options:
 - [Notifications Service](#).
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.
- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.

- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
 - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
 - **Topic Description:** Description of the new topic.
 - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com OR
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

Slack subscription



Note

See the following [known issue](#) for up-to-date information about creating



Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The `<webhook-token>` portion of the URL contains two slashes (/).

Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
 - **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
- **Notification Interval:** The period of time to wait before resending the notification.

- **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
 - **Suppression Description**
 - **Start Time**
 - **End Time**
5. If you want to disable the new alarm, clear **Enable This Alarm?**
 6. Click **Save alarm**.

The new alarm is listed on the **Alarm Definitions** page.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To create a query, use the [SummarizeMetricsData](#) operation.

To create an alarm, use the [CreateAlarm](#) operation.

Service Limits

This topic describes the service limits for Oracle Cloud Infrastructure and the process for requesting a service limit increase.

About Service Limits and Usage

When you sign up for Oracle Cloud Infrastructure, a set of service limits are configured for your tenancy. The service limit is the quota or allowance set on a resource. For example, your tenancy is allowed a maximum number of compute instances per availability domain. These limits are generally established with your Oracle sales representative when you purchase Oracle Cloud Infrastructure. If you did not establish limits with your Oracle sales

representative, or, if you signed up through the Oracle Store, default or trial limits are set for your tenancy. These limits may be increased for you automatically based on your Oracle Cloud Infrastructure resource usage and account standing. You can also [request a service limit increase](#).

Compartment Quotas

Compartment quotas are similar to Service Limits; the biggest difference is that service limits are set by Oracle, and compartment quotas are set by administrators, using policies that allow them to allocate resources with a high level of flexibility. Compartment quotas are set using *policy statements* written in a simple declarative language that is similar to the IAM policy language.

To learn more, see [Compartment Quotas](#).

Viewing Your Service Limits, Quotas, and Usage

You can view your tenancy's limits, quotas, and usage in the Console. Be aware that:

- The Console may not yet display limits and usage information for all of the Oracle Cloud Infrastructure services or resources.
- The usage level listed for a given resource type could be greater than the limit if the limit was reduced after the resources were created.
- If all the resource limits are listed as 0, this means your account has been suspended. For help, [contact Oracle Support](#).

If you don't yet have a tenancy or a user login for the Console, or if you don't find a particular limit listed in the Console, see [Limits by Service](#) for the default tenancy limits.

To view your tenancy's limits and usage (by region)



Note

Required Permission

If you're in the [Administrators group](#), you have permission to view the limits and usage. If you're not, here's an example [IAM policy](#) that grants the required permission to users in a group called `LimitsAndUsageViewers`:

```
Allow group LimitsAndUsageViewers to inspect resource-availability in tenancy
```

1. Open the Console. Open the navigation menu, under Governance and Administration, go to **Governance**, and then click **Limits, Quotas and Usage**.

Your resource limits, quotas, and usage for the specific region are displayed, broken out by service. You can use the filter drop-down lists at the top of the list to filter by service, scope, resource, and compartment.

When You Reach a Service Limit

When you reach the service limit for a resource, you receive an error when you try to create a new resource of that type. You are then prompted to submit a request to increase your limit. You cannot create a new resource until you are granted an increase to your service limit or you terminate an existing resource. Note that service limits apply to a specific scope, and when the service limit in one scope is reached you may still have resources available to you in other scopes (for example, other availability domains).

Requesting a Service Limit Increase

You can submit a request to increase your service limits from within the Console. If you try to create a resource for which limit has been met, you'll be prompted to submit a limit increase request. Additionally, you can launch the request from the service limits page or at any time by clicking the link under the **Help** menu (?).

Note that the service limit increase is not immediate.

To request a service limit increase

1. Open the **Help** menu (?), go to **Support** and click **Request service limit increase**.
2. Enter the following:
 - **Primary Contact Details:** Enter the name and email address of the person making the request. Enter one email address only. A confirmation will be sent to this address.
 - **Service Category:** Select the appropriate category for your request.
 - **Resource:** Select the appropriate resource.
Depending on your selection for resource, additional fields might display for more specific information.
 - **Reason for Request:** Enter a reason for your request. If your request is urgent or unusual, please provide details here.
3. Click **Submit Request**.

After you submit the request, the request is reviewed. If your request is awarded, a confirmation email is sent to the address provided in the primary contact details.

If we need additional information about your request, a follow-up email is sent to the address provided in the primary contact details.

Limits by Service

The following tables list the default limits for each service. Note the scope that each limit applies to (for example, per availability domain, per region, per tenant, etc.).



Note

Some services have additional limits. For more information, see the overview of each service.

Analytics Cloud Limits

For Analytics Cloud limits, see [Service Limits](#).

Block Volume Limits

Volume limits apply to each availability domain. Volume backup limits apply to each region.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Block Volumes aggregated size	100 TB	30 TB
Backups	1000	500

Compute Limits

Limits apply to each availability domain, unless otherwise noted.

Bare Metal Servers

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
BM.Standard.B1.44	Contact Us	Contact Us
BM.Standard2.52	5 (52 cores) (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 2 (52 cores) (other regions)	Contact Us
BM.DenseIO2.52	2 (52 cores)	Contact Us
BM.GPU2.2	5 (28 cores) (US East (Ashburn), Germany Central (Frankfurt))	Contact Us
BM.GPU3.8	Contact Us	Contact Us
BM.HPC2.36	Contact Us	Contact Us
BM.Standard.E2.64	5 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 2 (other regions)	Contact Us

Virtual Machines

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
VM.Standard2.1	100 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 50 (other regions)	2 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix))
VM.Standard2.2	80 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 40 (other regions)	2 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix))
VM.Standard2.4	80 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 40 (other regions)	Contact Us
VM.Standard.E2.1	40 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 20 (other regions)	2 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix))
VM.Standard.E2.2	30 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 15 (other regions)	2 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix))

CHAPTER 4 Service Essentials

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
VM.Standard.E2.4	30 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 15 (other regions)	Contact Us
VM.Standard.E2.8	10 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 5 (other regions)	Contact Us
VM.Standard2.8	40 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 20 (other regions)	Contact Us
VM.Standard2.16	40 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 20 (other regions)	Contact Us
VM.Standard2.24	40 (US East (Ashburn), Germany Central (Frankfurt), UK South (London), US West (Phoenix)) 20 (other regions)	Contact Us
VMDenseIO2.8	5	Contact Us

CHAPTER 4 Service Essentials

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
VMDenseIO2.16	5	Contact Us
VMDenseIO2.24	5	Contact Us
VM.GPU3.1	Contact Us	Contact Us
VM.GPU3.2	Contact Us	Contact Us
VM.GPU3.4	Contact Us	Contact Us

Other Compute Resources

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Custom images	100 per region	25 per region
Cluster networks	5 per tenancy	Contact Us

Container Engine for Kubernetes Limits

Container Engine for Kubernetes limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Clusters	3 clusters per OCI region	1 cluster per OCI region
Nodes	1000 nodes per cluster	1000 nodes per cluster

Data Safe Limits

To register an Oracle Database with [Data Safe](#), you must be using a [paid account](#).

Data Transfer Limits

Data Transfer limits are **regional**.

Disk-Based Data Transfer

Resource	Monthly Universal Credits	Pay-As-You-Go
Transfer package	Contact Us	Contact Us

File a service request at [My Oracle Support](#) to increase the service limits for Disk-Based Data Transfer. See [Requesting a Service Limit Increase](#) for details.

Appliance-Based Data Transfer

Resource	Monthly Universal Credits	Pay-As-You-Go
Transfer appliances	Contact Your CSM	Contact Your CSM

Contact your Oracle Account Manager or Customer Success Manager (CSM) to place an order for transfer appliances. Your service limits are automatically set to the number of transfer appliances ordered. You do not need to file a service request to increase the service limits for Appliance-Based Data Transfer.

Database Limits

Database limits are per [availability domain](#).

See [Data Safe limits](#) for information on [Data Safe](#).

CHAPTER 4 Service Essentials

Resources	Monthly Flex	Pay-as-You-Go or Promo
Autonomous Database Serverless Deployment - Total OCPUs	128 cores	8 cores
Always Free Autonomous Database	2 instances	2 instances
Always Free Autonomous Database - Total OCPUs	1 core	1 core
Always Free Autonomous Database - Total Block Storage	20 GB	20 GB
VM.Standard1 -Total OCPUs	10 cores	2 cores
VM.Standard2 -Total OCPUs	100 cores (US West (Phoenix), US East (Ashburn)) 50 cores (Germany Central (Frankfurt), uk-london-1)	2 cores
Total VM DB Block Storage (see note)	10TB	2TB
BM.DenseIO1.36 (see availability note)	1 instance	1 instance
BM.DenseIO2.52	1 instance	1 instance
BM.HighIO1.36	1 instance	1 instance
Exadata.Base.48	Contact Us	not available
Exadata.Quarter1.84 - X6	Contact Us	not available

CHAPTER 4 Service Essentials

Resources	Monthly Flex	Pay-as-You-Go or Promo
Exadata.Half1.168 - X6	Contact Us	not available
Exadata.Full1.336 - X6	Contact Us	not available
Exadata.Quarter2.92 - X7	Contact Us	not available
Exadata.Half2.184 - X7	Contact Us	not available
Exadata.Full2.368 - X7	Contact Us	not available
BM.RACLocalStorage1.72	Contact Us	Contact Us



Note

- **Autonomous Exadata Infrastructure** - Exadata X7 limits (Exadata.Quarter2.92, Exadata.Half2.184, and Exadata.Full2.368) cover the corresponding Autonomous Exadata Infrastructure shapes.
- **Total VM DB Block Storage** - covers block storage for all VM.Standard1 and VM.Standard2 virtual machine databases.
- **BM.DenseIO1.36** - This DB system shape is available only to monthly universal credit customers with tenancies existing on or before November 9th, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.
- Each of the two Always Free Autonomous Database available in your tenancy can be provisioned with your choice of Autonomous Transaction Processing or Autonomous Data Warehouse workload types.

DNS Limits

DNS limits are **global**.

Resources	Monthly Universal Credits	Pay-as-You-Go or Promo
Zones	1,000 zones	1,000 zones
Records	25,000 per zone	25,000 per zone
Zone File Size	1 MB	1 MB

Email Delivery Limits

Limits apply to each tenant or availability domain, as specified.

Resource	Monthly Universal Credits	Pay-As-You-Go or Promo
Email volume	50,000 emails per day	200 emails per day
Maximum approved senders	10,000	2,000
SMTP credentials	2 per user	2 per user
Sending rate	18,000 emails per minute	10 emails per minute



Note

Message size is limited to 2 MB, inclusive of message headers, body, and attachments.

Events Limits

Events limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Rules	50	50

File Storage Limits

Limits apply to each tenant or availability domain, as specified.

Resource	Pre-Paid	Pay-As-You-Go
File systems	100 per tenant per availability domain	100 per tenant per availability domain
Mount targets	2 per tenant per availability domain	2 per tenant per availability domain
Maximum file system size	8 exabytes	8 exabytes

Functions Limits

Limits apply to each tenancy.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Applications	10	10
Functions	20	20

Health Checks Limits

Health Checks limits are **global**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Endpoint tests	1000 per account	1000 per account

IAM Limits

IAM limits are **global**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Users in a tenancy	2000	2000
Groups in a tenancy	250	250
Compartments in a tenancy	50	50
Policies in a tenancy	100	100
Statements in a policy	50	50
Users per group in a tenancy	2000	2000
Groups per user in a tenancy	250	250
Identity providers in a tenancy	3	3
Group mappings for an identity provider	250	250

Key Management Limits

Key Management limits are **global**.

Resources	Monthly Universal Credits	Pay-as-You-Go or Promo
Vaults in a tenancy	Contact Us	Contact Us
Keys in a vault (Key versions, whether enabled or disabled, count against your limits.)	Contact Us	Contact Us

The Key Management service is not available to promo customers.

Load Balancing Limits

Load Balancing limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
LB-Capacity-10Mbps	1 Load Balancer	1 Load Balancer
LB-Capacity-100Mbps	3 Load Balancers	1 Load Balancer
LB-Capacity-400Mbps	3 Load Balancers	1 Load Balancer
LB-Capacity-8000Mbps	Contact Us	Contact Us

Monitoring Limits

CHAPTER 4 Service Essentials

Monitoring limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Alarms	50	50
Metrics (posted by services)	Unlimited	Unlimited

Networking Limits

Networking service limits apply to different scopes, depending on the resource.

VCN and Subnet Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
VCN	Region	10	10
Subnets	VCN	300	300

Gateway Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Dynamic routing gateways (DRGs)	Region	5	5
Internet gateways	VCN	1*	1*

CHAPTER 4 Service Essentials

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Local peering gateways (LPGs)	VCN	10	10
NAT gateways	VCN	1	1
Service gateways	VCN	1	1
* Limit for this resource cannot be increased			

IP Address Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Reserved public IPs	Region	50	50
Ephemeral public IPs	Instance	2 per VM instance 16 per bare metal instance	2 per VM instance 16 per bare metal instance

DHCP Option Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
DHCP options	VCN	300	300

Route Table Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Route tables	VCN	300	300
Route rules	Route table	100*	100*
* Limit for this resource cannot be increased			

Network Security Group Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Network security groups	VCN	1000	1000
VNICs	Network security group	A given network security group can have as many VNICs as are in the VCN. A given VNIC can belong to a maximum of 5 network security groups.*	A given network security group can have as many VNICs as are in the VCN. A given VNIC can belong to a maximum of 5 network security groups.*
Security rules	Network security group	120 (total ingress plus egress)	120 (total ingress plus egress)
* Limit for this resource cannot be increased			

Security List Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Security lists	VCN	300	300
Security lists	Subnet	5*	5*
Security rules	Security list	200 ingress rules* and 200 egress rules*	200 ingress rules* and 200 egress rules*
* Limit for this resource cannot be increased			

IPSec VPN Connection Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
IPSec VPN connections	Region	4	4
Customer-premises equipment objects (CPEs)	Region	10	10
Dynamic routing gateways (DRGs)	Region	See Gateway Limits	See Gateway Limits

FastConnect Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Cross-connects	Region	Contact Us	Contact Us
Virtual circuits	Region	10	10
Dynamic routing gateways (DRGs)	Region	See Gateway Limits	See Gateway Limits

Notifications Limits

Notifications limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Topics	50 (Active or Creating*) per tenancy	Contact Us
Subscriptions	10 (Active or Pending*) per topic 100 (Pending*) per tenancy	Contact Us

* A lifecycle state. See [NotificationTopic Reference](#) and [Subscription Reference](#).

Object Storage and Archive Storage Limits

Object Storage and Archive Storage limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Buckets	10,000 per tenancy	10,000 per tenancy
Objects per bucket	Unlimited	Unlimited

Registry Limits

Registry limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Repositories	500 repositories per OCI region	500 repositories per OCI region
Images	500 images per repository	500 images per repository

Resource Manager Limits

Resource Manager limits are **regional**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Stacks	100 stacks per tenant	10 stacks per tenant
Variables per stack	100	100
Zip file per stack	11 MB	11 MB
Jobs (concurrent)	5 per tenant	2 per tenant
Job Duration	24 hours	24 hours

Traffic Management Steering Policies Limits

Traffic Management Steering Policies limits are **global**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Policies	100 per tenant	100 per tenant
Attachments	1,000 per tenant	1,000 per tenant

WAF Limits

WAF limits are **global**.

Resource	Monthly Universal Credits	Pay-as-You-Go or Promo
Policies	50 per tenant	50 per tenant
Access rules	100 per policy	100 per policy

Viewing All Resources in a Compartment

This topic describes how you can use the compartment explorer to get a cross-region view of all resources in a compartment.

Compartment Explorer Overview

The compartment explorer allows you to view all your resources in a specific compartment, across all regions. You can also easily navigate up and down the compartment hierarchy to get the complete view of the resources in a specific compartment tree. From the compartment explorer you can navigate directly to resources, providing an alternative interface to access the resources in each compartment. The compartment explorer also allows you to move a resource to a different compartment.

The following image highlights these features:

CHAPTER 4 Service Essentials

ORACLE Cloud US West (Phoenix)

Governance

- Audit
 - Compartment Explorer
 - Quota Policies
 - Limits, Quotas and Usage
 - Tag Names

Compartment Explorer

Name: [CompartmentABC](#)
Description: Company ABC's compartment

Don't see what you're looking for? These results include only resources supported by [Search](#). Updates made to resources might not immediately appear in your results.

Instance x Vcn x Bucket x

Name	Resource Type	Status	Region	
Bucket-Ashburn	Bucket	Active	us-ashburn-1	
BucketC	Bucket	Active	us-phoen	View Details
D_CompanyVCN	Vcn	Available	eu-fra	Move Resource
F_CompanyVCN	Vcn	Available	us-phoenix-1	
G_CompanyVCN	Vcn	Available	eu-frankfurt-1	
Instance_1	Instance		us-phoenix-1	
Instance_2	Instance		us-phoenix-1	
Instance_3	Instance	TERMINATED	ap-mumbai-1	
Instance_4	Vcn	Available	us-phoenix-1	

When using the compartment explorer, be aware of the following:

- If you recently created a resource, it might not show up in the compartment explorer immediately. Similarly, if you recently updated a resource, your changes might not immediately appear.
- To navigate to the details page of a resource or to move it to another compartment, you must be in the same region as the resource. The compartment explorer displays the resource's region. Use the region selector at the top of the Console to change to the same region as the resource to enable these actions.

Resources Supported by the Compartment Explorer

The compartment explorer is powered by the Search service and supports the same resource types. Most resources are supported.

Supported resources

Service	Resource Type
Block Volume	bootvolume
Block Volume	bootvolumebackup
Block Volume	volume
Block Volume	volumebackup
Compute	autoscalingconfiguration
Compute	consolehistory
Compute	image
Compute	instance
Database	autonomousdatabase
Database	database
Database	dbsystem
Events	eventrule
Functions	functionsapplication
Functions	functionsfunction
IAM	compartment
IAM	group
IAM	identityprovider

Service	Resource Type
IAM	user
Key Management	key
Key Management	vault
Monitoring	alarm
Networking	routetable
Networking	securitylist
Networking	subnet
Networking	vcn
Notifications	onssubscription
Notifications	onstopic
Object Storage	bucket
Resource Manager	ormjob
Resource Manager	ormstack
WAF	waascertificate
WAF	waaspolicy

Required IAM Policy to View Resources in the Compartment Explorer

The resources that you see in the compartment explorer depend on the permissions you have in place for the resource type. You do not necessarily see results for everything in the compartment. For example, if your user account is not associated with a policy that grants you the ability to, at a minimum, `inspect` the `dbssystem` resource type, then you will not be

able to view DB systems in the compartment explorer. (The verb `inspect` lets you list and get resources.) For more information about policies, see [How Policies Work](#). For information about the permissions required for the list API operation for a specific resource type, see the [Policy Reference](#) for the appropriate service.

Navigating to the Compartment Explorer and Viewing Resources

Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Compartment Explorer**.

When you open the compartment explorer, the list of all resources that you have permission to view is displayed. The compartment explorer opens with a view of the root compartment. The **Name** and **Description** of the compartment you are viewing are displayed at the top of the page.

To navigate to the compartment you are interested in, use the compartment picker on the left of the Console page.

Filtering Displayed Resources

To view only specific resource types, select the resource types you are interested in from the **Filter by resource type** menu. You can select multiple resources to include in the filtered list.

Opening the Resource Details Page

To open the details page for a resource:

1. Locate the resource in the list.
2. Verify that you are in the same region as the resource. The resource's region is listed in the compartment explorer results. If it is not the same as the region you are currently in (shown at the top of the Console), then select the appropriate region from the **Regions** menu.
3. To open the details page, you can either:

- Click the name.
- Click the the Actions icon (three dots) and select **View Details**.

Detail page navigation is not supported for all resource types. If detail page navigation is not supported, the resource name does not display as a link and the option is not available from the Action menu.

Moving a Resource to a Different Compartment

To move a resource to a different compartment:

1. Locate the resource in the list.
2. Verify that you are in the same region as the resource. The resource's region is listed in the compartment explorer results. If it is not the same as the region you are currently in (shown at the top of the Console), then select the appropriate region from the **Regions** menu.
3. Click the the Actions icon (three dots) and select **Move Resource**.
4. In the dialog, choose the destination compartment from the list.
5. Click **Move Resource**.



Important

Ensure that you understand the impact for a resource before you move it. See the resource's service documentation for details.

Not all resources can be moved to a different compartment. If the resource cannot be moved, the option is not available from the Action menu.

Compartment Quotas

This topic describes compartment quotas for Oracle Cloud Infrastructure.

Compartment quotas give tenant and compartment administrators better control over how resources are consumed in Oracle Cloud Infrastructure, enabling administrators to easily allocate resources to compartments using the Console. Along with [compartment budgets](#), compartment quotas create a powerful toolset to manage your spending in Oracle Cloud Infrastructure tenancies.

You can start using compartment quotas from any compartment detail page in the Console.

About Compartment Quotas

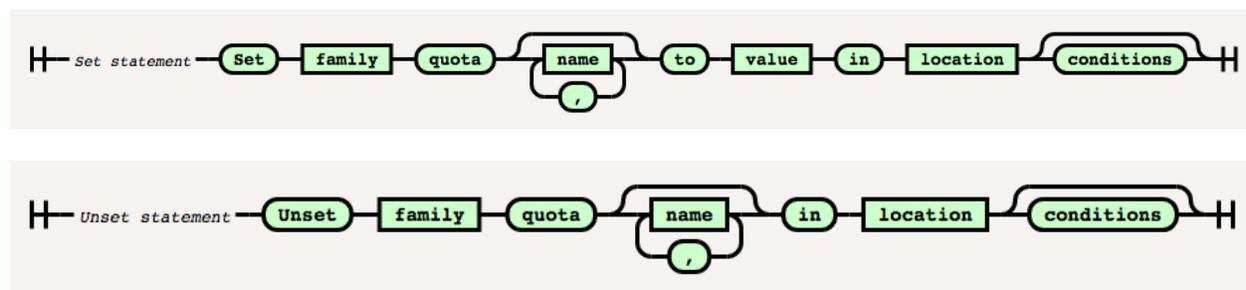
Compartment quotas are similar to [Service Limits](#); the biggest difference is that service limits are set by Oracle, and compartment quotas are set by administrators, using policies that allow them to allocate resources with a high level of flexibility.

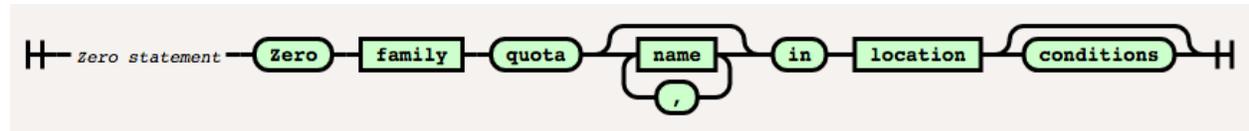
Compartment quotas are set using *policy statements* written in a simple declarative language that is similar to the IAM policy language.

There are three types of quota policy statements:

- `set` - sets the maximum number of a cloud resource that can be used for a compartment
- `unset` - resets quotas back to the default service limits
- `zero` - removes access to a cloud resource for a compartment

The quota policy statements look like this:





The language components for a quota policy statement are:

- The `action` keyword, which corresponds to the type of quota being defined. This can be `set`, `unset`, or `zero`.
- The name of the service family; for example: `compute`.
- The `quota` or `quotas` keyword
- The name of the quota, which varies by service family. For example, a valid quota in the `compute` family is `vm-standard2-16-count`.
 - You can also use wildcards to specify a range of names. For example, `"/vm-*/"` matches all Compute shapes that start with the letters "vm".
- For set statements, the value of the quota.
- The compartment that the quota covers.
- An optional condition. For example `where request.region = 'us-phoenix-1'`. Currently supported conditionals are `request.region` and `request.ad`.

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

For common policies used to authorize users, see [Common Policies](#).

To manage quotas in a compartment, you must belong to a group that has the correct permissions. For example:

```
allow group QuotaAdmins to { QUOTA_READ, QUOTA_CREATE, QUOTA_DELETE, QUOTA_UPDATE, QUOTA_INSPECT } in
tenancy
```

For in-depth information on granting users permissions for the Quotas service, see [Details for the Quotas Service](#) in the IAM policy reference.

Permissions and Nesting

Compartment quotas can be set on the root compartment. An administrator (who must be able to manage quotas on the root compartment) can set quotas on their own compartments and any child compartments. Quotas set on a parent compartment override quotas set on child compartments. This way, an administrator of a parent compartment can create a quota on a child compartment that cannot be overridden by the child.

Scope

Quotas can have different *scopes*, and work at the availability domain, the region, or globally.

There are a few important things to understand about scope when working with compartment quotas:

- When setting a quota at the availability domain (AD) level, the quota is allocated to each AD. So, for example, setting a quota of 2 X7 VMs on a compartment actually sets a limit of 2 VMs per AD. To target a specific AD, use the `request.ad` parameter in the `where` clause.
- Regional quotas apply to each region. For example, if a quota of 10 functions is set on a

compartment, 10 functions will be allocated per region. To target a specific region, use the `request.region` parameter in the `where` clause.

- Usage for sub-compartments counts towards usage for the main compartment.

For more information, see [Regions and Availability Domains](#).

Quota Evaluation and Precedence

The following rules apply when quota statements are evaluated:

- Within a policy, quota statements are evaluated in order, and later statements supersede previous statements that target the same resource.
- In cases where more than one policy is set for the same resource, the most restrictive policy is applied.
- Service limits always take precedence over quotas. Although it is possible to specify a quota for a resource that exceeds the service limit for that resource, the service limit will still be enforced.

Usage Examples

The following example sets the quota for `VM.DenseIO1.16` Compute shapes to 10 in each AD on compartment `MyCompartment` in the US West (Phoenix) region:

```
set compute quota vm-dense-io1-16-count to 10 in compartment MyCompartment
where request.region = us-phoenix-1
```

The next example shows how to make a whitelist, setting every quota in a family to zero and then explicitly allocating resources:

```
zero compute quotas in tenancy
set compute quota vm-dense-io1-16-count to 10 in tenancy
```

This example shows how to limit creating a bare metal compute resource to only one region:

```
zero compute quotas /*bm*/ in tenancy
set compute quota /*bm*/ to 5 in tenancy where request.region = us-phoenix-1
```

This example policy statement only allows one VM.Standard2.1 Compute instance in a single compartment in a single region:

```
zero compute quotas in tenancy
set compute quota vm-standard2-1-count to 10 in compartment sales_department
where request.region = us-phoenix-1
```

You can clear quotas by using an `unset` statement, which removes the quota for a resource - any limits on this resource will now be enforced by the service limits:

```
zero compute quotas in tenancy
unset compute quota vm-dense-iol-16-count in tenancy
```

Using the Console

To create a quota

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Quota Policies**. From the **Quota Policies** screen, click **Create Quota**.
2. Enter the following:
 - Enter a name for your quota in the **Name** field.
 - Enter a description for your quota in the **Description** field.
 - Enter a quota policy string in the **Quota Policy** field.
3. Click **Create Quota Policy**.



Note

New policies can take up to 10 minutes to start working.

To edit a quota

1. From the **Quota Policies** screen, click the quota you want to edit to display the quota policy details page, then click the **Edit Quota** button.
2. Edit the quota.
3. Click **Save Changes**.

To delete a quota

1. There are two ways to delete a quota from the console:
 - From the main **Quota Policies** page, click the context menu to the right of the quota you want to delete, then select **Delete**.
 - Click the quota you want to delete, then from the quota policy detail page click **Delete** .
2. From the **Confirm Delete** dialog, click **Delete** or **Cancel**.

Available Quotas by Service

Analytics Cloud

For Analytics Cloud quotas and examples, see [Service Quotas](#).

Block Volume Quotas

Family name: `block-storage`

Name	Scope	Description
<code>backup-count</code>	Regional	Total number of block and boot volume backups
<code>total-storage-gb</code>	Availability domain	Maximum storage space of block and boot volumes, in GB
<code>volume-count</code>	Availability domain	Total number of block and boot volumes

Example

```
set block-storage quota volume-count to 10 in compartment MyCompartment
```

Compute Quotas

Compute Shapes and Custom Images

Family name: `compute`

Name	Scope	Description
<code>custom-image-count</code>	Regional	Number of custom images
<code>bm-standard1-36-count</code>	Availability domain	Number of BM.Standard1.36 shapes
<code>bm-dense-io1-36-count</code>	Availability domain	Number of BM.DenseIO1.36 shapes

CHAPTER 4 Service Essentials

Name	Scope	Description
bm-standard-b1-44-count	Availability domain	Number of BM.Standard.B1.44 shapes
bm-standard2-52-count	Availability domain	Number of BM.Standard2.52 shapes
bm-dense-io2-52-count	Availability domain	Number of BM.DenseIO2.52 shapes
bm-gpu2-2-count	Availability domain	Number of BM.GPU2.2 shapes
bm-gpu3-8-count	Availability domain	Number of BM.GPU3.8 shapes
bm-standard-e2-64-count	Availability domain	Number of BM.Standard.E2.64 shapes
bm-hpc2-36-count	Availability domain	Number of BM.HPC2.36 shapes
vm-standard1-1-count	Availability domain	Number of VM.Standard1.1 shapes
vm-standard1-2-count	Availability domain	Number of VM.Standard1.2 shapes
vm-standard1-4-count	Availability domain	Number of VM.Standard1.4 shapes
vm-standard1-8-count	Availability domain	Number of VM.Standard1.8 shapes

CHAPTER 4 Service Essentials

Name	Scope	Description
vm-standard1-16-count	Availability domain	Number of VM.Standard1.16 shapes
vm-dense-io1-4-count	Availability domain	Number of VM.DenseIO1.4 shapes
vm-dense-io1-8-count	Availability domain	Number of VM.DenseIO1.8 shapes
vm-dense-io1-16-count	Availability domain	Number of VM.DenseIO1.16 shapes
vm-standard2-1-count	Availability domain	Number of VM.Standard2.1 shapes
vm-standard2-2-count	Availability domain	Number of VM.Standard2.2 shapes
vm-standard2-4-count	Availability domain	Number of VM.Standard2.4 shapes
vm-standard2-8-count	Availability domain	Number of VM.Standard2.8 shapes
vm-standard2-16-count	Availability domain	Number of VM.Standard2.16 shapes
vm-standard2-24-count	Availability domain	Number of VM.Standard2.24 shapes
vm-standard-e2-1-count	Availability domain	Number of VM.Standard.E2.1 shapes

CHAPTER 4 Service Essentials

Name	Scope	Description
vm-standard-e2-2-count	Availability domain	Number of VM.Standard.E2.2 shapes
vm-standard-e2-4-count	Availability domain	Number of VM.Standard.E2.4 shapes
vm-standard-e2-8-count	Availability domain	Number of VM.Standard.E2.8 shapes
vm-dense-io2-8-count	Availability domain	Number of VM.DenseIO2.8 shapes
vm-dense-io2-16-count	Availability domain	Number of VM.DenseIO2.16 shapes
vm-dense-io2-24-count	Availability domain	Number of VM.DenseIO2.24 shapes
vm-gpu2-1-count	Availability domain	Number of VM.GPU2.1 shapes
vm-gpu3-1-count	Availability domain	Number of VM.GPU3.1 shapes
vm-gpu3-2-count	Availability domain	Number of VM.GPU3.2 shapes
vm-gpu3-4-count	Availability domain	Number of VM.GPU3.4 shapes

Example

```
set compute quota vm-dense-io1-4-count to 10 in compartment MyCompartment where request.ad = 'us-phoenix-1-ad-2'
```

Instance Configurations and Instance Pools

Family name: `compute-management`

Name	Scope	Description
<code>config-count</code>	Regional	Number of instance configurations
<code>pool-count</code>	Regional	Number of instance pools

Example

```
set compute-management quota config-count to 10 in compartment MyCompartment
```

Autoscaling

Family name: `auto-scaling`

Name	Scope	Description
<code>config-count</code>	Regional	Number of autoscaling configurations

Example

```
Set auto-scaling quota config-count to 10 in compartment MyCompartment
```

Data Transfer Quotas

Family name: data-transfer

Name	Scope	Description
active-appliance-count	Regional	Number of approved transfer appliances
appliance-count	Regional	Number of transfer appliances
job-count	Regional	Number of transfer jobs

Example

```
zero data-transfer quota job-count in tenancy
set data-transfer quota job-count to 1 in compartment Finance
set data-transfer quota appliance-count to 3 in compartment Finance
```

Database Quotas

Family name: database

Name	Scope	Description
adw-ocpu-count	Regional	Number of Autonomous Data Warehouse OCPUs
atp-ocpu-count	Regional	Number of Autonomous Transaction Processing OCPUs
adb-free-count	Regional	Number of Always Free Autonomous Databases. Tenancies can have a total of two Always Free Autonomous Databases, and these resources must be provisioned in the home region. For each database, you can choose the workload type (Autonomous Transaction Processing or Autonomous Data Warehouse).

CHAPTER 4 Service Essentials

Name	Scope	Description
bm-dense-io1-36-count	Availability domain	Number of BM.DenseIO1.36 DB systems
bm-dense-io2-52-count	Availability domain	Number of BM.DenseIO2.52 DB systems
exadata-base-48-count	Availability domain	Number of Exadata.Base.48 DB systems
exadata-full1-336-x6-count	Availability domain	Number of Exadata.Full1.336 - X6 DB systems
exadata-full2-368-x7-count	Availability domain	Number of Exadata.Full2.368 - X7 DB systems and Autonomous Exadata Infrastructure
exadata-half1-168-x6-count	Availability domain	Number of Exadata.Half1.168 - X6 DB systems
exadata-half2-184-x7-count	Availability domain	Number of Exadata.Half2.184 - X7 DB systems and Autonomous Exadata Infrastructure
exadata-quarter1-84-x6-count	Availability domain	Number of Exadata.Quarter1.84 - X6 DB systems
exadata-quarter2-92-x7-count	Availability domain	Number of Exadata.Quarter2.92 - X7 DB systems and Autonomous Exadata Infrastructure
vm-block-storage-gb	Availability domain	Total size of block storage attachments across all virtual machine DB systems, in GB

CHAPTER 4 Service Essentials

Name	Scope	Description
vm-standard1-ocpu-count	Availability domain	Number of VM.Standard1.x OCPUs
vm-standard2-ocpu-count	Availability domain	Number of VM.Standard2.x OCPUs

For information about shapes that are not listed, including non-metered shapes, [contact Oracle Support](#).

Example

The following example shows how to limit the number of Autonomous Data Warehouse resources in a compartment:

```
#Limits the Autonomous Data Warehouse CPU core count to 2 in the MyCompartment compartment
set database quota adw-ocpu-count to 2 in compartment MyCompartment
```

To limit the number of virtual machine DB systems in a compartment, you must set a quota for the number of CPU cores and a separate quota for the block storage:

```
#Sets a quota for virtual machine Standard Edition OCPUs to 2 in the MyCompartment compartment
set database quota vm-standard1-ocpu-count to 2 in compartment MyCompartment

#Sets the virtual machine DB system block storage quota to 1024 GB in the same compartment
set database quota vm-block-storage-gb to 1024 in compartment MyCompartment
```

The following example shows how to prevent the usage of all database resources in the tenancy except for two Exadata full rack X7 resources in a specified compartment:

```
zero database quotas in tenancy
set database quota exadata-full12-368-x7-count to 2 in compartment MyCompartment
```

This example of *nested quotas* shows how to distribute limits for a resource type in a compartment among its subcompartments:

CHAPTER 4 Service Essentials

```
#Allows usage of 3 Autonomous Data Warehouse OCPUs in parent compartment Compartment1
set database quota adw-ocpu-count to 3 in compartment Compartment1

#Allows usage of 1 Autonomous Data Warehouse OCPU in child compartment Compartment1.1
set database quota adw-ocpu-count to 1 in compartment Compartment1.1

#Allows usage of 2 Autonomous Data Warehouse OCPUs in child compartment Compartment1.2
set database quota adw-ocpu-count to 2 in compartment Compartment1.2
```

This example shows how to set a quota for Autonomous Exadata Infrastructure quarter rack resources in a compartment:

```
#Limits the usage of Exadata.Quarter2.92 X7 shapes to 1 in the MyCompartment compartment
set database quota exadata-quarter2-92-x7-count to 1 in compartment MyCompartment
```

DNS Quotas

Family name: dns

Name	Scope	Description
global-zone-count	Global	Number of public DNS zones
steering-policy-count	Global	Number of traffic management steering policies
steering-policy-attachment-count	Global	Number of traffic management steering policy attachments

Example

```
zero dns quotas in compartment MyCompartment
zero dns quota global-zone-count in compartment MyCompartment
zero dns quota steering-policy-count in compartment MyCompartment
zero dns quota steering-policy-attachment-count in compartment MyCompartment
```

Email Delivery Quotas

Family name: email-delivery

Name	Scope	Description
approved-sender-count	Regional	Number of approved senders

Example

```
zero email-delivery quota approved-sender-count in compartment MyCompartment
```

Health Checks Quotas

Family name: health-checks

Name	Scope	Description
monitor-basic-count	Regional	Number of basic monitors
monitor-premium-count	Regional	Number of premium monitors

Example

```
zero health-checks quotas monitor-basic-count
```

Key Management Quotas

Family name: kms

Name	Scope	Description
virtual-private-vault-count	Regional	Number of virtual private vaults

Example

```
set kms quota virtual-private-vault-count to 10 in compartment MyCompartment
set kms quota virtual-vault-count to 10 in compartment MyCompartment
```

Notifications Quotas

Family name: notifications

Name	Scope	Description
topic-count	Regional	Number of topics

Example

```
set notifications quota topic-count to 10 in compartment MyCompartment
```

Resource Manager Quotas

Family name: `resource-manager`

Name	Scope	Description
<code>concurrent-job-count</code>	Regional	Number of concurrent Jobs per compartment
<code>stack-count</code>	Regional	Number of of stacks per compartment

Example

```
set resource-manager quota concurrent-job-count to 1 in compartment MyCompartment
zero resource-manager quota stack-count in compartment MyCompartment
```

Streaming Quotas

Family name: `streaming`

Name	Scope	Description
<code>partition-count</code>	Regional	Number of partitions

Example

```
set streaming quota partition-count to 10 in compartment MyCompartment
```

WAF Quotas

Family name: `waas`

Name	Scope	Description
<code>waas-policy-count</code>	Regional	Number of WAF policies

Example

```
zero waas quota waas-policy-count in compartment MyCompartment
```

Work Requests

This topic describes the work requests feature documented in the [Work Requests API](#). The following Oracle Cloud Infrastructure services are integrated with this API:

- [Compute](#)
- [Database](#)



Note

Some Oracle Cloud Infrastructure services offer work requests supported by the service API rather than the Work Requests API discussed in this topic. For information about work requests in these services, see the following topics:

- **Load Balancing:** [Viewing the State of a Work Request](#)
- **Object Storage:** [Copy Object Work Requests](#)
- **Identity:** [Deleting Compartments](#) and [Deleting Tag Key Definitions and Namespace](#)

Work requests allow you to monitor long-running operations such as Database backups or the provisioning of Compute instances. When you launch such an operation, the service spawns a *work request*. A work request is an activity log that enables you to track each step in the operation's progress. Each work request has an OCID that allows you to interact with it programmatically and use it for automation.

If an operation fails, a work request can help you determine which step of the process had an error.

Some operations affect multiple resources. For example, creating an instance pool also affects instances and instance configurations. A work request provides a list of the resources that an operation affects.

For workflows that require sequential operations, you can monitor each operation's work request and confirm that the operation has completed before proceeding to the next operation. For example, say that you want to create an instance pool with autoscaling enabled. To do this, you must first create the instance pool, and then configure autoscaling. You can monitor the work request for creating the instance pool to determine when that workflow is complete, and then configure autoscaling after it is done.

Work requests are retained for 12 hours.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: Work requests inherit the permissions of the operation that spawns the work request. To enable users to view the work requests, logs, and error messages for an operation, write a policy that grants users permission to do the operation. For example, to let users see the work requests associated with launching instances, write a policy that enables users to launch instances.

To enable users to list all work requests in a tenancy, use the following policy:

```
Allow group SupportTeam to inspect work-requests in tenancy
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Work Request States

Note: Work requests for some services or operations may support only a subset of the following statuses.

ACCEPTED

The request is in the work request queue to be processed.

IN_PROGRESS

A work request record exists for the specified request, but there is no associated WORK_COMPLETED record.

SUCCEEDED

A work request record exists for this request and an associated WORK_COMPLETED record has the state SUCCEEDED.

FAILED

A work request record exists for this request and an associated WORK_COMPLETED record has the state FAILED.

CANCELING

The work request is in the process of canceling.

CANCELED

The work request has been canceled.

Using the Console to View Work Requests

The steps to view a work request are similar for Oracle Cloud Infrastructure services that support work requests.

1. Navigate to resource whose work requests you want to see.
For example, to see the work requests for a Compute instance: Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. If the resource is displayed in a list view, click the resource name to view the resource details.
3. In the **Resources** section of the details page, click **Work Requests**. The status of all work requests appears on the page.
4. To see the log messages, error messages, and resources that are associated with a specific work request, click the operation name. Then, select an option in the **More information** section.
For associated resources, you can click the the Actions icon (three dots) next to a resource to copy the resource's OCID.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to monitor the state of work requests:

- [ListWorkRequests](#)
- [GetWorkRequest](#)
- [ListWorkRequestErrors](#)
- [ListWorkRequestLogs](#)

Console Announcements

This topic describes the announcements that Oracle Cloud Infrastructure displays in the Console. Console announcements appear at the top of the page to communicate timely, important information about service status. You can also view a list of past announcements.



Note

- Announcements is not available in Oracle Cloud Infrastructure Government Cloudrealms.
- If you use Oracle Platform Cloud Services or Oracle Cloud Applications and you have announcements about those service entitlements, the Console displays a banner with a link that you can use to access those announcements. For more information about these announcements, including how to set notification preferences, see [Monitoring Notifications](#).

Types of Announcements

There are different categories of announcements. An announcement's prefix helps you understand, at a glance, the type and relative severity of the information and whether there's anything you can or must do. Announcement types currently include the following, in order of most important to least:

- **Required action.** You must take specific action within your environment.
- **Emergency change.** There is a time period during which an unplanned, but urgent, change associated with your environment will take place.
- **Recommended action.** You have specific action to take within your environment, but the action is not required.

- **Planned change.** There is a time period during which a planned change associated with your environment will take place.
- **Planned change extended.** The scheduled change period has extended beyond what was previously communicated.
- **Planned change rescheduled.** The planned change to your environment has been postponed to a later time or date.
- **Production event.** An impactful change to your environment either recently occurred or is actively occurring.
- **Planned change completed.** The planned change to your environment has been completed and regular operations have resumed.
- **Information.** There is information that you might find useful, but is not urgent and does not require action on your part.

For announcements that require action and affect Oracle Cloud Infrastructure Compute instances, you will get 30 days of advance notice. If you need to delay the actions described in the announcement, [contact support](#) to request one of the alternate dates listed in the announcement. Critical vulnerabilities might not be eligible for delay.

Required IAM Policy

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

Depending on whether you have access, you might not see any announcements. With access to announcements, you can either see only the summary version of any given announcement or you can also view announcement details.

For administrators: for typical policies that give users access to announcements, see [Restrict user access to view only summary announcements](#) and [Let users view details of announcements](#). For more information, see [Details for the Announcements Service](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Email Delivery

As part of your service agreement, Oracle Cloud Infrastructure also contacts you about service status announcements through email. These emails help alert you to upcoming changes that will impact your tenancy, such as those involving data centers or instances you use, or about required action on your part. Oracle sends these announcements to the tenant administrator email address and you cannot opt out of receiving this operational information. Whenever possible, we try to provide advance notice of impactful events. If you want to change the tenant administrator email address, contact [Oracle Support](#). For more information, see [Contacting Support](#).

Viewing Announcements

This section describes how to view announcements. The Console displays announcements as banners that span the width of the top of your browser window. As long as an announcement remains in effect and you have the access to view announcements, the banner announcement displays each time you sign in to the Console until you mark it as read. You can also view all past announcements. The Announcements icon displays a green dot if you have any unread announcements.

To dismiss a banner announcement

- To close a banner announcement until the next time you sign in to the Console, click the X at the far right edge of the banner. If you want to stop seeing an announcement as a banner altogether, you must mark it as read. For more information, see [To mark an announcement as read](#).

To view the details of an announcement

1. Do one of the following:
 - If you are viewing a banner, click the **Show details** link near the far right edge of the banner.
 - If you are viewing a list of announcements, under the **Summary** column, click the announcement summary.
2. On the **Announcement Details** page, you can view the following information:
 - **Description.** This describes the issue or event in greater detail than the summary text of the announcement.
 - **OCID.** This is the announcement's unique, Oracle-assigned identifier.
 - **Reference Ticket Number.** You can use this number to refer to the issue when talking to Support.
 - **Type.** This is one of several predefined categories that helps to set expectations about the nature and severity of the issue described.
 - **Affected Service.** This indicates the Oracle Cloud Infrastructure services affected by the issue or event.
 - **Region.** This tells you what Oracle Cloud Infrastructure regions are impacted.
 - **Start Time.** This is when the issue or event was first detected.
 - **End Time.** This is when the issue or event was resolved.
 - **Required After.** This is the date after which you must address any required actions described in the announcement.
 - **Created.** This is when the announcement was created.
 - **Updated.** This is when the announcement was updated.
 - **Additional Information.** This includes information such as workarounds or background material.

- **Impacted Resources.** This shows the resources that were affected in some way by the event that prompted the announcement.
3. Optionally, if you want to refer to the list of impacted resources later, click **Download Impacted Resources List**.

To view a list of all announcements

1. Click the Announcements icon (△).
2. The **Announcements** page displays all announcements. From this page, you can do the following:
 - **Filter.** You can filter announcements by type or by start or end date.
 - **Sort.** You can sort announcements by summary, type, event start time, or publish time (which indicates when the announcement was last updated).
 - **Mark as read.** You can mark announcements as read if you want stop seeing them as banners in the Console in subsequent sessions.
 - **View announcement details.** You can view the details of an announcement.

To filter a list of announcements

1. Click the Announcements icon (△).
2. To filter the list, under **Filters**, do one of the following:
 - Click **Type**, and then click a type from the list.
 - Click **Start Date**, and then choose a date to see only events that started on that date.
 - Click **End Date**, and then choose a date to see only events that ended on that date.
3. To clear a filter on a date, click the X next to the date.

To sort a list of announcements

1. Click the Announcements icon (△).
2. By default, the list displays announcements according to the event start time, from most recent to least. To sort the list another way, do one of the following:
 - Click **Summary**. The list sorts alphabetically, according to the summary of the announcement.
 - Click **Type**. The list sorts according to the importance of the announcement.
 - Click **Start Time**. The list sorts according to the start time of the event described in the announcement. If you begin by viewing the default sort order, the sort order will change to show the oldest announcement at the beginning of the list.
 - Click **Publish Time**. The list sorts according to the time that an announcement was last updated. You might find it helpful to sort by this column if you want to track an ongoing issue or if an announcement requires action on your part.
3. To sort the list again, repeat the previous step.

To mark an announcement as read

1. Click the Announcements icon (△).
2. Find the announcement that you want to mark as read, click the Actions icon (three dots), and then click **Mark As Read**.

Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

To view the details of an announcement

Open a command prompt and run `oci announce announcements get` to view detailed information about an announcement:

CHAPTER 4 Service Essentials

```
oci announce announcements get --announcement-id <announcement_OCID>
```

For example:

```
oci announce announcements get --announcement-id  
ocid1.announcement.region1..examplear73oue4jdywjvietoc6im3cvb6xae4falm3faux5us3iwra3t6q
```

To view a list of all announcements

Open a command prompt and run `oci announce announcements list` to view a list of all announcements:

```
oci announce announcements list --compartment-id <compartment_OCID>
```

For example:

```
oci announce announcements list --compartment-id  
ocid1.tenancy.oc1..exampleati4wjo6cvbxq4iusld5ltpneskcfy7lr4a6wfauxuwrwed5bsdea
```

To filter a list of announcements

Open a command prompt and run `oci announce announcements list` to filter a list of announcements.

To filter a list of announcements by announcement type:

```
oci announce announcements list --compartment-id <compartment_OCID> --announcement-type <announcement_  
type>
```

For example:

```
oci announce announcements list --compartment-id  
ocid1.tenancy.oc1..exampleati4wjo6cvbxq4iusld5ltpneskcfy7lr4a6wfauxuwrwed5bsdea --announcement-type  
ACTION_REQUIRED
```

To sort a list of announcements

Open a command prompt and run `oci announce announcements list` to sort a list of announcements.

CHAPTER 4 Service Essentials

To sort a list of announcements in ascending order of time created, from oldest to newest:

```
oci announce announcements list --compartment-id <compartment_OCID> --sort-order ASC
```

For example:

```
oci announce announcements list --compartment-id  
ocid1.tenancy.oc1..exampleati4wjo6cvbxq4iusld5ltpneskcfy7lr4a6wfauxurwed5bsdea --sort-order ASC
```

To mark an announcement as read

Open a command prompt and run `oci announce user-status update` to mark an announcement as read:

```
oci announce user-status update --announcement-id <announcement_OCID> --user-status-announcement-id  
<announcement_OCID> --user-id <user_OCID> --time-acknowledged <date_and_time>
```

For example:

```
oci announce user-status update --announcement-id  
ocid1.announcement.region1..examplear73oue4jdywjvietoc6im3cvb6xae4falm3faux5us3iwra3t6g --user-status-  
announcement-id ocid1.announcement.region1..examplear73oue4jdywjvietoc6im3cvb6xae4falm3faux5us3iwra3t6g  
--user-id ocid1.user.region1..exampleaorxz3psplonigcvbzy5oaiwiubh7k7ip6zgzklfauxic67kksu4oq --time-  
acknowledged 2019-01-06T20:14:00+00:00
```

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage announcements:

- [GetAnnouncement](#)
- [GetAnnouncementUserStatus](#)
- [ListAnnouncements](#)
- [UpdateAnnouncementUserStatus](#)

Prerequisites for Oracle Platform Services on Oracle Cloud Infrastructure

This topic describes procedures that are required by some Oracle Platform Services before you can launch them on Oracle Cloud Infrastructure. The information in this topic applies only to the following services:

- Oracle Big Data Cloud
- Oracle Database Cloud Service
- Oracle Data Hub Cloud Service
- Oracle Event Hub Cloud Service
- Oracle Java Cloud Service
- Oracle MySQL Cloud Service
- Oracle SOA Cloud Service

For a list of all services supported on Oracle Cloud Infrastructure, see [Information About Supported Platform Services](#).

Accessing Oracle Cloud Infrastructure

Oracle Cloud Infrastructure has a different interface and credential set than your Oracle Platform Services. You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the REST API. Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser (Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

Required Identity and Access Management (IAM) Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

See [Common Policies](#) for more information and examples.

Resources Created in Your Tenancy by Oracle

Oracle creates a compartment in your tenancy for Oracle Platform Services. This compartment is specially configured by Oracle for the Oracle Cloud Infrastructure resources that you create through the Platform Services. You can't choose another compartment for Oracle to use.

Along with this compartment, Oracle creates the IAM policies to allow Oracle Platform Services access to the resources.

The compartment that Oracle creates for Oracle Platform Services is named:

`ManagedCompartmentForPaaS`.

The policies that Oracle creates for Oracle Platform Services are:

- `PSM-root-policy`

This policy is attached to the root compartment of your tenancy.

- `PSM-mgd-comp-policy`

This policy is attached to the `ManagedCompartmentForPaaS` compartment.



Warning

Do not make any changes to these resources. Editing or renaming the policies or the compartment can result in loss of functionality.

Prerequisites for Oracle Platform Services

Before you can create instances of an Oracle Platform Service on Oracle Cloud Infrastructure, you need to have the following resources in your Oracle Cloud Infrastructure tenancy:

- A compartment for your resources
- A virtual cloud network (VCN) with at least one public subnet
- IAM policies to allow Oracle Platform Services to access the VCN
- An Object Storage bucket
- Credentials to use with Object Storage

Some of the Platform Services automatically create some of these resources for you. See details about your service in the following sections.

Setting Up the Prerequisites



Note

To use **Autonomous Data Warehouse Cloud**, you don't need to set up any of the resources listed in this prerequisites section. However, if you optionally choose to use Oracle Cloud Infrastructure Object Storage for data loading, you need to perform these two tasks:

[Create a bucket](#)

[Create an auth token](#)

Following are two scenarios with procedure sets. If you need to set up all the required resources, follow Scenario 1. If you already have a VCN in your Oracle Cloud Infrastructure tenancy that you want to use for Oracle Platform Services, follow Scenario 2.

To follow a tutorial on how to set up the prerequisites for Scenario 1, see [Creating the Infrastructure Resources Required for Oracle Platform Services](#).

Scenario 1: I need to create all the prerequisite resources

Create a compartment



Important

You cannot use the `ManagedCompartmentForPaaS` for your VCN and bucket.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.
2. A list of the existing compartments in your tenancy is displayed.
3. Click **Create Compartment**.
4. Enter the following:
 - **Name:** For example, `PaaSResources`. Restrictions for compartment names are: Maximum 100 characters, including letters, numbers, periods, hyphens, and underscores. The name must be unique across all the compartments in your tenancy
 - **Description:** A friendly description.
5. Click **Create Compartment**.

Set up your virtual cloud network

This procedure creates a VCN with these characteristics:

- A VCN with CIDR 10.0.0.0/16.
- Three public subnets (10.0.0.0/24, 10.0.1.0/24, and 10.0.2.0/24) each using the VCN's [default security list](#), default [route table](#), and default [DHCP options](#).
- An internet gateway, with the required route rule in the default route table.

- Use of the [Internet and VCN Resolver](#) for DNS, so your instances can use their hostnames instead of their private IP addresses to communicate with each other.



Tip

This Quick VCN procedure is useful for getting started and trying out Oracle Platform Services on Oracle Cloud Infrastructure. For production, use the procedure in [VCNs and Subnets](#). That topic explains features such as how to specify the CIDR ranges for your VCN and subnets, and how to secure your network. When you use the advanced procedure, remember that the VCN that you create must have a public subnet for Oracle Platform Services to use.

1. Open the **Region** menu and select the region in which you want to create the Oracle PaaS service instance.
Select a region that's within the default data region of your account. For example, if your default data region is EMEA, then select Germany Central (Frankfurt) or UK South (London).
2. From the **Compartment** list, select the compartment you created.
3. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
4. Click **Create Virtual Cloud Network**.
5. In **Create in Compartment**, leave the default value (the compartment you're currently working in).
6. Enter a friendly name for the cloud network, for example: `PaaSVCN`. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API).

7. Select **Create Virtual Cloud Network Plus Related Resources**.
8. Scroll down to the bottom of the dialog box and click **Create Virtual Cloud Network**.

Permit Oracle Platform Services to access resources

1. In the Console, navigate to the root compartment of your tenancy by clicking your tenancy name in the **Compartment** list.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
3. Click **Create Policy**.
4. Enter the following:
 - **Name:** A unique name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later.
 - **Description:** A friendly description. You can change this later if you want to.
 - **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in format YYYY-MM-DD format. For more information, see [Policy Language Version](#).
 - **Statement:** To allow Oracle Platform Services access to use the network in your compartment, enter the following policy statements. Replace *<compartment_name>* with your compartment name. Click **+** after each statement to add another.

```
Allow service PSM to inspect vcns in compartment <compartment_name>
```

```
Allow service PSM to use subnets in compartment <compartment_name>
```

```
Allow service PSM to use vnics in compartment <compartment_name>
```

```
Allow service PSM to manage security-lists in compartment <compartment_name>
```

For more information about policies, see [Policy Basics](#) and also [Policy Syntax](#).

5. (Optional) If you want to enable the use of an Autonomous Transaction Processing or Oracle Cloud Infrastructure Database instance in your compartment as the infrastructure schema database for your Oracle Java Cloud Service instance, then add the following statements:

```
Allow service PSM to inspect autonomous-database in compartment <compartment_name>
```

```
Allow service PSM to inspect database-family in compartment <compartment_name>
```

6. Click **Create**.

Create a bucket

1. Open the **Region** menu and select the region in which you want to create the Oracle PaaS service instance.
Select a region that's within the default data region of your account. For example, if your default data region is EMEA, then select Germany Central (Frankfurt) or UK South (London).
2. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
3. Choose the compartment you created.
4. Click **Create Bucket**.
5. In the **Create Bucket** dialog, enter a bucket name, for example: `PaasBucket`.
Make a note of the name you enter. You will need it when you create an instance for your Oracle Platform Service later.
6. Click **Create Bucket**.

Set up credentials to use with Object Storage

For Big Data Cloud, set up an API signing key:

Set up an API signing key

Follow the instructions in this topic: [Required Keys and OCIDs](#).

For all other services, create an auth token. Note that your service might refer to this credential as a Swift password. Use the auth token wherever you are asked to provide a Swift password.

Create an auth token

1. View the user's details:
 - If you're creating an auth token for yourself: Open the **Profile** menu () and click **User Settings**.
 - If you're an administrator creating an auth token for another user: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Auth tokens**.
3. Click **Generate Token**.
4. Enter a friendly description for the token and click **Generate Token**.
The new token is displayed.
5. Copy the token immediately, because you can't retrieve it again after closing the dialog box. Also, make sure you have this token available when you create your Oracle Platform Services instance.

Scenario 2: I have an existing VCN in Oracle Cloud Infrastructure that I want to use for my Oracle Platform Services instance

You can use an existing VCN. The VCN must have at least one public subnet. Perform these tasks to complete the prerequisites:

Permit Oracle Platform Services to access resources

1. In the Console, navigate to the root compartment of your tenancy by clicking your tenancy name in the **Compartment** list.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
3. Click **Create Policy**.
4. Enter the following:
 - **Name:** A unique name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later.
 - **Description:** A friendly description. You can change this later if you want to.
 - **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in YYYY-MM-DD format. For more information, see [Policy Language Version](#).
 - **Statement:** To allow Oracle Platform Services access to use the network, enter the following policy. Click **+** after each statement to add another. In each statement, replace *<compartment_name>* with the name of the compartment where your VCN resides.

```
Allow service PSM to inspect vcns in compartment <compartment_name>
```

```
Allow service PSM to use subnets in compartment <compartment_name>
```

```
Allow service PSM to use vnics in compartment <compartment_name>
```

```
Allow service PSM to manage security-lists in compartment <compartment_name>
```

For more information about policies, see [Policy Basics](#) and also [Policy Syntax](#).

5. (Optional) If you want to enable the use of an Autonomous Transaction Processing or Oracle Cloud Infrastructure Database instance in your compartment as the infrastructure schema database for your Oracle Java Cloud Service instance, then add the following statements:

CHAPTER 4 Service Essentials

```
Allow service PSM to inspect autonomous-database in compartment <compartment_name>
```

```
Allow service PSM to inspect database-family in compartment <compartment_name>
```

6. Click **Create**.

Create a bucket

1. Open the **Region** menu and select the region in which you want to create the Oracle PaaS service instance.
Select a region that's within the default data region of your account. For example, if your default data region is EMEA, then select Germany Central (Frankfurt) or UK South (London).
2. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
3. Choose the compartment you want to create the bucket in.
4. Click **Create Bucket**.
5. In the **Create Bucket** dialog, enter a bucket name, for example: `PaasBucket`. Make a note of the name you enter. You will need it when you create an instance for your Oracle Platform Service later.
6. Click **Create Bucket**.

Set up credentials to use with Object Storage

For Big Data Cloud, set up an API signing key:

Set up an API signing key

Follow the instructions in this topic: [Required Keys and OCIDs](#).

For all other services, create an auth token. Note that your service might refer to this credential as a Swift password. Use the auth token wherever you are asked to provide a Swift password.

Create an auth token

1. View the user's details:
 - If you're creating an auth token for yourself: Open the **Profile** menu () and click **User Settings**.
 - If you're an administrator creating an auth token for another user: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Auth Tokens**.
3. Click **Generate Token**.
4. Enter a friendly description for the token and click **Generate Token**.
The new token is displayed.
5. Copy the auth token immediately, because you can't retrieve it again after closing the dialog box. Also, make sure you have this token available when you create your Oracle Platform Services instance.

Information About Supported Platform Services

The following table lists the services supported on Oracle Cloud Infrastructure and links to more information about using those services on Oracle Cloud Infrastructure:

Service	More Information
Autonomous Analytics Cloud	About Oracle Autonomous Analytics Cloud
Autonomous API Platform Cloud Service	Using Oracle Autonomous API Platform Cloud Service

CHAPTER 4 Service Essentials

Service	More Information
Autonomous Data Warehouse Cloud	About Autonomous Data Warehouse Cloud
Integration Cloud	Oracle Integration Cloud
Autonomous Mobile Cloud Enterprise	About Oracle Autonomous Mobile Cloud Enterprise
NoSQL Database Cloud	Oracle NoSQL Database Cloud
Oracle Visual Builder	Administering Oracle Visual Builder
Big Data Cloud	About Big Data Cloud Clusters in Oracle Cloud Infrastructure
Data Hub Cloud Service	About Oracle Data Hub Cloud Service Clusters in Oracle Cloud Infrastructure
Data Integration Platform Cloud	What is Oracle Data Integration Platform Cloud
Database Cloud Service	About Database Deployments in Oracle Cloud Infrastructure
Developer Cloud Service	About Oracle Developer Cloud Service in Oracle Cloud Infrastructure
Event Hub Cloud Service	About Oracle Event Hub Cloud Service - Dedicated Instances in Oracle Cloud Infrastructure
Java Cloud Service	About Java Cloud Service Instances in Oracle Cloud Infrastructure
MySQL Cloud Service	About MySQL Cloud Service Deployments in Oracle Cloud Infrastructure
Oracle SOA Cloud Service	About SOA Cloud Service Instances in Oracle Cloud Infrastructure Classic and Oracle Cloud Infrastructure

Billing and Payment Tools Overview

Oracle Cloud Infrastructure provides various billing and payment tools that make it easy to manage your service costs.

Budgets

Budgets can be used to set thresholds for your Oracle Cloud Infrastructure spending. You can set alerts on your budget to let you know when you might exceed your budget, and you can view all of your budgets and spending from one single place in the Oracle Cloud Infrastructure console.

See [Budgets Overview](#) for more information.

Cost Analysis

Cost Analysis provides easy-to-use visualization tools to help you track and optimize your Oracle Cloud Infrastructure spending. For more information, see [Checking Your Balance and Usage](#).

Usage Reports

A usage report is a comma-separated value (CSV) file that can be used to get a detailed breakdown of resources in Oracle Cloud Infrastructure for audit or invoice reconciliation.

For more information, see [Usage Reports Overview](#).

Invoices

You can view and download invoices for your Oracle Cloud Infrastructure usage. For more information, see [Viewing Your Subscription Invoice](#).

Payment Methods

The Payment Method section of the Oracle Cloud Infrastructure Console allows you to easily manage how you pay for your Oracle Cloud Infrastructure usage.

For more information, see [Changing Your Payment Method](#).

Built 11/26/2019 11:09 AM

Budgets Overview

A budget can be used to set soft limits on your Oracle Cloud Infrastructure spending. You can set alerts on your budget to let you know when you might exceed your budget, and you can view all of your budgets and spending from one single place in the Oracle Cloud Infrastructure console.

HOW BUDGETS WORK

Budgets are set on [cost-tracking tags](#) or on compartments (including the root compartment) to track all spending in that cost-tracking tag or for that compartment and its children.

All budgets alerts are evaluated every 15 minutes. To see the last time a budget was evaluated, open the details for a budget. You will see fields that show the current spend, the forecast and the "Spent in period" field which shows you the time period over which the budget was evaluated. When a budget alert fires, the email recipients configured in the budget alert receive an email.

BUDGET CONCEPTS

The following concepts are essential to working with budgets:

BUDGET

A monthly threshold you define for your Oracle Cloud Infrastructure spending. Budgets are set on cost-tracking tags or compartments and track all spending in the cost-tracking tag or compartment and any child compartments. Note: the budget tracks spending in the specified target compartment, but you need to have permissions to manage budgets in the root compartment of the tenancy to create and use budgets.

ALERT

You can define email alerts that get sent out for your budget. You can send a customized email message body with these alerts. Alerts are evaluated every 15 minutes, and can be

triggered when your actual or your forecasted spending hits either a percentage of your budget or a specified set amount.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

To use budgets, you must be in a group that can use "usage-budgets" in the tenancy (which is the root compartment) or be able to use all resources in the tenancy. All budgets are created in the root compartment, regardless of the compartment they are targeting, so IAM policies that grant budget permissions outside of the root will not be meaningful.

IAM Policy	Description
Allow group accountants to inspect usage-budgets in tenancy	Accountants can inspect budgets including spend.
Allow group accountants to read usage-budgets in tenancy	Accountants can read budgets including spend (same as list).
Allow group accountants to use usage-budgets in tenancy	Accountants can create and edit budgets and alerts rules.
Allow group accountants to manage usage-budgets in tenancy	Accountants can create, edit, and delete budgets and alerts rules.

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Managing Budgets

This topic discusses how to view and manage your budgets.

USING THE CONSOLE

To create a budget

1. Open the navigation menu. Under **Governance and Administration**, go to **Account Management** and click **Budgets**.
2. Click **Create Budget** at the top of the budgets list. The **Create Budget** dialog is displayed.
3. Select either **Compartment** or **Cost-Tracking Tag** to select the type of target for your budget.

4. Select the target for your budget:
 - For budgets targeting a compartment:
 - Select a target compartment for your budget from the **Target Compartment** drop-down list. Note that while the budget tracks spending in the specified target compartment, but you need to have permissions to manage budgets in the root compartment of the tenancy to create and use budgets.
 - For budgets targeting a cost-tracking tag:
 - Select a tag namespace
 - Select a target cost-tracking tag key.
 - Enter a value for the cost-tracking tag.
5. Enter a name for your budget in the **Name** text field. The name can only contain alphanumeric characters, dashes, and the underscore character, and can't begin with a number.
6. Enter a monthly amount for your budget in the **Monthly Budget Amount** field. The minimum allowed value for your monthly budget is 1; the maximum allowed value is 999,999,999,999.
7. You can optionally create an alert for your budget by creating a budget alert rule. In the Budget Alert Rule panel on the Create Budget dialog, configure your alert rule:
 - a. Select a threshold for your alert from the **Threshold Metric** drop-down list. There are two possible values:
 - Actual Spend** will watch the actual amount you spend in your compartment per month;
 - Forecast Spend** will watch your resource usage and alert you when it appears that you'll exceed your budget. The forecast algorithm is linear extrapolation and requires at least 3 days of consumption to trigger
 - b. Select a threshold type from the **Threshold Type** drop-down list. You can select either a percentage of your monthly budget (which must be greater than 0 and no greater than 10,000) or a fixed amount.

- c. The label of the next text field changes depending on what type of threshold you selected. Enter either a **Threshold %** or a **Threshold Amount**.
 - d. In the **Email Recipients** field, enter one or more email addresses to receive the alerts. Multiple addresses can be separated using a comma, semicolon, space, tab, or new line.
 - e. Enter the body of your email alert in the **Email Message** field. The text of the email message cannot exceed 1000 characters. This message will be included with metadata about your budget, including the budget name, the compartment, and the amount of your monthly budget. You can use this message to for things like providing instructions to the recipient that explain how to request a budget increase or reminding users about corporate policies.
8. **Advanced Options (optional):** Click the **Show Advanced Options** link to add **Tags** to your Budget. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
 9. Click the **Create** button to create your budget.

To view or edit a budget

1. Open the navigation menu. Under **Governance and Administration**, go to **Account Management** and click **Budgets**.
2. From the list of budgets, click on the budget you want to edit. The budget detail screen will appear.
3. Click the **Edit** button. The **Edit Budget** dialog will appear.
4. You can edit the name of your budget or the budget amount.
5. When you are finished, click **Save Changes**.

To delete a budget

1. From the list of budgets, select **Delete** from the context menu, or click the **Delete** button at the top of budget detail screen. The **Confirm Delete** dialog will appear.
2. Click the **Confirm** button to delete the budget, or cancel by clicking **Cancel**.

To manage tags for a budget

1. Open the navigation menu. Under **Governance and Administration**, go to **Account Management** and click **Budgets**.
2. From the list of budgets, click on the budget you want to tag. The budget detail screen will appear.
3. Click the **Add tag(s)** button to add a tag.
4. Click the **Tags** tab and then click on the pencil icon next to a tag you want to edit or remove.

USING THE API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operation to manage budgets:

- [ListBudgets](#)
- [GetBudget](#)
- [CreateBudget](#)
- [DeleteBudget](#)
- [UpdateBudget](#)

Managing Budget Alert Rules

You can set email alerts on your budgets. You can set alerts that are based on a percentage of your budget or an absolute amount, and on your actual spending or your forecast spending.

This topic covers how to view and manage your budget alert rules.

REQUIRED IAM POLICY

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

USING THE CONSOLE

To create a budget alert rule

1. Click the budget that you want to create an alert for from the budgets list.
2. In the **Budget Alert Rules** panel at the bottom of the screen, click the **Create Budget Alert Rule** button.
3. Configure your alert rule:
 - a. Select a threshold for your alert from the **Threshold Metric** drop-down list. There are two possible values:
 - Actual Spend** will watch the actual amount you spend in your compartment per month;
 - Forecast Spend** will watch your resource usage and alert you when it appears that you'll exceed your budget. The forecast algorithm is linear extrapolation and requires at least 3 days of consumption to trigger
 - b. Select a threshold type from the **Threshold Type** drop-down list. You can select either a percentage of your monthly budget (which must be greater than 0 and no greater than 10,000) or a fixed amount.
 - c. The label of the next text field changes depending on what type of threshold you selected. Enter either a **Threshold %** or a **Threshold Amount**.

- d. In the **Email Recipients** field, enter one or more email addresses to receive the alerts. Multiple addresses can be separated using a comma, semicolon, space, tab, or new line.
 - e. Enter the body of your email alert in the **Email Message** field. The text of the email message cannot exceed 1000 characters. This message will be included with metadata about your budget, including the budget name, the compartment, and the amount of your monthly budget. You can use this message for things like providing instructions to the recipient that explain how to request a budget increase or reminding users about corporate policies.
4. Click the **Create** button to create your alert.

To view or edit a budget alert rule

1. In the list of budget alert rules, click the menu icon at the right side of the list and select **View/Edit** from the context menu.
2. Edit your alert rule.
3. Confirm your changes by clicking **Save Changes**, or dismiss the dialog without saving by clicking the **Cancel** button.

To delete a budget alert rule

1. In the list of budget alert rules, click the menu icon at the right side of the list and select **Delete** from the context menu.
2. Confirm or cancel the delete operation in the **Confirm Delete** dialog by clicking either the **Confirm** or **Cancel** button.

USING THE API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operation to manage budget alert rules:

- [ListAlertRules](#)
- [GetAlertRule](#)
- [CreateAlertRule](#)
- [DeleteAlertRule](#)
- [UpdateAlertRule](#)

Usage Reports Overview

A usage report is a comma-separated value (CSV) file that can be used to get a detailed breakdown of resources in Oracle Cloud Infrastructure for audit or invoice reconciliation.

HOW USAGE REPORTS WORK

The usage report is automatically generated daily, and is stored in an Oracle-owned Object Storage bucket. It contains one row per each Oracle Cloud Infrastructure resource (such as instance, Object Storage bucket, VNIC) per hour along with consumption information, metadata, and tags. Usage reports generally contain 24 hours of usage data, although occasionally a usage report may contain late-arriving data that is older than 24 hours.

Usage reports are retained for one year.

The file name for each usage report is appended with an automatically incrementing numerical value.

The report may contain corrections. Corrections are added as new rows to the report, with the `lineItem/isCorrection` column set and the `referenceNo` value of the corrected line populated in the `lineItem/backReference` column.

USAGE REPORT SCHEMA

The following table shows the usage report schema.

CHAPTER 4 Service Essentials

Field Name	Description
lineItem/referenceNo	Line identifier. Used for debugging and corrections.
lineItem/TenantId	The identifier (OCID) for the Oracle Cloud Infrastructure tenant.
lineItem/intervalUsageStart	The start time of the usage interval for the resource in UTC.
lineItem/intervalUsageEnd	The end time of the usage interval for the resource in UTC.
product/service	The service that the resource is in.
product/resource	The resource name used by the metering system.
product/compartmentId	The ID of the compartment that contains the resource.
product/compartmentName	The name of the compartment that contains the resource.
product/region	The region that contains the resource.

CHAPTER 4 Service Essentials

Field Name	Description
product/availabilityDomain	The availability domain that contains the resource.
product/resourceId	The identifier for the resource.
usage/consumedQuantity	The quantity of the resource that has been consumed over the usage interval.
usage/billedQuantity	The quantity of the resource that has been billed over the usage interval.
usage/consumedQuantityUnits	The unit for the consumed quantity and billed quantity.
usage/consumedQuantityMeasure	The measure for the consumed quantity and billed quantity.

Field Name	Description
<code>lineItem/backreference</code>	Data amendments and corrections reference. If a correction to an existing line item is needed, a new row is added with the corrected values and a reference to the original line. Used with the <code>lineItem/isCorrection</code> field.
<code>lineItem/isCorrection</code>	Used if the current line is a correction. See the <code>lineitem/backreference</code> column for a reference to the corrected line item.
<code>tags/</code>	The usage report contains one column per tag definition (includes all tag definitions, not just cost tracking tags).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Accessing Usage Reports

A usage report is a comma-separated value (CSV) file that is generated daily and stored in an Object Storage bucket. This topic describes how to access usage reports.

REQUIRED IAM POLICY

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Reports are generated in another tenancy and stored in an Oracle-owned Object Storage bucket. You must set up a cross-tenancy IAM policy to access your usage reports as shown below, changing the group name as appropriate:

```
define tenancy usage-report as
ocid1.tenancy.oc1..aaaaaaaaaned4fcpkispbjlr56u7cj63lf3wffbilvqknstgtvzub7vhqkggq
```

CHAPTER 4 Service Essentials

```
endorse group MyGroupName to read objects in tenancy usage-report
```

USING THE CONSOLE

To download a usage report

1. Open the navigation menu. Under **Governance and Administration**, go to **Account Management** and select **Usage Report**.
2. Click the report you want to download from the list, and follow your browser's instructions for downloading.

USING THE API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To download a usage report, use the Object Storage APIs. The reports are stored in the tenancy's home region. The Object Storage namespace used for the reports is `bling`; the bucket name is the tenancy OCID.

The following example shows how to download a usage report using a Python script:

```
import oci
import os

# This script downloads all of the usage reports for a tenancy (specified in the config file)
#
# Pre-requisites: Create an IAM policy to endorse users in your tenancy to read usage reports from the
OCI tenancy
#
# Example policy:
# define tenancy usage-report as
ocid1.tenancy.oc1..aaaaaaaaaned4fcpkisbwjlr56u7cj631f3wffbilvqknstgtvzub7vhqkggq
# endorse group group_name to read objects in tenancy usage-report
#
# Note - the only value you need to change is group name. Do not change the OCID in the first statement

usage_report_namespace = 'bling'
```

CHAPTER 4 Service Essentials

```
# Update these values
destination_path = 'downloaded_reports'

# Make a directory to receive reports
if not os.path.exists(destination_path):
    os.mkdir(destination_path)

# Get the list of usage reports
config = oci.config.from_file('config/config', 'DEFAULT')
usage_report_bucket = config['tenancy']
object_storage = oci.object_storage.ObjectStorageClient(config)
report_bucket_objects = object_storage.list_objects(usage_report_namespace, usage_report_bucket)

for o in report_bucket_objects.data.objects:
    print('Found file ' + o.name)
    object_details = object_storage.get_object(usage_report_namespace, usage_report_bucket, o.name)
    filename = o.name.rsplit('/', 1)[-1]

    with open(destination_path + '/' + filename, 'wb') as f:
        for chunk in object_details.data.raw.stream(1024 * 1024, decode_content=False):
            f.write(chunk)

print('Finished downloading ' + o.name + '\n')
```

Console Cookies and Local Storage

The Oracle Cloud Infrastructure console uses browser cookies and local storage as detailed below.

Oracle Cloud Infrastructure Console Login Page Cookies

Name	Purpose	Duration	Impact of Disabling
bmc_tenancy	Tracks the default tenancy for the OCI Console login page	30 days	The last-used tenancy is not tracked, and users are asked to specify a tenancy when logging into the OCI Console

Oracle Cloud Infrastructure Console Local Storage

Name	Purpose	Duration	Impact of Disabling
hg-session- <userid> :<session>	UI State information, includes selected region, active compartment and other UI state	Never expires	Console UI may not work optimally
recorded-events	Temporary cache of console usage data. Does not include any sensitive information.	Never expires	Console usage not recorded for analysis and product improvement purposes
recorded-metrics	Temporary cache of console metrics (for example, page load time)	Never expires	Console metrics not recorded for analysis and product improvement purposes

Oracle Cloud Infrastructure Console Cookies

Name	Purpose	Duration	Impact of Disabling
s_fid	Adobe Analytics unique visitor information, anonymous	2 years	Cannot track users across different Oracle products (OCI, Cloud Marketplace)
s_nr	Adobe Analytics unique visitor information, anonymous	2 years	Cannot track users across different Oracle products (OCI, Cloud Marketplace)
gpw_e24	Records and saves the previously accessed URL (anonymous)	Never expires	Console metrics not recorded for analysis and product improvement purposes

Oracle Cloud Infrastructure Console Local Storage - Indexed DB

Table	Field	Purpose
duplo - <tenancy ocid>	activeCompartmentId	Stores the compartment a user has selected in the UI
duplo - <tenancy ocid>	activeRegionId	Stores the region a user has selected in the UI
duplo - <tenancy ocid>	selectedLocale	Stores the user's current locale
opc-key-store	key	Signed ID token (carries user identity information) and a signed security token (carries transient user public key as a JSON Web Key) used for signed request calls to API end points. The security token expires after 24 hours, at which point the user will be prompted to log in again.
opc-key-store-v2	key	Signed ID token (carries user identity information) and a signed security token (carries transient user public key as a JSON Web Key) used for Signed Request calls to API end points. The security token expires after 24 hours, at which point the user will be prompted to log in again.

My Services Use Cases

**Important**

The My Services dashboard and APIs are deprecated.

To interact programmatically with My Services, you can use the [Oracle Cloud My Services API](#). To help you get started, here are some use cases:

- [Service Discovery Use Case](#)
- [Exadata Use Cases](#)
- [Managing Exadata Instances](#)
- [Using Access Token Authorization with My Services API](#)

Service Discovery Use Case

This use case shows how you can get the list of your service entitlement IDs.



Important

The My Services dashboard and APIs are deprecated.

Discover Current Service Entitlement IDs

Many of the My Services API operations require you to specify the `serviceEntitlementId`. To get the list of all your service entitlement IDs, use the [GET ServiceEntitlements](#) operation. This operation returns information that you can use to make more specific requests using the [Oracle Cloud My Services API](#).

Example:

```
GET /itas/<domain>/myservices/api/v1/serviceEntitlements
```



Note

In the examples, *<domain>* is the identity domain ID. An identity domain ID can be either the *IDCS GUID* that identifies the identity domain for the users within Identity Cloud Service (IDCS) or the *Identity Domain name* for a traditional Cloud Account.

To obtain the IDCS GUID

Go to the Users page in My Services dashboard and click **Identity Console**. The URL in the browser address field displays the IDCS GUID for your identity domain. For example:

```
https://idcs-
105bbbdfe5644611bf7ce04496073adf.identity.oraclecloud.com/ui/v
1/adminconsole/?root=users
```

In the above URL, *idcs-105bbbdfe5644611bf7ce04496073adf* is the IDCS GUID for your identity domain.

Example payload returned for this request:

```
{
  "items": [
    {
      "id": "cesi-511202718", // Unique ServiceEntitlementId
      "purchaseEntitlement": { // Purchase Entitlement is the entity
bought by a customer
        "subscriptionId": "511203590",
        "id": "511203590",
        "canonicalLink": "/itas/<domain>/myservices/api/v1/purchaseEntitlements/511203590"
      },
      "serviceDefinition": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceDefinitions/500089778",
```

CHAPTER 4 Service Essentials

```
        "id": "500089778",
        "name": "Storage"                                     // The customer is entitled to use the
Storage Service
    },
    "createdOn": "2017-12-20T16:23:23.326Z",
    "createdBy": "paul.smith@oracle.com",
    "modifiedOn": "2017-12-20T18:35:40.628Z",
    "modifiedBy": "paul.smith@oracle.com",
    "identityDomain": {                                     // Identity Domain to which the Service
Entitlement is associated
        "id": "511203592",
        "name": "myenvironment",
        "displayName": "myenvironment"
    },
    "cloudAccount": {                                     // Cloud Account to which the Service
Entitlement is associated
        "id": "cacct-be7475efc2c54995bc842d3379d35812",
        "name": "myenvironment",
        "canonicalLink": "/itas/<domain>/myservices/api/v1/cloudAccounts/cacct-
be7475efc2c54995bc842d3379d35812"
    },
    "status": "ACTIVE",                                   // Current Status
    "serviceConfigurations": {                             // Specific configuration information such
as Exadata configuration
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-
511202718/serviceConfigurations"
    },
    "canonicalLink": "/itas/{domain}/myservices/api/v1/serviceEntitlements/cesi-511202718"
},
{
    "id": "cesi-511202719",
    "purchaseEntitlement": {
        "subscriptionId": "511203590",
        "id": "511203590",
        "canonicalLink": "/itas/<domain>/myservices/api/v1/purchaseEntitlements/511203590"
    },
    "serviceDefinition": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceDefinitions/500123193",
        "id": "500123193",
        "name": "Compute"                                 // The customer is entitled to use the
Compute Service
    },
}
```

CHAPTER 4 Service Essentials

```
"createdOn": "2017-12-20T16:23:23.326Z",
"createdBy": "paul.smith@oracle.com",
"modifiedOn": "2017-12-20T18:35:40.628Z",
"modifiedBy": "paul.smith@oracle.com",
"identityDomain": {
  "id": "511203592",
  "name": "myenvironment",
  "displayName": "myenvironment"
},
"cloudAccount": {
  "id": "cacct-be7475efc2c54995bc842d3379d35812",
  "name": "myenvironment",
  "canonicalLink": "/itas/<domain>/myservices/api/v1/cloudAccounts/cacct-
be7475efc2c54995bc842d3379d35812"
},
"status": "ACTIVE",
"serviceConfigurations": {
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-
511202719/serviceConfigurations"
},
"canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-511202719"
},
... // More Service Entitlements could be
displayed
},
"canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements",
"hasMore": false,
"limit": 25,
"offset": 0
}
```

Exadata Use Cases



Important

The My Services dashboard and APIs are deprecated.

The following use case examples can get you started working with the Exadata operations available in the [Oracle Cloud My Services API](#).



Important

These procedures are for use with Oracle Database Exadata Cloud at Customer **ONLY**. For more information, see [Administering Oracle Database Exadata Cloud at Customer](#). These procedures **DO NOT** apply to Oracle Database Exadata Cloud Service available in Oracle Cloud Infrastructure.

Exadata Firewall Whitelisting

To enable access to your Exadata Cloud Service instance, you can configure security rules and associate them with your instance. The security rules define a whitelist of allowed network access points.

The firewall provides a system of rules and groups. By default, the firewall denies network access to the Exadata Cloud Service instance. When you enable a security rule, you enable access to the Exadata Cloud Service instance. To enable access you must:

- Create a security group and create security rules that define specific network access allowances.
- Assign the security group to your Exadata Cloud Service instance.

You can define multiple security groups, and each security group can contain multiple security rules. You can associate multiple security groups with each Exadata Cloud Service instance, and each security group can be associated with multiple Exadata Cloud Service instances. You can dynamically enable and disable security rules by modifying the security groups that are associated with each Exadata Cloud Service instance.

To enable access to an Exadata Cloud Service instance:



Note

In the following examples, *<domain>* is the identity domain ID. An identity domain ID can be either the *IDCS GUID* that identifies the identity domain for the users within Identity Cloud Service (IDCS) or the *Identity Domain name* for a traditional Cloud Account.

To obtain the IDCS GUID

Go to the Users page in My Services dashboard and click **Identity Console**. The URL in the browser address field displays the IDCS GUID for your identity domain.

For example:

```
https://idcs-  
105bbbdfe5644611bf7ce04496073adf.identity.oraclecloud.com/ui/v  
1/adminconsole/?root=users
```

In the above URL, *idcs-105bbbdfe5644611bf7ce04496073adf* is the IDCS GUID for your identity domain.

1. Get the service instance IDs.

Operation: [GET ServiceInstances](#)

Example

Example request:

```
GET  
/itas/<domain>/myservices/api/v1/serviceInstances?serviceDefinitionNames=Exadata&statuses=ACTIVE
```

Example payload returned for this request:

```

{
  "items": [
    {
      "id": "csi-585928949",           // Unique ServiceInstanceId
      "serviceEntitlement": {
        "id": "cesi-585927251",
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-585927251"
      },
      "serviceDefinition": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceDefinitions/502579309",
        "id": "502579309",
        "name": "Exadata"           // The customer is entitled to use the Exadata Service
      },
      "cloudAccount": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/cloudAccounts/cacct-fd7a122448aaaa",
        "id": "cacct-fd7a122448aaaa",
        "name": "myAccountName"
      },
      ...
      "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949"
    }
    ...           // More Service Instances could be displayed
  ],
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances",
  "hasMore": false,
  "limit": 25,
  "offset": 0
}

```

This example payload returns the service instance ID csi-585928949, which is part of the service entitlement ID cesi-585927251.

2. Get the service configuration IDs.

Operation: [GET SIServiceConfigurations](#)

Example

Example request, using the service instance ID csi-585928949:

```
GET /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations
```

Example payload returned for this request:

```
{
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations",
  "items": [
    {
      "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata",
      "exadata": {
        "bursting": {
          "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/bursting"
        },
        "id": "Exadata",
        "securityGroupAssignments": {
          "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/securityGroupAssignments"
        }
      },
      "id": "Exadata"
    }
  ]
}
```

This example payload shows that `/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/securityGroupAssignments` is used for Exadata Firewall.

3. Get the current security groups for the service entitlement.
Operation: [GET SEExadataSecurityGroups](#)

Example

Example request, using the service entitlement ID `cesi-585927251`:

```
GET /itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-585927251/serviceConfigurations/Exadata/securityGroups
```

Example payload returned for this request:

```
{
  "items": [
    {
      "id": "1",
      "customerId": "585927251",
      "name": "SecGroup 1",
      "description": "My first Security group",
      "version": 10,
      "rules": [
        {
          "direction": "ingress",
          "proto": "tcp",
          "startPort": 1159,
          "endPort": 1159,
          "ipSubnet": "0.0.0.0/0",
          "ruleInterface": "data"
        }
      ],
      "canonicalLink":
"/itas/
<domain>
/myservices/api/v1/serviceEntitlements/585927251/serviceConfigurations/Exadata/securityGroups/1"
    },
    {
      "id": "2",
      "customerId": "585927251",
      "name": " SecGroup 2",
      "description": "My second Security group",
      "version": 3,
      "rules": [
        {
          "direction": "egress",
          "proto": "tcp",
          "startPort": 8123,
          "endPort": 8123,
```

```
        "ipSubnet": "192.168.1.0/28",
        "ruleInterface": "data"
      }
    ],
    "canonicalLink":
"/itas/
<domain>
/myservices/api/v1/serviceEntitlements/585927251/serviceConfigurations/Exadata/securityGroups/2"
  }
],
"canonicalLink":
"/itas/
<domain>
/myservices/api/v1/serviceEntitlements/585927251/serviceConfigurations/Exadata/securityGroups"
}
```

This example payload shows two security groups defined for the specified service entitlement ID.

4. Get the current security group assignments for the service instance

Operation: [GET SIExadataSecurityGroupAssignments](#)

Example

Example request, using the service instance ID csi-585928949:

```
GET /itas/<domain>/myservices/api/v1/serviceInstances/csi-
585928949/serviceConfigurations/Exadata/securityGroupAssignments
```

Example payload returned for this request:

```
{
  "items": [
    {
      "id": "11",
      "securityGroup":
      {
        "id": "1",
        "canonicalLink":
"/itas/
```

```
<domain>
/myServices/api/v1/serviceEntitlements/585927251/serviceConfigurations/Exadata/securityGroups/1"
  },
  "canonicalLink": "/itas/<domain>/myServices/api/v1/serviceInstances/csi-
585928949/serviceConfigurations/Exadata/securityGroupAssignments/11"
}
],
"canonicalLink": "/itas/<domain>/myServices/api/v1/serviceInstances/csi-
585928949/serviceConfigurations/Exadata/securityGroupAssignments"
}
```

This example payload shows one security group assigned to the service instance csi-585928949.

5. Create a security group with security rules.

Operation: [POST SEExadataSecurityGroups](#)

Example

Example request, using the service entitlement ID csi-585927251:

```
POST /itas/<domain>/myServices/api/v1/serviceEntitlements/csi-
585927251/serviceConfigurations/Exadata/securityGroups
{
  "customerId": "585927251",
  "name": "SecGroup 1",
  "description": "My third Security group",
  "version": 1,
  "rules": [
    {
      "direction": "ingress",
      "proto": "tcp",
      "startPort": 30,
      "endPort": 31,
      "ipSubnet": "100.100.100.255",
      "ruleInterface": "admin"
    },
    {
      "direction": "egress",
```

CHAPTER 4 Service Essentials

```
"proto": "tcp",
"startPort": 32,
"endPort": 32,
"ipSubnet": "100.100.255.0/16",
"ruleInterface": "admin"
}
]
```

Attributes:

Name	Description
customerId	Required: Yes String This must be the same as the <i><serviceEntitlementId></i>
direction	Required: Yes String Allowed values: [ingress egress] for inbound or outbound.
proto	Required: Yes String Allowed values: [tcp udp].
startPort	Required: Yes Integer startPort defines the beginning of a range of ports to open/white-list [0 - 65535].

Name	Description
endPort	Required: Yes Integer endPort defines the ending of a range of ports to open/white-list [0 - 65535].
ipSubnet	Required: Yes String Single IP address or range specified in CIDR notation.
ruleInterface	Required: Yes String Allowed values: [admin client backup] where: <ul style="list-style-type: none">• admin — specifies that the rule applies to network communications over the administration network interface. The administration network is typically used to support administration tasks by using terminal sessions, monitoring agents, and so on.• client — specifies that the rule applies to network communications over the client access network interface, which is typically used by Oracle Net Services connections.• backup — specifies that the rule applies to network communications over the backup network interface, which is typically used to transport backup information to and from network-based storage that is separate from Exadata Cloud Service.

If successful, the POST request will return the unique ID of the newly created security group. For the next step, we'll assume that the newly created security group ID is 3.



Note

A security group can also be modified or deleted. See [Oracle Cloud My Services API](#).

6. Assign the security group to a service instance.

Operation: [POST SIExadataSecurityGroupAssignments](#)

Example

Example request, using the service instance csi-585928949 and the security group ID 3:

```
POST /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/securityGroupAssignments

{
  "securityGroup": {
    "id": "3",
    "customerId": "585927251",
    "canonicalLink":
"/itas/
<domain>
/myservices/api/v1/serviceEntitlements/585927251/serviceConfigurations/Exadata/securityGroups/3"
  }
}
```

Attributes:

Name	Description
customerId	Required: Yes String This must be the same as the serviceEntitlementId.

If successful, the POST request will return the unique Id of the newly created security group assignment.



Note

A security group assignment can also be deleted. See [Oracle Cloud My Services API](#).

You can now verify all your security groups and assignments. See:

- [Get the current security groups for the service entitlement.](#)
- [Get the current security group assignments for the service instance .](#)

Exadata Scaling with Bursting

You can temporarily modify the capacity of your Exadata environment by configuring bursting. Bursting is a method you can use to scale Exadata Cloud Service non-metered instances within an Exadata system.

To scale up your non-metered instances, increase the number of compute nodes by modifying the `burstOcpu` attribute of the host. When you no longer need the additional nodes, update the `burstOcpu` attribute back to its original setting.



Note

In the following examples, *<domain>* is the identity domain ID. An identity domain ID can be either the *IDCS GUID* that identifies the identity domain for the users within Identity Cloud Service (IDCS) or the *Identity Domain name* for a traditional Cloud Account.

To obtain the IDCS GUID

Go to the Users page in My Services dashboard and click **Identity Console**. The URL in the browser address field displays the IDCS GUID for your identity domain.

For example:

```
https://idcs-  
105bbbdfe5644611bf7ce04496073adf.identity.oraclecloud.com/ui/v  
1/adminconsole/?root=users
```

In the above URL, *idcs-105bbbdfe5644611bf7ce04496073adf* is the IDCS GUID for your identity domain.

1. Get the service instance IDs.

Operation: [GET ServiceInstances](#)

Example

Example request:

```
GET  
/itas/<domain>/myservices/api/v1/serviceInstances?serviceDefinitionNames=Exadata&statuses=ACTIVE
```

Example payload returned for this request:

```

{
  "items": [
    {
      "id": "csi-585928949",           // Unique ServiceInstanceId
      "serviceEntitlement": {
        "id": "cesi-585927251",
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements/cesi-585927251"
      },
      "serviceDefinition": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceDefinitions/502579309",
        "id": "502579309",
        "name": "Exadata"           // The customer is entitled to use the Exadata Service
      },
      "cloudAccount": {
        "canonicalLink": "/itas/<domain>/myservices/api/v1/cloudAccounts/cacct-fd7a122448aaaa",
        "id": "cacct-fd7a122448aaaa",
        "name": "myAccountName"
      },
      ...
      "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949"
    }
    ...           // More Service Instances could be displayed
  ],
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances",
  "hasMore": false,
  "limit": 25,
  "offset": 0
}

```

This example payload returns the service instance ID csi-585928949.

2. Get the service configuration IDs.

Operation: [GET SIServiceConfigurations](#)

Example

Example request, using the service instance ID csi-585928949:

```
GET /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations
```

Example payload returned for this request:

```
{
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations",
  "items": [
    {
      "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata",
      "exadata": {
        "bursting": {
          "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/bursting"
        },
        "id": "Exadata",
        "securityGroupAssignments": {
          "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/securityGroupAssignments"
        }
      },
      "id": "Exadata"
    }
  ]
}
```

This example payload shows that `/itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/securityGroupAssignments` is used for Bursting.

3. Get the current compute node configuration.

Operation: [GET SIExadataBursting](#)

Example

Example request, using the service instance ID `csi-585928949`:

```
GET /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/bursting
```

Example payload returned for this request:

```
{
  "ocpuOpInProgress": false,
  "exaunitId": 50,
  "ocpuAllocations": [
    {
      "hostName": "host1.oraclecloud.com",
      "subscriptionOcpu": 11,
      "meteredOcpu": 0,
      "burstOcpu": 0, // Current Burst value
      "minOcpu": 11,
      "maxOcpu": 42,
      "maxBurstOcpu": 11,
      "maxSubOcpu": 38,
      "maxMetOcpu": 0
    },
    {
      "hostName": "host2.oraclecloud.com",
      "subscriptionOcpu": 11,
      "meteredOcpu": 0,
      "burstOcpu": 0, // Current Burst value
      "minOcpu": 11,
      "maxOcpu": 42,
      "maxBurstOcpu": 11,
      "maxSubOcpu": 38,
      "maxMetOcpu": 0
    }
  ],
  "status": 200,
  "op": "exaunit_coreinfo",
  "additionalNumOfCores": "0",
  "additionalNumOfCoresHourly": "0",
  "coreBursting": "Y"
}
```

4. Modify the values for `burstOcpu`.

Operation: [PUT SIExadataBursting](#)

You can modify `burstOcpu` to a value that is up to the value of `maxBurstOcpu`. This example adds two compute nodes to each host.

Example

Example request, using the service instance csi-585928949:

```
PUT /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/bursting/
{
  "ocpuOpInProgress": false,
  "exaunitId": 50,
  "ocpuAllocations": [
    {
      "hostName": "host1.oraclecloud.com",
      "subscriptionOcpu": 11,
      "meteredOcpu": 0,
      "burstOcpu": 2,
      "minOcpu": 11,
      "maxOcpu": 42,
      "maxBurstOcpu": 11,
      "maxSubOcpu": 38,
      "maxMetOcpu": 0
    },
    {
      "hostName": "host2.oraclecloud.com",
      "subscriptionOcpu": 11,
      "meteredOcpu": 0,
      "burstOcpu": 2,
      "minOcpu": 11,
      "maxOcpu": 42,
      "maxBurstOcpu": 11,
      "maxSubOcpu": 38,
      "maxMetOcpu": 0
    }
  ]
}
```

Attributes:

Name	Description
burstOcpu	Required: Yes Type: Integer, Minimum Value: 0, Maximum Value: maxBurstOcpu Number of additional cores



Note

This action may take a few minutes to complete.

5. Verify the new compute node configuration.

Operation: [GET SIExadataBursting](#)

Example

Example request, using the service instance ID csi-585928949:

```
GET /itas/<domain>/myservices/api/v1/serviceInstances/csi-585928949/serviceConfigurations/Exadata/bursting
```

Example payload returned for this request:

```
{
  "ocpuOpInProgress": false,
  "exaunitId": 50,
  "ocpuAllocations": [
    {
      "hostName": "host1.oraclecloud.com",
      "subscriptionOcpu": 11,
      "meteredOcpu": 0,
      "burstOcpu": 2, // New Burst value
      "minOcpu": 11,
      "maxOcpu": 42,
      "maxBurstOcpu": 11,
      "maxSubOcpu": 38,
      "maxMetOcpu": 0
    }
  ]
}
```

```
    },  
    {  
      "hostName": "host2.oraclecloud.com",  
      "subscriptionOcpu": 11,  
      "meteredOcpu": 0,  
      "burstOcpu": 2,                // New Burst value  
      "minOcpu": 11,  
      "maxOcpu": 42,  
      "maxBurstOcpu": 11,  
      "maxSubOcpu": 38,  
      "maxMetOcpu": 0  
    }  
  ],  
  "status": 200,  
  "op": "exaunit_coreinfo",  
  "additionalNumOfCores": "0",  
  "additionalNumOfCoresHourly": "0",  
  "coreBursting": "Y"  
}
```

Managing Exadata Instances



Important

The My Services dashboard and APIs are deprecated.

The following procedures walk you through creating, modifying, and deleting Exadata instances used with the [Oracle Cloud My Services API](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.



Important

These procedures are for use with Oracle Database Exadata Cloud at Customer **ONLY**. For more information, see [Administering Oracle Database Exadata Cloud at Customer](#). These procedures **DO NOT** apply to Oracle Database Exadata Cloud Service available in Oracle Cloud Infrastructure.

Prerequisites

Before you can manage Exadata instances, you need to:

- Subscribe to an Oracle Cloud service
- Obtain account credentials with required roles assigned
- Determine your API endpoint

To subscribe to an Oracle Cloud service

To access [Oracle Cloud My Services API](#), you must request a trial or paid subscription to an Oracle Cloud service.

To obtain account credentials and role assignments

Ask your account administrator for the following items to access [Oracle Cloud My Services API](#):

- Account credentials:
 - User name and password
 - Identity domain ID

CHAPTER 4 Service Essentials

An identity domain ID can be either the *IDCS GUID* that identifies the identity domain for the users within Identity Cloud Service (IDCS) or the *Identity Domain name* for a traditional Cloud Account.

- Required roles assigned to above user name

To determine your API endpoint

Insert the identity domain ID provided by the account administrator (*<domain>*) between `/itas/` and `/myservices/`.

Example:

```
https://itracore.oraclecloud.com/itas/<domain>/myservices/api/v1/serviceEntitlements
```

Creating Exadata Instances

This section covers how to create a basic Exadata instance, an instance with custom IP network configuration, and an instance with multi-VM support.

To create a basic Exadata instance

Post a request with the required payload to create a new instance for a given service entitlement (Exadata in our case).

In the following example, *<domain>* is the identity domain ID.

```
POST /itas/<domain>/myservices/api/v1/operations
{
  "operationItems": [
    {
      "attributes": [
        {
          "name": "requestPayload.name",
          "value": "newinstanceName"
        },
        {
          "name": "requestPayload.serviceEntitlementId",
```

```
    "value": "500073421"
  },
  {
    "name": "requestPayload.size",
    "value": "CUSTOM"
  },
  {
    "name": "requestPayload.serviceType",
    "value": "Exadata"
  },
  {
    "name": "requestPayload.adminUserName",
    "value": "john.smith@example.com"
  },
  {
    "name": "requestPayload.adminEmail",
    "value": "john.smith@example.com"
  },
  {
    "name": "requestPayload.adminFirstName",
    "value": "John"
  },
  {
    "name": "requestPayload.adminLastName",
    "value": "Smith"
  },
  {
    "name": "requestPayload.invokerAdminUserName",
    "value": "john.smith@example.com"
  },
  {
    "name": "requestPayload.invokerAdminEmail",
    "value": "john.smith@example.com"
  },
  {
    "name": "requestPayload.invokerAdminFirstName",
    "value": "John"
  },
  {
    "name": "requestPayload.invokerAdminLastName",
    "value": "Smith"
  }
}
```

CHAPTER 4 Service Essentials

```
    },
    {
      "name": "requestPayload.customAttributes.ExaUnitName",
      "value": "systemname"
    },
    {
      "name": "requestPayload.customAttributes.CreateSparse",
      "value": "N"
    },
    {
      "name": "requestPayload.customAttributes.BackupToDisk",
      "value": "N"
    },
    {
      "name": "requestPayload.customAttributes.isBYOL",
      "value": "N"
    },
    {
      "name": "requestPayload.customAttributes.PickRackSize",
      "value": "Quarter Rack"
    },
    {
      "name": "requestPayload.customAttributes.SELECTED_DC_ID",
      "value": "US001"
    }
  ],
  "operationItemDefinition": {
    "id": "CIM-Exadata-CUSTOM-PRODUCTION-CREATE"
  }
}
]
```

ATTRIBUTES

Name	Description
requestPayload.name	<p>Required: Yes</p> <p>Type: String</p> <p>Name of the Exadata instance. This name:</p> <ul style="list-style-type: none"> • Must not exceed 25 characters. • Must start with a letter. • Must contain only lower case letters and numbers. • Must not contain spaces or any other special characters. • Must be unique within the identity domain.
requestPayload.serviceEntitlementId	<p>Required: Yes</p> <p>Type: String</p> <p>Service Entitlement for the Exadata instance. See "Exadata Service Entitlement discovery". Note that any "cesi-" or "sub-" prefix should not be included.</p>
requestPayload.customAttributes.ExaUnitName	<p>Required: Yes</p> <p>Type: String</p> <p>A name for your Exadata Database Machine environment. This name is also used as the cluster name for the Oracle Grid Infrastructure installation.</p>

Name	Description
<p>requestPayload. customAttributes. CreateSparse</p>	<p>Required: Yes Type: String</p> <p>"Y" to create a disk group that is based on sparse grid disks, else "N".</p> <p>You must select this option to enable Exadata Cloud Service snapshots. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily.</p>
<p>requestPayload. customAttributes. BackupToDisk</p>	<p>Required: Yes Type: String</p> <p>"Y" to use "Database backups on Exadata Storage", else "N".</p> <p>This option configures the Exadata storage to enable local database backups on Exadata storage.</p>
<p>requestPayload. customAttributes. isBYOL</p>	<p>Required: Yes Type: String</p> <p>"Y" to indicate that the Exadata Cloud Service instance uses Oracle Database licenses that are provided by you rather than licenses that are provided are part of the service subscription, else "N".</p> <p>This option only affects the billing that is associated with the service instance. It has no effect on the technical configuration of the Exadata Cloud Service instance.</p>

Name	Description
requestPayload. customAttributes. PickRackSize	Required: Yes Type: String Specify the rack configuration for your service instance. Exact allowed values depend on your purchase. Typical values are like "Full Rack", "Half Rack", "Quarter Rack" or "Eighth Rack".
requestPayload. customAttributes. SELECTED_DC_ID	Required: Yes Type: String Data center that will host your Exadata Cloud Service instance. See "Exadata Service Entitlement discovery" to obtain the Eligible Data Center IDs.

To create an Exadata instance with custom IP network configuration

Post a request with the attributes ClientNetwork and BackupNetwork as part of the payload. The following example includes these optional attributes as well as required attributes.

In the following example, *<domain>* is the identity domain ID.

```
POST /itas/<domain>/myservices/api/v1/operations
{
  "operationItems": [
    {
      "attributes": [
        {
          "name": "requestPayload.name",
          "value": "newinstanceName"
        },
        {
          "name": "requestPayload.serviceEntitlementId",
          "value": "500073421"
        }
      ]
    }
  ]
}
```

```
{
  "name": "requestPayload.size",
  "value": "CUSTOM"
},
{
  "name": "requestPayload.serviceType",
  "value": "Exadata"
},
{
  "name": "requestPayload.adminUserName",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.adminEmail",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.adminFirstName",
  "value": "John"
},
{
  "name": "requestPayload.adminLastName",
  "value": "Smith"
},
{
  "name": "requestPayload.invokerAdminUserName",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.invokerAdminEmail",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.invokerAdminFirstName",
  "value": "John"
},
{
  "name": "requestPayload.invokerAdminLastName",
  "value": "Smith"
},
{
```

CHAPTER 4 Service Essentials

```
    "name": "requestPayload.customAttributes.ExaUnitName",
    "value": "systemname"
  },
  {
    "name": "requestPayload.customAttributes.CreateSparse",
    "value": "N"
  },
  {
    "name": "requestPayload.customAttributes.BackupToDisk",
    "value": "N"
  },
  {
    "name": "requestPayload.customAttributes.isBYOL",
    "value": "N"
  },
  {
    "name": "requestPayload.customAttributes.PickRackSize",
    "value": "Quarter Rack"
  },
  {
    "name": "requestPayload.customAttributes.SELECTED_DC_ID",
    "value": "US001"
  }
  {
    "name": "requestPayload.customAttributes.ClientNetwork",
    "value": "/root/root/1/ipnetwork1"
  },
  {
    "name": "requestPayload.customAttributes.BackupNetwork",
    "value": "/root/root/1/ipnetwork2"
  }
],
"operationItemDefinition": {
  "id": "CIM-Exadata-CUSTOM-PRODUCTION-CREATE"
}
}
]
```

ATTRIBUTES

Name	Description
requestPayload.	Required: Yes
customAttributes.	Type: Url
ClientNetwork	IP network definitions for the network that is primarily used for client access to the database servers. Applications typically access databases on Exadata Cloud Service through this network using Oracle Net Services in conjunction with Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) interfaces.
requestPayload.	Required: Yes
customAttributes.	Type: Url
BackupNetwork	IP network definitions for the network that is typically used to access the database servers for various purposes, including backups and bulk data transfers.

To create an Exadata instance with multi-VM support

If your Exadata system environment is enabled to support multiple virtual machine (VM) clusters, then you can define up to eight clusters and specify how the overall Exadata system resources are allocated to them.

In a configuration with multiple VM clusters, each VM cluster is allocated a dedicated portion of the overall Exadata system resources, with no over-provisioning or resource sharing. On the compute nodes, a separate VM is defined for each VM cluster, and each VM is allocated a dedicated portion of the available compute node CPU, memory, and local disk resources. Each VM cluster is also allocated a dedicated portion of the overall Exadata storage.

CHAPTER 4 Service Essentials

Post a request with the attributes EXAUNIT_ALLOCATIONS and MULTIVM_ENABLED as part of the payload. The following example includes these optional attributes as well as required attributes.

In the following example, *<domain>* is the identity domain ID and *<base64_encoded_string>* is a base64 encoding of the payload following the example.

Example payload for request:

```
POST /itas/<domain>/myservices/api/v1/operations
{
  "operationItems": [
    {
      "attributes": [
        {
          "name": "requestPayload.name",
          "value": "newinstanceName"
        },
        {
          "name": "requestPayload.serviceEntitlementId",
          "value": "500073421"
        },
        {
          "name": "requestPayload.size",
          "value": "CUSTOM"
        },
        {
          "name": "requestPayload.serviceType",
          "value": "Exadata"
        },
        {
          "name": "requestPayload.adminUserName",
          "value": "john.smith@example.com"
        },
        {
          "name": "requestPayload.adminEmail",
          "value": "john.smith@example.com"
        },
        {
          "name": "requestPayload.adminFirstName",
          "value": "John"
        }
      ]
    }
  ]
}
```

```
},
{
  "name": "requestPayload.adminLastName",
  "value": "Smith"
},
{
  "name": "requestPayload.invokerAdminUserName",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.invokerAdminEmail",
  "value": "john.smith@example.com"
},
{
  "name": "requestPayload.invokerAdminFirstName",
  "value": "John"
},
{
  "name": "requestPayload.invokerAdminLastName",
  "value": "Smith"
},
{
  "name": "requestPayload.customAttributes.ExaUnitName",
  "value": "systemname"
},
{
  "name": "requestPayload.customAttributes.CreateSparse",
  "value": "N"
},
{
  "name": "requestPayload.customAttributes.BackupToDisk",
  "value": "N"
},
{
  "name": "requestPayload.customAttributes.isBYOL",
  "value": "N"
},
{
  "name": "requestPayload.customAttributes.PickRackSize",
  "value": "Quarter Rack"
},
}
```

CHAPTER 4 Service Essentials

```
{
  "name": "requestPayload.customAttributes.SELECTED_DC_ID",
  "value": "US001"
}
{
  "name": "requestPayload.customAttributes.EXAUNIT_ALLOCATIONS",
  "value": "<base64_encoded_string>"
},
{
  "name": "requestPayload.customAttributes.MULTIVM_ENABLED",
  "value": "true"
}
],
"operationItemDefinition": {
  "id": "CIM-Exadata-CUSTOM-PRODUCTION-CREATE"
}
}
]
```

Payload for *<base64_encoded_string>*:

```
{
  ExaunitProperties: [
    {name:requestId, value:27ac0ee3-0c72-4493-b02b-40038f07d2a0},
    {name:Operation, value:AddCluster},
    {name:TotalNumOfCoresForCluster, value:4},
    {name:TotalMemoryInGb, value:30},
    {name:StorageInTb, value:3},
    {name:OracleHomeDiskSizeInGb, value:60},
    {name:ClientNetwork, value:/root/root/1/ipnetwork1}, // Only if Higgs is also required
    {name:BackupNetwork, value:/root/root/1/ipnetwork2}, // Only if Higgs is also required
    {name:ExaUnitName, value:systemname},
    {name:CreateSparse, value:N},
    {name:BackupToDisk, value:N}
  ]
}
```

ATTRIBUTES

Name	Description
requestId	Required: Optional Type: String Unique UUID
TotalNumOfCores ForCluster	Required: Yes Type: String The number of CPU cores that are allocated to the VM cluster. This is the total number of CPU cores that are allocated evenly across all of the compute nodes in the VM cluster. Must be a multiple of numComputes as returned by a call to ecra/endpoint/clustershapes.
TotalMemoryInGb	Required: Yes Type: String The amount of memory (in GB) that is allocated to the VM cluster. This is the total amount of memory that is allocated evenly across all of the compute nodes in the VM cluster. Must be a multiple of numComputes as returned by a call to ecra/endpoint/clustershapes.

Name	Description
StorageInTb	<p>Required: Yes</p> <p>Type: String</p> <p>The total amount of Exadata storage (in TB) that is allocated to the VM cluster. This storage is allocated evenly from all of the Exadata Storage Servers.</p>
OracleHomeDiskSize InGb	<p>Required: Yes</p> <p>Type: String</p> <p>The amount of local disk storage (in GB) that is allocated to each database server in the first VM cluster.</p>

Modifying Exadata Instances

This section covers how to add a cluster to an existing instance, reshape a cluster, and delete a cluster.

To add a cluster to an existing instance

Post a request with the operationItemDefinition of CIM-Exadata-CUSTOM-PRODUCTION-UPDATE and a base64 encoding of a payload that includes the Operation value of AddCluster.

In the following example, *<domain>* is the identity domain ID, *<instanceId>* and *<serviceEntitlementId>* are returned from iTAS serviceInstances, and *<base64_encoded_string>* is a base64 encoding of the payload following the example.

Example payload for request:

```
POST /itas/<domain>/myservices/api/v1/operations HTTP/1.1
{
  "operationItems": [
    {
      "attributes": [
```

```
{
  "name": "instanceId",
  "value": "<instanceId>"
},
{
  "name": "requestPayload.serviceEntitlementId",
  "value": "<serviceEntitlementId>"
},
{
  "name": "requestPayload.size",
  "value": "CUSTOM"
},
{
  "name": "requestPayload.serviceType",
  "value": "Exadata"
},
{
  "name": "requestPayload.customAttributes.EXAUNIT_ALLOCATIONS",
  "value": "<base64_encoded_string>"
},
{
  "name": "requestPayload.customAttributes.MULTIVM_ENABLED",
  "value": "true"
}
],
"operationItemDefinition": {
  "id": "CIM-Exadata-CUSTOM-PRODUCTION-UPDATE"
}
}
]
```

Payload for *<base64_encoded_string>*:

```
{
  ExaunitProperties: [
    {name:requestId, value:27ac0ee3-0c72-4493-b02b-40038f07d2a0},
    {name:Operation, value:AddCluster},
    {name:TotalNumOfCoresForCluster, value:4},
    {name:TotalMemoryInGb, value:30},
    {name:StorageInTb, value:3},
    {name:OracleHomeDiskSizeInGb, value:60},
```

CHAPTER 4 Service Essentials

```
{name:ClientNetwork, value:/root/root/1/ipnetwork1}, // Only if Higgs is also required
{name:BackupNetwork, value:/root/root/1/ipnetwork2}, // Only if Higgs is also required
{name:ExaUnitName, value:Cluster2},
{name:CreateSparse, value:N},
{name:BackupToDisk, value:N}
]
}
```

To reshape a cluster

Post a request with the operationItemDefinition of CIM-Exadata-CUSTOM-PRODUCTION-UPDATE and a base64 encoding of a payload that includes the Operation value of ReshapeCluster.

In the following example, *<domain>* is the identity domain ID and *<base64_encoded_string>* is a base64 encoding of the payload following the example.

Example payload for request:

```
POST /itas/<domain>/myservices/api/v1/operations HTTP/1.1
{
  "operationItems": [
    {
      "attributes": [
        {
          "name": "instanceId",
          "value": "500076173"
        },
        {
          "name": "requestPayload.serviceEntitlementId",
          "value": "500073421"
        },
        {
          "name": "requestPayload.size",
          "value": "CUSTOM"
        },
        {
          "name": "requestPayload.serviceType",
          "value": "Exadata"
        }
      ]
    }
  ]
}
```

CHAPTER 4 Service Essentials

```
{
  {
    "name": "requestPayload.customAttributes.EXAUNIT_ALLOCATIONS",
    "value": "<base64_encoded_string>"
  },
  {
    "name": "requestPayload.customAttributes. MULTIVM_ENABLED",
    "value": "true"
  }
],
"operationItemDefinition": {
  "id": "CIM-Exadata-CUSTOM-PRODUCTION-UPDATE"
}
}
]
```

Payload for *<base64_encoded_string>*:

```
{
  ExaunitProperties: [
    {name:requestId, value:27ac0ee3-0c72-4493-b02b-40038f07d2a0},
    {name:ExaunitID, value:1, // From ecra/endpoint/exaservice/{serviceInstance}/resourceinfo
    {name:Operation, value:ReshapeCluster},
    {name:TotalNumOfCoresForCluster, value:10},
    {name:TotalMemoryInGb, value:10},
    {name:StorageInTb, value:4},
    {name:OhomePartitionInGB, value:100},
    {name:ClientNetwork, value:/root/root/1/ipnetwork1}, // Only if Higgs is also required
    {name:BackupNetwork, value:/root/root/1/ipnetwork2} // Only if Higgs is also required
  ]
}
```



Important

- Only one attribute can be modified per Reshape request. The payload should contain only the modified attribute. Example:

```
{ExaunitProperties:
  [{name:Operation,value:ReshapeCluster},
  {name:ExaunitID,value:5},
  {name:TotalNumOfCoresForCluster,value:6}]}
```

- When doing a Reshape with the OracleHomeDiskSizeInGb attribute, use the name OhomePartitionInGB.
- The value for TotalNumOfCoresForCluster must be a multiple of numComputes as returned by a call to ecra/endpoint/clustershapes.
- The value for TotalMemoryInGb must be a multiple of numComputes as returned by a call to ecra/endpoint/clustershapes.

To delete a cluster

Post a request with the operationItemDefinition of CIM-Exadata-CUSTOM-PRODUCTION-UPDATE and a base64 encoding of a payload that includes the Operation value of DeleteCluster.

In the following example, *<domain>* is the identity domain ID and *<base64_encoded_string>* is a base64 encoding of the payload following the example.

Example payload for request:

```
POST /itas/<domain>/myservices/api/v1/operations HTTP/1.1
{
```

```
"operationItems": [
  {
    "attributes": [
      {
        "name": "instanceId",
        "value": "500076173"
      },
      {
        "name": "requestPayload.serviceEntitlementId",
        "value": "500073421"
      },
      {
        "name": "requestPayload.size",
        "value": "CUSTOM"
      },
      {
        "name": "requestPayload.serviceType",
        "value": "Exadata"
      },
      {
        "name": "requestPayload.customAttributes.EXAUNIT_ALLOCATIONS",
        "value": "<base64_encoded_string>"
      },
      {
        "name": "requestPayload.customAttributes.MULTIVM_ENABLED",
        "value": "true"
      }
    ],
    "operationItemDefinition": {
      "id": "CIM-Exadata-CUSTOM-PRODUCTION-UPDATE"
    }
  }
]
```

Payload for *<base64_encoded_string>*:

```
{
  ExaunitProperties: [
    {name:requestId, value:27ac0ee3-0c72-4493-b02b-40038f07d202}, // Optional
    {name:ExaunitID, value:2},
    {name:Operation, value>DeleteCluster}
  ]
}
```

```
]
}
```

Deleting Exadata Instances

This section covers how to delete Exadata instances.



Important

Delete all existing multi-VM clusters before deleting the Exadata instance. Following this guidance prevents the instance ending up in an invalid state.

To delete an instance

Post a request with the operationItemDefinition of CIM-Exadata-CUSTOM-PRODUCTION-DELETE.

In the following example, *<domain>* is the identity domain ID.

Example payload for request:

```
POST /itas/<domain>/myservices/api/v1/operations HTTP/1.1
{
  "operationItems": [
    {
      "attributes": [
        {
          "name": "instanceId",
          "value": "500076173"
        },
        {
          "name": "requestPayload.serviceEntitlementId",
          "value": "500073421"
        },
        {
          "name": "requestPayload.serviceType",
```

```
    "value": "Exadata"
  }
],
"operationItemDefinition": {
  "id": "CIM-Exadata-CUSTOM-PRODUCTION-DELETE"
}
}
]
}
```

Discovering Entitlements and Instances

This section describes how to discover service entitlements and service instances.

To discover service entitlements

Send the following request:

```
GET /itas/<domain>/myservices/api/v1/serviceEntitlements?serviceDefinitionNames=Exadata
```

Example payload returned for this request:

```
{
  "items": [
    {
      "id": "cesi-585927251",           // Unique ServiceEntitlementId
      "serviceDefinition": {
        "canonicalLink": "/itas/a517289/myservices/api/v1/serviceDefinitions/502579309",
        "id": "502579309",
        "name": "Exadata"           // The customer is entitled to use the Exadata Service
      },
      "status": "ACTIVE",
      ...
      "canonicalLink": "/itas/a517289/myservices/api/v1/serviceInstances/cesi-585928949"
    }
    ...                               // More Service Entitlements could be displayed
  ],
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements",
  "hasMore": false,
  "limit": 25,
}
```

CHAPTER 4 Service Essentials

```
"offset": 0
}
```

Eligible Data Centers:

Use:

```
/itas/<domain>/myservices/api/v1/serviceEntitlements/
{ServiceEntitlementId}?expands=serviceInstancesEligibleDataCenters
```

where {ServiceEntitlementId} is a service entitlement ID such as cesi-500074601. This will provide additional information such as:

```
"serviceInstancesEligibleDataCenters": [
  {
    "id": "US001"
  }
],
```

To discover service instances

Send the following request:

```
GET /<domain>/myservices/api/v1/serviceInstances?serviceDefinitionNames=Exadata
```

Example payload returned for this request:

```
{
  "items": [
    {
      "id": "csi-585928949", // Unique ServiceInstanceId
      "serviceEntitlement": {
        "id": "cesi-585927251", // Related ServiceEntitlementId
        "canonicalLink": "/itas/a517289/myservices/api/v1/serviceEntitlements/cesi-585927251"
      },
      "serviceDefinition": {
        "canonicalLink": "/itas/a517289/myservices/api/v1/serviceDefinitions/502579309",
        "id": "502579309",
        "name": "Exadata" // The customer is entitled to use the Exadata Service
      },
      ...
      "canonicalLink": "/itas/a517289/myservices/api/v1/serviceInstances/csi-585928949"
    }
  ]
}
```

```
    }  
    ... // More Service Entitlements could be displayed  
  ],  
  "canonicalLink": "/itas/<domain>/myservices/api/v1/serviceEntitlements",  
  "hasMore": false,  
  "limit": 25,  
  "offset": 0  
}
```

Using Access Token Authorization with My Services API



Important

The My Services dashboard and APIs are deprecated.

This topic explains how to set up and use access token authorization with the Oracle Cloud My Services API. Access token authorization allows a developer to access programmatic endpoints (APIs) to obtain some information (for example, entitlements, instances, or metering data) for your cloud account.

About Access Tokens

An access token contains the information required to allow a developer to access information on your cloud account. A developer presents the token when making API calls. The allowed actions and endpoints depend on the scopes (permissions) that you select when you generate the token. An access token is valid for about an hour.

A refresh token allows the developer to generate a new access token without having to contact an administrator. A refresh token is valid for about one year.

Process Overview

Setup steps for the Administrator:

1. Create an Identity Cloud Service client application with the specific privileges you want to grant to developers.
2. Generate an access token that contains the required privileges for the intended developer.
3. Provide the access token and required information to the developer.
4. Configure Identity Cloud Service for access token validation.

Steps for developer to use the token:

1. Issue requests against My Services API endpoints. Include the access token for the authorization parameter.
2. When the access token expires, refresh the access token without administrator intervention until the privilege is terminated.

Administrator Tasks to Set Up Token Validation

Perform the following tasks to enable developer access with an access token:

Create the IDCS client application

1. Sign in to Identity Cloud Services as an Administrator and go to the administration console. See [How to Access Oracle Identity Cloud Service](#) if you need help signing in.
2. Click the **Applications** tile. A list of the applications is displayed.
3. Click **+ Add** to create a new application.
4. Click **Trusted Application** as the type of application.
5. In the **App Details** section, enter a **Name** and **Description** and then click **Next**.
6. In the **Client** section:

- a. Select **Configure this application as a client now**.
- b. Under **Authorization**, for **Allowed grant types**, select the following options:
 - **JWT Assertion**
 - **Refresh Token**
7. Under **Accessing APIs from Other Applications**, from the **Trust Scope** list, select **Allowed scopes**.
8. Under **Allowed Scopes** click **+ Add**.
9. In the **Add Scope** dialog, click the arrow next to **CloudPortalResourceApp** in the list of App.
10. Select the box next to each authorization that you might want to give the developers to whom you will provide an Access Token. (The permissions are assigned in another step.)
11. Click **Add** to close the dialog. Your selections are displayed.
12. Click **Next**.
13. In the **Resources** section, accept the default and click **Next**.
14. In the **Web Tier Policy** section, accept the default and click **Next**.
15. In the **Authorization** section, click **Finish**.

The **Application Added** notification displays the new Client ID and Client Secret for the application.



Important

Copy and store the Client ID and Client Secret in a safe place and then click **Close**. The Client ID and Client Secret are credentials that are specific to the application that you just created. You will need these credentials later.

16. To complete the creation process, click **Activate** at the top of the page.

Generate an access token

1. Navigate to the IDCS application that you created in the preceding task and select the **Details** tab.
2. Click **Generate Access Token**.
3. On the **Generate Token** dialog, select **Customized Scopes**, then select **Invokes Other APIs**.
4. Select the scopes that you want to give to the developer who will receive this access token.



Note

Oracle recommends that you provide only the minimum required privileges.

5. Select **Include Refresh Token**.
6. Click **Download Token**. Your browser will prompt you to download a token file (.tok). The token file contains an access token and a refresh token.
7. Provide this file to the developer.

Send the access information to a developer

To call API endpoints, the developer needs:

- A token file that you generated.
- The Client ID and Client Secret for the IDCS application used to generate the token file. The Client ID and Secret are required for the developer to generate a new access token from the refresh token.

- The endpoints for the APIs.
 - End points related to the itas:myservices scopes are:
`https://itra.oraclecloud.com/itas/<tenant-IDCS-ID>/myservices/api/v1`
 - End points related to the itas:metering scopes are:
`https://itra.oraclecloud.com/metering/api/v1`

Make sure that you send the above information in a secure way. If you think that this information has been compromised, see [Revoking a Developer's Ability to Refresh Access Tokens](#).

Configure Identity Cloud Service for access token validation

To allow clients to access the tenant signing certificate without logging in to Oracle Identity Cloud Service:

1. Sign in to the Oracle Identity Cloud Services admin console. See [How to Access Oracle Identity Cloud Service](#) if you need help signing in.
2. Open the navigation menu. Under **Settings** select **Default Settings**.
3. Set the **Access Signing Certificate** toggle button to on.

Using the Access Token

The token file has a .tok extension. The file contains the access token and the refresh token. The content looks like:

```
{"app_access_token":"eyJ4N...aabb...CpNwA","refresh_token":"AQID...9NCA="}
```

To use the token with the My Services API:

1. Open the token file.
2. Issue a request to a valid endpoint, inserting the access token for the `Authorization` parameter.
For example:

CHAPTER 4 Service Essentials

```
curl -X GET https://itra.oraclecloud.com/itas/<tenant-IDCS-ID>/myservices/api/v1/serviceEntitlements -H 'Authorization: Bearer eyJ4N...aabb...CpNwA'
```

REQUESTING A NEW ACCESS TOKEN FROM A REFRESH TOKEN

An access token is valid for about one hour. When the token is no longer valid you will get a 401 response code and an Error Message ("errorMessage") value containing "Expired".

You can generate a new short-lived access token from the refresh token. You'll need the Client ID and Client Secret to generate the new token. You can only generate tokens with the same or lower access (scopes) as your original token.

Example using the curl command:

```
curl -i -H 'Authorization: Basic <base64Encoded clientid:secret>' -H 'Content-Type: application/x-www-form-urlencoded;charset=UTF-8' --request POST https://<tenant-IDCS-ID>/oauth2/v1/token -d 'grant_type=refresh_token&refresh_token=<refresh-token>'
```

Using the sample token file from the previous section, the value for *<refresh-token>* would be AQID...9NCA=.

Sample response:

```
{ "access_token": "eyJraWQiO...2nqA", "token_type": "Bearer", "expires_in": 3600, "refresh_token": "AQIDBAUn...VxxNCB7djF9NCA=" }
```



Note

When a developer generates a new access token and refresh token, the previous refresh token becomes invalid.

Revoking a Developer's Ability to Refresh Access Tokens

If you need to revoke a developer's ability to refresh access tokens, you can either invalidate the existing refresh token by generating a new Client Secret for the token; or, you can temporarily revoke access by deactivating the application.



Important

Taking either of these actions will terminate or suspend the ability of all developers using the current Client Secret or application. When generating tokens for multiple developers, consider creating more than one IDCS application to isolate developers from each other.

To terminate a developer's ability to refresh their access token

1. Sign in to Identity Cloud Services as an Administrator and go to the administration console. See [How to Access Oracle Identity Cloud Service](#) if you need help signing in.
2. Click the **Applications** tile. A list of the applications is displayed.
3. Click the application used to generate the token to view its details.
4. Click **Configuration**.
5. Under **General Information**, next to **Client Secret**, click **Regenerate** to generate a new Client Secret.

To restore the ability for the developer to generate an access token from a refresh token, [generate a new access token](#). Then [provide the token](#) along with the new Client Secret to the developer.

To temporarily suspend a developer's ability to refresh their access token

1. Sign in to Identity Cloud Services as an Administrator and go to the administration console. See [How to Access Oracle Identity Cloud Service](#) if you need help signing in.
2. Click the **Applications** tile. A list of the applications is displayed.
3. Click the application used to generate the token to view its details.
4. In the upper right corner of the page, click **Deactivate**.
5. At the prompt, click **Deactivate Application**.

To re-enable developers to use the same tokens, click **Activate**.

CHAPTER 5 Archive Storage

This chapter explains how to upload, manage, and access data using Archive Storage.

Overview of Archive Storage

Oracle Cloud Infrastructure offers two distinct storage class tiers to address the need for both performant, frequently accessed "hot" storage, and less frequently accessed "cold" storage. Storage tiers help you maximize performance where appropriate and minimize costs where possible.

- Use **Archive Storage** for data to which you seldom or rarely access, but that must be retained and preserved for long periods of time. The cost efficiency of the Archive Storage offsets the long lead time required to access the data.
- Use **Object Storage** for data to which you need fast, immediate, and frequent access. Data accessibility and performance justifies a higher price point to store data in the Object Storage. For more information, see [Overview of Object Storage](#).

About Archive Storage

Archive Storage is ideal for storing data that is accessed infrequently and requires long retention periods. Archive Storage is more cost effective than Object Storage for preserving cold data for:

- Compliance and audit mandates
- Retroactively analyzing log data to determine usage pattern or debug problems
- Historical or infrequently accessed content repository data
- Application generated data that requires archival for future analysis or legal purposes

Unlike Object Storage, Archive Storage data retrieval is **not** instantaneous.

Archive Storage is Always Free eligible. For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).

Using Archive Storage



Important

You interact with the data stored in the Archive Storage using the same resources and management interfaces that you use for data stored in Object Storage.

The following summarizes the Object Storage resources you use to store and manage Archive Storage data:

Buckets

Buckets are logical containers for storing objects. A bucket is associated with a single compartment that has policies that determine what actions a user can perform on a bucket and on all the objects in the bucket.

You decide which storage tier (Archive Storage or standard Object Storage) is appropriate for your data when you initially create the bucket container for your data. The storage tier is expressed as a property of the bucket. A bucket's storage tier property sets the initial storage tier of objects added to the bucket. However, objects placed in standard tier buckets can be archived automatically by Object Storage (while remaining in the standard tier bucket) if they meet the criteria of an [object lifecycle policy rule](#) in effect for the bucket.

Once set, you cannot change the storage tier property for a bucket:

- An existing Object Storage bucket cannot be downgraded to an Archive Storage bucket.
- An Archive Storage bucket cannot be upgraded to an Object Storage bucket.

In addition to the inability to change the storage tier designation of a bucket, there are other reasons why storage tier selection for buckets requires careful consideration:

- The minimum retention requirement for Archive Storage is 90 days. If you delete objects from Archive Storage before the minimum retention requirements are met, you

are charged a deletion penalty. The deletion penalty is the prorated cost of storing the data for the full 90 days.

- While Archive Storage is more cost effective than Object Storage for cold storage, understand that when you restore objects, you are returning those objects to Object Storage. You are billed for that storage service class while the objects reside in that tier.

You can use object lifecycle policy rules to automatically delete objects in an Archive Storage bucket based on the age of the object.



Important

You *cannot* use object lifecycle policy rules to automatically restore archived objects to the regular Object Storage tier. See [Restoring and Downloading Objects](#) for information on restoring objects.

See [Managing Buckets](#) for detailed instructions on creating an Archive Storage bucket.

Objects

Any type of data, regardless of content type, is stored as an object. The object is composed of the object itself and metadata about the object. Each object is stored in a bucket.

You upload objects to an Archive Storage bucket the same way you upload objects to a standard Object Storage bucket. The difference is that when you upload an object to an Archive Storage bucket, the object is immediately archived. You must first restore the object before you can download it.

Archived objects are displayed in the object listing of a bucket. You can also display the details of each object.

See [Managing Objects](#) for detailed instructions on uploading objects to an Archive Storage bucket.

Restoring and Downloading Objects

To download an object from Archive Storage, you must first restore the object. Restoration takes about four hours from the time an Archive Storage restore request is made, to the time the first byte of data is retrieved. The retrieval time metric is measured by Time To First Byte (TTFB). How long the full restoration takes, depends on the size of the object. You can determine the status of the restoration by looking at the object **Details**. Once the status shows as **Restored**, you can then download the object.

After an object is restored, you have a window of time to download the object. By default, you have 24 hours to download an object, but you can alternatively specify a time from 1 to 240 hours. You can find out how much of the download time is remaining by looking at **Available for Download** in object **Details**. After the allotted download time expires, the object returns to Archive Storage. You always have access to the metadata for an object, regardless of whether the object is in an archived or restored state.

See [Managing Objects](#) for detailed instructions on restoring, checking status of, and downloading Archive Storage objects.

Ways to Access Archive Storage

Archive Storage and Object Storage share the *same* management interfaces:

- The Console is an easy-to-use, browser-based interface. To access Archive Storage in the console, do the following:
 - [Sign in](#) to the Console.
 - Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**. A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
 - Click the name of the Archive Storage tier bucket you want to manage.
- The command line interface (CLI) provides both quick access and full functionality without the need for programming. For more information, see [Command Line Interface \(CLI\)](#).

CHAPTER 5 Archive Storage

The syntax for the CLI commands include specifying a service. You will use the Object Storage service designation: `oci os` to manage Archive Storage using the CLI.

- The REST API provides the most functionality, but requires programming expertise. [API Reference and Endpoints](#) provides endpoint details and links to the available API reference documents. For general information about using the API, see [REST APIs](#). Archive Storage is accessible with the following APIs:
 - Object Storage Service API
 - Amazon S3 Compatibility API
 - Swift API (for use with Oracle RMAN)
- Oracle Cloud Infrastructure provides SDKs that interact with Archive Storage and Object Storage without you having to create a framework. For general information about using the SDKs, see [Software Development Kits and Command Line Interface](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API). IAM also manages user credentials for things like API signing keys, auth tokens, and customer secret keys for Amazon S3 Compatibility API. See [User Credentials](#) for details.

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see the [Policy Reference](#). For specific details about writing policies for Archive Storage, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

For administrators:

CHAPTER 5 Archive Storage

- The policy [Let Object Storage admins manage buckets and objects](#) lets the specified group do everything with buckets and objects.
- Users that need to restore archived objects require the OBJECT_RESTORE permission.

WORM Compliance

You can achieve WORM compliance with Archive Storage by applying IAM policy permissions so that data once written, cannot be overwritten.

For administrators: There is not a direct way to disallow OBJECT_OVERWRITE. To achieve WORM compliance, you must specifically grant groups OBJECT_CREATE, OBJECT_READ, and OBJECT_INSPECT permissions to keep the data from being overwritten. For example, you can allow groups to inspect objects using a policy like the following:

```
Allow group <group_name> to inspect in compartment <compartment_name>
```

See [Details for Object Storage, Archive Storage, and Data Transfer](#) for more information. If you are new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Limits on Archive Storage Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

Additional limits include:

- Number of namespaces per root compartment: 1
- Maximum object size: 10 TiB
- Maximum object part size in a multipart upload: 50 GiB
- Maximum number of parts in a multipart upload: 10,000
- Maximum size of object metadata: 2 K

CHAPTER 6 Audit

This chapter explains how to work with audit logs.

Overview of Audit

The Oracle Cloud Infrastructure Audit service automatically records calls to all supported Oracle Cloud Infrastructure public application programming interface (API) endpoints as log events. Currently, all services support logging by Audit. Object Storage service supports logging for bucket-related events, but not for object-related events. Log events recorded by the Audit service include API calls made by the Oracle Cloud Infrastructure Console, Command Line Interface (CLI), Software Development Kits (SDK), your own custom clients, or other Oracle Cloud Infrastructure services. Information in the logs includes the following:

- Time the API activity occurred
- Source of the activity
- Target of the activity
- Type of action
- Type of response

Each log event includes a header ID, target resources, timestamp of the recorded event, request parameters, and response parameters. You can view events logged by the Audit service by using the Console, API, or the SDK for Java. Data from events can be used to perform diagnostics, track resource usage, monitor compliance, and collect security-related events.

Version 2 Audit Log Schema

On October 8, 2019, Oracle introduced the Audit version 2 schema, which provides the following benefits:

- Captures state changes of resources
- Better tracking of long running APIs
- Provides troubleshooting information in logs

The new schema is being implemented over time. Oracle continues to provide Audit logs in the version 1 format, but you cannot access version 1 format logs from the Console. The Console displays only the version 2 format logs. However, not all resources are emitting logs using the version 2 schema. For those services that are not emitting in the version 2 format, Oracle converts version 1 logs to version 2 logs, leaving fields blank if information for the version 2 schema cannot be determined.

Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see

[Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Administrators: For an example of policy that gives groups access to audit logs, see [Required IAM Policy](#). To modify the Audit log retention period, you must be a member of the Administrators group. See [The Administrators Group and Policy](#).

Contents of an Audit Log Event

The following explains the contents of an Audit log event. Every audit log event includes two main parts:

- Envelopes that act as a container for all event messages
- Payloads that contain data from the resource emitting the event message

Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

Event Envelope

These attributes for an event envelope are the same for all events. The structure of the envelope follows the [CloudEvents](#) industry standard format hosted by the [Cloud Native Computing Foundation \(CNCF\)](#).

Property	Description
cloudEventsVersion	<p>The version of the CloudEvents specification.</p> <div data-bbox="768 495 1302 762">Note<p>Audit uses version 0.1 specification of the CloudEvents event envelope.</p></div>
contentType	<p>Set to <code>application/json</code>. The content type of the data contained in the <code>data</code> attribute.</p>
data	<p>The payload of the event. Information within <code>data</code> comes from the resource emitting the event.</p>
eventID	<p>The UUID of the event. This identifier is not an OCID, but just a unique ID for the event.</p>
eventTime	<p>The time of the event, expressed in RFC 3339 timestamp format.</p>

Property	Description
eventType	<p>The type of event that happened.</p> <div data-bbox="769 495 1300 961" style="border: 1px solid #0070c0; background-color: #e1f5fe; padding: 10px;">  <p>Note</p> <p>The service that produces the event can also add, remove, or change the meaning of a field. A service implementing these type changes would publish a new version of an <code>eventType</code> and revise the <code>eventTypeVersion</code> field.</p> </div>
eventTypeVersion	<p>The version of the event type. This version applies to the payload of the event, not the envelope. Use <code>cloudEventsVersion</code> to determine the version of the envelope.</p>
source	<p>The resource that produced the event. For example, an Autonomous Database or an Object Storage bucket.</p>

Payload

The data in these fields depends on which service produced the event log and the event type it defines.

Data

The data object contains the following attributes.

CHAPTER 6 Audit

Property	Description
<code>data.additionalDetails</code>	A container object for attributes unique to the resource emitting the event.
<code>data.availabilityDomain</code>	The availability domain where the resource resides.
<code>data.compartmentId</code>	The OCID of the compartment of the resource emitting the event.
<code>data.compartmentName</code>	The name of the compartment of the resource emitting the event.
<code>data.definedTags</code>	Defined tags added to the resource emitting the event.
<code>data.eventGroupingId</code>	This value links multiple audit events that are part of the same API operation. For example, a long running API operation that emits an event at the start and the end of the operation.
<code>data.eventName</code>	Name of the API operation that generated this event. Example: <code>LaunchInstance</code>
<code>data.freeformTags</code>	Free-form tags added to the resource emitting the event.
<code>data.identity</code>	A container object for identity attributes. See Identity .
<code>data.request</code>	A container object for request attributes. See Request .
<code>data.resourceId</code>	An OCID or an ID for the resource emitting the event.
<code>data.resourceName</code>	The name of the resource emitting the event.
<code>data.response</code>	A container object for response attributes. See Response .
<code>data.stateChange</code>	A container object for state change attributes. See State Change .

Identity

The identity object contains the following attributes.

Property	Description
<code>data.identity.authType</code>	The type of authentication used.
<code>data.identity.callerId</code>	The OCID of the caller. The caller that made a request on behalf of the principal.
<code>data.identity.callerName</code>	The name of the user or service issuing the request. This value is the friendly name associated with <code>callerId</code> .
<code>data.identity.consoleSessionId</code>	This value identifies any Console session associated with this request.
<code>data.identity.credentials</code>	The credential ID of the user.
<code>data.identity.ipAddress</code>	The IP address of the source of the request.
<code>data.identity.principalId</code>	The OCID of the principal.
<code>data.identity.principalName</code>	The name of the user or service. This value is the friendly name associated with <code>principalId</code> .
<code>data.identity.tenantId</code>	The OCID of the tenant.
<code>data.identity.userAgent</code>	The user agent of the client that made the request.

Request

The request object contains the following attributes.

CHAPTER 6 Audit

Property	Description
<code>data.request.action</code>	The HTTP method of the request. Example: <code>GET</code>
<code>data.request.headers</code>	The HTTP header fields and values in the request.
<code>data.request.id</code>	The unique identifier of a request.
<code>data.request.parameters</code>	All the parameters supplied by the caller during this operation.
<code>data.request.path</code>	The full path of the API request. Example: <code>/20160918/instances/ocid1.instance.oc1.phx. <unique_ID></code>

Response

The response object contains the following attributes.

Property	Description
<code>data.response.headers</code>	The headers of the response.
<code>data.response.message</code>	A friendly description of what happened during the operation.
<code>data.response.payload</code>	This value is included for backward compatibility with the Audit version 1 schema, where it contained metadata of interest from the response payload.

CHAPTER 6 Audit

Property	Description
<code>data.response.responseTime</code>	The time of the response to the audited request, expressed in RFC 3339 timestamp format.
<code>data.response.status</code>	The status code of the response.

State Change

The state change object contains the following attributes.

Property	Description
<code>data.stateChange.current</code>	Provides the current state of fields that may have changed during an operation. To determine how the current operation changed a resource, compare the information in this attribute to <code>data.stateChange.previous</code> .
<code>data.stateChange.previous</code>	Provides the previous state of fields that may have changed during an operation. To determine how the current operation changed a resource, compare the information in this attribute to <code>data.stateChange.current</code> .

An Example Audit Log

The following is an example an event recorded by the Audit service.

```
{
  "eventType": "com.oraclecloud.ComputeApi.GetInstance",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "ComputeApi",
```

```

"eventId": "<unique_ID>",
"eventTime": "2019-09-18T00:10:59.252Z",
"contentType": "application/json",
"data": {
  "eventGroupingId": null,
  "eventName": "GetInstance",
  "compartmentId": "ocidl.tenancy.oc1..<unique_ID>",
  "compartmentName": "compartmentA",
  "resourceName": "my_instance",
  "resourceId": "ocidl.instance.oc1.phx.<unique_ID>",
  "availabilityDomain": "<availability_domain>",
  "freeformTags": null,
  "definedTags": null,
  "identity": {
    "principalName": "ExampleName",
    "principalId": "ocidl.user.oc1..<unique_ID>",
    "authType": "natv",
    "callerName": null,
    "callerId": null,
    "tenantId": "ocidl.tenancy.oc1..<unique_ID>",
    "ipAddress": "172.24.80.88",
    "credentials": null,
    "userAgent": "Jersey/2.23 (URLConnection 1.8.0_212)",
    "consoleSessionId": null
  },
  "request": {
    "id": "<unique_ID>",
    "path": "/20160918/instances/ocidl.instance.oc1.phx.<unique_ID>",
    "action": "GET",
    "parameters": {},
    "headers": {
      "opc-principal": [
        {"tenantId": "ocidl.tenancy.oc1..<unique_ID>", "subjectId": "ocidl.user.oc1.
ID>", "claims": [{"key": "pstype", "value": "natv", "issuer": "authService.oracle.com"},
{"key": "h_host", "value": "iaas.r2.oracleiaas.com", "issuer": "h"}, {"key": "h_opc-request-
id", "value": "<unique_ID>", "issuer": "h"},
{"key": "ptype", "value": "user", "issuer": "authService.oracle.com"}, {"key": "h_
date", "value": "Wed, 18 Sep 2019 00:10:58 UTC", "issuer": "h"}, {"key": "h_
accept", "value": "application/json", "issuer": "h"},
{"key": "authorization", "value": "Signature headers=\\\"date (request-target) host accept opc-
request-id\\\",keyId=\\\"ocidl.tenancy.oc1..<unique_ID>/ocidl.user.oc1..<unique_
ID>/8c:b4:5f:18:e7:ec:db:08:b8:fa:d2:2a:7d:11:76:ac\\\",algorithm=\\\"rsa-pss-

```



```
    "Connection": [
      "close"
    ],
    "Content-Length": [
      "1828"
    ],
    "opc-request-id": [
      "<unique_ID>"
    ],
    "Date": [
      "Wed, 18 Sep 2019 00:10:59 GMT"
    ],
    "Content-Type": [
      "application/json"
    ]
  },
  "payload": {
    "resourceName": "my_instance",
    "id": "ocidl.instance.oc1.phx.<unique_ID>"
  },
  "message": null
},
"stateChange": {
  "previous": null,
  "current": null
},
"additionalDetails": {
  "imageId": "ocidl.image.oc1.phx.<unique_ID>",
  "shape": "VM.Standard1.1",
  "type": "CustomerVmi"
}
}
```

Viewing Audit Log Events

Audit provides records of API operations performed against supported services as a list of log events. The service logs events at both the tenant and compartment level. By default, audit logs are maintained for 90 days. You can configure [audit log retention](#) for up to 365 days.

When viewing events logged by Audit, you might be interested in specific activities that happened in the tenancy or compartment and who was responsible for the activity. You will need to know that the approximate time and date something happened and the compartment in which it happened to display a list of log events that includes the activity in question. List log events by specifying a time range on the 24-hour clock in Greenwich Mean Time (GMT), calculating the offset for your local time zone, as appropriate. New activity is appended to the existing list, usually within 15 minutes of the API call, though processing time can vary.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The following policy statement gives the specified group (Auditors) the ability to view all the Audit event logs in the tenancy:

```
Allow group Auditors to read audit-events in tenancy
```

To give the group access to the Audit event logs in a specific compartment only (ProjectA), write a policy like the following:

```
Allow group Auditors to read audit-events in compartment ProjectA
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for the Audit, see [Details for the Audit Service](#).

Searching and Filtering in the Console

When you navigate to Audit in the Console, a list of results is generated for the current compartment. Audit logs are organized by compartment, so if you are looking for a particular event, you must know which compartment the event occurred in. You can filter the list in all the following ways:

- Date and time
- Request Action Types (operations)
- Keywords

For example, users begin to report that their attempts to log in are failing. You want to use Audit to research the problem. Adjust the date and time to search for corresponding failures during a window of time that starts a little before the events were reported. Look for corresponding failures and similar operations preceding the failures to correlate a reason for the failures.



Note

The service logs events at the time they are processed. There can be a delay between the time an operation occurs and when it is processed.

You can filter results by request actions to zero in on only the events with operations that interest you. For example, say that you only want to know about instances that were deleted during a specific time frame. Select a delete request action filter to see only the events with delete operations.

You can also filter by keywords. Keyword filters are powerful when combined with the values from audit event fields. For example, say that you know the user name of an account and want a list of all activity by that account in a particulate time frame. Do a search using the user name as a keyword filter.

Every audit event contains the same fields, so search for values from those fields. To get a better understanding of what values are available, see [Contents of an Audit Log Event](#).

Using the Console

To search log events

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Audit**.
The list of events that occurred in the current compartment is displayed.
2. Click one of the compartments under **Compartment**.
Audit organizes logs by compartment, so if you are looking for a particular event, you must know which compartment the event occurred in.
3. Click in the **Start Date** box to choose the start date and time for the range of results you want to see. You can click the arrows on either side of the month to go backward or forward.
4. (Optional) Specify a time by doing one of the following:
 - a. Click **Time** and specify an exact start time in thirty-minute increments.
 - b. Type an exact time in the **Start Date** box.
The service uses a 24-hour clock, so you must provide a number between 0 and 23 for the hour. Also remember to calculate the offset between Greenwich Mean Time (GMT) and your local time.
5. Repeat step 3 and 4 to choose an end date and time.



Note

The age of the results available can be from 90 and 365 days, depending on your tenancy's setting for [audit log retention](#) period.

6. (Optional) In **Request Action Types**, specify one or more operations with which to filter results.

- GET
 - POST
 - PUT
 - PATCH
 - DELETE
7. (Optional) In the **Keywords** box, type the text you want to find and click **Search**.
- Tip:** If you want to find log events with a specific status code, include quotes (") around the code to avoid results that have those numbers embedded in a longer string.

The results are updated to include only log events that were processed within the time range and filters you specified. If an event occurred in the recent past, you might have to wait to see it in the list. The service typically requires up to 15 minutes for processing.

If there are more than 100 results for the specified time range, you can click the right arrow next to the page number at the bottom of the page to advance to the next page of log events.



Tip

If you get fewer than 100 results on the last page of a results list, you might still have more results, which you can access by clicking the right arrow. If there are more results, Audit prompts you.

If you want to view all the key-value pairs in a log event, see [To view the details of a log event](#).

To view the details of a log event

View the details of your event:

- To see only the top-level details, click the down arrow to the right of an event.
- To see lower-level details, click { . . . } to the right of the collapsed parameter.

To copy the details of a log event

The following assumes that you have expanded a row in your results.

- To copy an entire event, click the clipboard icon to the right of the `event` parameter.
- To copy a portion of an event, click the clipboard icon to the right of the nested parameter or value you want to copy.

The log event is copied to your clipboard. The Audit service logs events in JSON format. You can paste the log event details into a text editor to save and review later or to use with standard log analysis tools.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operation to list audit log events:

- [ListEvents](#)



Note

This API is not intended for bulk-export operations. For bulk export, see [Bulk Export of Audit Log Events](#).

Setting Audit Log Retention Period

By default, Audit logs are retained for 90 days. You can configure log retention for up to 365 days. You can edit the log retention period in the [tenancy details](#) page.

Retention period is a tenancy-level setting. The value of the retention period setting affects all regions and all compartments. You can't set different retention periods for different regions or compartments.

Required IAM Policy

To modify the Audit log retention period, you must be a member of the Administrators group. See [The Administrators Group and Policy](#)

Using the Console

To modify the Audit log retention period

1. Open the **Profile** menu () and click **Tenancy: <your_tenancy_name>**.
The tenancy details are displayed. The **Audit Retention Period** is displayed under **Tenancy Information**.
2. Click **Edit Audit Retention Policy**. Enter the number of days you want to retain the audit logs for. The minimum you can set the value to is 90 and the maximum is 365. This value is enforced for all regions and all compartments in the tenancy.
3. Click **Submit**.



Note

You don't see the new value immediately

It may take several minutes for the new setting to display in the Console.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage the log retention period configuration:

- [GetConfiguration](#)
- [UpdateConfiguration](#)

Bulk Export of Audit Log Events

You can request a bulk export of audit logs, and within 3–4 business days Oracle support will begin making copies of the logs and adding them to buckets in your tenancy. The export includes logs for the specified regions, beginning after you make the request and continuing into the future.

Highlights

- Administrators have full control of the buckets and can provide access to others with IAM policy statements.
- Exported logs remain available indefinitely.



Tip

You can automatically manage archiving and deleting logs using Object Storage. See [Using Object Lifecycle Management](#).

- Specify all the regions you want exported in your request. If you only request some regions, then decide later you want to add other regions, you must make another request.

- To disable your bulk export, contact Oracle support. New logs will stop being added to the bucket, and audit logs will only be available through the Console, based on the retention period you have defined.

Required IAM Policy

To access the bucket where Oracle exports the audit logs, you must be a member of the Administrators group. See [The Administrators Group and Policy](#)

Requesting an Export of Audit Logs

A member of the Administrators group for your tenancy must create a ticket at [My Oracle Support](#) and provide the following information:

- Ticket name: Export Audit Logs - *<your_company_name>*
- Tenancy OCID
- Regions

For example:

- Ticket name: Export Audit Logs - ACME
- Tenancy OCID: ocid1.tenancy.oc1.*<unique_ID>*
- Regions: US East (Ashburn), region identifier= us-ashburn-1; (US West (Phoenix)), region identifier = us-phoenix-1



Note

It can take 3–4 business days before your My Oracle Support ticket is complete and the logs are available to you.

Bucket and Object Details

This section specifies the naming conventions of the bucket and objects you receive.

Bucket Name Format

Oracle support creates buckets for audit log exports using the following naming format:

`oci-logs._audit.<compartment_OCID>`

- `oci-logs` identifies that Oracle created this bucket.
- `_audit` identifies that the bucket contains audit events.
- `<compartment_OCID>` identifies the compartment where the audit events were generated.

For example:

```
oci-logs._audit.ocidlcompartment.oc1..<unique_ID>
```



Important

If the OCID of the compartment that generated the audit log contains a colon, your bucket name will not match the OCID. To create a bucket, Oracle must substitute colon characters (:) from the OCID with dot characters (.) in the bucket name.

Object Name Format

Objects use the following naming format:

`<region>/<ad>/<YYYY-MM-DDTHH:MMZ>[_<seqNum>].log.gz`

CHAPTER 6 Audit

- *<region>* identifies the region where the audit events were generated.
- *<ad>* identifies the availability domain where the audit events were generated.
- *<YYYY-MM-DDTHH:MMZ>* identifies the start time of the earliest audit event listed in the object.
- [*_<seqNum>*] identifies a conditional sequence number. If present, this number means that either an event came in late or the object became too large to write. Sequence numbers start at two. Apply multiple sequence numbers to the original object in the order listed.

For example:

```
us-phoenix-1/ad1/2019-03-21T00:00Z.log.gz
us-phoenix-1/ad1/2019-03-21T00:00Z_2.log.gz
```

File Format

Files list a single audit event per line. For more information, see [Contents of an Audit Log Event](#).



Note

Audit introduced a version 2 schema of Audit logs but bulk export is currently only available for version 1 schema logs.

CHAPTER 7 Block Volume

This chapter explains how to create storage volumes and attach them to instances.

Overview of Block Volume

The Oracle Cloud Infrastructure Block Volume service lets you dynamically provision and manage block storage volumes. You can create, attach, connect, and move volumes, as well as change volume performance, as needed, to meet your storage, performance, and application requirements. After you attach and connect a volume to an instance, you can use the volume like a regular hard drive. You can also disconnect a volume and attach it to another instance without the loss of data.

These components are required to create a volume and attach it to an instance:

- **Instance:** A bare metal or virtual machine (VM) host running in the cloud.
- **Volume attachment:** There are two types of volume attachments:
 - [iSCSI](#): A TCP/IP-based standard used for communication between a volume and attached instance.
 - [Paravirtualized](#): A virtualized attachment available for VMs.
- **Volume:** There are two types of volumes:
 - **Block volume:** A detachable block storage device that allows you to dynamically expand the storage capacity of an instance.
 - **Boot volume:** A detachable boot volume device that contains the image used to boot a Compute instance. See [Boot Volumes](#) for more information.

For additional Oracle Cloud Infrastructure terms, see the [Glossary](#).

Block Volume is Always Free eligible. For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).

Typical Block Volume Scenarios

Scenario A: Expanding an Instance's Storage

A common usage of Block Volume is adding storage capacity to an Oracle Cloud Infrastructure instance. After you have [launched an instance](#) and [set up your cloud network](#), you can create a block storage volume through the Console or API. Then, you attach the volume to an instance using a volume attachment. After the volume is attached, you connect to the volume from your instance's guest OS using iSCSI. The volume can then be mounted and used by your instance.

Scenario B: Persistent and Durable Storage

A Block Volume volume can be detached from an instance and moved to a different instance without the loss of data. This data persistence enables you to migrate data between instances and ensures that your data is safely stored, even when it is not connected to an instance. Any data remains intact until you reformat or delete the volume.

To move your volume to another instance, unmount the drive from the initial instance, terminate the iSCSI connection, and attach the volume to the second instance. From there, you connect and mount the drive from that instance's guest OS to have access to all of your data.

Additionally, Block Volume volumes offer a high level of data durability compared to standard, attached drives. All volumes are automatically replicated for you, helping to protect against data loss.

Scenario C: Instance Scaling

When you terminate an instance, you can keep the associated boot volume and use it to launch a new instance with a different instance type or shape. This allows you to easily switch from a bare metal instance to a VM instance and vice versa, or scale up or scale down the number of cores for an instance. See [Creating an Instance](#) for steps to launch an instance based on a boot volume.

Volume Attachment Types

When you attach a block volume to a VM instance, you have two options for attachment type, iSCSI or paravirtualized. Paravirtualized attachments simplify the process of configuring your block storage by removing the extra commands that are required before connecting to an iSCSI-attached volume. The trade-off is that IOPS performance for iSCSI attachments is greater than that for paravirtualized attachments. You should consider your requirements when selecting a volume's attachment type.



Important

Connecting to Volumes on Linux Instances

When connecting to volumes on Linux instances, if you want to automatically mount these volumes on instance boot, you need to use some specific options in the `/etc/fstab` file, or the instance may fail to launch. See [Traditional fstab Options](#) and [fstab Options for Block Volumes Using Consistent Device Paths](#) for more information.

iSCSI

iSCSI attachments are the only option when connecting a block volume to any of the following types of instances:

- Bare metal instances
- VM instances based on Windows images that were published before February 2018
- VM instances based on Linux images that were published before December 2017

After the volume is attached, you need to log in to the instance and use the `iscsiadm` command-line tool to configure the iSCSI connection. For more information about the additional configuration steps required for iSCSI attachments, see [iSCSI Commands and Information](#), [Connecting to a Volume](#), and [Disconnecting From a Volume](#).

IOPS performance is better with iSCSI attachments compared to paravirtualized attachments. For more information about iSCSI-attached volume performance, see [Block Volume Performance](#).

Paravirtualized

Paravirtualized attachments are an option when attaching volumes to the following types of VM instances:

- For VM instances launched from [Oracle-provided images](#), you can select this option for Linux-based images published in December 2017 or later, and Windows images published in February 2018 or later.
- For VM instances launched from custom images, the volume attachment type is based on the volume attachment type from the VM the custom image was created from.

After you attach a volume using the paravirtualized attachment type, it is ready to use, and you do not need to run any additional commands. However, because of the overhead of virtualization, this reduces the maximum IOPS performance for larger block volumes. See [Block Volume Performance](#) for more information.

Volume Access Types

When you attach a block volume, you can specify one of the following options for access type:

- **Read/write:** This is the default option for volume attachments. With this option, an instance can read and write data to the volume.
- **Read-only:** With this option, an instance can only read data on the volume. It cannot update data on the volume. Specify this option to safeguard data against accidental or malicious modifications.

To change the access type for a block volume, you need to detach the volume and specify the new access type when you reattach the volume. For more information, see [Detaching a Volume](#) and [Attaching a Volume](#).

The access type for boot volumes is always read/write. If you want to change the access type, you need to stop the instance and detach the boot volume. You can then reattach it to another

instance as a block volume, with read-only specified as the access type. For more information, see [Detaching a Boot Volume](#) and [Attaching a Volume](#).

Device Paths

When you attach a block volume to a compatible Linux-based instance, you can select a device path that remains consistent between instance reboots. This enables you to refer to the volume using a consistent device path. For example, you can use the device path when you set options in the `/etc/fstab` file to automatically mount the volume on instance boot.

Consistent device paths are supported on instances when all of the following things are true:

- The instance was created using an [Oracle-provided image](#).
- The image is a Linux-based image.
- The image was released in November 2018 or later. For specific version numbers, see [Oracle-Provided Image Release Notes](#).
- The instance was launched after January 11, 2019.

For instances launched using the image OCID or an existing boot volume, if the source image supports consistent device paths, the instance supports device paths.

Consistent device paths are not supported on Linux-based partner images or custom images that are created from other sources. This feature does not apply to Windows-based images.



Important

You must select a device path when you attach a volume using the Console, it is required. Specifying a device path is optional when you attach a volume using the CLI, REST APIs, or SDK.

For more information about consistent device paths, see [Connecting to Volumes With Consistent Device Paths](#).

Regions and Availability Domains

Volumes are only accessible to instances in the same availability domain . You cannot move a volume between availability domains or regions, they are only accessible within the region or availability domain they were created in. However volume backups are not limited to the availability domain of the source volume, you can restore them to any availability domain within that region, see [Restoring a Backup to a New Volume](#). You can also copy a volume backup to a new region and restore the backup to a volume in any availability domain in the new region, for more information see [Copying a Volume Backup Between Regions](#).

For more information, see [Regions and Availability Domains](#).

Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

Moving Resources

You can move Block Volume resources such as block volumes, boot volumes, volume backups, volume groups, and volume group backups from one compartment to another. For more information, see [Move Block Volume Resources Between Compartments](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

The following Block Volume resources emit events:

- Block volumes and block volume backups
- Boot volumes and boot volume backups
- Volume groups and volume group backups



Note

For troubleshooting, see [Known Issues - Block Volume](#) for a list of known issues related to Block Volume events.

Block Volume Encryption

The Oracle Cloud Infrastructure Block Volume service always encrypts all block volumes, boot volumes, and volume backups at rest by using the Advanced Encryption Standard (AES) algorithm with 256-bit encryption. By default all volumes and their backups are encrypted using the Oracle-provided encryption keys. Each time a volume is cloned or restored from a backup the volume is assigned a new unique encryption key.

You have the option to encrypt all of your volumes and their backups using the keys that you own and manage using the Key Management service, for more information see [Overview of Key Management](#). If you do not configure a volume to use the Key Management service or you later unassign a key from the volume, the Block Volume service uses the Oracle-provided encryption key instead. This applies to both encryption at-rest and in-transit encryption.

For how to use your own key for new volumes, see [Creating a Volume](#). See [To assign a key to an existing Block Volume](#) for how to assign or change the key for an existing volume.

All the data moving between the instance and the block volume is transferred over an internal and highly secure network. If you have specific compliance requirements related to the encryption of the data while it is moving between the instance and the block volume, the Block Volume service provides the option to enable in-transit encryption for paravirtualized volume attachments on virtual machine (VM) instances.



Important

In-transit encryption for boot and block volumes is only available for virtual machine (VM) instances launched from Oracle-provided images, it is not supported on bare metal instances. It is also not supported in most cases for instances launched from custom images imported for "bring your own image" (BYOI) scenarios. To confirm support for certain Linux-based custom images and for more information contact Oracle support, see [Contacting Support](#).

Block Volume Data Eradication

The Oracle Cloud Infrastructure Block Volume service uses eventual-overwrite data eradication, which guarantees that block volumes you delete cannot be accessed by anyone else and that the deleted data is eventually overwritten. When you terminate a volume, its associated data is overwritten in the storage infrastructure before any future volume allocations.

Block Volume Performance

Block Volume performance varies with volume size, see [Block Volume Performance](#) for more information.

The Block Volume service's elastic performance feature enables you to dynamically change the volume performance. You can select one of the following volume performance options for your block volumes:

- Balanced
- Higher Performance
- Lower Cost

For more information about this feature and the performance options, see [Block Volume Elastic Performance](#) and [Changing the Performance of a Volume](#)

Block Volume Capabilities and Limits

Block Volume volumes can be created in sizes ranging from **50 GB** to **32 TB** in **1 GB** increments. By default, Block Volume volumes are **1 TB**.

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

Additional limits include:

- **Volumes per instance:** 32
- **Number of backups**
 - Monthly universal credits: 1000
 - Pay-as-you-go: 500

iSCSI Commands and Information

Block volumes attached with the [iSCSI](#) attachment type use the iSCSI protocol to connect a volume to an instance. See [Volume Attachment Types](#) for more information about volume attachment options.

Once the volume is attached, you need to log on to the instance and use the `iscsiadm` command-line tool to configure the iSCSI connection. After you configure the volume, you can mount it and use it like a normal hard drive.

To enhance security, Oracle enforces an iSCSI security protocol called CHAP that provides authentication between the instance and volume.

Accessing a Volume's iSCSI Information

When you successfully attach a volume to an instance, Block Volume provides a list of iSCSI information. You need the following information from the list when you connect the instance to the volume.

- IP address
- Port
- CHAP user name and password (if enabled)
- IQN



Note

The CHAP credentials are auto-generated by the system and cannot be changed. They are also unique to their assigned volume/instance pair and cannot be used to authenticated another volume/instance pair.

The Console provides this information on the details page of the volume's attached instance. Click the Actions icon (three dots) on your volume's row, and then click **iSCSI Information**. The system also returns this information when the [AttachVolume](#) API operation completes successfully. You can re-run the operation with the same parameter values to review the information.

See [Attaching a Volume](#) and [Connecting to a Volume](#) for step-by-step instructions.

Additional Reading

There is a wealth of information on the internet about iSCSI and CHAP. If you need more information on these topics, try the following pages:

- [Oracle Linux Administrator's Guide for Release 7 - About iSCSI Storage](#)
- [Oracle Linux Administrator's Guide for Release 6 - About iSCSI Storage](#)
- [Troubleshooting iSCSI Configuration Problems](#)

Volume Groups

The Oracle Cloud Infrastructure Block Volume service provides you with the capability to group together multiple volumes in a volume group. A volume group can include both types of volumes, boot volumes, which are the system disks for your Compute instances, and block volumes for your data storage. You can use volume groups to create volume group backups and clones that are point-in-time and crash-consistent.

This simplifies the process to create time-consistent backups of running enterprise applications that span multiple storage volumes across multiple instances. You can then restore an entire group of volumes from a volume group backup.

Similarly, you can also clone an entire volume group in a time-consistent and crash-consistent manner. A deep disk-to-disk and fully isolated clone of a volume group, with all the volumes associated in it, becomes available for use within a matter of seconds. This speeds up the process of creating new environments for development, quality assurance, user acceptance testing, and troubleshooting.

For more information about Block Volume-backed system disks, see [Boot Volumes](#). For more information about Block Volume backups see [Overview of Block Volume Backups](#). See [Cloning a Volume](#) for more information about Block Volume clones.

This capability is available using the Console, command line interface (CLI), SDKs, or REST APIs.

Volume groups and volume group backups are high-level constructs that allow you to group together multiple volumes. When working with volume groups and volume group backups, keep the following in mind:

- You can only add a volume to a volume group when the volume status is available.
- You can add up to 32 volumes in a volume group, up to a maximum size limit of 128 TB. For example, if you wanted to add 32 volumes of equal size to a volume group, the maximum size for each volume would be 4 TB. Or you could add volumes that vary in size, however the overall combined size of all the block and boot volumes in the volume group must be 128 TB or less. Make sure you account for the size of any boot volumes in your volume group when considering volume group size limits.
- Each volume may only be in one volume group.
- When you clone a volume group, a new group with new volumes are created. For example, if you clone a volume group containing three volumes, once this operation is complete, you will now have two separate volume groups and six different volumes with nothing shared between the volume groups.
- When you update a volume group using the CLI, SDKs, or REST APIs you need to specify all the volumes to include in the volume group each time you use the update operation. If you do not include a volume ID in the update call, that volume will be removed from the volume group.
- When you delete a volume group the individual volumes in the group are not deleted, only the volume group is deleted.
- When you delete a volume that is part of a volume group you must first remove it from the volume group before you can delete it.
- When you delete a volume group backup, all the volume backups in the volume group backup are deleted.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes, backups, and volume groups.

See the following policy examples for working with volume groups:

- [Let users create a volume group](#) lets the specified group create a volume group from a set of volumes.
- [Let users clone a volume group](#) lets the specified group clone a volume group from an existing volume group.
- [Let users create a volume group backup](#) lets the specified group create a volume group backup.
- [Let users restore a volume group backup](#) lets the specified group create a volume group by restoring a volume group backup.



Tip

When users create a backup from a volume or restore a volume from a backup, the volume and backup don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Using the Console

To create a volume group from existing volumes

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. Click **Create Volume Group**.
3. Fill in the required volume information:
 - **Name:** A user-friendly name or description.
 - **Compartment:** The compartment for the volume group.
 - **Availability Domain:** The availability domain for the volume group.

- **Volumes:** For each volume you want to add, select the compartment containing the volume and then the volume to add. Click **+ Volume** to add additional volumes.
4. Click **Create Volume Group**.

To view the volumes in a volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to view the volumes for.
3. To view the block volumes for the volume group, in **Resources**, click **Block Volumes**.
4. To view the boot volumes for the volume group, in **Resources**, click **Boot Volumes**.

To add block volumes to an existing volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to add the volume to.
3. In **Resources**, click **Block Volumes**.
4. Click **Add Block Volumes**.
5. For each block volume you want to add, select the compartment containing the volume and then select the volume to add. Click **+ Volume** to add additional volumes.
6. Once you have selected all the block volumes to add to the volume group, click **Add**.

To remove block volumes from an existing volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to add the volume to.

3. In **Resources**, click **Block Volumes**.
4. In **Actions** menu for the block volume you want to remove, click **Remove**.
5. In the **Confirm** dialog, click **Remove**.



Note

When you remove the last volume in a volume group the volume group is terminated.

To add boot volumes to an existing volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to add the volume to.
3. In **Resources**, click **Boot Volumes**.
4. Click **Add Boot Volumes**.
5. For each boot volume you want to add, select the compartment containing the volume and then select the volume to add. Click **+ Volume** to add additional volumes.
6. Once you have selected all the boot volumes to add to the volume group, click **Add**.

To remove boot volumes from an existing volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to add the volume to.
3. In **Resources**, click **Boot Volumes**.
4. In **Actions** menu for the boot volume you want to remove, click **Remove**.
5. In the **Confirm** dialog, click **Remove**.

To create a backup of the volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click **Create Volume Group Backup** in the **Actions** menu for the volume group you want to create a backup for.

To create a clone of the volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click **Create Volume Group Clone** in the **Actions** menu for the volume group you want to clone.

To delete the volume group

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Volume Groups** list, click the volume group you want to delete.
3. On the **Volume Group Details** page, click **Terminate**.
4. On the **Terminate Volume Group** dialog, click **Terminate**.



Note

When you delete a volume group the individual volumes in the group are not deleted, only the volume group is deleted.

To restore a volume group from a volume group backup

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Group Backups**.
2. In the **Volume Group Backups** list, click the volume group backup you want to restore.
3. Click **Create Volume Group**.
4. Fill in the required volume information:
 - **Name:** A user-friendly name or description.
 - **Compartment:** The compartment for the volume group.
 - **Availability Domain:** The availability domain for the volume group.
5. Click **Create Volume Group**.

Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

To retrieve information about the supported operations

Open a command prompt and run the one of the following commands to retrieve the information.

- To retrieve the supported operations for volume groups:

```
oci bv volume-group --help
```

- To retrieve the supported operations for volume group backups:

```
oci bv volume-group-backup --help
```

- To retrieve help for a specific volume group operation:

```
oci bv volume-group <operation_name> --help
```

CHAPTER 7 Block Volume

- To retrieve help for a specific volume group backup operation:

```
oci bv volume-group-backup <operation_name> --help
```

Volume Group Operations

To list the volume groups in a specified compartment

Open a command prompt and run:

```
oci bv volume-group list --compartment-id <compartment_ID>
```

For example:

```
oci bv volume-group list --compartment-id ocid1.compartment.oc1..<unique_ID>
```

To create a volume group from existing volumes

Open a command prompt and run:

```
oci bv volume-group create --compartment-id <compartment_ID> --availability-domain <external_AD> --  
source-details <Source_details_JSON>
```

Volume status must be available to add it to a volume group.

For example:

```
oci bv volume-group create --compartment-id ocid1.compartment.oc1..<unique_ID> --availability-domain  
ABbv:PHX-AD-1 --source-details '{"type": "volumeIds", "volumeIds":["ocid1.volume.oc1.phx.<unique_ID_1>",  
"ocid1.volume.oc1.phx.<unique_ID_2>
```

To clone a volume group from another volume group

Open a command prompt and run:

```
oci bv volume-group create --compartment-id <compartment_ID> --availability-domain <external_AD> --  
source-details <Source_details_JSON>
```

For example:

CHAPTER 7 Block Volume

```
oci bv volume-group create --compartment-id ocidl.compartment.oc1..<unique_ID> --availability-domain  
ABbv:PHX-AD-1 --source-details '{"type": "volumeGroupId", "volumeGroupId":  
"ocidl.volumegroup.oc1.phx.<unique_ID>"}'
```

To restore a volume group from a volume group backup

Open a command prompt and run:

```
oci bv volume-group create --compartment-id <compartment_ID> --availability-domain <external_AD> --  
source-details <Source_details_JSON>
```

For example:

```
oci bv volume-group create --compartment-id ocidl.compartment.oc1..<unique_ID> --availability-domain  
ABbv:PHX-AD-1 --source-details '{"type": "volumeGroupBackupId", "volumeGroupBackupId":  
"ocidl.volumegroup.oc1.sea.<unique_ID>"}'
```

To retrieve a volume group

Open a command prompt and run:

```
oci bv volume-group get --volume-group-id <volume-group-ID>
```

For example:

```
oci bv volume-group get --volume-group-id ocidl.volumegroup.oc1.phx.<unique_ID>
```

To update display name or add/remove volumes from a volume group

Open a command prompt and run:

```
oci bv volume-group update --volume-group-id <volume-group_ID> --volume-ids <volume_ID_JSON>
```

You can update the volume group display name along with adding or removing volumes from the volume group. The volume group is updated to include only the volumes specified in the update operation. This means that you need to specify the volume IDs for all of the volumes in the volume group each time you update the volume group.

CHAPTER 7 Block Volume

The following example changes the volume group's display name for a volume group with two volumes:

```
oci bv volume-group update --volume-group-id ocidl.volumegroup.oc1.phx.<unique_ID> --volume-ids '
["ocidl.volume.oc1.phx.<unique_ID_1>","ocidl.volume.oc1.phx.<unique_ID_2>"]' --display-name "new display
name"
```

If you specify volumes in the command that are not part of the volume group they are added to the group. Any volumes not specified in the command are removed from the volume group.

To delete a volume group

Open a command prompt and run:

```
oci bv volume-group delete --volume-group-id <volume-group_ID>
```

When you delete a volume group, the individual volumes in the group are not deleted, only the volume group is deleted.

For example:

```
oci bv volume-group delete --volume-group-id ocidl.volumegroup.oc1.phx.<unique_ID>
```

Volume Group Backup Operations

To list volume backup groups

Open a command prompt and run:

```
oci bv volume-group-backup list --compartment-id <compartment_ID>
```

For example:

```
oci bv volume-group-backup list --compartment-id ocidl.compartment.oc1..<unique_ID>
```

To create a volume group backup

Open a command prompt and run:

CHAPTER 7 Block Volume

```
oci bv volume-group-backup create --volume-group-id <volume-group_ID>
```

For example:

```
oci bv volume-group-backup create --volume-group-id ocid1.volumegroup.oc1.phx.<unique_ID>
```

To retrieve a volume group backup

Open a command prompt and run:

```
oci bv volume-group-backup get --volume-group-backup-id <volume-group-backup_ID>
```

For example:

```
oci bv volume-group-backup get --volume-group-backup-id ocid1.volumegroupbackup.oc1.phx.<unique_ID>
```

To update display name for a volume group backup

Open a command prompt and run:

```
oci bv volume-group-backup update --volume-group-backup-id <volume-group-backup_ID> --display-name <new_display_name>
```

You can only update the display name for the volume group backup.

For example:

```
oci bv volume-group-backup update --volume-group-backup-id ocid1.volumegroupbackup.oc1.phx.<unique_ID> -  
-display-name "new display name"
```

To delete a volume group backup

Open a command prompt and run:

```
oci bv volume-group-backup delete --volume-group-backup-id <volume-group-backup_ID>
```

When you delete a volume group backup, all volume backups in the group are deleted.

For example:

```
oci bv volume-group-backup delete --volume-group-backup-id ocid1.volumegroupbackup.oc1.phx.<unique_ID>
```

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations for working with volume groups:

- [ListVolumeGroups](#)
- [CreateVolumeGroup](#)
- [DeleteVolumeGroup](#)
- [GetVolumeGroup](#)
- [UpdateVolumeGroup](#)

Use the following operations for working with volume group backups:

- [ListVolumeGroupBackups](#)
- [CreateVolumeGroupBackup](#)
- [DeleteVolumeGroupBackup](#)
- [GetVolumeGroupBackup](#)
- [UpdateVolumeGroupBackup](#)

Creating a Volume

You can create a volume using Block Volume.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Click **Create Block Volume**.

3. Fill in the required volume information:

- **Name:** A user-friendly name or description.
- **Domain:** Must be in the same **availability domain** as the instance.
- **Size:** Must be between **50 GB** and **32 TB**. You can choose in 1 GB increments within this range. The default is 1024 GB. If you choose a size outside of your service limit, you may be prompted to request an increase. For more information, see [Service Limits](#).
- **Backup Policy:** Optionally, you can select the appropriate backup policy for your requirements. See [Policy-Based Backups](#) for more information about backup policies.
- **Volume Performance:** Optionally, you can select the appropriate performance setting for your requirements. See [Block Volume Elastic Performance](#) for more information about volume performance options. The default option is **Balanced**.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Encryption:** Optionally, you can encrypt the data in this volume using your own Key Management encryption key. To use Key Management for your encryption needs, select the **Encrypt using customer-managed keys** radio button. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#).

4. Click **Create Block Volume**.

The volume will be ready to attach once its icon no longer lists it as **PROVISIONING** in the volume list. For more information, see [Attaching a Volume](#).

Using the API

To create a volume, use the following operation:

- [CreateVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Attaching a Volume

You can attach a volume to an instance in order to expand the available storage on the instance. If you specify [iSCSI](#) as the volume attachment type, you must also connect and mount the volume from the instance for the volume to be usable. For more information, see [Volume Attachment Types](#) and [Connecting to a Volume](#).



Note

You should only attach Linux volumes to Linux instances and Windows volumes to Windows instances.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to attach/detach existing block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the **Instances** list, click the instance that you want to attach a volume to.
3. In the **Resources** section, click **Attached Block Volumes**.
4. Click **Attach Block Volume**.
5. Select the volume attachment type, **iSCSI** or **Paravirtualized**.
For more information, see [Volume Attachment Types](#).
6. In the **Block Volume Compartment** drop-down list, select the compartment.
7. Specify the volume you want to attach to. To use the volume name, choose **SELECT VOLUME** and then select the volume from the **Block Volume** drop-down list. To specify the volume OCID, choose **ENTER VOLUME OCID** and then enter the OCID into the **Block Volume OCID** field.
8. If the instance supports consistent device paths, and the volume you are attaching is not a boot volume, select a path from the **Device Path** drop-down list when attaching. This is required and enables you to specify a device path for the volume attachment that remains consistent between instance reboots.
For more information about this feature and the instances that support it, see [Connecting to Volumes With Consistent Device Paths](#)



Tip

You must select a device path when you attach a volume from the Console, it is not optional. Specifying a device path is optional when you attach a volume using the CLI, REST APIs, or SDK.

9. Select the access type, **Read/Write** or **Read-only**.
For more information, see [Volume Access Types](#).
10. For paravirtualized volume attachments on virtual machine (VM) instances, you can optionally encrypt data that is transferred between the instance and the Block Volume service storage servers. To do this, select the **Use in-transit encryption** check box. If you configured the volume to use an encryption key that you manage using the Key Management service, this key is used for in-transit encryption. Otherwise, the Oracle-provided encryption key is used. See [Block Volume Encryption](#) for more information.
11. Click **Attach**.
When the volume's icon no longer lists it as **Attaching**, if the attachment type is [Paravirtualized](#), you can use the volume. If the attachment type is [iSCSI](#), you need to connect to the volume first. For more information, see [Connecting to a Volume](#).
On Linux-based instances, if you want to automatically mount volumes on instance boot, you need to set some specific options in the `/etc/fstab` file, or the instance may fail to launch. This applies to both iSCSI and paravirtualized attachment types. For volumes using consistent device paths, see [fstab Options for Block Volumes Using Consistent Device Paths](#). For all other volumes, see [Traditional fstab Options](#).

Using the API

To attach a volume to an instance, use the following operation:

- [AttachVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Connecting to Volumes With Consistent Device Paths

Oracle Cloud Infrastructure supports consistent device paths for block volumes that are attached to compatible Linux-based instances. When you attach a block volume to an instance, you must select a device path that remains consistent between instance reboots. This enables you to use a consistent device path when you refer to the volume to perform tasks such as:

- Creating partitions.
- Creating file systems.
- Mounting file systems.
- Specifying options in the `/etc/fstab` file to ensure that volumes are mounted properly when automatically mounting volumes on instance boot. For more information, see [fstab Options for Block Volumes Using Consistent Device Paths](#).

When you use consistent device paths on compatible Linux-based instances, the boot volume's device path is:

```
/dev/oracleoci/oraclevd
```



Note

Device paths are not available when you attach a boot volume as a data volume to a second instance.

Images that Support Consistent Device Paths

Consistent device paths are supported on instances when all of the following things are true:

- The instance was created using an [Oracle-provided image](#).
- The image is a Linux-based image.
- The image was released in November 2018 or later. For specific version numbers, see [Oracle-Provided Image Release Notes](#).
- The instance was launched after January 11, 2019.

For instances launched using the image OCID or an existing boot volume, if the source image supports consistent device paths, the instance supports device paths.

Consistent device paths are not supported on Linux-based partner images or custom images that are created from other sources. This feature does not apply to Windows-based images.



Important

You must select a device path when you attach a volume using the Console, it is required. Specifying a device path is optional when you attach a volume using the CLI, REST APIs, or SDK.

Device Paths in the Console

You select a device path when you [attach a block volume to an instance](#).

If you specify a device path, the path appears in the **Attached Block Volumes** list for an instance, in the **Device Path** field. An example is shown in the following screenshot.

Attached Block Volumes

Displaying 1 Attached Block Volumes

Attach Block Volume

 ATTACHED	block-volume-1 OCID: ...sbaosq Show Copy	Attachment Type: iscsi Block Volume Compartment: blockstoragebetatenant (root)	Size: 50.0 GB Device Path: /dev/oracleoci/oraclevdb	Created: Tue, 11 Dec 2018 15:10:15 GMT Availability Domain: uSXE:SEA-AD-1	...
---	---	---	--	--	-----

Device Paths on the Instance

Use the following sample commands to perform various configuration tasks on the attached volume. Commands are provided for volumes that use consistent device paths and for volumes that don't.

Creating a partition with `fdisk`

- **No device path specified:**

```
fdisk /dev/sdb
```

- **Device path specified:**

```
fdisk /dev/oracleoci/oracleovdb
```

Creating an ext3 file system

- **No device path specified:**

```
/sbin/mkfs.ext3 /dev/sdb1
```

- **Device path specified:**

```
/sbin/mkfs.ext3 /dev/oracleoci/oracleovdb1
```

Updating the `/etc/fstab` file

- **No device path specified:**

```
UUID=84dc162c-43dc-429c-9ac1-b511f3f0e23c /oradiskvdb1 xfs defaults,_netdev,noatime 0 2
```

- **Device path specified:**

```
/dev/oracleoci/oracleovdb1 /oradiskvdb1 ext3 defaults,_netdev,noatime 0 2
```

Mounting the file system

- **No device path specified:**

```
mount /dev/sdb1 /oradiskvdb1
```

- **Device path specified:**

```
mount /dev/oracleoci/oracleovdb1 /oradiskvdb1
```

Connecting to a Volume

For volumes attached with [Paravirtualized](#) as the volume attachment type, you do not need to perform any additional steps after [Attaching a Volume](#), the volumes are connected automatically. However, for Linux-based images, if you want to mount these volumes on instance boot, you need to perform additional configuration steps. If you specified a device path when you attached the volume, see [fstab Options for Block Volumes Using Consistent Device Paths](#). If you did not specify a device path or if your instance was created from an image that does not support device paths, see [Traditional fstab Options](#).

For volumes attached with [iSCSI](#) as the volume attachment type, you need to connect and mount the volume from the instance for the volume to be usable. For more information about attachment type options, see [Volume Attachment Types](#). In order to connect the volume, you must first attach the volume to the instance, see [Attaching a Volume](#).

Connecting to iSCSI-Attached Volumes

Required IAM Policy

Connecting a volume to an instance does not require a specific IAM policy. However, you may need permission to run the necessary commands on the instance's guest OS. Contact your system administrator for more information.

Prerequisites

You must attach the volume to the instance before you can connect the volume to the instance's guest OS. For details, see [Attaching a Volume](#).

To connect the volume, you need the following information:

- iSCSI IP Address
- iSCSI Port numbers
- CHAP credentials (if you enabled CHAP)
- IQN

The Console provides the commands required to configure, authenticate, and log on to iSCSI.

Connecting to a Volume on a Linux Instance

1. Use the Console to obtain the iSCSI data you need to connect the volume:
 - a. Log on to Oracle Cloud Infrastructure.
 - b. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
 - c. Click the name of the instance to display the instance details.
 - d. In the **Resources** section on the **Instance Details** page, click **Attached Block Volumes** to view the attached block volume.
 - e. Click the Actions icon (three dots) next to the volume you're interested in, and then click **iSCSI Commands and Information**.

The **iSCSI Commands and Information** dialog box displays specific identifying information about your volume and the iSCSI commands you'll need. The commands are ready to use with the appropriate information included. You can copy and paste the commands into your instance session window for each of the following steps.
2. Log on to your instance's guest OS.

3. Register the volume with the `iscsiadm` tool.

```
iscsiadm -m node -o new -T <volume IQN> -p <iSCSI IP address>:<iSCSI port>
```

A successful registration response resembles the following:

```
New iSCSI node [tcp:[hw=,ip=,net_if=,iscsi_if=default] 169.254.0.2,3260,-1 iqn.2015-12.us.oracle.com:c6acda73-90b4-4bbb-9a75-faux09015418] added
```

4. Configure iSCSI to automatically connect to the authenticated block storage volumes after a reboot:

```
iscsiadm -m node -T <volume IQN> -o update -n node.startup -v automatic
```

Note: All command arguments are essential. Success returns no response.

5. Skip this step if CHAP is not enabled. If you enabled CHAP when you attached the volume, authenticate the iSCSI connection by providing the volume's CHAP credentials as follows:

```
iscsiadm -m node -T <volume IQN> -p <iSCSI IP address>:<iSCSI port> -o update -n node.session.auth.authmethod -v CHAP
```

```
iscsiadm -m node -T <volume IQN> -p <iSCSI IP address>:<iSCSI port> -o update -n node.session.auth.username -v <CHAP user name>
```

```
iscsiadm -m node -T <volume's IQN> -p <iSCSI IP address>:<iSCSI port> -o update -n node.session.auth.password -v <CHAP password>
```

Success returns no response.

6. Log in to iSCSI:

```
iscsiadm -m node -T <volume's IQN> -p <iSCSI IP Address>:<iSCSI port> -l
```

A successful login response resembles the following:

```
Logging in to [iface: default, target: iqn.2015-12.us.oracle.com:c6acda73-90b4-4bbb-9a75-faux09015418, portal: 169.254.0.2,3260] (multiple)
Login to [iface: default, target: iqn.2015-12.us.oracle.com:c6acda73-90b4-4bbb-9a75-faux09015418, portal: 169.254.0.2,3260] successful.
```

7. You can now format (if needed) and mount the volume. To get a list of mountable iSCSI devices on the instance, run the following command:

```
fdisk -l
```

The connected volume listing resembles the following:

```
Disk /dev/sdb: 274.9 GB, 274877906944 bytes, 536870912 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```



Tip

If you have multiple volumes that do not have CHAP enabled, you can log in to them all at once by using the following commands:

```
iscsiadm -m discovery -t sendtargets -p <iSCSI IP
address>:<iSCSI port>
iscsiadm -m node -l
```

Connecting to a Volume on a Windows Instance



Warning

When connecting to a Windows boot volume as a data volume from a second instance, you need to append `-IsMultipathEnabled $True` to the `Connect-IscsiTarget` command. See [Attaching a Windows boot volume as a data volume to another instance fails](#) for more information.

1. Use the Console to obtain the iSCSI data you need to connect the volume:
 - a. Log on to Oracle Cloud Infrastructure.
 - b. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
 - c. Click your instance's name to display the instance details.
 - d. In the **Resources** section on the **Instance Details** page, click **Attached Block Volumes** to view the attached block volume.
 - e. Click the Actions icon (three dots) next to the volume you're interested in, and then click **iSCSI Commands and Information**.

The **iSCSI Commands and Information** dialog box displays your volume's IP address and port, which you'll need to know later in this procedure.
2. Log in to your instance using a Remote Desktop client.
3. On your Windows instance, open the iSCSI Initiator. The steps to open the iSCSI Initiator may vary depending on the version of Windows.

For example: Open **Server Manager**, click **Tools**, and then select **iSCSI Initiator**.
4. In the iSCSI Initiator Properties dialog box, click the **Discovery** tab, and then click **Discover Portal**.
5. Enter the block volume **IP Address** and **Port**, and then click **OK**.
6. Click the **Targets** tab.
7. Under **Discovered targets**, select the volume IQN.
8. Click **Connect**.
9. Make sure that the **Add this connection to the list of favorite targets** check box is selected, and then click **OK**.
10. You can now format (if needed) and mount the volume. To view a list of mountable iSCSI devices on your instance, in **Server Manager**, click **File and Storage Services**, and then click **Disks**.

The disk is displayed in the list.

fstab Options for Block Volumes Using Consistent Device Paths

On Linux instances, if you want to automatically mount volumes on instance boot, you need to set some specific options in the `/etc/fstab` file, or the instance may fail to launch.



Note

These steps are for block volumes that are attached with [consistent device paths](#) enabled. If the block volume does not have consistent device paths enabled, use the [legacy etc/fstab options](#) instead.

Prerequisites

Before using a consistent device path, you should confirm that the [instance supports consistent device paths](#) and is correctly configured.

To verify that the volume is attached to a supported instance, connect to the instance and run the following command:

```
ll /dev/oracleoci/oraclevd*
```

The output will look similar to the following:

```
lrwxrwxrwx. 1 root root 6 Feb  7 21:02 /dev/oracleoci/oraclevda -> ../sda
lrwxrwxrwx. 1 root root 7 Feb  7 21:02 /dev/oracleoci/oraclevda1 -> ../sda1
lrwxrwxrwx. 1 root root 7 Feb  7 21:02 /dev/oracleoci/oraclevda2 -> ../sda2
lrwxrwxrwx. 1 root root 7 Feb  7 21:02 /dev/oracleoci/oraclevda3 -> ../sda3
```

If you don't see this output and instead see the following error message:

```
cannot access /dev/oracleoci/oraclevd*: No such file or directory
```

there may be a problem with the instance configuration for device paths. For assistance with this, [contact Support](#).

Use the `_netdev` and `nofail` Options

By default, the `/etc/fstab` file is processed before the initiator starts. To configure the mount process to initiate before the volumes are mounted, specify the `_netdev` option on each line of the `/etc/fstab` file.

When you create a custom image of an instance where the volumes, excluding the root volume, are listed in the `/etc/fstab` file, instances will fail to launch from the custom image. To prevent this issue, specify the `nofail` option in the `/etc/fstab` file.

In the example scenario with three volumes, the `/etc/fstab` file entries for the volumes with the `_netdev` and `nofail` options are as follows:

```
/dev/oracleoci/oracleovdb /mnt/vol1 xfs defaults,_netdev,nofail 0 2
/dev/oracleoci/oracleovdc /mnt/vol2 xfs defaults,_netdev,nofail 0 2
/dev/oracleoci/oracleovdd /mnt/vol3 xfs defaults,_netdev,nofail 0 2
```

After you have updated the `/etc/fstab` file, use the following command to mount the volumes:

```
bash-4.2$ sudo mount -a
```

Reboot the instance to confirm that the volumes are mounted properly on reboot with the following command:

```
bash-4.2$ sudo reboot
```

Troubleshooting Issues with the `/etc/fstab` File

If the instance fails to reboot after you update the `/etc/fstab` file, you may need to undo the changes to the `/etc/fstab` file. To update the file, first [connect to the serial console for the instance](#). When you have access to the instance using the serial console connection, you can remove, comment out, or fix the changes that you made to the `/etc/fstab` file.

Traditional `fstab` Options

On Linux instances, if you want to automatically mount volumes on instance boot, you need to set some specific options in the `/etc/fstab` file, or the instance may fail to launch.



Note

These steps are for block volumes that do not have [consistent device paths](#) enabled. If consistent device paths are enabled for the block volume, use the [/etc/fstab options for block volumes using consistent device paths](#) instead.

Volume UUIDs

On Linux operating systems, the order in which volumes are attached is non-deterministic, so it can change with each reboot. If you refer to a volume using the device name, such as `/dev/sdb`, and you have more than one non-root volume, you can't guarantee that the volume you intend to mount for a specific device name will be the volume mounted.

To prevent this issue, specify the volume UUID in the `/etc/fstab` file instead of the device name. When you use the UUID, the mount process matches the UUID in the superblock with the mount point specified in the `/etc/fstab` file. This process guarantees that the same volume is always mounted to the same mount point.

DETERMINING THE UUID FOR A VOLUME

1. Follow the steps to [attach a volume](#) and [connect to the volume](#).
2. After the volumes are connected, create the file system of your choice on each volume using standard Linux tools.

The remaining steps assume that three volumes were connected, and that an XFS file system was created on each volume.

3. Run the following command to use the **blkid** utility to get the UUIDs for the volumes:

```
sudo blkid
```

The output will look similar to the following:

```
{{ /dev/sda3: UUID="1701c7e0-7527-4338-ae9f-672fd8d24ec7" TYPE="xfs" PARTUUID="82d2ba4e-4d6e-4a33-9c4d-ba52db57ea61"}}
{{ /dev/sda1: UUID="5750-10A1" TYPE="vfat" PARTLABEL="EFI System Partition" PARTUUID="082c26fd-85f5-4db2-9f4e-9288a3f3e784"}}
{{ /dev/sda2: UUID="1aad7aca-689d-4f4f-aff0-e0d46fc1b89f" TYPE="swap" PARTUUID="94ee5675-a805-49b2-aaf5-2fa15aade8d5"}}
{{ /dev/sdb: UUID="699a776a-3d8d-4c88-8f46-209101f318b6" TYPE="xfs"}}
{{ /dev/sdd: UUID="85566369-7148-4ffc-bf97-50954cae7854" TYPE="xfs"}}
{{ /dev/sdc: UUID="ba0ac1d3-58cf-4ff0-bd28-f2df532f7de9" TYPE="xfs"}}
```

The root volume in this output is `/dev/sda*`. The additional remote volumes are:

- `/dev/sdb`
 - `/dev/sdc`
 - `/dev/sdd`
4. To automatically attach the volumes at `/mnt/vol1`, `/mnt/vol2`, and `/mnt/vol3` respectively, create the three directories using the following commands:

```
bash-4.2$ sudo mkdir /mnt/vol1
{{ bash-4.2$ sudo mkdir /mnt/vol2}}
{{ bash-4.2$ sudo mkdir /mnt/vol3}}
```

Use the `_netdev` and `nofail` Options

By default, the `/etc/fstab` file is processed before the initiator starts. To configure the mount process to initiate before the volumes are mounted, specify the `_netdev` option on each line of the `/etc/fstab` file.

When you create a custom image of an instance where the volumes, excluding the root volume, are listed in the `/etc/fstab` file, instances will fail to launch from the custom image. To prevent this issue, specify the `nofail` option in the `/etc/fstab` file.

CHAPTER 7 Block Volume

In the example scenario with three volumes, the `/etc/fstab` file entries for the volumes with the `_netdev` and `nofail` options are as follows:

```
UUID=699a776a-3d8d-4c88-8f46-209101f318b6 /mnt/vol1 xfs defaults, _netdev, nofail 0 2
UUID=ba0ac1d3-58cf-4ff0-bd28-f2df532f7de9 /mnt/vol2 xfs defaults, _netdev, nofail 0 2
UUID=85566369-7148-4ffc-bf97-50954cae7854 /mnt/vol3 xfs defaults, _netdev, nofail 0 2
```

After you have updated the `/etc/fstab` file, use the following command to mount the volumes:

```
bash-4.2$ sudo mount -a
```

Reboot the instance to confirm that the volumes are mounted properly on reboot with the following command:

```
bash-4.2$ sudo reboot
```

Troubleshooting Issues with the `/etc/fstab` File

If the instance fails to reboot after you update the `/etc/fstab` file, you may need to undo the changes to the `/etc/fstab` file. To update the file, first [connect to the serial console for the instance](#). When you have access to the instance using the serial console connection, you can remove, comment out, or fix the changes that you made to the `/etc/fstab` file.

Listing Volumes

You can list all Block Volume volumes in a specific compartment, as well as detailed information on a single volume.

Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to list volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**. A detailed list of volumes in your current compartment is displayed.

- To view the volumes in a different compartment, change the compartment in the **Compartment** drop-down menu.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

List Volumes:

Get a list of volumes within a compartment.

- [ListVolumes](#)

Get a Single Volume:

Get detailed information on a single volume:

- [GetVolume](#)

Listing Volume Attachments

You can use the API to list all Block Volume volume attachments in a specific compartment, as well as detailed information on a single volume attachment.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to list volume attachments. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

List Attachments:

Get information on all volume attachments in a specific compartment.

- [ListVolumeAttachments](#)

Get a Single Attachment:

Get detailed information on a single attachment.

- [GetVolumeAttachment](#)

Listing Boot Volume Attachments

You can use the API to list all the boot volume attachments in a specific compartment. You can also use the API to retrieve detailed information on a single boot volume attachment.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to list volume attachments. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

List Boot Volume Attachments:

Get information on all boot volume attachments in a specific compartment.

- [ListBootVolumeAttachments](#)

Get a Single Boot Volume Attachment:

Get detailed information on a single boot volume attachment.

- [GetBootVolumeAttachment](#)

Renaming a Volume

You can use the API to change the display name of a Block Volume volume.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to rename block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the API

To update a volume's display name, use the following operation:

- [UpdateVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Resizing a Volume

The Oracle Cloud Infrastructure Block Volume service lets you expand the size of block volumes and boot volumes. You have three options to increase the size of your volumes:

- Expand an existing volume in place with offline resizing. See [Resizing a Volume Using the Console](#) for the steps to do this.
- Restore from a volume backup to a larger volume. See [Restoring a Backup to a New Volume](#) and [Restoring a Boot Volume](#).
- Clone an existing volume to a new, larger volume. See [Cloning a Volume](#) and [Cloning a Boot Volume](#).

For more information about the Block Volume service, see the [Block Volume FAQ](#).

You can only increase the size of the volume, you cannot decrease the size. You can attach a volume and start using it as soon as it's resized and becomes available.

This topic describes how to expand your volume in place with offline resizing.



Warning

Before you resize a boot or block volume, you should create a backup of the volume.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to attach/detach existing block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Considerations When Resizing a Volume

Whenever you detach and re-attach volumes, there are complexities and risks for both Linux-based and Windows-based instances. This applies to both paravirtualized and iSCSI attachment types. You should keep the following in mind when resizing volumes:

- When you re-attach a volume to an instance after resizing, if you are not using consistent device paths, or the instance does not support consistent device paths, device order and path may change. If you are using a tool such as Logical Volume Manager (LVM), you may need to fix the device mappings. For more information about consistent device paths, see [Connecting to Volumes With Consistent Device Paths](#).

- When you detach and then re-attach an iSCSI-attached volume to an instance, the volume's IP address will increment.
- Before you resize a volume, you should create a full backup of the volume.

Resizing a Volume Using the Console

Resizing a Block Volume

1. Detach the block volume, see [Detaching a Volume](#).
2. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
3. In the **Block Volumes** list, click the block volume you want to resize.
4. Click **Resize**.
5. Specify the new size and click **Resize**. You must specify a larger value than the block volume's current size.
6. Reattach the block volume, see [Attaching a Volume](#).
7. Extend the partition, see [Extending the Partition for a Block Volume](#).

Resizing a Boot Volume for a Windows Instance

1. Stop the instance, see [Stopping and Starting an Instance](#).
2. Detach the boot volume, see [Detaching a Boot Volume](#).
3. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
4. In the **Boot Volumes** list, click the boot volume you want to resize.
5. Click **Resize**.
6. Specify the new size and click **Resize**. You must specify a larger value than the boot volume's current size.

7. Reattach the boot volume, see [Attaching a Boot Volume](#).
8. Restart the instance, see [Stopping and Starting an Instance](#).
9. Extend the partition, see [Extending the System Partition on a Windows-Based Image](#).

Resizing a Boot Volume for a Linux Instance

1. Stop the instance, see [Stopping and Starting an Instance](#).
2. Detach the boot volume, see [Detaching a Boot Volume](#).
3. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
4. In the **Boot Volumes** list, click the boot volume you want to resize.
5. Click **Resize**.
6. Specify the new size and click **Resize**. You must specify a larger value than the boot volume's current size.
7. Attach the boot volume to a second instance as a data volume. See [Attaching a Volume](#) and [Connecting to a Volume](#).
8. Extend the partition and grow the file system, see [Extending the Root Partition on a Linux-Based Image](#).
9. Reattach the boot volume, see [Attaching a Boot Volume](#).
10. Restart the instance, see [Stopping and Starting an Instance](#).

Extending the Partition for a Block Volume

The Oracle Cloud Infrastructure Block Volume service lets you expand the size of block volumes with offline volume resizing. For more information, see [Resizing a Volume](#). In order to take advantage of the larger volume size, you need to extend the partition for the block volume.

Required IAM Policy

Extending a partition on an instance does not require a specific IAM policy. However, you may need permission to run the necessary commands on the instance's guest OS. Contact your system administrator for more information.

Extending a Partition on a Linux-Based Image

On Linux-based images, use the following steps to extend the partition for a block volume.

PREREQUISITES

After you have resized a volume, you need to attach it to an instance before you can extend the partition and grow the file system. See [Attaching a Volume](#) and [Connecting to a Volume](#) for more information.

EXTENDING THE LINUX PARTITION

Extending a partition

1. To identify the volume that you want to extend the partition for, run the following command to list the attached block volumes:

```
lsblk
```

2. Run the following command to edit the volume's partition table with `parted`:

```
parted <volume_id>
```

`<volume_id>` is the volume identifier, for example `/dev/sdc`.

3. When you run `parted`, you may encounter the following error message:

```
Warning: Not all of the space available to <volume_id> appears to be used,  
you can fix the GPT to use all of the space (an extra <volume_size> blocks)  
or continue with the current setting?
```

You are then prompted to fix the error or ignore the error and continue with the current setting. Specify the option to fix the error.

CHAPTER 7 Block Volume

4. Run the following command to change the display units to sectors so that you can see the precise start position for the volume:

```
(parted) unit s
```

5. Run the following command to display the current partitions in the partition table:

```
(parted) print
```

Make note of the values in the **Number**, **Start**, and **File system** columns for the root partition.

6. Run the following command to remove the existing root partition:

```
(parted) rm <partition_number>
```

<partition_number> is the value from the **Number** column.

7. Run the following command to recreate the partition:

```
(parted) mkpart
```

At the `Start?` prompt, specify the value from the **Start** column. At the `File system type?` prompt, specify the value from the **File system** column. Specify `100%` for the `End?` prompt.

8. Run the following command to exit `parted`:

```
(parted) quit
```

This command forces a rewrite of the partition table with the new partition settings that you specified.

9. To verify that the root partition was extended, run the following command to list the attached block volumes:

```
lsblk
```

After you extend the root partition you need to grow the file system. The steps in the following procedure apply only to xfs file systems.

Growing the file system for a partition

1. Before you grow the file system, repair any issues with the file system on the extended partition by running the following command:

```
xfs_repair <partition_id>
```

<partition_id> is the partition identifier, for example `/dev/sdc1`. See [Checking and Repairing an XFS File System](#) for more information.

2. After you have confirmed that there are no more issues to repair, you need to create a mount point to run the `xfs_growfs` against. To do this, create a directory and mount the partition to that directory by running the following commands:

```
mkdir <directory_name>
mount <partition_id> <directory_name> -o nouuid
```

<partition_id> is the partition identifier, for example `/dev/sdc1`, and *<directory_name>* is the directory name, for example `data`.

3. After you have created the mount point run the following command to grow the file system:

```
xfs_growfs -d <directory_name>
```

<directory_name> is the name for the directory you created in the previous step, for example `data`.

4. To verify that the file system size is correct, run the following command to display the file system details:

```
df -lh
```

Extending a Partition on a Windows-Based Image

On Windows-based images, you can extend a partition using the Windows interface or from the command line using the DISKPART utility.

WINDOWS SERVER 2016 AND WINDOWS SERVER 2012

The steps to extend a partition for a block volume attached to an instance running Windows 2012 or Windows 2016 are the same, and are described in the following procedures.

Extending a partition using the Windows interface

1. Open the [Disk Management](#) system utility on the instance.
2. Right-click the expanded block volume and select **Extend Volume**.
3. Follow the instructions in the **Extend Volume Wizard**:
 - a. Select the disk that you want to extend, enter the size, and then click **Next**.
 - b. Confirm that the disk and size settings are correct, and then click **Finish**.
4. Verify that the block volume's disk has been extended in Disk Management.

Extending a partition using the command line with DISKPART

1. Open a command prompt as administrator on the instance.
2. Run the following command to start the DISKPART utility:

```
diskpart
```

3. At the DISKPART prompt, run the following command to display the instance's volumes:

```
list volume
```

4. Run the following command to select the expanded block volume:

```
select volume <volume_number>
```

<volume_number> is the number associated with the block volume that you want to extend the partition for.

5. Run the following command to extend the partition:

```
extend size=<increased_size_in_MB>
```

<increased_size_in_MB> is the size in MB that you want to extend the partition to.



Warning

When using the DISKPART utility, do not overextend the partition beyond the current available space. Overextending the partition could result in data loss.

6. To confirm that the partition was extended, run the following command and verify that the block volume's partition has been extended:

```
list volume
```

WINDOWS SERVER 2008

Use the steps described in the following procedures to extend a partition on instances running Windows 2008.

Extending the system partition using the Windows interface

1. Open the [Server Manager](#) on the instance.
2. Expand the **Storage** node in the left navigation pane and click **Disk Management**.
3. Right-click the expanded block volume and select **Extend Volume**.
4. Follow the instructions in the **Extend Volume Wizard**:
 - a. Select the disk that you want to extend, enter the size, and then click **Next**.
 - b. Confirm that the disk and size settings are correct, and then click **Finish**.
5. Verify that the block volume's disk has been extended in the Server Manager's **Disk Management** node.

Extending a partition using the command line with DISKPART

1. Open a command prompt as administrator on the instance.
2. Run the following command to start the DISKPART utility:

```
diskpart
```

3. At the DISKPART prompt, run the following command to display the instance's volumes:

```
list volume
```

4. Run the following command to select the expanded block volume:

```
select volume <volume_number>
```

<volume_number> is the number associated with the block volume that you want to extend the partition for.

5. Run the following command to extend the partition:

```
extend size=<increased_size_in_MB>
```

<increased_size_in_MB> is the size in MB that you want to extend the partition to.



Warning

When using the DISKPART utility, do not overextend the partition beyond the current available space. Overextending the partition could result in data loss.

6. To confirm that the partition was extended, run the following command and verify that the boot volume's partition has been extended:

```
list volume
```

Overview of Block Volume Backups

The backups feature of the Oracle Cloud Infrastructure Block Volume service lets you make a point-in-time backup of data on a block volume. These backups can then be restored to new volumes either immediately after a backup or at a later time that you choose.

Backups are encrypted and stored in [Oracle Cloud Infrastructure Object Storage](#), and can be restored as new volumes to any availability domain within the same region they are stored. This capability provides you with a spare copy of a volume and gives you the ability to successfully complete disaster recovery within the same region.

There are two ways you can initiate a backup, either by manually starting the backup, or by assigning a policy which defines a set backup schedule.

Manual Backups

These are on-demand one-off backups that you can launch immediately by following the steps described in [Backing Up a Volume](#). When launching a manual backup, you can specify whether an incremental or a full backup should be performed. See [Volume Backup Types](#) for more information about backup types.

Policy-Based Backups

These are automated scheduled backups. Each backup policy has a set backup frequency and retention period. There are three predefined policies, Bronze, Silver, and Gold.

See [Policy-Based Backups](#) for more information.

Volume Backup Types

There are two backup types available in the Block Volume service:

- **Incremental:** This backup type includes only the changes since the last backup.
- **Full:** This backup type includes all changes since the volume was created.



Note

Backup Details

Backups are not an identical copy of the volume being backed up. For incremental backups, they are a record of all the changes since the last backup. For full backups, they are a record of all the changes since the volume was created. For example, in a scenario where you create a 16 TB block volume, modify 40 GB on the volume, and then launch a full backup, upon completion the volume backup size is 40 GB.

Planning Your Backup

The primary use of backups is to support business continuity, disaster recovery, and long-term archiving requirements. When determining a backup schedule, your backup plan and goals should consider the following:

- **Frequency:** How often you want to back up your data.
- **Recovery time:** How long you can wait for a backup to be restored and accessible to the applications that use it. The time for a backup to complete varies on several factors, but it will generally take a few minutes or longer, depending on the size of the data being backed up and the amount of data that has changed since your last backup.
- **Number of stored backups:** How many backups you need to keep available and the deletion schedule for those you no longer need. You can only create one backup at a time, so if a backup is underway, it will need to complete before you can create another one. For details about the number of backups you can store, see [Block Volume Capabilities and Limits](#).

The common use cases for using backups are:

- Needing to create multiple copies of the same volume. Backups are highly useful in cases where you need to create many instances with many volumes that need to have the same data formation.
- Taking a snapshot of your work that you can restore to a new volume at a later time.
- Ensuring you have a spare copy of your volume in case something goes wrong with your primary copy.

Copying Block Volume Backups Across Regions

You can copy block volume backups between regions using the Console, command line interface (CLI), SDKs, or REST APIs. For steps, see [Copying a Volume Backup Between Regions](#). This capability enhances the following scenarios:

- **Disaster recovery and business continuity:** By copying block volume backups to another region at regular intervals, it makes it easier for you to rebuild applications and data in the destination region if a region-wide disaster occurs in the source region.
- **Migration and expansion:** You can easily migrate and expand your applications to another region.

To copy volume backups between regions, you must have permission to read and copy volume backups in the source region, and permission to create volume backups in the destination region. For more information see [Required IAM Policy](#).

Once you have copied the volume backup to the new region you can then restore from that backup by creating a new volume from the backup using the steps described in [Restoring a Backup to a New Volume](#).

Volume Backup Encryption

The Oracle Cloud Infrastructure Block Volume service always encrypts all block volumes, boot volumes, and volume backups at rest by using the Advanced Encryption Standard (AES) algorithm with 256-bit encryption.

The Oracle Cloud Infrastructure Key Management service enables you to bring and manage your own keys to use for encrypting volumes and their backups. When you create a volume backup, the encryption key used for the volume is also used for the volume backup. When you restore the backup to create a new volume you configure a new key, see [Restoring a Backup to a New Volume](#). See also [Overview of Key Management](#).

If you do not configure a volume to use the Key Management service, the Block Volume service uses the Oracle-provided encryption key instead. This applies to both encryption at-rest and in-transit encryption.

Best Practices When Creating Block Volume Backups

When creating and restoring from backups, keep in mind the following:

- Before creating a backup, you should ensure that the data is consistent: Sync the file system, unmount the file system if possible, and save your application data. Only the data on the disk will be backed up. When creating a backup, after the backup state changes from REQUEST_RECEIVED to CREATING, you can return to writing data to the volume. While a backup is in progress, the volume that is being backed up cannot be deleted.
- If you want to attach a restored volume that has the original volume attached, be aware that some operating systems do not allow you to restore identical volumes. To resolve this, you should change the partition IDs before restoring the volume. The steps to change an operating system's partition ID vary by operating system. For instructions, see your operating system's documentation.
- You should not delete the original volume until you have verified that the backup you created of it completed successfully.

See [Backing Up a Volume](#) and [Restoring a Backup to a New Volume](#) for more information.

Differences Between Block Volume Backups and Clones

Consider the following criteria when you decide whether to create a backup or a clone of a volume.

	Volume Backup	Volume Clone
Description	Creates a point-in-time backup of data on a volume. You can restore multiple new volumes from the backup later in the future.	Creates a single point-in-time copy of a volume without having to go through the backup and restore process.
Use case	<p>Retain a backup of the data in a volume, so that you can duplicate an environment later or preserve the data for future use.</p> <p>Meet compliance and regulatory requirements, because the data in a backup remains unchanged over time.</p> <p>Support business continuity requirements.</p> <p>Reduce the risk of outages or data mutation over time.</p>	Rapidly duplicate an existing environment. For example, you can use a clone to test configuration changes without impacting your production environment.
Speed	Slower (minutes or hours)	Faster (seconds)
Cost	Lower cost	Higher cost
Storage location	Object Storage	Block Volume

	Volume Backup	Volume Clone
Retention policy	Policy-based backups expire, manual backups do not expire	No expiration
Volume groups	Supported. You can back up a volume group.	Supported. You can clone a volume group.

For background information and steps to clone a block volume, see [Cloning a Volume](#).

Using the CLI or REST APIs to Customize and Manage the Lifecycle of Volume Backups

You can use the CLI, REST APIs, or the SDKs to automate, script, and manage volume backups and their lifecycle.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Using the CLI

This section provides basic sample CLI commands that you can use in a script, such as a cron job run by the cron utility on Linux-based operating systems, to perform automatic backups at specific times. For information about using the CLI, see [Command Line Interface \(CLI\)](#).

To create a manual backup of the specified block volume

Open a command prompt and run:

CHAPTER 7 Block Volume

```
oci bv backup create --volume-id <block_volume_OCID> --display-name <Name> --type <FULL|INCREMENTAL>
```

For example:

```
oci bv backup create --volume-id ocid1.volume.oc1..<unique_ID> --display-name "backup display name" --type FULL
```

To delete a block volume backup

Open a command prompt and run:

```
oci bv backup delete --volume-backup-id <volume_backup_OCID>
```

For example:

```
oci bv backup delete --volume-backup-id ocid1.volume.oc1..<unique_ID>
```

To create a manual backup of the specified boot volume

Open a command prompt and run:

```
oci bv boot-volume-backup create --volume-id <boot_volume_OCID> --display-name <Name> --type <FULL|INCREMENTAL>
```

For example:

```
oci bv boot-volume-backup create --volume-id ocid1.volume.oc1..<unique_ID> --display-name "backup display name" --type FULL
```

To delete a boot volume backup

Open a command prompt and run:

```
oci bv backup delete --boot-volume-backup-id <boot_volume__backup_OCID>
```

For example:

```
oci bv backup delete --boot-volume-backup-id ocid1.volume.oc1..<unique_ID>
```

CHAPTER 7 Block Volume

To list the Oracle-defined backup policies

Open a command prompt and run:

```
oci bv volume-backup-policy list
```

To assign an Oracle-defined backup policy to a boot or block volume

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment create --asset-id <volume_OCID> --policy-id <policy_OCID>
```

For example:

```
oci bv volume-backup-policy-assignment create --asset-id ocid1.volume.oc1..<unique_ID> --policy-id ocid1.volumebackuppolicy.oc1..<unique_ID>
```

To un-assign an Oracle-defined backup policy from a boot or block volume

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment delete --policy-assignment-id <policy_assignment_OCID>
```

For example:

```
oci bv volume-backup-policy-assignment delete --policy-assignment-id ocid1.volumebackuppolicyassign.oc1..<unique_ID>
```

To retrieve the backup policy assignment ID for a boot or block volume

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment get-volume-backup-policy-asset-assignment --asset-id <volume_OCID>
```

For example:

```
oci bv volume-backup-policy-assignment get-volume-backup-policy-asset-assignment --asset-id ocid1.volume.oc1..<unique_ID>
```

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations for working with block volume backups, boot volume backups, and backup policies.

BLOCK VOLUME BACKUPS

- [CreateVolumeBackup](#)
- [DeleteVolumeBackup](#)
- [GetVolumeBackup](#)
- [ListVolumeBackups](#)
- [UpdateVolumeBackup](#)

BOOT VOLUME BACKUPS

- [CreateBootVolumeBackup](#)
- [DeleteBootVolumeBackup](#)
- [GetBootVolumeBackup](#)
- [ListBootVolumeBackups](#)
- [UpdateBootVolumeBackup](#)

VOLUME BACKUP POLICIES AND POLICY ASSIGNMENTS

- [GetVolumeBackupPolicy](#)
- [ListVolumeBackupPolicies](#)
- [CreateVolumeBackupPolicyAssignment](#)
- [DeleteVolumeBackupPolicyAssignment](#)
- [GetVolumeBackupPolicyAssetAssignment](#)
- [GetVolumeBackupPolicyAssignment](#)

Backing Up a Volume

You can create a backup of a volume using Block Volume.

For information to help you decide whether to create a backup or a clone of a boot volume, see [Differences Between Block Volume Backups and Clones](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups. The policy in [Let volume backup admins manage only backups](#) further restricts access to just creating and managing backups.



Tip

When users create a backup from a volume or restore a volume from a backup, the volume and backup don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Click the block volume that you want to create a backup for.
3. Click **Create Manual Backup**.
4. Enter a name for the backup.
5. Select the backup type, either incremental or full. See [Volume Backup Types](#) for information about backup types.
6. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Backup**.
The backup will be completed once its icon no longer lists it as **CREATING** in the volume list.

Using the API

To back up a volume, use the following operation:

- [CreateVolumeBackup](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about backups, see [Overview of Block Volume Backups](#) and [Restoring a Backup to a New Volume](#).

Policy-Based Backups

The Oracle Cloud Infrastructure Block Volume service provides you with the capability to perform volume backups automatically on a schedule and retain them based on the selected backup policy. This allows you to adhere to your data compliance and regulatory requirements.



Warning

Deleting Block Volumes with Policy-Based Backups

All policy-based backups will eventually expire, so if you want to keep a volume backup indefinitely, you need to create a manual backup.

There are two kinds of backup policies:

- **Oracle defined:** Predefined backup policies that have a set backup frequency and retention period. You cannot modify these policies.
- **User defined:** Custom backup policies that you create and configure schedules for.

Oracle Defined Backup Policies

There are three Oracle defined backup policies, Bronze, Silver, and Gold. Each backup policy is comprised of schedules with a set backup frequency and a retention period that you cannot modify. If the backup policy settings for Oracle defined policies don't meet your requirements, you should use [User Defined Backup Policies](#) instead. With user defined backup policies you define and control the schedules.

Bronze Policy

The bronze policy includes monthly incremental backups, run on the first day of the month. These backups are retained for twelve months. This policy also includes a full backup, run yearly on January 1st. Full backups are retained for five years.

Silver Policy

The silver policy includes weekly incremental backups that run on Sunday. These backups are retained for four weeks. This policy also includes monthly incremental backups, run on the first day of the month and are retained for twelve months. Also includes a full backup, run yearly on January 1st. Full backups are retained for five years.

Gold Policy

The gold policy includes daily incremental backups. These backups are retained for seven days. This policy also includes weekly incremental backups that run on Sunday and are retained for four weeks. Also includes monthly incremental backups, run on the first day of the month, retained for twelve months, and a full backup, run yearly on January 1st. Full backups are retained for five years.

User Defined Backup Policies

Oracle Cloud Infrastructure enables you to customize your backup schedules with user defined policies. These are backup policies that you define the backup frequency and retention period for. There are two parts to user defined backup policies, the backup policy itself, and then one or more schedules in the policy.

To get started with user defined backup policies, you need to first create the backup policy, see [To create a user defined backup policy](#). After this step, you have an empty backup policy, so the next step is to define and add schedules to the policy.

Schedules

Schedules define the backup frequency and retention period for a user defined backup policy, just like Oracle defined backup policies. The difference is that you can customize the

schedules associated with user defined policies. This gives you control over the backup frequency and retention period.

When defining a schedule for a user defined backup policy, the first thing you configure is the schedule type, this specifies the backup frequency. Oracle Cloud Infrastructure provides the following schedule types:

- **Daily:** Backups are generated daily. You specify the hour of the day for the backup.
- **Weekly:** Backups are generated weekly. You specify the day of the week, and the hour of that day for the backup.
- **Monthly:** Backups are generated monthly. You specify the day of the month, and the hour of that day for the backup.
- **Yearly:** Backups are generated yearly. You specify the month, the day of that month, and the hour of that day for the backup.

In addition to frequency, you also configure the following:

- **Retention time:** The amount of time to keep the backup, in days, weeks, months, or years. The time period is based the schedule type.
- **Backup type** Options are full or incremental, see [Volume Backup Types](#) for more information.
- **Timezone** The time zone to use for the backup schedule. Options are UTC or the regional data center time zone.

For more information, see [To add a schedule to a user defined backup policy](#).

You can also edit or remove schedules for a user defined policy at any time, see [To edit a schedule for a user defined backup policy](#) and [To delete a schedule for a user defined backup policy](#).

Duplicating Existing Backup Policies

You can create a new backup policy by duplicating any of the existing backup policies.

If one of the Oracle defined policies is close to meeting your volume backup requirements, but with some changes, you can create a new backup policy by duplicating the Oracle defined

policy. This creates a new user defined backup policy with schedules already assigned, enabling you to use the Oracle defined policy's settings as a starting point to save time and simplify the process.

You can also duplicate an existing user defined policy. For more information, see [To duplicate a backup policy](#). You can then add, edit, or delete schedules for the new backup policy.

Working with Backup Policies

There are two types of tasks when working with backup policies:

- [Creating and Configuring User Defined Backup Policies](#)
- [Managing Backup Policy Assignments to Volumes](#)

The linked sections listed above provide information for working with backup policies using the Console, CLI, and REST APIs.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.



Important

To view or work with backup policies, you need access to the root compartment, which is where the predefined backup policies are located.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups. The policy in [Let volume backup admins manage only backups](#) further restricts access to just creating and managing backups.



Tip

When users create a backup from a volume or restore a volume from a backup, the volume and backup don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Creating and Configuring User Defined Backup Policies

Using the Console

You can use the Console to create and update user defined backup policies.

To create a user defined backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Backup Policies**.
2. Click **Create Backup Policy**.
3. Specify a name for the backup policy.
4. Select the compartment to create the backup policy in.

While you select a compartment for the backup policy, it is accessible across your tenancy.

5. Click **Create Backup Policy** to create the backup policy.

To add a schedule to a user defined backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Backup Policies**.
2. Click the backup policy you want to add the schedule to.
3. Click **Add Schedule**.
4. Specify the backup frequency by selecting from the **Schedule Type** options: **Daily**, **Weekly**, **Monthly**, or **Yearly**, and then configure the additional schedule options. Depending on the schedule type, the additional schedule options will include one or more of the following:
 - Hour of the day
 - Day of the week
 - Day of the month
 - Month of the year
5. Specify the **Retention Time**, which will be in days, weeks, months, or years, depending on the schedule type you selected in the previous step.
6. Select **Full** or **Incremental** for **Backup Type**.
7. Select the **Timezone** to base the schedule settings on, either **UTC** or **Regional Data Center Time**.
8. Click **Add Schedule**.

To duplicate a backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and

click **Backup Policies**.

2. Click the backup policy that you want to duplicate. Both Oracle defined and user defined backup policies can be duplicated.
3. Click **Duplicate**.
4. Specify a name for the policy.
5. Select the compartment to create the backup policy in. It does not need to be the same compartment as the backup policy you are duplicating.
6. Click **Duplicate Backup Policy**.

To edit a schedule for a user defined backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Backup Policies**.
2. Click the backup policy that you want to edit a schedule for.
3. In **Schedules**, for the schedule you want to edit, click the Actions icon (three dots), and then click **Edit**.
4. After making your changes to the schedule, click **Update**.

To delete a schedule for a user defined backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Backup Policies**.
2. Click the user defined backup policy that you want to delete a schedule for.
3. In **Schedules**, for the schedule you want to delete, click the Actions icon (three dots), and then click **Delete**.
4. Click **Delete** in the confirmation dialog.

To delete a user defined backup policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Backup Policies**.
2. Click the user defined backup policy you want to delete.
3. Click **Delete**.
4. Enter the name of the backup policy and click **Delete**.

Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

Use the following operations to work with backup policies:

To create a user defined backup policy

Open a command prompt and run:

```
oci bv volume-backup-policy create --compartment-id <compartment_ID> --schedules  
file//<path>/<scheduleJSON>.json
```

For example:

```
oci bv volume-backup-policy create --compartment-id ocid1.compartment.oc1..  
<unique_ID> --schedules  
file//~/input.json
```

To list the backup policies in a specified compartment

Open a command prompt and run:

```
oci bv volume-backup-policy list --compartment-id <compartment_ID>
```

For example:

```
oci bv volume-backup-policy list --compartment-id ocid1.compartment.oc1..  
<unique_ID>
```

CHAPTER 7 Block Volume

To retrieve a specific backup policy

Open a command prompt and run:

```
oci bv volume-backup-policy get --backup-policy-id <backup-policy-ID>
```

For example:

```
oci bv volume-backup-policy get --backup-policy-id ocid1.volumebackuppolicy.oc1.phx.<unique_ID>
```

To update the display name for a user defined backup policy

Open a command prompt and run:

```
oci bv volume-backup-policy update --backup-policy-id <backup-policy_ID> --display-name <backup-policy_name>
```

For example:

```
oci bv volume-backup-policy update --backup-policy-id ocid1.volumebackuppolicy.oc1.phx.<unique_ID> --display-name "new display name"
```

To update the schedules for a user defined backup policy

Open a command prompt and run:

```
oci bv volume-backup-policy update --backup-policy-id <backup-policy_ID> --schedules file//<path>/<scheduleJSON>.json
```

For example:

```
oci bv volume-backup-policy update --volume-group-id ocid1.volumebackuppolicy.oc1.phx.<unique_ID> --schedules file//~/input.json
```

To delete a user defined backup policy

Open a command prompt and run:

```
oci bv volume-backup-policy delete --backup-policy-id <backup-policy_ID>
```

CHAPTER 7 Block Volume

You can only delete a user defined backup policy if it is not assigned to any volumes. You cannot delete Oracle defined backup policies.

For example:

```
oci bv volume-backup-policy delete --backup-policy-id ocid1.volumebackuppolicy.oc1.phx.<unique_ID>
```

Using the API

Use the following operations to work with backup policies:

- [CreateVolumeBackupPolicy](#)
- [DeleteVolumeBackupPolicy](#)
- [UpdateVolumeBackupPolicy](#)
- [ListVolumeBackupPolicies](#)
- [GetVolumeBackupPolicy](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about backups, see [Overview of Block Volume Backups](#) and [Restoring a Backup to a New Volume](#).

Managing Backup Policy Assignments to Volumes

Using the Console

You can use the Console to assign, change, or remove both user defined and Oracle defined backup policies for existing volumes.

To assign a backup policy to a volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and

click **Block Volumes**.

2. Click the volume for which you want to assign a backup policy to.
3. On the **Block Volume Information** tab click **Assign** for **Backup Policy**.
4. Select the compartment containing the backup policies.
5. Select the appropriate backup policy for your requirements.
6. Click **Assign**.

To change a backup policy assigned to a volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Click the volume for which you want to change the backup policy for.
3. On the **Block Volume Information** tab click **Edit** for **Backup Policy**.
4. Select the compartment containing the backup policies.
5. Select the backup policy you want to switch to.
6. Click **Save Changes**.

To remove a backup policy assigned to a volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Click the volume for which you want to remove the backup policy for.
3. On the **Block Volume Information** tab click **Unassign** for **Backup Policy**.
4. Click **Unassign** to confirm the backup policy removal.

Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

Use the following operations to work with volume backup policy assignments to volumes:

To assign a backup policy to a volume

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment create --asset-id <volume_ID> --policy-id <policy_ID>
```

For example:

```
oci bv volume-backup-policy-assignment create --asset-id ocid1.volume.oc1..<unique_ID> --policy-id ocid1.volumebackuppolicy.oc1..<unique_ID>
```

To get the backup policy assigned to a volume

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment get-volume-backup-policy-asset-assignment --asset-id <volume_ID>
```

For example:

```
oci bv volume-backup-policy-assignment get-volume-backup-policy-asset-assignment --asset-id ocid1.volume.oc1..<unique_ID>
```

To retrieve a specific backup policy assignment

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment get --policy-assignment-id <backup-policy-ID>
```

For example:

```
oci bv volume-backup-policy-assignment get --policy-assignment-id ocid1.volumebackuppolicyassignment.oc1.phx..<unique_ID>
```

To delete a backup policy assignment

Open a command prompt and run:

```
oci bv volume-backup-policy-assignment delete ----policy-assignment-id <backup-policy_ID>
```

You can only delete a user defined backup policy if it is not assigned to any volumes. You cannot delete Oracle defined backup policies.

For example:

```
oci bv volume-backup-policy-assignment delete ----policy-assignment-id  
ocid1.volumebackuppolicyassignment.oc1.phx.<unique_ID>
```

Using the API

Use the following operations to manage backup policy assignments to volumes:

- [CreateVolumeBackupPolicyAssignment](#)
- [DeleteVolumeBackupPolicyAssignment](#)
- [GetVolumeBackupPolicyAssetAssignment](#)
- [GetVolumeBackupPolicyAssignment](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about backups, see [Overview of Block Volume Backups](#) and [Restoring a Backup to a New Volume](#).

Copying a Volume Backup Between Regions

You can copy volume backups from one region to another region using the Oracle Cloud Infrastructure Block Volume service. For more information, see [Copying Block Volume Backups Across Regions](#).



Note

Limitations for Copying Volume Backups Across Regions

Copying boot volume backups across regions is not supported.

When copying block volume backups across regions in your tenancy, you can only copy one backup at a time from a specific source region.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The first two statements listed in the [Let volume admins manage block volumes, backups, and volume groups](#) policy lets the specified group do everything with block volumes and backups with the exception of copying volume backups across regions. The aggregate resource type `volume-family` does not include the `VOLUME_BACKUP_COPY` permission, so to enable copying volume backups across regions you need to ensure that you include the third statement in that policy, which is:

CHAPTER 7 Block Volume

```
Allow group VolumeAdmins to copy volume-backups in tenancy
```

To restrict access to just creating and managing volume backups, including copying volume backups between regions, use the policy in [Let volume backup admins manage only backups](#). The individual resource type `volume-backups` includes the `VOLUME_BACKUP_COPY` permission, so you do not need to specify it explicitly in this policy.

If you are copying volume backups encrypted using Key Management between regions or you want the copied volume backup to use Key Management for encryption in the destination region, you need to use a policy that allows the Block Volume service to perform cryptographic operations with keys in the destination region. For a sample policy showing this, see [Let Block Volume, Object Storage, File Storage services encrypt and decrypt volumes, volume backups, buckets, and file systems](#).

Restricting Access

The specific permissions needed to copy volume backups across regions are:

- **Source region:** `VOLUME_BACKUP_READ`, `VOLUME_BACKUP_COPY`
- **Destination region:** `VOLUME_BACKUP_CREATE`

Sample Policies

To restrict a group to specific source and destination regions for copying volume backups

In this example, the group is restricted to copying volume backups from the UK South (London) region to the Germany Central (Frankfurt) region.

```
Allow group MyTestGroup to read volume-backups in tenancy where all {request.region='lhr'}
Allow group MyTestGroup to use volume-backups in tenancy where all {request.permission='VOLUME_BACKUP_COPY', request.region = 'lhr'},
Allow group MyTestGroup to manage volume-backups in tenancy where all {request.permission='VOLUME_BACKUP_CREATE', request.region = 'fra'}
```

To restrict some source regions to specific destination regions while enabling all destination regions for other source regions

In this example, the following is enabled for the group:

- Manage volume backups in all regions.
- Copy volume backups from the US West (Phoenix) and US East (Ashburn) regions to any destination regions.
- Copy volume backups from the Germany Central (Frankfurt) and UK South (London) regions only to the Germany Central (Frankfurt) or UK South (London) regions.

```
Allow group MyTestGroup to read volume-backups in tenancy where all {request.region='lhr'}
Allow group MyTestGroup to manage volume-backups in tenancy where any {request.permission!='VOLUME_
BACKUP_COPY'}
Allow group MyTestGroup to use volume-backups in tenancy where all {request.permission='VOLUME_BACKUP_
COPY', any {request.region='lhr', request.region='fra'}, any{target.region='fra', target.region='lhr'}}
Allow group MyTestGroup to use volume-backups in tenancy where all {request.permission='VOLUME_BACKUP_
COPY', any {request.region='phx', request.region='iad'}}
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Block Storage**, click **Block Volume Backups**. A list of the block volume backups in the compartment you're viewing is displayed. If you don't see the one you're looking for, make sure you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the Actions icon (three dots) for the block volume backup you want to copy to another region.
3. Click **Copy to Another Region**.
4. Enter a name for the backup and choose the region to copy the backup to.
5. In the **Encryption** section select whether you want the volume backup to use the

Oracle-provided encryption key or your own Key Management encryption key. If you select the option to use your own key, paste the OCID for encryption key from the destination region.

6. Click **Copy Block Volume Backup**.
7. Confirm that the source and destination region details are correct in the confirmation dialog and then click **OK**.

Using the API

To copy a volume backup to another region, use the following operation:

- [CopyVolumeBackup](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Next Steps

After copying the block volume backup, switch to the destination region in the Console and verify that the copied backup appears in the list of block volume backups for that region. You can then restore the backup by creating a new block volume from it using the steps in [Restoring a Backup to a New Volume](#).

For more information about backups, see [Overview of Block Volume Backups](#).

Restoring a Backup to a New Volume

You can restore a backup of a volume as a new volume using Block Volume.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups.



Tip

When users create a backup from a volume or restore a volume from a backup, the volume and backup don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Block Storage**, click **Block Volume Backups**.
A list of the block volume backups in the compartment you're viewing is displayed. If you don't see the one you're looking for, make sure you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the Actions icon (three dots) for the block volume backup you want to restore.
3. Click **Create Block Volume**.
4. Enter a name for the block volume and choose the availability domain in which you want to restore it.
5. You can restore a block volume backup to a larger volume size. To do this, check **Custom Block Volume Size (GB)**, and then specify the new size. You can only increase the size of the volume, you cannot decrease the size. If you restore the block volume backup to a larger size volume, you need to extend the volume's partition, see [Extending the Partition for a Block Volume](#) for more information.
6. Optionally, you can select the appropriate backup policy for your requirements. See [Policy-Based Backups](#) for more information about backup policies.
7. Optionally, you can encrypt the data in this volume using your own Key Management encryption key. To use Key Management for your encryption needs, select the **Encrypt using Key Management** check box. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#).
8. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
9. Click **Create**.
The volume will be ready to attach once its icon no longer lists it as **PROVISIONING** in the volume list. For more information, see [Attaching a Volume](#).



Warning

If you want to attach a restored volume that has the original volume attached, be aware that some operating systems do not allow you to restore identical volumes. To resolve this, you should change the partition IDs before restoring the volume. How to change an operating system's partition ID varies by operating system; for instructions, see your operating system's documentation.

Using the API

The API used to restore a backup is [CreateVolume](#). The API has an optional `volumeBackupId` parameter that you can use to define the backup from which the data should be restored on the newly created volume. For details, see [CreateVolumeDetails Reference](#).

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about backups, see [Overview of Block Volume Backups](#) and [Backing Up a Volume](#).

Cloning a Volume

You can create a clone from a volume using the Block Volume service. Cloning enables you to make a copy of an existing block volume without needing to go through the backup and restore process.

A cloned volume is a point-in-time direct disk-to-disk deep copy of the source volume, so all the data that is in the source volume when the clone is created is copied to the clone volume. Any subsequent changes to the data on the source volume are not copied to the clone. Since

the clone is a copy of the source volume it will be the same size as the source volume unless you specify a larger volume size when you create the clone.

The clone operation occurs immediately, and you can attach and use the cloned volume as a regular volume as soon as the state changes to available. At this point, the volume data is being copied in the background, and can take up to thirty minutes depending on the size of the volume.

There is a single point-in-time reference for a source volume while it is being cloned, so if the source volume is attached when a clone is created, you need to wait for the first clone operation to complete from the source volume before creating additional clones. If the source volume is detached, you can create up to ten clones from the same source volume simultaneously.

You can only create a clone for a volume within the same region, availability domain and tenant. You can create a clone for a volume between compartments as long as you have the required access permissions for the operation.

For more information about the Block Volume service and cloned volumes, see the [Block Volume FAQ](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Differences Between Block Volume Clones and Backups

Consider the following criteria when you decide whether to create a backup or a clone of a volume.

	Volume Backup	Volume Clone
Description	Creates a point-in-time backup of data on a volume. You can restore multiple new volumes from the backup later in the future.	Creates a single point-in-time copy of a volume without having to go through the backup and restore process.
Use case	<p>Retain a backup of the data in a volume, so that you can duplicate an environment later or preserve the data for future use.</p> <p>Meet compliance and regulatory requirements, because the data in a backup remains unchanged over time.</p> <p>Support business continuity requirements.</p> <p>Reduce the risk of outages or data mutation over time.</p>	Rapidly duplicate an existing environment. For example, you can use a clone to test configuration changes without impacting your production environment.
Speed	Slower (minutes or hours)	Faster (seconds)
Cost	Lower cost	Higher cost
Storage location	Object Storage	Block Volume

	Volume Backup	Volume Clone
Retention policy	Policy-based backups expire, manual backups do not expire	No expiration
Volume groups	Supported. You can back up a volume group.	Supported. You can clone a volume group.

For more information about block volume backups, see [Overview of Block Volume Backups](#) and [Backing Up a Volume](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. In the **Block Volumes** list, click the volume that you want to clone.
3. In **Resources**, click **Clones**.
4. Click **Create Clone**.
5. Specify a name for the clone.
6. If you want to clone the block volume to a larger size volume, check **Custom Block Volume Size (GB)** and then specify the new size. You can only increase the size of the volume, you cannot decrease the size. If you clone the block volume to a larger size volume, you need to extend the volume's partition. See [Extending the Partition for a Block Volume](#) for more information.
7. If you want to change the elastic performance setting when cloning the volume, check **Custom Block Volume Performance** and select the elastic performance setting you want the volume clone to use. See [Block Volume Elastic Performance](#) for more information. You can also change the elastic performance setting after you have cloned the volume, see [Block Volume Elastic Performance](#). If you leave **Custom Block Volume Performance** unchecked, the cloned volume will use the same elastic performance setting as the source volume.
8. Click **Create Clone**.

The volume is ready use when its icon lists it as **AVAILABLE** in the volume list. At this point, you can perform various actions on the volume such as creating a clone from the volume, attaching it to an instance, or deleting the volume.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To create a clone from a volume, use the [CreateVolume](#) operation and specify [VolumeSourceFromVolumeDetails](#) for [CreateVolumeDetails](#).

Disconnecting From a Volume

For volumes attached with [iSCSI](#) as the volume attachment type you need to disconnect the volume from an instance before you detach the volume. For more information about attachment type options, see [Volume Attachment Types](#).

Required IAM Policy

Disconnecting a volume from an instance does not require a specific IAM policy. Don't confuse this with detaching a volume (see [Detaching a Volume](#)).

Disconnecting from a Volume on a Linux Instance



Warning

We recommend that you unmount and disconnect the volume from the instance using `iscsiadm` before you detach the volume. Failure to do so may lead to loss of



data.

1. Log on to your instance's guest OS and unmount the volume.
2. Run the following command to disconnect the instance from the volume:

```
iscsiadm -m node -T <IQN> -p <ISCSI IP ADDRESS>:<ISCSI PORT> -u
```

A successful logout response resembles the following:

```
Logging out of session [sid: 2, target: iqn.2015-12.us.oracle.com:c6acda73-90b4-4bbb-9a75-faux09015418, portal: 169.254.0.2,3260]
Logout of [sid: 2, target: iqn.2015-12.us.oracle.com:c6acda73-90b4-4bbb-9a75-faux09015418, portal: 169.254.0.2,3260] successful.
```

3. You can now detach the volume without the risk of losing data.

Disconnecting from a Volume on a Windows Instance

1. Use a Remote Desktop client to log on to your Windows instance, and then open **Disk Management**.
2. Right-click the volume you want to disconnect, and then click **Offline**.
3. Open **iSCSI Initiator**, select the target, and then click **Disconnect**.
4. Confirm the session termination. The status should show as **Inactive**.
5. In **iSCSI Initiator**, click the **Favorite Targets** tab, select the target you are disconnecting, and then click **Remove**.
6. Click the **Volumes and Devices** tab, select the volume from the **Volume List**, and then click **Remove**.
7. You can now detach the volume without the risk of losing data.

Detaching a Volume

When an instance no longer needs access to a volume, you can detach the volume from the instance without affecting the volume's data.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to attach/detach existing block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console



Warning

For volumes attached using [iSCSI](#), we recommend that you unmount and disconnect the volume from the instance using `iscsiadm` before you detach the volume. Failure to do so may lead to loss of data. See [Disconnecting From a Volume](#) for more information.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the **Instance** list locate the instance. Click its name to display the instance details.
3. In the **Resources** section on the **Instance Details** page, click **Attached Block Volumes**

4. Click **Detach** next to the volume you want to detach and confirm the selection when prompted.

Using the API

To delete an attachment, use the following operation:

- [DetachVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Deleting a Volume

You can delete a volume that is no longer needed.



Warning

- You cannot undo this operation. Any data on a volume will be permanently deleted once the volume is deleted.
- All policy-based backups will eventually expire, so if you want to keep a volume backup indefinitely, you need to create a manual backup. See [Overview of Block Volume Backups](#) for information about policy-based and manual backups.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. In the **Block Volumes** list, find the volume you want to delete.
3. Click **Terminate** next to the volume you want to delete and confirm the selection when prompted.

Using the API

To delete a volume, use the following operation:

- [DeleteVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Move Block Volume Resources Between Compartments

You can move Block Volume resources such as block volumes, boot volumes, volume backups, volume groups, and volume group backups from one compartment to another. When you move a Block Volume resource to a new compartment, associated resources are not moved. After you move the resource to the new compartment, inherent policies apply

immediately and affect access to the resource through the Console. For more information, see [Managing Compartments](#).



Important

When moving Block Volume resources between compartments you need to ensure that the resource users have sufficient access permissions on the compartment the resource is being moved to.

Using the Console

To move a block volume to a new compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. In the **Scope** section, select a compartment.
3. Find the block volume in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

To move a block volume backup to a new compartment

1. Open the navigation menu. Under **Block Storage**, click **Block Volume Backups**.
2. In the **Scope** section, select a compartment.
3. Find the block volume backup in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.

5. Click **Move Resource**.

To move a volume group to a new compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Groups**.
2. In the **Scope** section, select a compartment.
3. Find the volume group in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

To move a volume group backup to a new compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Volumes Group Backups**.
2. In the **Scope** section, select a compartment.
3. Find the volume group backup in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

To move a boot volume to a new compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. In the **Scope** section, select a compartment.
3. Find the boot volume in the list, click the the Actions icon (three dots), and then click **Move Resource**.

4. Choose the destination compartment from the list.
5. Click **Move Resource**.

To move a boot volume backup to a new compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volume Backups**.
2. In the **Scope** section, select a compartment.
3. Find the boot volume backup in the list, click the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

To move a block volume to a new compartment

Open a command prompt and run:

```
oci bv volume change-volume-compartment --volume-id <volume_OCID> --compartment-id <destination_compartment_OCID>
```

To move a block volume backup to a new compartment

Open a command prompt and run:

```
oci bv volume-backup change-volume-backup-compartment --volume-backup-id <volume_backup_OCID> --compartment-id <destination_compartment_OCID>
```

CHAPTER 7 Block Volume

To move a volume group to a new compartment

Open a command prompt and run:

```
oci bv volume-group change-volume-group-compartment --volume-group-id <volume_group_OCID> --compartment-id <destination_compartment_OCID>
```

To move a volume group backup to a new compartment

Open a command prompt and run:

```
oci bv volume-group-backup change-volume-group-backup-compartment --volume-group-backup-id <volume_group_backup_OCID> --compartment-id <destination_compartment_OCID>
```

To move a boot volume to a new compartment

Open a command prompt and run:

```
oci bv boot-volume change-boot-volume-compartment --boot-volume-id <boot_volume_OCID> --compartment-id <destination_compartment_OCID>
```

To move a boot volume backup to a new compartment

Open a command prompt and run:

```
oci bv boot-volume-backup change-boot-volume-backup-compartment --boot-volume-backup-id <boot_volume_backup_OCID> --compartment-id <destination_compartment_OCID>
```

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations for moving Block Volume resources between compartments:

- [ChangeVolumeCompartment](#)
- [ChangeVolumeBackupCompartment](#)

- [ChangeVolumeGroupCompartment](#)
- [ChangeVolumeGroupBackupCompartment](#)
- [ChangeBootVolumeCompartment](#)
- [ChangeBootVolumeBackupCompartment](#)

Block Volume Performance

The content in the sections below apply to **Category 7** and **Section 3.b** of the [Oracle PaaS and IaaS Public Cloud Services Pillar documentation](#).

The Oracle Cloud Infrastructure Block Volume service lets you dynamically provision and manage block storage volumes. You can create, attach, connect and move volumes as needed to meet your storage and application requirements. The Block Volume service uses NVMe-based storage infrastructure, and is designed for consistency. You just need to provision the capacity needed and performance scales with the performance characteristics of the elastic performance option selected up to the service maximums. See [Block Volume Elastic Performance](#) for specific details about the elastic performance options.

The Block Volume service supports creating volumes sized from 50 GB to a maximum size of 32 TB, in 1 GB increments. You can attach up to 32 volumes to an instance, with a maximum of 1 PB of attached volumes per instance. Latency performance is independent of the instance shape or volume size, and is always sub-millisecond at the 95th percentile for the **Balanced** and **Higher Performance** elastic performance options.



Note

Block Volume performance may be limited by the network bandwidth of the instance shape, for more information see [Compute Shapes](#).

Higher Performance

The **Higher Performance** elastic performance option is recommended for workloads with the highest I/O requirements, requiring the best possible performance, such as large databases. This option provides the best linear performance scale with 75 IOPS/GB up to a maximum of 35,000 IOPS per volume. Throughput also scales at the highest rate at 600 KBPS/GB up to a maximum of 480 MBPS per volume.

The following table lists the Block Volume service's throughput and IOPS performance numbers based on volume size for this option. IOPS and KPBS performance scales linearly per GB volume size up to the service maximums so you can predictably calculate the performance numbers for a specific volume size. If you're trying to achieve certain performance targets for volumes configured to use the **Higher Performance** elastic performance option you can provision a minimum volume size using this table as a reference.

Volume Size	Max Throughput (1 MB block size)	Max Throughput (8 KB block size)	Max IOPS (4 KB block size)
50 GB	30 MB/s	30 MB/s	3750
100 GB	60 MB/s	60 MB/s	7500
200 GB	120 MB/s	96 MB/s	15,000
300 GB	180 MB/s	180 MB/s	22,500
400 GB	240 MB/s	240 MB/s	30,000
500 GB	300 MB/s	280 MB/s	35,000
800 GB - 32 TB	480 MB/s	280 MB/s	35,000

Balanced Performance

The **Balanced** elastic performance option provides a good balance between performance and cost savings for most workloads, including those that perform random I/O such as boot

CHAPTER 7 Block Volume

volumes. This option provides linear performance scaling with 60 IOPS/GB up to 25,000 IOPS per volume. Throughput scales at 480 KBPS/GB up to a maximum of 480 MBPS per volume.

The following table lists the Block Volume service's throughput and IOPS performance numbers based on volume size for this option. IOPS and KPBS performance scales linearly per GB volume size up to the service maximums so you can predictably calculate the performance numbers for a specific volume size. If you're trying to achieve certain performance targets for volumes configured to use the **Balanced** elastic performance option you can provision a minimum volume size using this table as a reference.

Volume Size	Max Throughput (1 MB block size)	Max Throughput (8 KB block size)	Max IOPS (4 KB block size)
50 GB	24 MB/s	24 MB/s	3000
100 GB	48 MB/s	48 MB/s	6000
200 GB	96 MB/s	96 MB/s	12,000
300 GB	144 MB/s	144 MB/s	18,000
400 GB	192 MB/s	192 MB/s	24,000
500 GB	240 MB/s	200 MB/s	25,000
750 GB	360 MB/s	200 MB/s	25,000
1 TB - 32 TB	480 MB/s	200 MB/s	25,000

Lower Cost

The **Lower Cost** elastic performance option is recommended for throughput intensive workloads with large sequential I/O, such as streaming, log processing, and data warehouses. This option gives you linear scaling 2 IOPS/GB up to a maximum of 3000 IOPS per volume.

The following table lists the Block Volume service's throughput and IOPS performance numbers based on volume size for this option. IOPS and KPBS performance scales linearly per

CHAPTER 7 Block Volume

GB volume size up to the service maximums so you can predictably calculate the performance numbers for a specific volume size. If you're trying to achieve certain performance targets for volumes configured to use the **Lower Cost** elastic performance option you can provision a minimum volume size using this table as a reference.

Volume Size	Max Throughput (1 MB block size)	Max Throughput (8 KB block size)	Max IOPS (4 KB block size)
50 GB	12 MB/s	0.8 MB/s	100
100 GB	24 MB/s	1.6 MB/s	200
200 GB	48 MB/s	3.2 MB/s	400
300 GB	72 MB/s	4.8 MB/s	600
400 GB	96 MB/s	6.4 MB/s	800
500 GB	120 MB/s	8 MB/s	1000
750 GB	180 MB/s	12 MB/s	1500
1 TB	240 MB/s	16 MB/s	2000
1.5 TB - 32 TB	480 MB/s	23 MB/s	3000

For more information about FIO command samples you can use for performance testing see [Sample FIO Commands for Block Volume Performance Tests on Linux-based Instances](#).

See [Using Block Volumes Service Metrics to Calculate Block Volume Throughput and IOPS](#) for a walkthrough of a performance testing scenario with FIO that shows how you can use Block Volume metrics to determine the performance characteristics of your block volume.

Limitations and Considerations

- Block Volume performance SLA for IOPS per volume and IOPS per instance applies to the **Balanced** and **Higher Performance** elastic performance settings only, not to the **Lower Cost** setting.
- The throughput performance results are for bare metal Compute instances. Throughput performance on virtual machine (VM) Compute instances is dependent on the network bandwidth that is available to the instance, and further limited by that bandwidth for the volume. For details about the network bandwidth available for VM shapes, see the **Network Bandwidth** column in the [VM Shapes](#) table.
- IOPS performance is independent of the instance type or shape, so is applicable to all bare metal and VM shapes, for iSCSI attached volumes. For VM shapes with paravirtualized attached volumes, see [Block Volume Performance](#).
- For the **Lower Cost** option you may not see the same latency performance that you see with the **Balanced** or **Higher Performance** elastic performance options. You may also see a greater variance in latency with the **Lower Cost** option.
- Windows Defender Advanced Threat Protection (Windows Defender ATP) is enabled by default on all Oracle-provided Windows images. This tool has a significant negative impact on disk I/O performance. The IOPS performance characteristics described in this topic are valid for Windows bare metal instances with Windows Defender ATP disabled for disk I/O. Customers must carefully consider the security implications of disabling Windows Defender ATP. See [Windows Defender Advanced Threat Protection](#).
- The IOPS performance characteristics described in this topic are for volumes with iSCSI attachments. Block Volume performance SLA for IOPS per volume and IOPS per instance applies to iSCSI volume attachments only, not to paravirtualized attachments. Paravirtualized attachments simplify the process of configuring your block storage by removing the extra commands needed before accessing a volume. However, due to the overhead of virtualization, this reduces the maximum IOPS performance for larger block volumes. If storage IOPS performance is of paramount importance for your workloads, you can continue to experience the guaranteed performance Oracle Cloud Infrastructure Block Volume offers by using iSCSI attachments.

Testing Methodology and Performance for Balanced Elastic Performance Option



Warning

- Before running any tests, protect your data by making a backup of your data and operating system environment to prevent any data loss.
- Do not run FIO tests directly against a device that is already in use, such as /dev/sdX. If it is in use as a formatted disk and there is data on it, running FIO with a write workload (readwrite, randrw, write, trimwrite) will overwrite the data on the disk, and cause data corruption. Run FIO only on unformatted raw devices that are not in use.

This section describes the setup of the test environments, the methodology, and the observed performance for the Balanced elastic performance configuration option. Some of the sample volume sizes tested were:

- 50 GB volume - 3,000 IOPS @ 4K
- 1 TB volume - 25,000 IOPS @ 4K
- Host maximum, Ashburn (IAD) region, twenty 1 TB volumes - 400,000 IOPS @ 4K

These tests used a wide range of volume sizes and the most common read and write patterns and were generated with the [Gartner Cloud Harmony test suite](#). To show the throughput performance limits, 256k or larger block sizes should be used. For most environments, 4K, 8K, or 16K blocks are common depending on the application workload, and these are used specifically for IOPS measurements.

In the observed performance images in this section, the X axis represents the volume size tested, ranging from 4KB to 1MB. The Y axis represents the IOPS delivered. The Z axis represents the read/write mix tested, ranging from 100% read to 100% write.



Note

Performance Notes for Instance Types

- The throughput performance results are for bare metal instances. Throughput performance on VM instances is dependent on the network bandwidth that is available to the instance, and further limited by that bandwidth for the volume. For details about the network bandwidth available for VM shapes, see the Network Bandwidth column in the VM Shapes table.
- IOPS performance is independent of the instance type or shape, so is applicable to all bare metal and VM shapes, for iSCSI attached volumes. For VM shapes with paravirtualized attached volumes, see [Block Volume Performance](#).

1 TB Block Volume

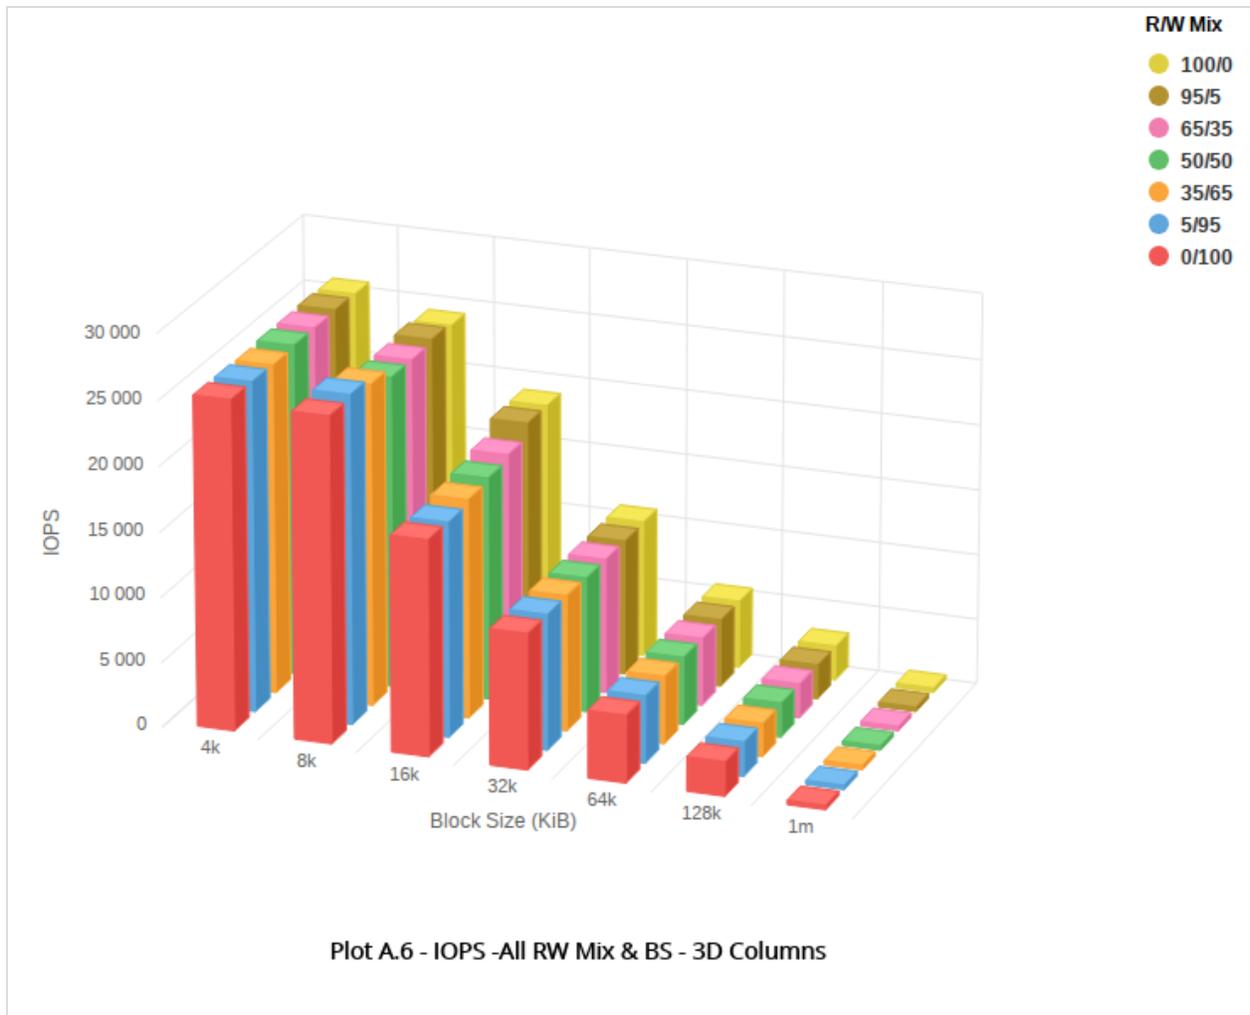
A 1 TB volume was mounted to a bare metal instance running in the Phoenix region. The instance shape was dense, workload was direct I/O with 10GB working set. The following command was run for the Gartner Cloud Harmony test suite:

```
~/block-storage/run.sh --nopurge --noprecondition --fio_direct=1 --fio_size=10g --target /dev/sdb --  
test iops --skip_blocksize 512b
```

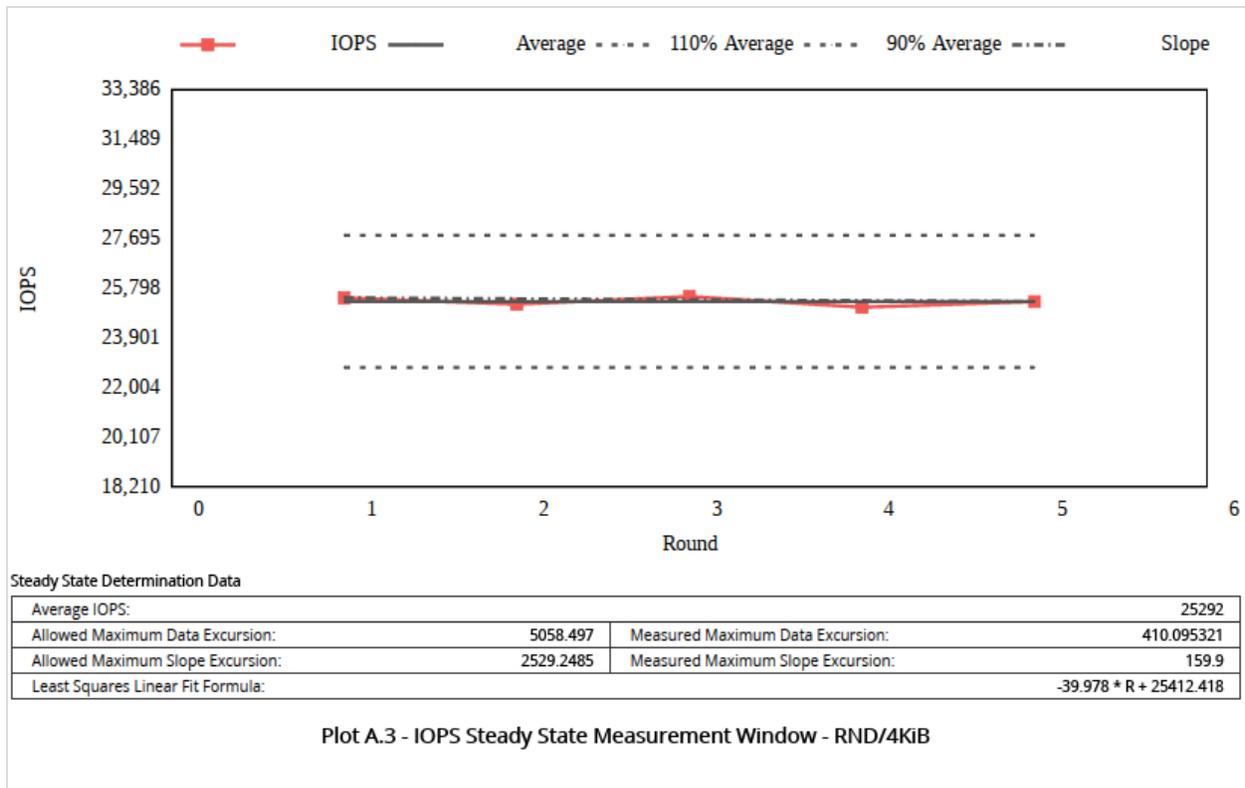
The results showed that for 1 TB, the bandwidth limit for the larger block size test occurs at 320MBS.

The following images show the observed performance for 1 TB:

CHAPTER 7 Block Volume



CHAPTER 7 Block Volume



50 GB Block Volume

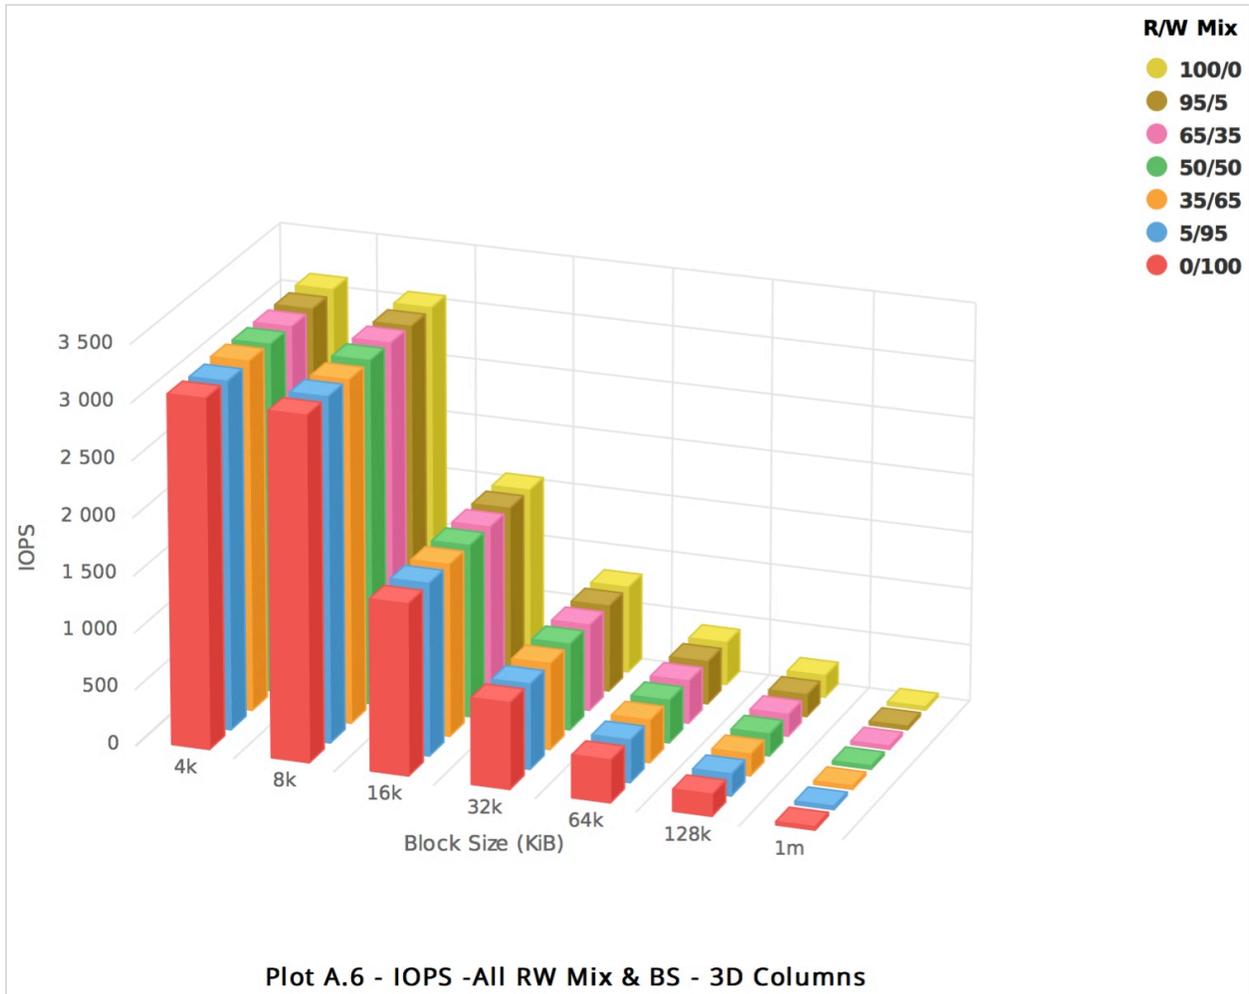
A 50 GB volume was mounted to a bare metal instance running in the Phoenix region. The instance shape was dense, workload was direct I/O with 10GB working set. The following command was run for the Gartner Cloud Harmony test suite:

```
~/block-storage/run.sh --nopurge --noprecondition --fio_direct=1 --fio_size=10g --target /dev/sdb --test iops --skip_blocksize 512b
```

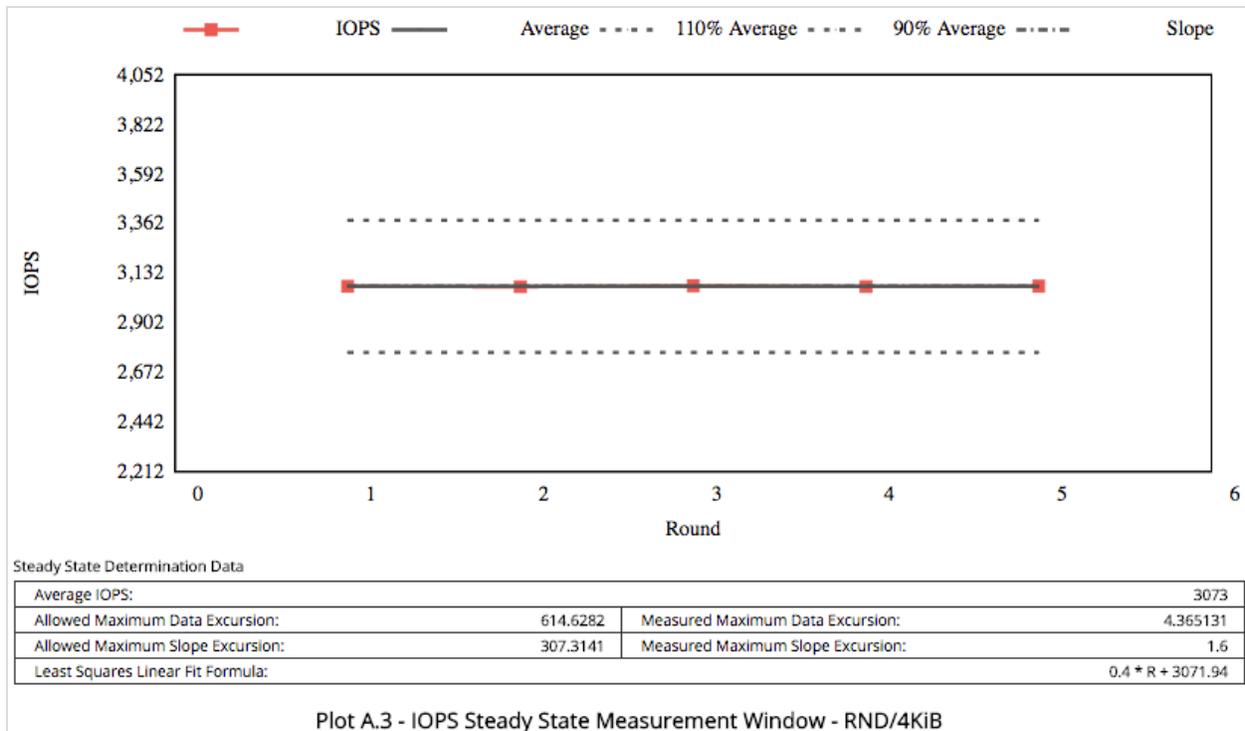
The results showed that for the 50 GB volume, the bandwidth limit is confirmed as 24,000 KBPS for the larger block size tests (256 KB or larger block sizes), and the maximum of 3,000 IOPS at 4K block size is delivered. For small volumes, a 4K block size is common.

The following images show the observed performance for 50 GB:

CHAPTER 7 Block Volume



CHAPTER 7 Block Volume



Host Maximum - Twenty 1 TB Volumes

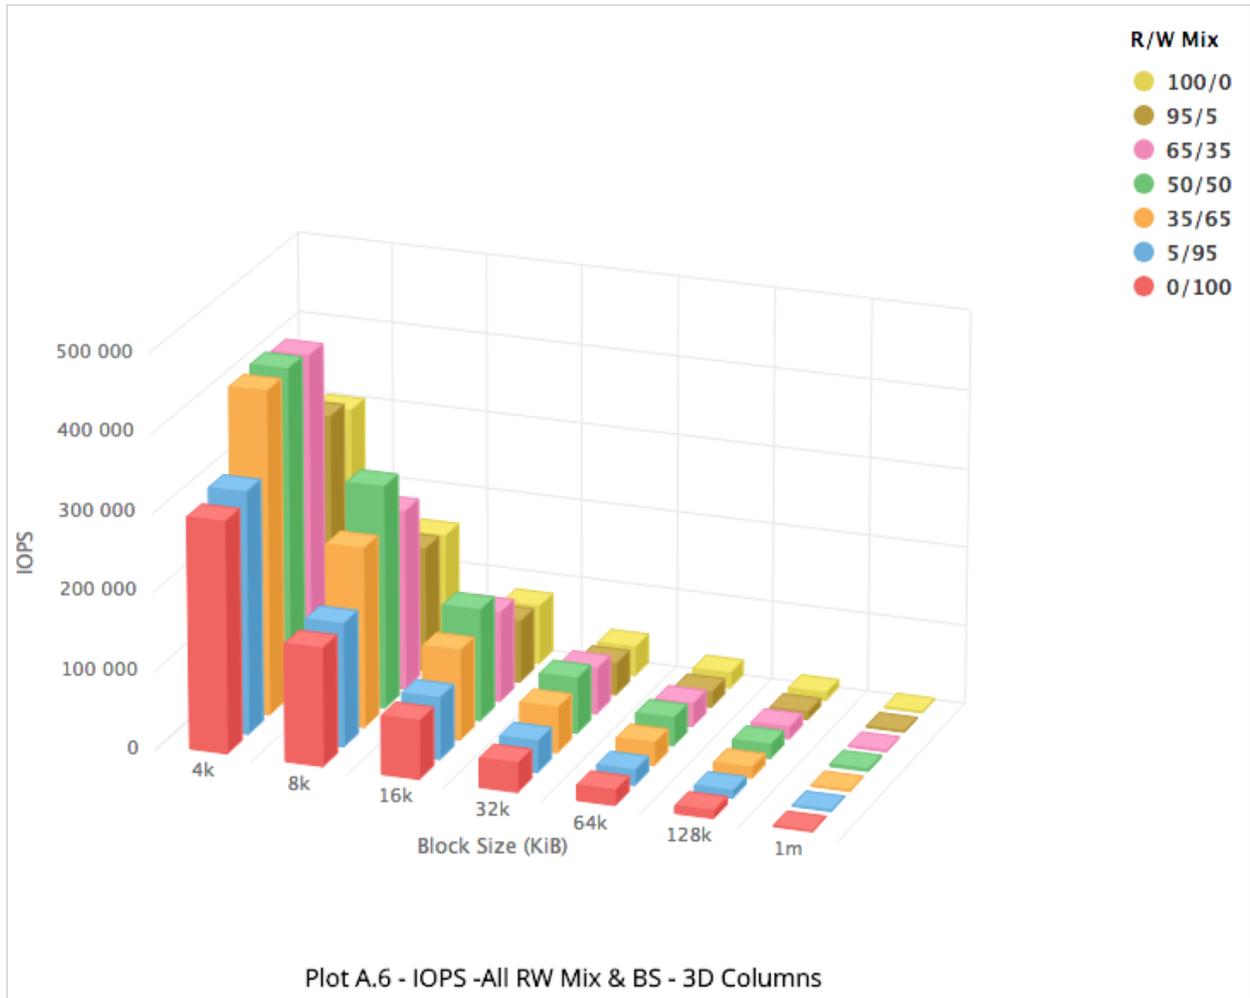
Twenty 1 TB volumes were mounted to a bare metal instance running in the Ashburn region. The instance shape was dense, workload was direct I/O with 10GB working set. The following command was run for the Gartner Cloud Harmony test suite:

```
~/block-storage/run.sh --nopurge --noprecondition --fio_direct=1 --fio_size=10g --target
/dev/sdb,/dev/sdc,/dev/sdd,/dev/sde,/dev/sdf,/dev/sdg,/dev/sdh,/dev/sdi,/dev/sdj,/dev/sdk,/dev/sdl,/dev
/sdm,/dev/sdn,/dev/sdo,/dev/sdp,/dev/sdq,/dev/sdr,/dev/sds,/dev/sdt,/dev/sdu --test iops --skip_
blocksize 512b
```

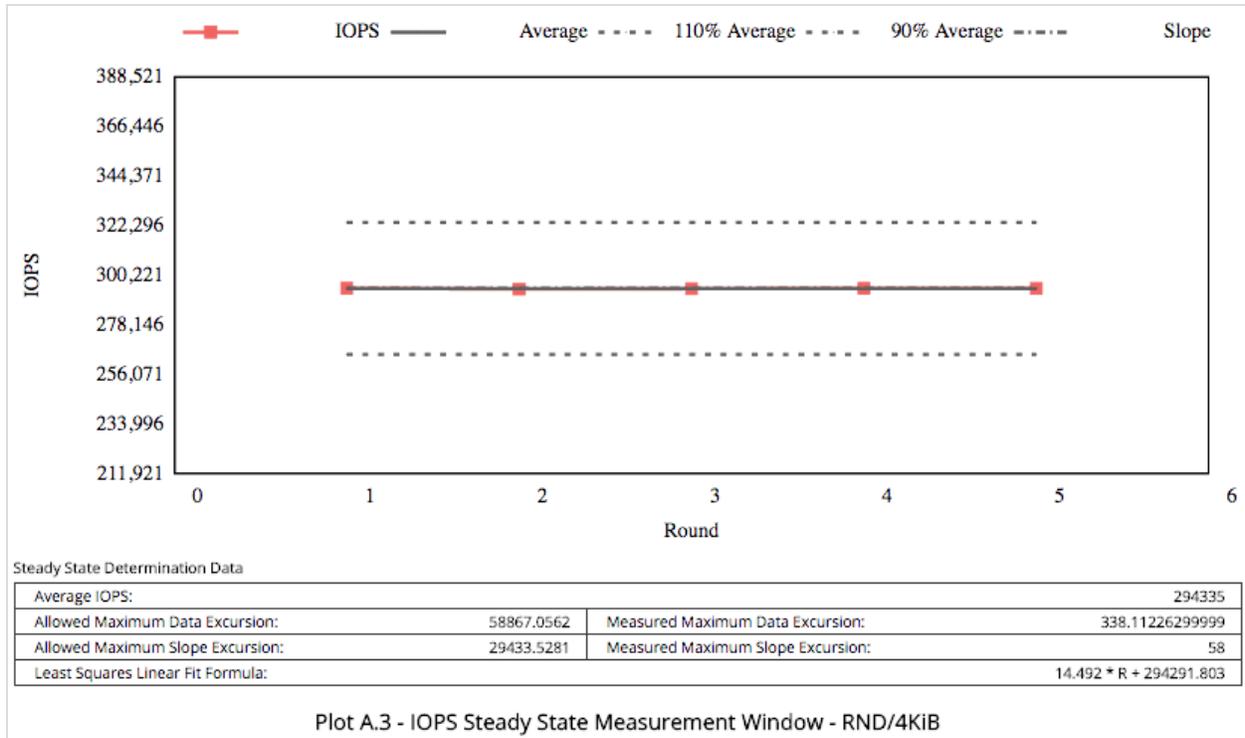
The results showed that for the host maximum test of twenty 1 TB volumes, the average is 2.1GBPS, and 400,000 IOPS to the host for the 50/50 read/write pattern.

The following images show the observed performance for 50 GB:

CHAPTER 7 Block Volume



CHAPTER 7 Block Volume



Sample FIO Commands for Block Volume Performance Tests on Linux-based Instances

This topic describes sample FIO commands you can use to run performance tests for the Oracle Cloud Infrastructure Block Volume service on instances created from Linux-based images.

Installing FIO

To install and configure FIO on your instances with Linux-based operating systems, run the commands applicable to the operating system version for your instance.

Oracle Linux and CentOS

Run the following command to install and configure FIO for your Oracle Linux CentOS systems:

```
sudo yum install fio -y
```

This applies to Oracle Linux 6.x, Oracle Linux 7.x, CentOS 6.x, and CentOS 7.x.

Ubuntu

Run the following commands to install and configure FIO for your Ubuntu systems:

```
sudo apt-get update && sudo apt-get install fio -y
```

This applies to Ubuntu 16.04, 18.04, and Ubuntu Minimal 16.04, 18.04.

FIO Commands

IOPS PERFORMANCE TESTS

Use the following FIO example commands to test IOPS performance. You can run the commands directly or create a job file with the command and then run the job file.

Test random reads

Run the following command directly to test random reads:

```
sudo fio --filename=device name --direct=1 --rw=randread --bs=4k --ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based --group_reporting --name=iops-test-job --eta-newline=1 --readonly
```

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomread.fio`, with the following:

```
[global]
bs=4K
iodepth=256
```

```
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-randread
rw=randread

[job1]
filename=device name
```

2. Run the job using the following command:

```
 fio randomread.fio
```

Test file random read/writes

Run the following command against the mount point to test file read/writes:

```
sudo fio --filename=/custom mount point/file --size=500GB --direct=1 --rw=randrw --bs=4k --
ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based --group_reporting --name=iops-test-
job --eta-newline=1
```

Add both the read IOPS and the write IOPS returned.

Test random read/writes



Warning

Do not run FIO tests with a write workload (`readwrite`, `randrw`, `write`, `trimwrite`) directly against a device that is in use.

Run the following command to test random read/writes:

CHAPTER 7 Block Volume

```
sudo fio --filename=device name --direct=1 --rw=randrw --bs=4k --ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based --group_reporting --name=iops-test-job --eta-newline=1
```

Add both the read IOPS and the write IOPS returned.

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomreadwrite.fio`, with the following:

```
[global]
bs=4K
iodepth=256
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-randreadwrite
rw=randrw

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio randomreadwrite.fio
```

Test sequential reads

For workloads that enable you to take advantage of sequential access patterns, such as database workloads, you can confirm performance for this pattern by testing sequential reads.

Run the following command to test sequential reads:

```
sudo fio --filename=device name --direct=1 --rw=read --bs=4k --ioengine=libaio --iodepth=256 --runtime=120 --numjobs=4 --time_based --group_reporting --name=iops-test-job --eta-newline=1 --readonly
```

In some cases you may see more consistent results if you use a job file instead of running the command directly. Use the following instructions for this approach:

1. Create a job file, `fioread.fio`, with the following:

```
[global]
bs=4K
iodepth=256
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-read
rw=read

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio read.fio
```

THROUGHPUT PERFORMANCE TESTS

Use the following FIO example commands to test throughput performance.

Test random reads

Run the following command to test random reads:

```
sudo fio --filename=device name --direct=1 --rw=randread --bs=64k --ioengine=libaio --iodepth=64 --
runtime=120 --numjobs=4 --time_based --group_reporting --name=throughput-test-job --eta-newline=1 --
readonly
```

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomread.fio`, with the following:

```
[global]
bs=64K
iodepth=64
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-randread
rw=randread

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio randomread.fio
```

Test file random read/writes

Run the following command against the mount point to test file read/writes:

```
sudo fio --filename=/custom mount point/file --size=500GB --direct=1 --rw=randrw --bs=64k --
ioengine=libaio --iodepth=64 --runtime=120 --numjobs=4 --time_based --group_reporting --name=throughput-
test-job --eta-newline=1
```

Add both the read MBPs and the write MBPs returned.

Test random read/writes



Warning

Do not run FIO tests with a write workload (`readwrite`, `randrw`, `write`, `trimwrite`) directly against a device



that is in use.

Run the following command to test random read/writes:

```
sudo fio --filename=device name --direct=1 --rw=randrw --bs=64k --ioengine=libaio --iodepth=64 --runtime=120 --numjobs=4 --time_based --group_reporting --name=throughput-test-job --eta-newline=1
```

Add both the read MBPs and the write MBPs returned.

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomread.fio`, with the following:

```
[global]
bs=64K
iodepth=64
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-randreadwrite
rw=randrw

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio randomreadwrite.fio
```

Test sequential reads

For workloads that enable you to take advantage of sequential access patterns, such as database workloads, you can confirm performance for this pattern by testing sequential reads.

CHAPTER 7 Block Volume

Run the following command to test sequential reads:

```
sudo fio --filename=device name --direct=1 --rw=read --bs=64k --ioengine=libaio --iodepth=64 --runtime=120 --numjobs=4 --time_based --group_reporting --name=throughput-test-job --eta-newline=1 --readonly
```

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fioread.fio`, with the following:

```
[global]
bs=64K
iodepth=64
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=4
name=raw-read
rw=read

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio read.fio
```

LATENCY PERFORMANCE TESTS

Use the following FIO example commands to test latency performance. You can run the commands directly or create a job file with the command and then run the job file.

Test random reads for latency

Run the following command directly to test random reads for latency:

CHAPTER 7 Block Volume

```
sudo fio --filename=device name --direct=1 --rw=randread --bs=4k --ioengine=libaio --iodepth=1 --
numjobs=1 --time_based --group_reporting --name=readlatency-test-job --runtime=120 --eta-newline=1 --
readonly
```

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomreadlatency.fio`, with the following:

```
[global]
bs=4K
iodepth=1
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=1
name=readlatency-test-job
rw=randread

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio fiorandomreadlatency.fio
```

Test random read/writes for latency



Warning

Do not run FIO tests with a write workload (`readwrite`, `randrw`, `write`, `trimwrite`) directly against a device that is in use.

Run the following command directly to test random read/writes for latency:

CHAPTER 7 Block Volume

```
sudo fio --filename=device name --direct=1 --rw=randrw --bs=4k --ioengine=libaio --iodepth=1 --numjobs=1
--time_based --group_reporting --name=rwlatency-test-job --runtime=120 --eta-newline=1 --readonly
```

In some cases you might see more consistent results if you use a job file instead of running the command directly. Use the following steps for this approach.

1. Create a job file, `fiorandomrwlatency.fio`, with the following:

```
[global]
bs=4K
iodepth=1
direct=1
ioengine=libaio
group_reporting
time_based
runtime=120
numjobs=1
name=rwlatency-test-job
rw=randrw

[job1]
filename=device name
```

2. Run the job using the following command:

```
fio fiorandomrwlatency.fio
```

Block Volume Elastic Performance

The elastic performance feature of the Oracle Cloud Infrastructure Block Volume service allows you to dynamically change the volume performance, along with enabling you to pay for the performance characteristics you require independently from the size of your block volumes and boot volumes.

This feature includes the concept of volume performance units (VPUs). You can purchase more VPUs to allocate more resources to a volume, increasing IOPS/GB and throughput per GB. You also have the flexibility to purchase fewer VPUs, which reduces the performance characteristics for a volume, however it can also provide cost savings. You can also choose not to purchase any VPUs which can provide significant cost savings for volumes that don't require the increased performance characteristics.

For specific pricing details, see [Oracle Storage Cloud Pricing](#).

Elastic Performance Configuration Options

There are three elastic performance configuration options, as described below.

- **Balanced:** This is the default setting for new and existing block and boot volumes. It provides a good balance between performance and cost savings for most workloads, including those that perform random I/O such as boot volumes. This option provides linear performance scaling with 60 IOPS/GB up to 25,000 IOPS per volume. Throughput scales at 480 KBPS/GB up to a maximum of 480 MBPS per volume. With this option you are purchasing 10 VPUs per GB/month.
- **Higher Performance:** Recommended for workloads with the highest I/O requirements, requiring the best possible performance, such as large databases. This option provides the best linear performance scale with 75 IOPS/GB up to a maximum of 35,000 IOPS per volume. Throughput also scales at the highest rate at 600 KBPS/GB up to a maximum of 480 MBPS per volume. With this option you are purchasing 20 VPUs per GB/month.
- **Lower Cost:** Recommended for throughput intensive workloads with large sequential I/O, such as streaming, log processing, and data warehouses. The cost is only the storage cost, there is no additional VPU cost. This option gives you linear scaling 2 IOPS/GB up to a maximum of 3000 IOPS per volume. This option is only available for block volumes, it is not available for boot volumes.

The following table lists the performance characteristics for each elastic performance level.

Performance Level	IOPS/GB	Max IOPS/Volume	Throughput/GB	Max Throughput/Volume	VPUs/GB ¹
Lower Cost	2	3000	240	Up to 480	0
Balanced	60	25,000	480	480	10
Higher Performance	75	35,000	600	480	20

See [Block Volume Performance](#) for additional performance details for the Block Volume service.

1: Volume performance units, see [Oracle Storage Cloud Pricing](#) for specific pricing details.

Configuring Volume Performance

You can configure the volume performance for a block volume when you create a volume, see [Creating a Volume](#). You can also change the volume performance for an existing block volume, see [To change the volume performance for an existing block volume](#).

When you create a Compute instance, the volume performance for the instance's boot volume is set to **Balanced** by default. You can change this setting after the instance has launched, see [To change the volume performance for an existing boot volume](#).

Changing the Performance of a Volume

The Block Volume service's elastic performance feature enables you to dynamically configure the volume performance for block volumes and boot volumes, for more information, see [Block Volume Elastic Performance](#).

REQUIRED IAM SERVICE POLICY

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

LIMITATIONS

- You can only change the elastic performance configuration on three volumes concurrently per tenancy.
- When changing volume performance for boot volumes, you can only select the **Balanced** or **Higher Performance** options.

USING THE CONSOLE

The default elastic performance setting for existing block volumes or when you create a new block volume is **Balanced**. You can change the default setting when you create a new block volume, see [Creating a Volume](#). You can also change the elastic performance setting for an existing block volume using the steps in the following procedure.

To change the volume performance for an existing block volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Click the block volume that you want to change the performance for.
3. Click **Change Performance**.
4. Click the performance option you want to change to.
5. Click **Change Performance**.

When you create an instance, the elastic performance setting for the instance's boot volume is set to **Balanced**. You can change this setting to **Higher Performance** after the instance has been launched.

To change the volume performance for an existing boot volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Click the boot volume that you want to change the performance for.
3. Click **Change Performance**.

4. Click the performance option you want to change to.
5. Click **Change Performance**.

USING THE CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

Use the `volume update` operation or the `boot-volume update` operation with `vpus-per-gb` parameter to update a block volume's elastic performance setting. The `vpus-per-gb` parameter is where you specify the volume performance units (VPUs). VPUs represent the elastic performance settings, with the following allowed values:

- 0: Represents **Lower Cost** setting, applies to block volumes only.
- 10: Represents **Balanced** setting, applies to both block volumes and boot volumes.
- 20: Represents **Higher Performance** setting, applies to both block volumes and boot volumes.

For example:

```
oci bv volume update --volume-id <volume_ID> --vpus-per-gb 20
```

USING THE API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

BLOCK VOLUMES

To update a block volume's elastic performance setting, use the following operation:

- [UpdateVolume](#)

The elastic performance setting is specified in the `vpusPerGB` attribute of [UpdateVolumeDetails](#). Allowed values are 0, 10, and 20.

BOOT VOLUMES

To update a block volume's elastic performance setting, use the following operation:

- [UpdateBootVolume](#)

The elastic performance setting is specified in the `vpusPerGB` attribute of [UpdateBootVolumeDetails](#). Allowed values are 10 and 20.

Block Volume Metrics

You can monitor the health, capacity, and performance of your block volumes and boot volumes by using metrics, alarms, and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_blockstore` (the Block Volume service).

Resources: Block volumes and boot volumes.

See [Using Block Volumes Service Metrics to Calculate Block Volume Throughput and IOPS](#) for a walkthrough of a performance testing scenario with FIO that shows how you can use these metrics to determine the performance characteristics of your block volume.

Overview of Metrics for an Instance and Its Storage Devices

If you're not already familiar with the different types of metrics available for an instance and its storage and network devices, see [Compute Instance Metrics](#).

Available Metrics: `oci_blockstore`

The Block Volume service metrics help you measure volume operations and throughput related to Compute instances.

The metrics listed in the following table are automatically available for any block volume or boot volume, regardless of whether the attached instance has [monitoring enabled](#). You do not need to enable monitoring on the volumes to get these metrics.

You also can use the Monitoring service to create [custom queries](#).

Each metric includes the following dimensions:

CHAPTER 7 Block Volume

ATTACHMENTID

The OCID of the volume attachment.

RESOURCEID

The OCID of the volume.

Metric	Metric Display Name	Unit	Description	Dimensions
VolumeReadThroughput*	Volume Read Throughput	bytes	Read throughput. Expressed as bytes read per interval.	attachmentId resourceId
VolumeWriteThroughput*	Volume Write Throughput	bytes	Write throughput. Expressed as bytes written per interval.	
VolumeReadOps*	Volume Read Operations	reads	Activity level from I/O reads. Expressed as reads per interval.	
VolumeWriteOps*	Volume Write Operations	writes	Activity level from I/O writes. Expressed as writes per interval.	

* The Compute service separately reports network-related metrics *as measured on the instance itself and aggregated across all the attached volumes*. Those metrics are available in the `oci_computeagent` metric namespace. For more information, see [Compute Instance Metrics](#).

Using the Console

To view default metric charts for a single volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance to view its details.
3. Click either **Attached Block Volumes** or **Boot Volume** to view the volume you're interested in.
4. Click the volume to view its details.
5. Under **Resources**, click **Metrics**.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).
For information about notifications for alarms, see [Notifications Overview](#).

To view default metric charts for multiple volumes

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that contains the volumes you're interested in.
3. For **Metric Namespace**, select **oci_blockstore**.
The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).
For information about notifications for alarms, see [Notifications Overview](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line](#)

[Interface.](#)

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

CHAPTER 8 Compute

This chapter explains how to launch, access, rename, and terminate compute instances.

Overview of the Compute Service

Oracle Cloud Infrastructure Compute lets you provision and manage compute hosts, known as instances. You can launch instances as needed to meet your compute and application requirements. After you launch an instance, you can access it securely from your computer, restart it, attach and detach volumes, and terminate it when you're done with it. Any changes made to the instance's local drives are lost when you terminate it. Any saved changes to volumes attached to the instance are retained.

Oracle Cloud Infrastructure offers both bare metal and virtual machine instances:

- **Bare Metal:** A bare metal compute instance gives you dedicated physical server access for highest performance and strong isolation.
- **Virtual Machine:** A virtual machine (VM) is an independent computing environment that runs on top of physical bare metal hardware. The virtualization makes it possible to run multiple VMs that are isolated from each other. VMs are ideal for running applications that do not require the performance and resources (CPU, memory, network bandwidth, storage) of an entire physical machine.

An Oracle Cloud Infrastructure VM compute instance runs on the same hardware as a bare metal instance, leveraging the same cloud-optimized hardware, firmware, software stack, and networking infrastructure.

Be sure to review [Best Practices for Your Compute Instance](#) for important information about working with your Oracle Cloud Infrastructure Compute instance.

Oracle Cloud Infrastructure uses [Oracle Ksplice](#) to apply important security and other critical kernel updates to the hypervisor hosts without a reboot.

Compute is Always Free eligible. For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).

Instance Types

When you create a Compute instance, you can select the most appropriate type of instance for your applications based on characteristics such as the number of CPUs, amount of memory, and network resources. Oracle Cloud Infrastructure offers a variety of shapes that are designed to meet a range of compute and application requirements:

- **Standard shapes:** Designed for general purpose workloads and suitable for a wide range of applications and use cases. Standard shapes provide a balance of cores, memory, and network resources. Standard shapes are available with Intel or AMD processors.
- **DenseIO shapes:** Designed for large databases, big data workloads, and applications that require high-performance local storage. DenseIO shapes include locally-attached NVMe-based SSDs.
- **GPU shapes:** Designed for hardware-accelerated workloads. GPU shapes include Intel CPUs and NVIDIA graphics processors.
- **High performance computing (HPC) shapes:** Designed for high-performance computing workloads that require high frequency processor cores and cluster networking for massively parallel HPC workloads. HPC shapes are available for bare metal instances only.

For more information about the available bare metal and VM shapes, see [Compute Shapes](#), [Bare Metal Instance Types](#), [Virtual Machine Shapes](#), and [Accelerated GPU Instance Types](#).

Components for Launching Instances

The components required to launch an instance are:

AVAILABILITY DOMAIN

The Oracle Cloud Infrastructure data center within your geographical region that hosts cloud resources, including your instances. You can place instances in the same or different availability domains, depending on your performance and redundancy requirements. For more information, see [Regions and Availability Domains](#).

VIRTUAL CLOUD NETWORK

A virtual version of a traditional network—including subnets, route tables, and gateways—on which your instance runs. At least one cloud network has to be set up before you launch instances. For information about setting up cloud networks, see [Overview of Networking](#).

KEY PAIR (FOR LINUX INSTANCES)

A security mechanism required for Secure Shell (SSH) access to an instance. Before you launch an instance, you'll need at least one key pair. For more information, see [Managing Key Pairs on Linux Instances](#).

TAGS

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

PASSWORD (FOR WINDOWS INSTANCES)

A security mechanism required to access an instance that uses an Oracle-provided Windows image. The first time you launch an instance using a Windows image, Oracle Cloud Infrastructure will generate an initial, one-time password that you can retrieve using the console or API. This password must be changed after you initially log on.

IMAGE

A template of a virtual hard drive that determines the operating system and other software for an instance. For details about Oracle Cloud Infrastructure platform images, see [Oracle-Provided Images](#). You can also launch instances from:

- Trusted third-party images published by Oracle partners from the Partner Image catalog. For more information about partner images, see [Overview of Marketplace](#) and [Working with Listings](#).
- Pre-built Oracle enterprise images and solutions enabled for Oracle Cloud Infrastructure

- [Custom images](#), including [bring your own image scenarios](#).
- [Boot Volumes](#).

SHAPE

A template that determines the number of CPUs, amount of memory, and other resources allocated to a newly created instance. You choose the most appropriate shape when you launch an instance. See [Compute Shapes](#) for a list of available bare metal and VM shapes.

You can optionally attach volumes to an instance. For more information, see [Overview of Block Volume](#).

Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

The following Compute resources emit events:

- Autoscaling configurations and autoscaling policies
- Cluster networks
- Console histories
- Images
- Instances and instance attachments
- Instance configurations
- Instance console connections
- Instance pools

Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

Work Requests

Compute is one of the Oracle Cloud Infrastructure services that is integrated with the Work Requests API. For general information on using work requests in Oracle Cloud Infrastructure, see [Work Requests](#) in the user guide, and the [Work Requests API](#).

Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Limits on Compute Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

Additional limits include:

- To attach a volume to an instance, both the instance and volume must be within the same availability domain.
- Many Compute operations are subject to [throttling](#).

Metadata Key Limits

Custom metadata keys (any key you define that is not `ssh_authorized_keys` or `user_data`) have the following limits:

- Max number of metadata keys: 128
- Max size of key name: 255 characters
- Max size of key value: 255 characters

`ssh_authorized_keys` is a special key that does not have these limits, but its value is validated to conform to a public key in the OpenSSH format.

`user_data` has a maximum size of 16KB. For Linux instances with cloud-init configured, you can populate the `user_data` field with a Base64-encoded string of cloud-init user data. For more information on formats that cloud-init accepts, see [cloud-init formats](#). On Windows instances, the `user_data` field can be provided but isn't used by Oracle-provided images.

Best Practices for Your Compute Instance

Oracle Cloud Infrastructure Compute provides bare metal compute capacity that delivers performance, flexibility, and control without compromise. It is powered by Oracle's next generation, internet-scale infrastructure designed to help you develop and run your most demanding applications and workloads in the cloud.

You can provision compute capacity through an easy-to-use web console or an API. The bare metal compute instance, once provisioned, provides you with access to the host. This gives you complete control of your instance.

While you have full management authority for your instance, Oracle recommends a variety of best practices to ensure system availability and top performance.

IP Addresses Reserved for Use by Oracle

Certain IP addresses are reserved for Oracle Cloud Infrastructure use and may not be used in your address numbering scheme.

169.254.0.0/16

These addresses are used for iSCSI connections to the boot and block volumes, instance metadata, and other services.

Three IP Addresses in Each Subnet

These addresses consist of:

- The first IP address in the CIDR (the network address)
- The last IP address in the CIDR (the broadcast address)
- The first host address in the CIDR (the subnet default gateway address)

For example, in a subnet with CIDR 192.168.0.0/24, these addresses are reserved:

- 192.168.0.0 (the network address)
- 192.168.0.255 (the broadcast address)
- 192.168.0.1 (the subnet default gateway address)

The remaining addresses in the CIDR (192.168.0.2 to 192.168.0.254) are available for use.

Essential Firewall Rules



Warning

Windows 2008 Server R2 images do not support restricting certain firewall rules for local principals, such as "Administrators", so any authenticated user on an instance can make outgoing connections to the iSCSI network endpoints (169.254.0.2:3260, 169.254.2.0/24:3260) that serve the instance's boot and block volumes.

All Oracle-provided images include rules that allow only "root" on Linux instances or "Administrators" on Windows Server 2012 R2 and Windows Server 2016 instances to make outgoing connections to the iSCSI network endpoints (169.254.0.2:3260, 169.254.2.0/24:3260) that serve the instance's boot and block volumes.

- We recommend that you do not reconfigure the firewall on your instance to remove these rules. Removing these rules allows non-root users or non-administrators to access the instance's boot disk volume.
- We recommend that you do not create custom images without these rules unless you understand the security risks.
- Running Uncomplicated Firewall (UFW) on Ubuntu images might cause issues with these rules. Because of this, we recommend that you do not enable UFW on your instances. See [Ubuntu Instance fails to reboot after enabling Uncomplicated Firewall \(UFW\)](#) for more information.

System Resilience

Oracle Cloud Infrastructure runs on Oracle's high-quality Sun servers. However, any hardware can experience a failure. Follow industry-wide hardware failure best practices to ensure the resilience of your solution. Some best practices include:

- Design your system with redundant compute nodes in different availability domains to support failover capability.
- Create a [custom image](#) of your system drive each time you change the image.
- [Back up](#) your data drives, or sync to spare drives, regularly.

If you experience a hardware failure and have followed these practices, you can terminate the failed instance, launch your custom image to create a new instance, and then apply the backup data.

Uninterrupted Access to the Instance

Make sure to keep the DHCP client running so you can always access the instance. If you stop the DHCP client manually or disable NetworkManager (which stops the DHCP client on Linux instances), the instance can't renew its DHCP lease and will become inaccessible when the lease expires (typically within 24 hours). Do not disable NetworkManager unless you use another method to ensure renewal of the lease.

Stopping the DHCP client might remove the host route table when the lease expires. Also, loss of network connectivity to your iSCSI connections might result in loss of the boot drive.

User Access

If you created your instance using an Oracle-provided Linux image, you can use SSH to access your instance from a remote host as the `opc` user. After logging in, you can add users on your instance.

If you do not want to share SSH keys, you can [create additional SSH-enabled users](#).

If you created your instance using an Oracle-provided Windows image, you can access your instance using a Remote Desktop client as the `opc` user. After logging in, you can add users on your instance.

For more information about user access, see [Adding Users on an Instance](#).

NTP Service

Oracle Cloud Infrastructure offers a fully managed, secure, and highly available NTP service that you can use to set the date and time of your Compute and Database instances from within your virtual cloud network (VCN). Oracle recommends that you configure your instances to use the Oracle Cloud Infrastructure NTP service. For information about how to configure instances to use this service, see [Configuring the Oracle Cloud Infrastructure NTP Service for an Instance](#).

Fault Domains

A fault domain is a grouping of hardware and infrastructure that is distinct from other fault domains in the same availability domain. Each availability domain has three fault domains. By properly leveraging fault domains you can increase the availability of applications running on Oracle Cloud Infrastructure. See [Fault Domains](#) for more information.

Your application's architecture will determine whether you should separate or group instances using fault domains.

Scenario 1: Highly Available Application Architecture

In this scenario you have a highly available application, for example you have two web servers and a clustered database. In this scenario you should group one web server and one database node in one fault domain and the other half of each pair in another fault domain. This ensures that a failure of any one fault domain does not result in an outage for your application.

Scenario 2: Single Web Server and Database Instance Architecture

In this scenario your application architecture is not highly available, for example you have one web server and one database instance. In this scenario both the web server and the database instance must be placed in the same fault domain. This ensures that your application will only be impacted by the failure of that single fault domain.

Customer-Managed Virtual Machine (VM) Maintenance

When an underlying infrastructure component needs to undergo maintenance, we notify you in advance of the planned maintenance downtime. To avoid this planned downtime, you have the options to reboot, or stop and restart your instances prior to the scheduled maintenance. This makes it easy for you to control your instance downtime during the notification period. The reboot, or stop and restart of a VM instance during the notification period is different than a normal reboot. The reboot, or stop and start workflow will stop your instance on the existing VM host that needs maintenance and restart it on a healthy VM host. If you choose not to reboot during the notification period, then Oracle Cloud Infrastructure will reboot your VM instance before proceeding with the planned infrastructure maintenance. For information on rebooting or restarting your instance prior to planned maintenance, see [Rebooting Your Virtual Machine \(VM\) Instance During Planned Maintenance](#).

Configuring the Oracle Cloud Infrastructure NTP Service for an Instance

Oracle Cloud Infrastructure offers a fully managed, secure, and highly available NTP service that you can use to set the date and time of your Compute and Database instances from within your virtual cloud network (VCN). The Oracle Cloud Infrastructure NTP service uses redundant Stratum 1 devices in every availability domain. The Stratum 1 devices are synchronized to dedicated Stratum 2 devices that every host synchronizes against. The service is available in every region.

This topic describes how to configure your Compute instances to use this NTP service.

You can also choose to configure your instances to use a public NTP service or use FastConnect to leverage an on-premises NTP service.

Oracle Linux 6.x

Use the following steps to configure your Oracle Linux 6.x instances to use the Oracle Cloud Infrastructure NTP service.

1. Configure IPTables to allow connections to the Oracle Cloud Infrastructure NTP service, using the following commands:

```
sudo iptables -I BareMetalInstanceServices 8 -d 169.254.169.254/32 -p udp -m udp --dport 123 -m comment --comment "Allow access to OCI local NTP service" -j ACCEPT
```

```
sudo service iptables save
```

2. Install the NTP service with the following command:

```
sudo yum install ntp
```

3. Set the date of your instance with the following command:

```
sudo ntpdate 169.254.169.254
```

4. Configure the instance to use the Oracle Cloud Infrastructure NTP service for iburst. To configure, modify the `/etc/ntp.conf` file as follows:

- a. In the `server` section, comment out the lines specifying the RHEL servers:

```
#server 0.rhel.pool.ntp.org iburst
#server 1.rhel.pool.ntp.org iburst
#server 2.rhel.pool.ntp.org iburst
#server 3.rhel.pool.ntp.org iburst
```

- b. Add an entry for the Oracle Cloud Infrastructure NTP server:

```
server 169.254.169.254 iburst
```

The modified `server` section now contains the following:

```
# Please consider joining the pool (http://www.pool.ntp.org/join.html).
#server 0.rhel.pool.ntp.org iburst
#server 1.rhel.pool.ntp.org iburst
#server 2.rhel.pool.ntp.org iburst
#server 3.rhel.pool.ntp.org iburst
server 169.254.169.254 iburst
```

CHAPTER 8 Compute

5. Set the NTP service to launch automatically when the instance boots with the following command:

```
sudo chkconfig ntpd on
```

6. Start the NTP service with the following command:

```
sudo /etc/init.d/ntpd start
```

7. Confirm that the NTP service is configured correctly with the following command:

```
ntpq -p
```

The output will be similar to the following:

```
remote          refid          st t when poll reach  delay  offset  jitter
=====
169.254.169.254 192.168.32.3  2 u  2   64   1   0.338  0.278  0.187
```

Oracle Linux 7.x



Note

Oracle-provided Oracle 7.x and CentOS 7.x images released after February 2018 include the Chrony service by default, so you do not need to configure the Oracle Cloud Infrastructure NTP service for these instances.

Use the following steps to configure your Oracle Linux 7.x instances to use the Oracle Cloud Infrastructure NTP service.

1. Run commands in this section as root with the following command:

```
sudo su -
```

2. Install the NTP service with the following command:

```
yum -y install ntp
```

3. Change the firewall rules to allow inbound and outbound traffic with the Oracle Cloud Infrastructure NTP server, at 169.254.169.254, on UDP port 123 with the following command:

```
awk -v n=13 -v s=' <passthrough ipv="ipv4">-A OUTPUT -d 169.254.169.254/32 -p udp -m udp --dport 123 -m comment --comment "Allow access to OCI local NTP service" -j ACCEPT </passthrough>' 'NR == n {print s} {print}' /etc/firewalld/direct.xml > tmp && mv tmp /etc/firewalld/direct.xml
```

At the prompt:

```
mv: overwrite '/etc/firewalld/direct.xml'?
```

enter *y*

4. Restart the firewall with the following command:

```
service firewalld restart
```

5. Set the date of your instance with the following command:

```
ntpdate 169.254.169.254
```

6. Configure the instance to use the Oracle Cloud Infrastructure NTP service for iburst. To configure, modify the `/etc/ntp.conf` file as follows:

- a. In the `server` section comment out the lines specifying the RHEL servers:

```
#server 0.rhel.pool.ntp.org iburst
#server 1.rhel.pool.ntp.org iburst
#server 2.rhel.pool.ntp.org iburst
#server 3.rhel.pool.ntp.org iburst
```

- b. Add an entry for the Oracle Cloud Infrastructure NTP service:

```
server 169.254.169.254 iburst
```

The modified `server` section should now contain the following:

```
# Please consider joining the pool (http://www.pool.ntp.org/join.html).
#server 0.rhel.pool.ntp.org iburst
#server 1.rhel.pool.ntp.org iburst
#server 2.rhel.pool.ntp.org iburst
#server 3.rhel.pool.ntp.org iburst
server 169.254.169.254 iburst
```

CHAPTER 8 Compute

7. Start and enable the NTP service with the following commands:

```
systemctl start ntpd
systemctl enable ntpd
```

You also need disable the chrony NTP client to ensure that the NTP service starts automatically after a reboot, using the following commands:

```
systemctl stop chronyd
systemctl disable chronyd
```

8. Confirm that the NTP service is configured correctly with the following command:

```
ntpq -p
```

The output will be similar to the following:

```
remote          refid          st t when poll reach  delay  offset  jitter
=====
169.254.169.254 192.168.32.3  2 u   2   64   1   0.338  0.278  0.187
```

Windows Server 2016, Windows Server 2012 R2 and Windows Server 2008 R2

You can configure your Windows Server instances to use the Oracle Cloud Infrastructure NTP service by running the following commands in [Windows Powershell](#) as Administrator.

Windows 2012 and Windows 2016:

```
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Parameters' -Name 'Type' -Value
NTP -Type String
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name 'AnnounceFlags' -
Value 5 -Type DWord
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer' -Name
'Enabled' -Value 1 -Type DWord
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Parameters' -Name 'NtpServer' -
Value '169.254.169.254,0x9' -Type String
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient' -Name
'SpecialPollInterval' -Value 900 -Type DWord
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name
'MaxPosPhaseCorrection' -Value 1800 -Type DWord
```

CHAPTER 8 Compute

```
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name  
'MaxNegPhaseCorrection' -Value 1800 -Type DWord
```

Windows 2008:

```
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Parameters' -Name 'Type' -Value  
NTP -Type String  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name 'AnnounceFlags' -  
Value 5 -Type DWord  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer' -Name  
'Enabled' -Value 1 -Type DWord  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Parameters' -Name 'NtpServer' -  
Value '169.254.169.254,0x9' -Type String  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient' -Name  
'SpecialPollInterval' -Value 900 -Type DWord  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name  
'MaxPosPhaseCorrection' -Value 1800 -Type DWord  
Set-ItemProperty -Path 'HKLM:\System\CurrentControlSet\Services\W32Time\Config' -Name  
'MaxNegPhaseCorrection' -Value 1800 -Type DWord  
Set-ItemProperty -Path 'HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\DateTime\Servers' -Name 0 -Value  
"169.254.169.254"  
Set-ItemProperty -Path 'HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\DateTime\Servers' -Name "  
(Default)" -Value "0"
```

Steps 1 - 7 below walk you through these registry changes, you can use these steps to manually edit the registry instead of using PowerShell. If you use the PowerShell commands, you can skip steps 1 - 7, and proceed with steps 8 and 9 to complete the process of configuring your Windows instance to use the Oracle Cloud Infrastructure NTP service.

1. Change the server type to NTP:

- a. From Registry Editor, navigate to:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\
```

- b. Click **Type**.

- c. Change the value to NTP and click **OK**.

2. Configure the Windows Time service to enable the `Timeserv_Announce_Yes` and `Reliable_Timeserv_Announce_Auto` flags.

To configure, set the `AnnounceFlags` parameter to 5:

CHAPTER 8 Compute

- a. From Registry Editor, navigate to:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config\`
 - b. Click **AnnounceFlags**.
 - c. Change the value to 5 and click **OK**.
3. Enable the NTP server:
- a. From Registry Editor, navigate to:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpServer\`
 - b. Click **Enabled**.
 - c. Change the value to 1 and click **OK**.
4. Set the time sources:
- a. From Registry Editor, navigate to:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\`
 - b. Click **NtpServer**.
 - c. Change the value to 169.254.169.254,0x9 and click **OK**.
5. Set the poll interval:
- a. From Registry Editor, navigate to:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\TimeProviders\NtpClient\`
 - b. Click **SpecialPollInterval**.
 - c. Set the value to the interval that you want the time service to synchronize on. The value is in seconds. To set it for 15 minutes, set the value to 900, and click **OK**.
6. Set the phase correction limit settings to restrict the time sample boundaries:
- a. From Registry Editor, navigate to:
`HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W32Time\Config\`
 - b. Click **MaxPosPhaseCorrection**.
 - c. Set the value to the maximum time offset in the future for time samples. The

CHAPTER 8 Compute

value is in seconds. To set it for 30 minutes, set the value to 1800 and click **OK**.

- d. Click **MaxNegPhaseCorrection**.
 - e. Set the value to the maximum time offset in the past for time samples. The value is in seconds. To set it for 30 minutes, set the value to 1800 and click **OK**.
7. For Windows 2008 only, you need to add the Oracle NTP service to the list of available servers and then configure the Oracle NTP service to be the active server:
- a. From Registry Editor, navigate to:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\DateTime\Servers\
```

- b. Right-click **Servers**, and select **New, String Value**.
 - c. Set the **Name** to 0.
 - d. Set the **Value data** to 169.254.169.254.
 - e. In the same path, find the string value named (Default) and double-click it to open the **Edit String** dialog.
 - f. Set the **Value data** to 0 and click **OK**.
8. Restart the time service by running the following command from a command prompt:

```
net stop w32time && net start w32time
```

9. Test the connection to the NTP service by running the following command from a command prompt:

```
w32tm /query /peers
```

The output will be similar to the following:

```
#Peer: 1
Peer: 169.254.169.254,0x9
State: Active
Time Remaining: 22.1901786s
Mode: 3 (Client)
Stratum: 0 (unspecified)
PeerPoll Interval: 10 (1024s)
HostPoll Interval: 10 (1024s)
```

After the time specified in the poll interval has elapsed, `State` will change from `Pending` to `Active`.

Protecting Data on NVMe Devices

Some instance shapes in Oracle Cloud Infrastructure include locally attached NVMe devices. These devices provide extremely low latency, high performance block storage that is ideal for big data, OLTP, and any other workload that can benefit from high-performance block storage.

Note that these devices are not protected in any way; they are individual devices locally installed on your instance. Oracle Cloud Infrastructure does not take images, back up, or use RAID or any other methods to protect the data on NVMe devices. It is your responsibility to protect and manage the durability the data on these devices.

Oracle Cloud Infrastructure offers high-performance remote block (iSCSI) LUNs that are redundant and can be backed up using an API call. See [Overview of Block Volume](#) for more information.

See [Compute Shapes](#) for information about which instance types support local NVMe storage.

Finding the NVMe devices on your instance

You can identify the NVMe devices by using the `lsblk` command. The response returns a list. NVMe devices begin with "nvme", as shown in the following example for a BM.DenseIO1.36 instance:

```
[opc@somehost ~]$ lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
sda          8:0    0 46.6G  0 disk
├─sda1       8:1    0  512M  0 part /boot/efi
├─sda2       8:2    0    8G   0 part [SWAP]
└─sda3       8:3    0   38G  0 part /
nvme0n1     259:6   0  2.9T  0 disk
nvme1n1     259:8   0  2.9T  0 disk
nvme2n1     259:0   0  2.9T  0 disk
nvme3n1     259:1   0  2.9T  0 disk
nvme4n1     259:7   0  2.9T  0 disk
nvme5n1     259:4   0  2.9T  0 disk
```

```
nvme6n1 259:5    0 2.9T 0 disk
nvme7n1 259:2    0 2.9T 0 disk
nvme8n1 259:3    0 2.9T 0 disk
[opc@somehost ~]$
```

Failure Modes and How to Protect Against Them

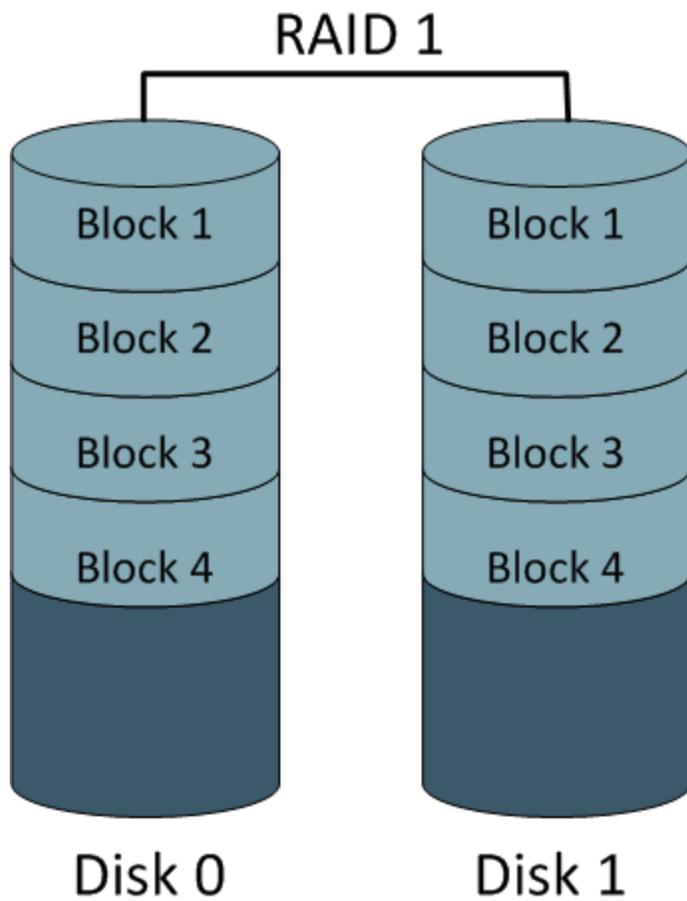
There are three primary failure modes you should plan for:

- [Protecting against the failure of an NVMe device](#)
- [Protecting Against the Loss of the Instance or Availability Domain](#)
- [Protecting Against Data Corruption or Loss from Application or User Error](#)

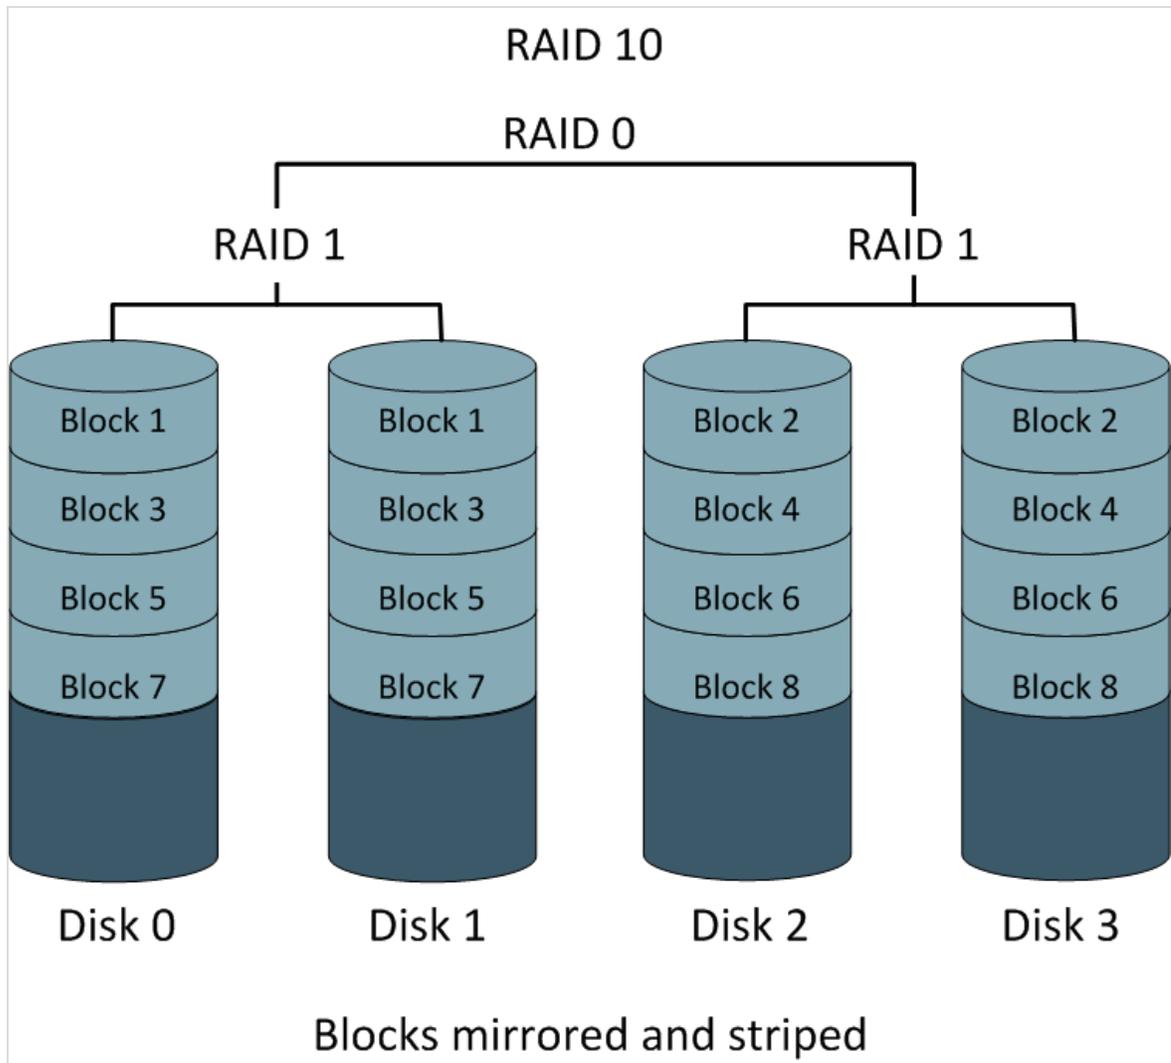
Protecting against the failure of an NVMe device

A protected RAID array is the most recommended way to protect against an NVMe device failure. There are three RAID levels that can be used for the majority of workloads:

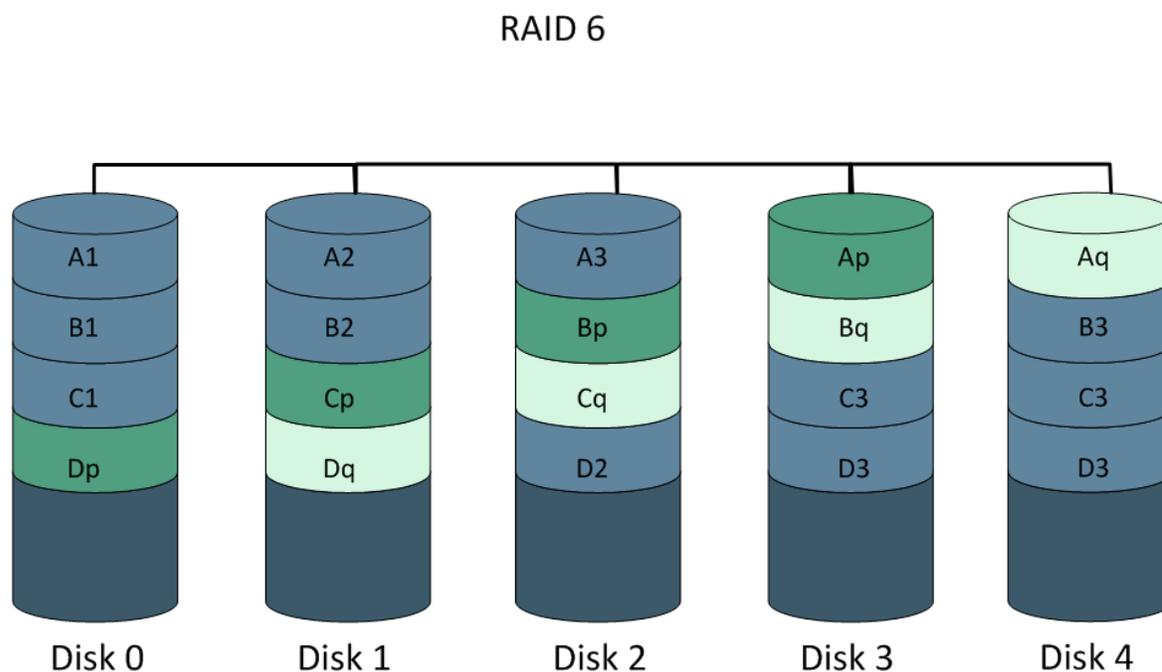
- RAID 1: An exact copy (or mirror) of a set of data on two or more disks; a classic RAID 1 mirrored pair contains two disks, as shown:



- RAID 10: Stripes data across multiple mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved, as shown:



- RAID 6: Block-level striping with two parity blocks distributed across all member disks, as shown.



Blocks striped with dual parity across drives

For more information about RAID and RAID levels, see [RAID](#).

Because the appropriate RAID level is a function of the number of available drives, the number of individual LUNs needed, the amount of space needed, and the performance requirements, there isn't one correct choice. You must understand your workload and design accordingly.

OPTIONS FOR USING A `BM.DenseIO1.36` SHAPE

There are several options for `BM.DenseIO1.36` instances with nine NVMe devices:

Create a single RAID 6 device across all nine devices. This array is redundant, performs well, will survive the failure of any two devices, and will be exposed as a single LUN with about 23.8TB of usable space.

CHAPTER 8 Compute

Use the following commands to create a single RAID 6 device across all nine devices:

```
$ sudo yum install mdadm -y
```

```
$ sudo mdadm --create /dev/md0 --raid-devices=9 --level=6 /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1 /dev/nvme4n1 /dev/nvme5n1 /dev/nvme6n1 /dev/nvme7n1 /dev/nvme8n1
```

```
$ sudo mdadm --detail --scan | sudo tee -a /etc/mdadm.conf >> /dev/null
```

Create a four device RAID 10 and a five device RAID 6 array. These arrays would be exposed as two different LUNs to your applications. This is a recommended choice when you need to isolate one type of I/O from another, such as log and data files. In this example, your RAID 10 array would have about 6.4TB of usable space and the RAID 6 array would have about 9.6TB of usable space.

Use the following commands to create a four-device RAID 10 and a five-device RAID 6 array:

```
$ sudo yum install mdadm -y
```

```
$ sudo mdadm --create /dev/md0 --raid-devices=4 --level=10 /dev/nvme5n1 /dev/nvme6n1 /dev/nvme7n1 /dev/nvme8n1
```

```
$ sudo mdadm --create /dev/md1 --raid-devices=5 --level=6 /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1 /dev/nvme4n1
```

```
$ sudo mdadm --detail --scan | sudo tee -a /etc/mdadm.conf >> /dev/null
```

If you need the best possible performance and can sacrifice some of your available space, then an eight-device RAID 10 array is an option. Because RAID 10 requires an even number of devices, the ninth device is left out of the array and serves as a hot spare in case another device fails. This creates a single LUN with about 12.8 TB of usable space.

Use the following commands to create an eight-device RAID 10 array:

```
$ sudo yum install mdadm -y
```

```
$ sudo mdadm --create /dev/md0 --raid-devices=8 --level=10 /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1 /dev/nvme4n1 /dev/nvme5n1 /dev/nvme6n1 /dev/nvme7n1
```

The following command adds `/dev/nvme8n1` as a hot spare for the `/dev/md0` array:

```
$ sudo mdadm /dev/md0 --add /dev/nvme8n1
```

```
$ sudo mdadm --detail --scan | sudo tee -a /etc/mdadm.conf >> /dev/null
```

CHAPTER 8 Compute

For the best possible performance and I/O isolation across LUNs, create two four-device RAID 10 arrays. Because RAID 10 requires an even number of devices, the ninth device is left out of the arrays and serves as a global hot spare in case another device in either array fails. This creates two LUNS, each with about 6.4 TB of usable space.

Use the following commands to create two four-device RAID 10 arrays with a global hot spare:

```
$ sudo yum install mdadm -y
```

```
$ sudo mdadm --create /dev/md0 --raid-devices=4 --level=10 /dev/nvme4n1 /dev/nvme5n1 /dev/nvme6n1 /dev/nvme7n1
```

```
$ sudo mdadm --create /dev/md1 --raid-devices=4 --level=10 /dev/nvme0n1 /dev/nvme1n1 /dev/nvme2n1 /dev/nvme3n1
```

Creating a global hot spare requires the following two steps:

1. Add the spare to either array (it does not matter which one) by running these commands:

```
$ sudo mdadm /dev/md0 --add /dev/nvme8n1
```

```
$ sudo mdadm --detail --scan | sudo tee -a /etc/mdadm.conf >> /dev/null
```

2. Edit `/etc/mdadm` to put both arrays in the same spare-group. Add `spare-group=global` to the end of the line that starts with `ARRAY`, as follows:

```
$ sudo vi /etc/mdadm.conf
```

```
ARRAY /dev/md0 metadata=1.2 spares=1 name=mdadm.localdomain:0  
UUID=43f93ce6:4a19d07b:51762f1b:250e2327 spare-group=global
```

```
ARRAY /dev/md1 metadata=1.2 name=mdadm.localdomain:1 UUID=7521e51a:83999f00:99459a19:0c836693  
spare-group=global
```

Monitoring Your Array

It's important for you to be notified if a device in one of your arrays fails. Mdadm has built-in tools that can be utilized for monitoring, and there are two options you can use:

CHAPTER 8 Compute

- Set the `MAILADDR` option in `/etc/mdadm.conf` and then run the `mdadm` monitor as a daemon
- Run an external script when `mdadm` detects a failure

SET THE `MAILADDR` OPTION IN `/ETC/MDADM.CONF` AND RUN THE `MDADM` MONITOR AS A DAEMON

The simplest method is to set the `MAILADDR` option in `/etc/mdadm.conf`, and then run the `mdadm` monitor as a daemon, as follows:

1. The `DEVICE` partitions line is required for `MAILADDR` to work; if it is missing, you must add it, as follows:

```
$ sudo vi /etc/mdadm.conf
```

```
DEVICE partitions
```

```
ARRAY /dev/md0 level=raid1 UUID=1b70e34a:2930b5a6:016we78d:eese14532
```

```
MAILADDR <my.name@example.com>
```

2. Run the monitor using the following command:

```
$ sudo nohup mdadm --monitor --scan --daemonize &
```

3. To verify that the monitor runs at startup, run the following commands:

```
$ sudo chmod +x /etc/rc.d/rc.local
```

```
$ sudo vi /etc/rc.local
```

Add the following line to the end of `/etc/rc.local`:

```
nohup mdadm --monitor --scan --daemonize &
```

4. To verify that the email and monitor are both working run the following command:

```
$ sudo mdadm --monitor --scan --test -1
```

Note that these emails will likely be marked as spam. The `PROGRAM` option, described later in this topic, allows for more sophisticated alerting and messaging.

CHAPTER 8 Compute

RUN AN EXTERNAL SCRIPT WHEN A FAILURE IS DETECTED

A more advanced option is to create an external script that would run if the `mdadm` monitor detects a failure. You would integrate this type of script with your existing monitoring solution. The following is an example of this type of script:

```
$ sudo vi /etc/mdadm.events

#!/bin/bash
event=$1
device=$2
if [ $event == "Fail" ]
then
  <"do something">
else
  if [ $event == "FailSpare" ]
  then
    <"do something else">
  else
    if [ $event == "DegradedArray" ]
    then
      <"do something else else">
    else
      if [ $event == "TestMessage" ]
      then
        <"do something else else else">
      fi
    fi
  fi
fi
fi

$ sudo chmod +x /etc/mdadm.events
```

Next, add the `PROGRAM` option to `/etc/mdadm.conf`, as shown in the following example:

1. The `DEVICE` partitions line is required for `MAILADDR` to work; if it is missing, you must add it, as follows:

```
$ sudo vi /etc/mdadm.conf
```

```
DEVICE partitions
```

```
ARRAY /dev/md0 level=raid1 UUID=1b70e34a:2930b5a6:016we78d:eese14532
```

```
MAILADDR <my.name@example.com>
```

```
PROGRAM /etc/mdadm.events
```

2. Run the monitor using the following command:

```
$ sudo nohup mdadm --monitor --scan --daemonize &
```

CHAPTER 8 Compute

3. To verify that the monitor runs at startup, run the following commands:

```
$ sudo chmod +x /etc/rc.d/rc.local
```

```
$ sudo vi /etc/rc.local
```

Add the following line to the end of `/etc/rc.local`:

```
nohup mdadm --monitor --scan --daemonize &
```

4. To verify that the email and monitor are both working run the following command:

```
$ sudo mdadm --monitor --scan --test -1
```

Note that these emails will likely be marked as spam. The `PROGRAM` option, described later in this topic, allows for more sophisticated alerting and messaging.

SIMULATE THE FAILURE OF A DEVICE

You can use `mdadm` to manually cause a failure of a device to see whether your RAID array can survive the failure, as well as test the alerts you have set up.

1. Mark a device in the array as failed by running the following command:

```
$ sudo mdadm /dev/md0 --fail /dev/nvme0n1
```

2. Recover the device or your array might not be protected. Use the following command:

```
$ sudo mdadm /dev/md0 --add /dev/nvme0n1
```

Your array will automatically rebuild in order to use the "new" device. Performance will be decreased during this process.

3. You can monitor the rebuild status by running the following command:

```
$ sudo mdadm --detail /dev/md0
```

What To Do When an NVMe Device Fails

Compute resources in the cloud are designed to be temporary and fungible. If an NVMe device fails while the instance is in service, you should start another instance with the same amount of storage or more, and then copy the data onto the new instance, replacing the old instance. There are multiple toolsets for copying large amounts of data, with [rsync](#) being the most

popular. Since the connectivity between instances is a full 10Gb/sec, copying data should be quick. Remember that with a failed device, your array may no longer be protected, so you should copy the data off of the impacted instance as quickly as possible.

Using the Linux Logical Volume Manager

The Linux [Logical Volume Manager \(LVM\)](#) provides a rich set of features for managing volumes. If you need these features, we strongly recommend that you use `mdadm` as described in preceding sections of this topic to create the RAID arrays, and then use LVM's `pvcreate`, `vgcreate`, and `lvcreate` commands to create volumes on the `mdadm` LUNs. You should not use LVM directly against your NVMe devices.

Protecting Against the Loss of the Instance or Availability Domain

Once your data is protected against the loss of a NVMe device, you need to protect it against the loss of an instance or the loss of the availability domain. This type of protection is typically done by replicating your data to another availability domain or backing up your data to another location. The method you choose depends on your objectives. For details, see [Recovery Time Objective](#) (RTO) and [Recovery Point Objective](#) (RPO).

REPLICATION

Replicating your data from one instance in one availability domain to another has the lowest RTO and RPO at a significantly higher cost than backups; for every instance in one availability domain, you must have another instance in a different availability domain.

For Oracle database workloads, you should use the built-in Oracle Data Guard functionality to replicate your databases. Oracle Cloud Infrastructure availability domains are each close enough to each other to support high performance, synchronous replication. Asynchronous replication is also an option.

For general-purpose block replication, [DRBD](#) is the recommended option. You can configure DRBD to replicate, synchronously or asynchronously, every write in one availability domain to another availability domain.

BACKUPS

Traditional backups are another way to protect data. All commercial backup products are fully supported on Oracle Cloud Infrastructure. If you use backups, the RTO and RPO are significantly higher than using replication because you must recreate the compute resources that failed and then restore the most recent backup. Costs are significantly lower because you don't need to maintain a second instance. Do not store your backups in the same availability domain as their original instance.

Protecting Against Data Corruption or Loss from Application or User Error

The two recommended ways of protecting against data corruption or loss from application or user error are regularly taking snapshots or creating backups.

SNAPSHOTS

The two easiest ways to maintain snapshots are to either use a file system that supports snapshots, such as ZFS, or use LVM to create and manage the snapshots. Because of the way LVM has implemented [copy-on-write \(COW\)](#), performance may significantly decrease when a snapshot is taken using LVM.

BACKUPS

All commercial backup products are fully supported on Oracle Cloud Infrastructure. Make sure that your backups are not stored in the same availability domain as their original instance.

Boot Volumes

When you launch a virtual machine (VM) or bare metal instance based on an Oracle-provided image or custom image, a new boot volume for the instance is created in the same compartment. That boot volume is associated with that instance until you terminate the instance. When you terminate the instance, you can preserve the boot volume and its data. For more information, see [Terminating an Instance](#). This feature gives you more control and management options for your compute instance boot volumes, and enables:

- **Instance scaling:** When you terminate your instance, you can keep the associated boot volume and use it to launch a new instance using a different instance type or shape. See [Creating an Instance](#) for steps to launch an instance based on a boot volume. This allows you to switch easily from a bare metal instance to a VM instance and vice versa, or scale up or down the number of cores for an instance.
- **Troubleshooting and repair:** If you think a boot volume issue is causing a compute instance problem, you can stop the instance and detach the boot volume. Then you can attach it to another instance as a data volume to troubleshoot it. After resolving the issue, you can then reattach it to the original instance or use it to launch a new instance.

Boot volumes are encrypted by default, the same as other block storage volumes. For more information, see [Block Volume Encryption](#).



Important

In-transit encryption for boot and block volumes is only available for virtual machine (VM) instances launched from Oracle-provided images, it is not supported on bare metal instances. It is also not supported in most cases for instances launched from custom images imported for "bring your own image" (BYOI) scenarios. To confirm support for certain Linux-based custom images and for more information contact Oracle support, see [Contacting Support](#).

You can group boot volumes with block volumes into the same volume group, making it easy to create a group volume backup or a clone of your entire instance, including both the system disk and storage disks at the same time. See [Volume Groups](#) for more information.

You can move Block Volume resources such as boot volumes and boot volume backups between compartments. For more information, see [Move Block Volume Resources Between Compartments](#).

For more information about the Block Volume service and boot volumes, see the [Block Volume FAQ](#).

Custom Boot Volume Sizes

When you launch an instance, you can specify whether to use the selected image's default boot volume size, or you can specify a custom size up to 32 TB. This capability is available for the following image source options:

- Oracle-provided image
- Custom image
- Image OCID

See [Creating an Instance](#) for more information.

The specified size must be larger than the image's default boot volume size or 50 GB, whichever is higher. After you launch the instance, you can't change the boot volume size.



Important

The minimum size for Windows-based image boot volumes is 256 GB.

If you specify a custom boot volume size, you need to extend the volume to take advantage of the larger size. For steps, see [Extending the Partition for a Boot Volume](#).

Boot Volume Performance

Boot volume performance varies with volume size, see [Block Volume Performance](#) for more information.

The Block Volume service's elastic performance feature enables you to dynamically change the volume performance for boot volumes. Once an instance has been created, you can

change the volume performance of the boot volume to one of the following performance options:

- Balanced
- Higher Performance

For more information about this feature and the performance options, see [Block Volume Elastic Performance](#) and [Changing the Performance of a Volume](#)

Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to list boot volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes, boot volumes, and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

To access the Console, you must use a [supported browser](#).

See the following tasks for managing boot volumes:

- [Listing Boot Volumes](#)
- [Attaching a Boot Volume](#)
- [Detaching a Boot Volume](#)

- [Listing Boot Volume Attachments](#)
- [Deleting a Boot Volume](#)

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage boot volumes:

- [BootVolume](#)
- [ListBootVolumes](#)
- [GetBootVolume](#)
- [UpdateBootVolume](#)
- [DetachBootVolume](#)
- [DeleteVolume](#)
- [BootVolumeAttachment](#)
- [AttachBootVolume](#)
- [GetBootVolumeAttachment](#)
- [ListBootVolumeAttachments](#)

Extending the Partition for a Boot Volume

When you create a new virtual machine (VM) instance or bare metal instance based on an Oracle-provided image or custom image, you have the option of specifying a custom boot volume size. You can also expand the size of the boot volume for an existing instance; see [Resizing a Volume](#) for more information. In order to take advantage of the larger size, you need to extend the partition for the boot volume.

Required IAM Policy

Extending a partition on an instance does not require a specific IAM policy. However, you may need permission to run the necessary commands on the instance's guest OS. Contact your system administrator for more information.

Extending the Root Partition on a Linux-Based Image

For instances running Linux-based images, you need to extend the root partition and then grow the file system.

PREREQUISITES

Before you can extend the partition, you need to detach the boot volume from the instance and attach it to another instance as a data volume. To do this:

1. Stop the instance. For steps, see [Stopping and Starting an Instance](#).
2. After the instance has stopped, detach the boot volume. For steps, see [Detaching a Volume](#).
3. Attach the boot volume to a second instance as a data volume. For steps, see [Attaching a Volume](#) and [Connecting to a Volume](#).

EXTENDING THE LINUX PARTITION

After attaching the boot volume as a data volume to the second instance, connect to this instance and perform the following steps to extend the partition.

Extending the root partition

1. To identify the volume you want to extend the partition for, run the following command to list the attached block volumes:

```
lsblk
```

2. Run the following command to edit the volume's partition table with `parted`:

```
parted <volume_id>
```

`<volume_id>` is the volume identifier, for example `/dev/sdc`.

3. When you run `parted`, you may encounter the following error message:

```
Warning: Not all of the space available to <volume_id> appears to be used,
you can fix the GPT to use all of the space (an extra <volume_size> blocks)
or continue with the current setting?
```

You are then prompted to fix the error or ignore the error and continue with the current setting. Specify the option to fix the error.

4. Run the following command to change the display units to sectors so that you can see the precise start position for the volume:

```
(parted) unit s
```

5. Run the following command to display the current partitions in the partition table:

```
(parted) print
```

Make note of the values in the **Number**, **Start**, and **File system** columns for the root partition.

6. Run the following command to remove the existing root partition:

```
(parted) rm <partition_number>
```

`<partition_number>` is the value from the **Number** column.

7. Run the following command to recreate the partition:

```
(parted) mkpart
```

At the `Start?` prompt, specify the value from the **Start** column. At the `File system type?` prompt, specify the value from the **File system** column. Specify `100%` for the `End?` prompt.

8. Run the following command to exit `parted`:

```
(parted) quit
```

This command forces a rewrite of the partition table with the new partition settings that you specified.

9. To verify that the root partition was extended, run the following command to list the attached block volumes:

```
lsblk
```

After you extend the root partition, you need to grow the file system. The steps in the following procedure apply only to xfs file systems.

Growing the file system for the root partition

1. Before you grow the file system, repair any issues with the file system on the extended partition by running the following command:

```
xfs_repair <partition_id>
```

<partition_id> is the partition identifier, for example `/dev/sdc3`. See [Checking and Repairing an XFS File System](#) for more information.

2. After you have confirmed that there are no more issues to repair, you need to create a mount point to run the `xfs_growfs` against. To do this, create a directory and mount the partition to that directory by running the following commands:

```
mkdir <directory_name>  
mount <partition_id> <directory_name> -o nouuid
```

<partition_id> is the partition identifier, for example `/dev/sdc3`, and *<directory_name>* is the directory to create and mount.

CHAPTER 8 Compute

3. After you have created the mount point run the following command to grow the file system:

```
xfs_growfs -d <directory_name>
```

<directory_name> is the name for the directory you created in the previous step.

4. To verify that the file system size is correct, run the following command to display the file system details:

```
df -lh
```

5. Once you have verified that the file system size is correct, run the following command to unmount the partition:

```
umount <partition_id>
```

NEXT STEPS

After you have extended the partition and grown the file system, you can restart the original instance with the boot volume. To do this:

1. Disconnect the volume from the second instance. For steps, see [Disconnecting From a Volume](#).
2. Detach the volume from the second instance. For steps, see [Detaching a Volume](#).
3. Attach the volume to the original instance as a boot volume. For steps, see [Attaching a Boot Volume](#).
4. Restart the instance. For steps, see [Stopping and Starting an Instance](#).

Extending the System Partition on a Windows-Based Image

On Windows-based images, you can extend a partition using the Windows interface or from the command line using the DISKPART utility.

WINDOWS SERVER 2016 AND WINDOWS SERVER 2012

The steps for extending a system partition on instances running Windows 2012 or Windows 2016 are the same, and are described in the following procedures.

Extending the system partition using the Windows interface

1. Open the [Disk Management](#) system utility on the instance.
2. Right-click the boot volume and select **Extend Volume**.
3. Follow the instructions in the **Extend Volume Wizard**:
 - a. Select the disk that you want to extend, enter the size, and then click **Next**.
 - b. Confirm that the disk and size settings are correct, and then click **Finish**.
4. Verify that the boot volume's system disk has been extended in Disk Management.

Extending the system partition using the command line with DISKPART

1. Open a command prompt as administrator on the instance.
2. Run the following command to start the DISKPART utility:

```
diskpart
```

3. At the DISKPART prompt, run the following command to display the instance's volumes:

```
list volume
```

4. Run the following command to select the boot volume:

```
select volume <volume_number>
```

<volume_number> is the number associated with the boot volume that you want to extend the partition for.

5. Run the following command to extend the partition:

```
extend size=<increased_size_in_MB>
```

<increased_size_in_MB> is the size in MB that you want to extend the partition to.



Warning

When using the DISKPART utility, do not overextend the partition beyond the current available space. Overextending the partition could result in data loss.

6. To confirm that the partition was extended, run the following command and verify that the boot volume's partition has been extended:

```
list volume
```

WINDOWS SERVER 2008

Use the steps described in the following procedures to extend a system partition on instances running Windows 2008.

Extending the system partition using the Windows interface

1. Open the [Server Manager](#) on the instance.
2. Expand the **Storage** node in the left navigation pane and click **Disk Management**.
3. Right-click the boot volume and select **Extend Volume**.
4. Follow the instructions in the **Extend Volume Wizard**:
 - a. Select the disk that you want to extend, enter the size, and then click **Next**.
 - b. Confirm that the disk and size settings are correct, and then click **Finish**.
5. Verify that the boot volume's system disk has been extended in the Server Manager's **Disk Management** node.

Extending the system partition using the command line with DISKPART

1. Open a command prompt as administrator on the instance.
2. Run the following command to start the DISKPART utility:

```
diskpart
```

3. At the DISKPART prompt, run the following command to display the instance's volumes:

```
list volume
```

4. Run the following command to select the boot volume:

```
select volume <volume_number>
```

<volume_number> is the number associated with the boot volume that you want to extend the partition for.

5. Run the following command to extend the partition:

```
extend size=<increased_size_in_MB>
```

<increased_size_in_MB> is the size in MB that you want to extend the partition to.



Warning

When using the DISKPART utility, do not overextend the partition beyond the current available space. Overextending the partition could result in data loss.

6. To confirm that the partition was extended, run the following command and verify that the boot volume's partition has been extended:

```
list volume
```

Attaching a Boot Volume

If a boot volume has been detached from the associated instance, or if the instance is stopped or terminated, you can attach it to another instance as a data volume. The steps are the same as the steps for attaching a block volume, see [Attaching a Volume](#).

You can also reattach a boot volume to the associated instance. If you want to restart an instance with a detached boot volume, you must reattach the boot volume using the steps described in this topic.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to attach/detach existing block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the **Instances** list, select the instance you want to attach the boot volume to.
3. Click the name of the instance to display the instance details.
4. In the **Resources**, click **Boot Volume**.
5. Click the Actions icon (three dots) for the boot volume.

6. Click **Attach** and confirm the selection when prompted.

You can start the instance once the boot volume's icon no longer lists it as **ATTACHING**. For more information, see [Stopping and Starting an Instance](#).

Using the API

To attach a volume to an instance, use the following operation:

- [AttachBootVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Listing Boot Volumes

You can list all boot volumes in a specific compartment, or detailed information on a single boot volume.

Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to list volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Choose your **Compartment**.

A detailed list of volumes in the current compartment is displayed. To see detailed information for a specific volume, click the boot volume name.

The instance associated with the boot volume is listed in the **Attached Instance** field. If the value for this field displays:

```
None in this Compartment.
```

the boot volume has been detached from the associated instance, or the instance has been terminated while the boot volume was preserved.

To view the volumes in a different compartment, change the compartment in the **Compartment** drop-down menu.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

LIST BOOT VOLUMES:

Get a list of boot volumes within a compartment.

- [ListBootVolumes](#)

GET A SINGLE BOOT VOLUME:

Get detailed information on a single boot volume:

- [GetBootVolume](#)

Overview of Boot Volume Backups

The backups feature of the Oracle Cloud Infrastructure Block Volume service lets you make a point-in-time crash-consistent backup of a boot volume without application interruption or downtime. Boot volume backup capabilities are the same as block volume backup capabilities. See [Overview of Block Volume Backups](#) for more information.

There are two ways you can initiate a boot volume backup, the same as block volume backups. You can either manually start the backup, or assign a policy which defines a set backup schedule. See [Manual Backups](#) and [Policy-Based Backups](#) for more information.

Boot Volume Backup Types

The Block Volume service supports the same backups types for boot volumes as for block volumes:

- **Incremental:** This backup type includes only the changes since the last backup.
- **Full:** This backup type includes all changes since the volume was created.

Backing Up a Boot Volume

You can create boot volume backups using the Console or the REST APIs/command line interface (CLI). See [Backing Up a Boot Volume](#) and the [BootVolumeBackup](#) API for more information.

Restoring a Boot Volume

Before you can use a boot volume backup, you need to restore it. For steps, see [Restoring a Boot Volume](#).

Making a boot volume backup while an instance is running creates a crash-consistent backup, meaning the data is in the identical state it was in at the time the backup was made. This is the same state it would be in the case of a loss of power or hard crash. In most cases, you can restore a boot volume backup and use it to create an instance. Alternatively you can attach it to an instance as a data volume to repair it or recover data, see [Attaching a Volume](#). To

ensure a bootable image, you should create a custom image from your instance. For information about creating custom images, see [Managing Custom Images](#).

Differences Between Boot Volume Backups and Clones

Consider the following criteria when you decide whether to create a backup or a clone of a volume.

	Volume Backup	Volume Clone
Description	Creates a point-in-time backup of data on a volume. You can restore multiple new volumes from the backup later in the future.	Creates a single point-in-time copy of a volume without having to go through the backup and restore process.
Use case	<p>Retain a backup of the data in a volume, so that you can duplicate an environment later or preserve the data for future use.</p> <p>Meet compliance and regulatory requirements, because the data in a backup remains unchanged over time.</p> <p>Support business continuity requirements.</p> <p>Reduce the risk of outages or data mutation over time.</p>	Rapidly duplicate an existing environment. For example, you can use a clone to test configuration changes without impacting your production environment.
Speed	Slower (minutes or hours)	Faster (seconds)
Cost	Lower cost	Higher cost
Storage location	Object Storage	Block Volume

	Volume Backup	Volume Clone
Retention policy	Policy-based backups expire, manual backups do not expire	No expiration
Volume groups	Supported. You can back up a volume group.	Supported. You can clone a volume group.

Backing Up a Boot Volume

You can create a backup of a boot volume using the Oracle Cloud Infrastructure Block Volume service. This topic describes how to create a manual boot volume backup.

You can also configure a backup policy that creates backups automatically based on a specified schedule and retention policy. This works the same as block volumes. See [Policy-Based Backups](#) for more information.

For information to help you decide whether to create a backup or a clone of a boot volume, see [Differences Between Boot Volume Backups and Clones](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.



Tip

When users create a backup from a volume or restore a volume from a backup, the volume and backup don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Click the boot volume that you want to create a backup for.
3. Click **Create Manual Backup**.
4. Enter a name for the backup.
5. Select the backup type, either incremental or full. See [Boot Volume Backup Types](#) for information about backup types.
6. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Backup**.
The backup is completed when its icon no longer lists it as **CREATING** in the **Boot Volume Backup** list.

Using the API

To back up a boot volume, use the following operation:

- [CreateBootVolumeBackup](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about backups, see [Overview of Block Volume Backups](#) and [Restoring a Backup to a New Volume](#).

Cloning a Boot Volume

You can create a clone from a boot volume using the Oracle Cloud Infrastructure Block Volume service. Cloning enables you to make a copy of an existing boot volume without needing to go through the backup and restore process. For more information about the Block Volume service, see [Overview of Block Volume](#) and the [Block Volume FAQ](#).

A boot volume clone is a point-in-time direct disk-to-disk deep copy of the source boot volume, so all the data that is in the source boot volume when the clone is created is copied to the boot volume clone. Any subsequent changes to the data on the source boot volume are not copied to the boot volume clone. Since the clone is a copy of the source boot volume it will be the same size as the source boot volume unless you specify a larger volume size when you create the clone.

The clone operation occurs immediately and you can use the cloned boot volume as soon as the state changes to available.

There is a single point-in-time reference for a source boot volume while it is being cloned, so if you clone a boot volume while the associated instance is running, you need to wait for the first clone operation to complete from the source before creating additional clones. You also need to wait for any backup operations to complete as well.

You can only create a clone for a boot volume within the same region, availability domain, and tenant. You can create a clone for a boot volume between compartments as long as you have the required access permissions for the operation.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Differences Between Boot Volume Backups and Clones

Consider the following criteria when you decide whether to create a backup or a clone of a volume.

	Volume Backup	Volume Clone
Description	Creates a point-in-time backup of data on a volume. You can restore multiple new volumes from the backup later in the future.	Creates a single point-in-time copy of a volume without having to go through the backup and restore process.
Use case	<p>Retain a backup of the data in a volume, so that you can duplicate an environment later or preserve the data for future use.</p> <p>Meet compliance and regulatory requirements, because the data in a backup remains unchanged over time.</p> <p>Support business continuity requirements.</p> <p>Reduce the risk of outages or data mutation over time.</p>	<p>Rapidly duplicate an existing environment. For example, you can use a clone to test configuration changes without impacting your production environment.</p>

	Volume Backup	Volume Clone
Speed	Slower (minutes or hours)	Faster (seconds)
Cost	Lower cost	Higher cost
Storage location	Object Storage	Block Volume
Retention policy	Policy-based backups expire, manual backups do not expire	No expiration
Volume groups	Supported. You can back up a volume group.	Supported. You can clone a volume group.

For more information about boot volume backups, see [Overview of Boot Volume Backups](#) and [Backing Up a Boot Volume](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. In the **Boot Volumes** list, click the boot volume that you want to clone.
3. In **Resources**, click **Boot Volume Clones**.
4. Click **Create Clone**.
5. Specify a name for the clone.
6. If you want to clone the boot volume to a larger size volume, select **Custom Boot Volume Size (GB)** and then specify the new size. You can only increase the size of the volume, you cannot decrease the size. If you clone the boot volume to a larger size volume, you need to extend the volume's partition. See [Extending the Partition for a Boot Volume](#) for more information.
7. Click **Create Clone**.

The boot volume is ready use when its icon lists it as **AVAILABLE** in the **Boot Volumes** list.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To create a clone from a boot volume, use the [CreateBootVolume](#) operation and specify [BootVolumeSourceFromBootVolumeDetails](#) for [CreateBootVolumeDetails](#).

Next Steps

After you have cloned a boot volume backup, you can:

- Use the boot volume to create an instance. For more information, see [Creating an Instance](#).
- Attach the boot volume to an instance as a data volume. For more information, see [Attaching a Volume](#).

Making a boot volume clone while an instance is running creates a crash-consistent clone, meaning the data is in the identical state it was in at the time the clone was made. This is the same state it would be in the case of a loss of power or hard crash. In most cases you can use the cloned boot volume to create an instance, however to ensure a bootable image, you should create a custom image from your instance. For information about creating custom images, see [Managing Custom Images](#).

Detaching a Boot Volume

If you think a boot volume issue is causing a compute instance problem, you can stop the instance and detach the boot volume using the steps described in this topic. Then you can attach it to another instance as a data volume to troubleshoot it.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have

CHAPTER 8 Compute

permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to attach/detach existing block volumes. The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups, but not launch instances.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

You can only detach a boot volume from an instance when the instance is stopped. See [Stopping and Starting an Instance](#) for more information about managing an instance's state.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Choose your **Compartment**.
3. In the **Instances** list, select the instance you want to detach the boot volume from.
4. Click the name of the instance to display the instance details.
5. In the **Resources**, click **Boot Volume**.
6. Click the Actions icon (three dots), for the boot volume.
7. Click **Detach** and confirm the selection when prompted.

You can now attach the boot volume to another instance, for more information see [Attaching a Volume](#).

Using the API

To delete an attachment, use the following operation:

- [DetachBootVolume](#)

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Deleting a Boot Volume

When you terminate an instance, you choose to delete or preserve the associated boot volume. For more information, see [Terminating an Instance](#). You can also delete a boot volume if it has been detached from the associated instance. See [Detaching a Boot Volume](#) for how to detach a boot volume.



Warning

You cannot undo this operation. Any data on a volume will be permanently deleted once the volume is deleted. You will also not be able to restart the associated instance.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let volume admins manage block volumes, backups, and volume groups](#) lets the specified group do everything with block volumes and backups.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Choose your **Compartment**.
3. In the **Boot Volumes** list, find the volume you want to delete.
4. Click the Actions icon (three dots) for the boot volume.
5. Click **Terminate** and confirm the selection when prompted.

Using the API

Use the [DeleteBootVolume](#) operation to delete a boot volume.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Boot Volume Metrics

You can monitor the health, capacity, and performance of your Compute instances by using [metrics](#), [alarms](#), and [notifications](#).

The Block Volume service provides a set of metrics that apply to both boot volumes and block volumes. For more information, see [Block Volume Metrics](#).

Recovering a Corrupted Boot Volume for Linux-Based Instances

If your instance fails to boot successfully or boots with the boot volume set to read-only access, the instance's boot volume may be corrupted. While it is rare, boot volume corruption can occur in the following scenarios:

- When an instance experiences a forced shutdown using the API.
- When an instance experiences a system hang due to an operating system or software error and a graceful reboot or shutdown of the instance times out, and then a forced

shutdown occurs.

- When an error or outage occurs in the underlying infrastructure and there were critical disk writes pending in the system.



Important

In most cases a simple reboot will resolve boot volume corruption issues, so this is the first action you should take when troubleshooting this.

This topic describes how to determine if your Linux-based instance's boot volume is corrupted and what steps to take to troubleshoot and recover the corrupted boot volume. For Windows instances, see [Recovering a Corrupted Boot Volume for Windows Instances](#).

Detecting Boot Volume Corruption

Boot volume corruption can prevent an instance from booting successfully, so you may not be able to connect to the instance using SSH. Instead, you can use the instance console connection feature to connect to the malfunctioning instance. For more information about using this feature, see [Instance Console Connections](#).

This section describes how to use a serial console connection to detect if boot volume corruption has occurred.



Tip

If you have already confirmed your instance's boot volume is corrupted or if you are using an imported custom image, proceed to the [Recovering the Boot Volume](#) section, which describes how to use a second instance along with standard file system tools to both detect and repair boot volume corruption.

1. Create a serial console connection for the instance.

To create the serial console connection for an instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the list of instances, find the instance you think may have a corrupted boot volume, and then click the instance name.
- c. In the **Resources** section on the **Instance Details** page, click **Console Connections**, and then click **Create Console Connection**.
- d. Specify the public key portion for the SSH key, either by browsing and selecting a public key file, for example `id_rsa.pub`, or by pasting your public key into the text box, and then click **Create Console Connection**.

2. Connect to the instance through serial console.

To connect to the serial console for an instance using OpenSSH on Mac OS X or Linux

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
- b. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
- c. Click the Actions icon (three dots), and then click **Connect with SSH**.
- d. Select **LINUX/MAC OS** for **PLATFORM**.
- e. Click **Copy** to copy the string to the clipboard.
- f. Paste the connection string copied from the previous step to a terminal window on a Mac OS X or Linux system, and press Enter to connect to the console.
If you are not using the default SSH key or `ssh-agent`, you can modify the serial console connection string to include the identity file flag, `-i` to specify the SSH

key to use. You need to specify this for both the SSH connection and the SSH ProxyCommand, as shown in the following line:

```
ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W %h:%p -p 443...
```

- g. Press Enter again to activate the console.

At this point, it's normal for the serial console to appear to hang, as the system may have already crashed.

3. Reboot the instance from the Console, on the **Instances Details page**, click **Reboot**.
4. Once the reboot process starts, switch back to the terminal window, and you should see system messages from the instance start to appear in the window.
5. Monitor the messages that appear as the system is starting up. Most operating systems will set the boot volume to read-only as soon as disk corruption is detected to prevent writes from further corrupting the volume, so look for messages that indicate the boot volume is in read-only mode. Following are some examples:

- On an instance with iSCSI-attached boot volumes, the `iscsiadm` service will fail to attach a volume because the volume is in read-only mode. This will typically prevent instances from continuing to boot. The serial console may display a message similar to the following:

```
iscsiadm: Maybe you are not root?
iscsiadm: Could not lock discovery DB: /var/lock/iscsi/lock.write: Read-only file system
touch: cannot touch `/var/lock/subsys/iscsid': Read-only file system
touch: cannot touch `/var/lock/subsys/iscsi': Read-only file system
```

- On an instance with paravirtualized-attached boot volumes, the system may continue the boot process, but will be in a degraded state because nothing can be written to the boot drive. The serial console may display error messages similar to the following:

```
[FAILED] Failed to start Create Volatile Files and Directories.
See 'systemctl status systemd-tmpfiles-setup.service' for details.
...
[ 27.160070] cloud-init[819]: os.chmod(path, real_mode)
[ 27.166027] cloud-init[819]: OSError: [Errno 30] Read-only file system:
'/var/lib/cloud/data'
```

The error messages and system behavior described here are the most commonly seen for boot volume corruption, however depending on the operating system, you may see different error messages and system behavior. If you don't see the ones described here, consult the documentation for your operating system for additional troubleshooting information.

Recovering the Boot Volume

To troubleshoot and recover the corrupted boot volume, you need to detach the boot volume from the instance and then attach the boot volume to a second instance as a data volume.

DETACHING THE BOOT VOLUME

If you have detected that your instance's boot volume is corrupted, you need to detach the boot volume from the instance before you can begin troubleshooting and recovery steps.

1. Stop the instance. For more information, see [Stopping and Starting an Instance](#).

To stop the instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
 - b. In the list of instances, find the instance you have detected boot volume corruption for and then click the instance name to display the instance details.
 - c. Click **Stop**.
2. Detach the boot volume from the instance. For more information, see [Detaching a Boot Volume](#).

To detach the boot volume from the instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the **Instances** list, select the instance you want to detach the boot volume from.

- c. Click the name of the instance to display the instance details.
- d. In the **Resources**, click **Boot Volume**.
- e. Click the Actions icon (three dots), for the boot volume.
- f. Click **Detach** and confirm the selection when prompted.

ATTACHING THE BOOT VOLUME AS A DATA VOLUME TO A SECOND INSTANCE

For the second instance we recommend that you use an instance running an operating system that most closely matches the operating system for the boot volume's instance. You should only attach boot volumes for Linux-based instances to other Linux-based instances. The second instance must be in the same availability domain and region as the boot volume's instance. If no existing instance is available, create a new Linux instance using the steps described in [Creating an Instance](#). Once you have the second instance, make sure you can log into the instance and that it is functional before proceeding with the recovery steps. For steps to access the instance, see [Connecting to a Linux Instance](#). After you have confirmed that the instance is functional, perform the following steps.

1. Run the `lsblk` command and make note of the drives that are currently on the instance, for example:

```
lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sda   8:0    0 46.6G  0 disk
├─sda2 8:2    0   8G  0 part [SWAP]
├─sda3 8:3    0 38.4G  0 part /
└─sda1 8:1    0 200M  0 part /boot/efi
```

2. Attach the boot volume to the second instance as a data volume. For more information, see [Attaching a Volume](#).

To attach the boot volume as a data volume

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the **Instances** list, click the instance that you want to attach a volume to.

- c. In the **Resources** section, click **Attached Block Volumes**.
- d. Click **Attach Block Volume**.
- e. Select the volume attachment type. If **Paravirtualized** attachments are available for this instance, we recommend that you select this attachment type for this procedure.
If you select **iSCSI** as the volume attachment type, you need to connect to the volume, see [Connecting to a Volume](#) for more information.
- f. In the **Block Volume Compartment** drop-down list, select the compartment.
- g. Choose the **SELECT VOLUME** option and then select the volume from the **Boot Volume** section of the **Block Volume** drop-down list.
- h. Select **Read/Write** as the access type.
- i. Click **Attach**.
When the volume's icon no longer lists it as **Attaching**, proceed with the next steps.

3. Run the `lsblk` command again to confirm that the boot volume now shows up as a volume attached to the instance. In this sample output for the `lsblk`, the boot volume attached as a data volume shows up as `sdb`:

```
lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
sdb   8:16   0 46.6G 0 disk
├─sdb2 8:18   0   8G 0 part
├─sdb3 8:19   0 38.4G 0 part
└─sdb1 8:17   0 200M 0 part
sda   8:0    0 46.6G 0 disk
├─sda2 8:2    0   8G 0 part [SWAP]
├─sda3 8:3    0 38.4G 0 part /
└─sda1 8:1    0 200M 0 part /boot/efi
```

4. Run the `fsck` command on the volume's root partition. The root partition is usually the largest partition on the volume.

The following sample for the `fsck` command shows the output when there are no errors or corruption present on the partitions for an Oracle 7.6 instance:

CHAPTER 8 Compute

```
sudo fsck -V /dev/sdb1
fsck from util-linux 2.23.2
[/sbin/fsck.vfat (1) -- /boot/efi] fsck.vfat /dev/sdb1
fsck.fat 3.0.20 (12 Jun 2013)
/dev/sdb1: 17 files, 2466/51145 clusters

sudo fsck -V /dev/sdb2
fsck from util-linux 2.23.2

sudo fsck -V /dev/sdb3
fsck from util-linux 2.23.2
[/sbin/fsck.xfs (1) -- /] fsck.xfs /dev/sdb3
If you wish to check the consistency of an XFS filesystem or
repair a damaged filesystem, see xfs_repair(8).
```

If errors are present on a partition, you will usually be prompted to repair the errors. Following is an example of an interactive repair session of a corrupt ext4 boot volume for an Ubuntu instance:

```
sudo fsck -V /dev/sdb1
fsck from util-linux 2.31.1
[/sbin/fsck.ext4 (1) -- /] fsck.ext4 /dev/sdb1
e2fsck 1.44.1 (24-Mar-2018)
One or more block group descriptor checksums are invalid.  Fix<y> yes
Group descriptor 92 checksum is 0xe9a1, should be 0x1f53.  FIXED.
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure

Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
Block bitmap differences: Group 92 block bitmap does not match checksum.
FIXED.

cloudimg-rootfs: ***** FILE SYSTEM WAS MODIFIED *****
cloudimg-rootfs: 75336/5999616 files (0.1% non-contiguous), 798678/12181243 blocks
```



Note

XFS file systems will usually auto-repair their contents when the system boots up, fixing any corruption during the boot process. You can use the `xfstool` command to force a repair for scenarios where boot volume corruption is preventing the auto-repair functionality from working, as shown in the following example:

```
sudo xfstool /dev/sdb3
Phase 1 - find and verify superblock...
Phase 2 - using internal log
- zero log...
- scan filesystem freespace and inode maps...
...
Phase 7 - verify and correct link counts...
done
```

Recovering a Corrupted Boot Volume for Windows Instances

If your instance fails to boot successfully or boots with the boot volume set to read-only access, the instance's boot volume may be corrupted. While it is rare, boot volume corruption can occur in the following scenarios:

- When an instance experiences a forced shutdown using the API.
- When an instance experiences a system hang due to an operating system or software error and a graceful reboot or shutdown of the instance times out, and then a forced shutdown occurs.
- When an error or outage occurs in the underlying infrastructure and there were critical disk writes pending in the system.



Important

In most cases a simple reboot will resolve boot volume corruption issues, so this is the first action you should take when troubleshooting this.

This topic describes how to determine if your Windows instance's boot volume is corrupted and what steps to take to troubleshoot and recover the corrupted boot volume. For Linux-based instances, see [Recovering a Corrupted Boot Volume for Linux-Based Instances](#).

Detecting Boot Volume Corruption

When Windows operating systems detect boot volume corruption, the instance is usually able to recover from it by automatically repairing the file system. You can use a VNC console connection to verify that the instance isn't experiencing a system hang while repairing the file system, or to detect if there are other issues. VNC console connections enable you to see what's displayed through the VGA port, for more information about the VNC console, see [Instance Console Connections](#).



Important

VNC console connections only work for virtual machine (VM) instances launched on October 13, 2017 or later, and bare metal instances launched on February 21, 2019 or later. If your instance does not support VNC console connections, proceed to [Recovering the Boot Volume](#).

1. Create a VNC console connection for the instance.

To create the VNC console connection for an instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the list of instances, find the instance you think may have a corrupted boot volume, and then click the instance name.
- c. In the **Resources** section on the **Instance Details** page, click **Console Connections**, and then click **Create Console Connection**.
- d. Specify the public key portion for the SSH key, either by browsing and selecting a public key file, for example `id_rsa.pub`, or by pasting your public key into the text box, and then click **Create Console Connection**.

2. Connect to the instance through VNC console.

To set up a secure tunnel to the VNC server on the instance using PowerShell on Windows

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
- b. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
- c. Click the Actions icon (three dots), and then click **Connect with VNC**.
- d. Select **WINDOWS** for **PLATFORM**.
- e. Click **Copy** to copy the string to the clipboard.
- f. Paste the connection string copied from the previous step to [Windows Powershell](#) and hit enter to set up the secure connection.
- g. Once the connection has been established, open your VNC client and specify `localhost` as the host to connect to and `5900` as the port to use.

Check what is displayed in the VNC console to see if the instance is stuck in the boot process or if it is in the recovery partition.



Tip

For Windows 2012 and Windows 2016 operating systems, if the instance has booted into the recovery partition it may be possible to directly perform the steps to recover the boot volume in the recovery partition .

Detaching the Boot Volume

If you have detected that your instance's boot volume is corrupted, you need to detach the boot volume from the instance before you can begin troubleshooting and recovery steps.

1. Stop the instance. For more information, see [Stopping and Starting an Instance](#).

To stop the instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
 - b. In the list of instances, find the instance you have detected boot volume corruption for and then click the instance name to display the instance details.
 - c. Click **Stop**.
2. Detach the boot volume from the instance. For more information, see [Detaching a Boot Volume](#).

To detach the boot volume from the instance

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.

- b. In the **Instances** list, select the instance you want to detach the boot volume from.
- c. Click the name of the instance to display the instance details.
- d. In the **Resources**, click **Boot Volume**.
- e. Click the Actions icon (three dots), for the boot volume.
- f. Click **Detach** and confirm the selection when prompted.

Recovering the Boot Volume

To troubleshoot and recover the corrupted boot volume, you need to attach the boot volume to a second instance as a data volume. For the second instance we recommend that you use an instance running an operating system that most closely matches the operating system for the boot volume's instance, and you should only attach boot volumes for Windows instances to other Windows instances. The second instance must be in the same availability domain and region as the boot volume's instance. If no existing instance is available create a new Windows instance using the steps described in [Creating an Instance](#).

Once you have the second instance, make sure you can log in to the instance and that it is functional before proceeding with the recovery steps. After you have confirmed that the instance is functional perform the following steps.

1. Attach the boot volume to the second instance as a data volume. For more information, see [Attaching a Volume](#).

To attach the boot volume as a data volume

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the **Instances** list, click the instance that you want to attach a volume to.
- c. In the **Resources** section, click **Attached Block Volumes**.
- d. Click **Attach Block Volume**.

- e. Select **iSCSI** for the volume attachment type.
- f. In the **Block Volume Compartment** drop-down list, select the compartment.
- g. Choose the **SELECT VOLUME** option and then select the volume from the **Boot Volume** section of the **Block Volume** drop-down list.
- h. Select **Read/Write** as the access type.
- i. Click **Attach**.

When the volume's icon no longer lists it as **Attaching**, proceed with the next steps.

2. Connect to the second instance, see [Connecting to a Windows Instance](#) for more information.
3. Connect to the volume, see [Connecting to a Volume on a Windows Instance](#) for more information. Since you are attaching a boot volume as a data volume you must also run the `Connect-IscsiTarget` and set `IsMultiEnabled` to true. For example:

```
Set-Service -Name msiscsi -StartupType Automatic
Start-Service msiscsi
New-IscsiTargetPortal -TargetPortalAddress 169.254.2.4
Connect-IscsiTarget -NodeAddress iqn.2015-02.oracle.boot:uefi -TargetPortalAddress 169.254.2.4 -
IsPersistent $True -IsMultipathEnabled $True
```

4. Open **Computer Management** and navigate to **Storage**, and then **Disk Management**.
5. Select the new disk and mark it **Online**.
6. Click **This PC** and then right-click on the new disk and select **Properties**.
7. Navigate to **Tools, Error Checking**, and then **Check**.
8. Select **Scan Drive** and fix issues as they come up.
9. Mark the new disk **Offline**.
10. Open iscsi initiator with administrator privileges.
11. In **Favorite Targets**, remove the iscsi target of the attached volume.

Oracle-Provided Images

An image is a template of a virtual hard drive. The image determines the operating system and other software for an instance. The following table lists the images that are available in Oracle Cloud Infrastructure. For specific image and kernel version details, along with changes between versions, see [Oracle-Provided Image Release Notes](#).

Image	Name	Description
Oracle Linux 7 Unbreakable Enterprise Kernel Release 5	Oracle-Linux-7.x-<date>-<number>	The Unbreakable Enterprise Kernel (UEK) is Oracle's optimized operating system kernel for demanding Oracle workloads. GPU shapes are supported with this image.
Oracle Linux 6 Unbreakable Enterprise Kernel Release 4	Oracle-Linux-6.x-<date>-<number>	The Unbreakable Enterprise Kernel (UEK) is Oracle's optimized operating system kernel for demanding Oracle workloads.
CentOS 7	CentOS-7-<date>-<number>	CentOS is a free, open-source Linux distribution that is suitable for use in enterprise cloud environments. For more information, see https://www.centos.org/ .
CentOS 6	CentOS-6.x-<date>-<number>	CentOS is a free, open-source Linux distribution that is suitable for use in enterprise cloud environments. For more information, see https://www.centos.org/ .

Image	Name	Description
Ubuntu 18.04 LTS	Canonical-Ubuntu-18.04- <date>- <number>	<p>Ubuntu is a free, open-source Linux distribution that is suitable for use in the cloud. For more information, see https://www.ubuntu.com.</p> <p>Minimal Ubuntu is designed for automated use at scale. It uses a smaller boot volume, boots faster, and has a smaller surface for security patches than standard Ubuntu images. For more information, see https://wiki.ubuntu.com/Minimal.</p> <p>GPU shapes are supported with this image. You must install the appropriate GPU drivers from NVIDIA.</p>
Ubuntu 16.04 LTS	Canonical-Ubuntu-16.04- <date>- <number>	<p>Ubuntu is a free, open-source Linux distribution that is suitable for use in the cloud. For more information, see https://www.ubuntu.com.</p> <p>Minimal Ubuntu is designed for automated use at scale. It uses a smaller boot volume, boots faster, and has a smaller surface for security patches than standard Ubuntu images. For more information, see https://wiki.ubuntu.com/Minimal.</p> <p>GPU shapes are supported with this image. For Minimal Ubuntu, you must install the appropriate GPU drivers from NVIDIA.</p>
Windows Server 2016	Windows-Server-2016- <edition>- Gen2.<date>- <number>	<p>Windows Server 2016 supports running production Windows workloads on Oracle Cloud Infrastructure.</p> <p>GPU shapes are supported with this image. You must install the appropriate GPU drivers from NVIDIA.</p>

CHAPTER 8 Compute

Image	Name	Description
Windows Server 2012 R2	Windows-Server-2012-R2- <i><edition></i> - <i><gen></i> - <i><date></i> - <i><number></i>	Windows Server 2012 R2 supports running production Windows workloads on Oracle Cloud Infrastructure. GPU shapes are supported with this image. You must install the GPU drivers from NVIDIA.
Windows Server 2008 R2 - Virtual Machine (VM)	Windows-Server-2008-R2-Enterprise-Edition-VM- <i><date></i> - <i><number></i>	Windows Server 2008 R2 Enterprise Edition supports running production Windows workloads on Oracle Cloud Infrastructure.

You also can [create custom images](#) of your boot disk OS and software configuration for launching new instances.

Essential Firewall Rules



Warning

Windows 2008 Server R2 images do not support restricting certain firewall rules for local principals, such as "Administrators", so any authenticated user on an instance can make outgoing connections to the iSCSI network endpoints (169.254.0.2:3260, 169.254.2.0/24:3260) that serve the instance's boot and block volumes.

CHAPTER 8 Compute

All Oracle-provided images include rules that allow only "root" on Linux instances or "Administrators" on Windows Server 2012 R2 and Windows Server 2016 instances to make outgoing connections to the iSCSI network endpoints (169.254.0.2:3260, 169.254.2.0/24:3260) that serve the instance's boot and block volumes.

- We recommend that you do not reconfigure the firewall on your instance to remove these rules. Removing these rules allows non-root users or non-administrators to access the instance's boot disk volume.
- We recommend that you do not create custom images without these rules unless you understand the security risks.
- Running Uncomplicated Firewall (UFW) on Ubuntu images might cause issues with these rules. Because of this, we recommend that you do not enable UFW on your instances. See [Ubuntu Instance fails to reboot after enabling Uncomplicated Firewall \(UFW\)](#) for more information.

User Data

Oracle-provided images give you the ability to run custom scripts or supply custom metadata when the instance launches. To do this, you specify a custom startup script in the **User Data** field when you [create the instance](#). For more information about startup scripts, see [cloud-init](#) for Linux-based images and [cloudbase-init](#) for Windows-based images.

OS Updates for Linux Images

Oracle Linux and CentOS images are preconfigured to let you install and update packages from the repositories on the Oracle public yum server. The repository configuration file is in the `/etc/yum.repos.d` directory on your instance. You can install, update, and remove packages by using the yum utility.



Note

OS Security Updates for Oracle Linux and CentOS images

After you launch an instance using Oracle Linux or CentOS images, you are responsible for applying the required OS security updates published through the Oracle public yum server. For more information, see [Installing and Using the Yum Security Plugin](#).

The Ubuntu image is preconfigured with suitable repositories to allow you to install, update, and remove packages.



Note

OS Security Updates for the Ubuntu image

After you launch an instance using the Ubuntu image, you are responsible for applying the required OS security updates using the `sudo apt-get upgrade` command.

Linux Kernel Updates

Oracle Linux images on Oracle Cloud Infrastructure include Oracle Linux Premier Support at no extra cost. This gives you all the services included with Premier Support, including Oracle Ksplice. Ksplice enables you to apply important security and other critical kernel updates without a reboot. For more information, see [About Oracle Ksplice](#) and [Ksplice Overview](#).

Ksplice is only available for Linux instances launched on or after February 15, 2017. For instances launched before August 25, 2017, you must install Ksplice before running it. See [Installing and Running Oracle Ksplice](#) for more information.



Note

Ksplice Support

Oracle Ksplice is not supported for CentOS and Ubuntu images, or on Linux images launched before February 15, 2017.

Configuring Automatic Package Updating on Instance Launch

You can configure your instance to automatically update to the latest package versions when the instance first launches using a cloud-init startup script. To do this, add the following code to the startup script:

```
package_upgrade: true
```

The upgrade process starts when the instance launches and runs in the background until it completes. To verify that it completed successfully, check the cloud-init logs in `/var/log`.

See [User Data](#) and [Cloud config examples - Run apt or yum upgrade](#) for more information.

Linux Image Details

See [Lifetime Support Policy: Coverage for Oracle Linux and Oracle VM](#) for details about the Oracle Linux support policy.

Users

For instances created using Oracle Linux and CentOS images, the user name `opc` is created automatically. The `opc` user has `sudo` privileges and is configured for remote access over the SSH v2 protocol using RSA keys. The SSH public keys that you specify while creating instances are added to the `/home/opc/.ssh/authorized_keys` file.

For instances created using the Ubuntu image, the user name `ubuntu` is created automatically. The `ubuntu` user has `sudo` privileges and is configured for remote access over

the SSH v2 protocol using RSA keys. The SSH public keys that you specify while creating instances are added to the `/home/ubuntu/.ssh/authorized_keys` file.

Note that `root` login is disabled.

Remote Access

[Access to the instance](#) is permitted only over the SSH v2 protocol. All other remote access services are disabled.

Firewall Rules

Instances created using Oracle-provided images have a default set of firewall rules that allow only SSH access. Instance owners can modify those rules as needed, but must not restrict link local traffic to address 169.254.0.2 in accordance with the warning at the top of this page.

Be aware that the Networking service uses [network security groups](#) and [security lists](#) to control packet-level traffic in and out of the instance. When troubleshooting access to an instance, make sure all of the following items are set correctly: the network security groups that the instance is in, the security lists associated with the instance's subnet, and the instance's firewall rules.

Cloud-init Compatibility

Instances created using Oracle-provided images are compatible with cloud-init. When launching an instance with the Core Services API, you can pass cloud-init directives with the `metadata` parameter. For more information, see [LaunchInstance](#).

OCI Utilities

Instances created using Oracle Linux include a preinstalled set of utilities that are designed to make it easier to work with Oracle Linux images. These utilities consist of a service component and related command line tools.

The following table summarizes the components that are included in the OCI utilities.

CHAPTER 8 Compute

Name	Description
ocid	The service component of oci-utils. This normally runs as a daemon started via <code>systemd</code> . This service scans for changes in the iSCSI and VNIC device configurations and caches the OCI metadata and public IP address of the instance.
oci-iscsi-config	Used to display and configure iSCSI devices attached to a compute instance. If no command line options are specified, lists devices that need attention.
oci-metadata	Displays metadata for the compute instance. If no command line options are specified, lists all available metadata. Metadata includes the instance OCID, display name, compartment, shape, region, availability domain, creation date, state, image, and any custom metadata that you provide, such as an SSH public key.
oci-network-config	Lists or configures virtual network interface cards (VNICs) attached to the Compute instance. When a secondary VNIC is provisioned in the cloud, it must be explicitly configured on the instance using this script or similar commands.
oci-public-ip	Displays the public IP address of the current system in either human-readable or JSON format.

For more information, see the [OCI Utilities](#) reference.

Windows OS Updates for Windows Images

Windows images include the Windows Update utility, which you can run to get the latest Windows updates from Microsoft. You have to configure the instance's [network security group](#) or the [security list](#) used by the instance's subnet to allow instances to access Windows update servers.

Windows Image Details

Users

For instances created using Oracle-provided Windows images, the user name `opc` is created automatically. When you launch an instance using the Windows image, Oracle Cloud Infrastructure will generate an initial, one-time password that you can retrieve using the console or API. This password must be changed after you initially log on.

Remote Access

[Access to the instance](#) is permitted only through a Remote Desktop connection.

Firewall Rules

Instances created using the Windows image have a default set of firewall rules that allow Remote Desktop protocol or RDP access on port 3389. Instance owners can modify these rules as needed, but must not restrict link local traffic to 169.254.169.253 for the instance to activate with Microsoft Key Management Service (KMS). This is how the instance stays active and licensed.

Be aware that the Networking service uses [network security groups](#) and [security lists](#) to control packet-level traffic in and out of the instance. When troubleshooting access to an instance, make sure all of the following items are set correctly: the network security groups that the instance is in, the security lists associated with the instance's subnet, and the instance's firewall rules.

User Data on Windows Images

On Windows images custom user data scripts are executed using [cloudbase-init](#), which is the equivalent of [cloud-init](#) on Linux-based images. All Oracle-provided Windows images on Oracle Cloud Infrastructure include cloudbase-init installed by default. When an instance launches, cloudbase-init runs PowerShell, batch scripts, or additional user data content. See [cloudbase-init Userdata](#) for information about supported content types.

You can use user data scripts to perform various tasks, such as:

- Enable GPU support using a custom script to install the applicable GPU driver.
- Add or update local user accounts.
- Join the instance to a domain controller.
- Install certificates into the certificate store.
- Copy any required application workload files from the Object Storage service directly to the instance.



Warning

Do not include anything in the script that could trigger a reboot, because this could impact the instance launch, causing it to fail. Any actions requiring a reboot should only be performed after the instance state is **RUNNING**.

Windows Remote Management

[Windows Remote Management](#) (WinRM) is enabled by default on Oracle-provided Windows images. WinRM provides you with the capability to remotely manage the operating system.

To use WinRM you need to add a stateful ingress [security rule](#) for TCP traffic on destination port 5986. You can implement this security rule in either a [network security group](#) that the instance belongs to, or a [security list](#) that is used by the instance's subnet.



Warning

The following procedure allows WinRM connections from 0.0.0.0/0, which means any IP address, including public IP addresses. To allow access only from instances within the VCN, change the source CIDR value to the VCN's CIDR block. For more information, see [Security Recommendations](#).

To enable WinRM access

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. To add the rule to a network security group that the instance belongs to:
 - a. Under **Resources**, click **Network Security Groups**. Then click the network security group that you're interested in.
 - b. Click **Add Rules**.
 - c. Enter the following values for the rule:
 - **Stateless:** Leave the check box cleared
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** TCP
 - **Source Port Range:** All
 - **Destination Port Range:** 5986
 - d. When done, click **Add**.
4. Or, to add the rule to a security list that is used by the instance's subnet:

- a. Under **Resources**, click **Security Lists**. Then click the security list you're interested in.
- b. Click **Add Ingress Rules**.
- c. Enter the following values for the rule:
 - **Stateless:** Leave the check box cleared
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** TCP
 - **Source Port Range:** All
 - **Destination Port Range:** 5986
- d. When done, click **Add Ingress Rules**.

To use WinRM on an instance

1. [Get the instance's public IP address](#).
2. Open Windows PowerShell on the Windows client that you're using to connect to the instance.
3. Run the following command:

```
# Get the public IP from your OCI running windows instance
$ComputerName = Public IP Address

# Store your username and password credentials (default username is opc)
$c = Get-Credential

# Options
$opt = New-PSSessionOption -SkipCACheck -SkipCNCheck -SkipRevocationCheck

# Create new PSSession (Pre-requisite: ensure network security group or security list has Ingress
Rule for port 5986)
$PSSession = New-PSSession -ComputerName $ComputerName -UseSSL -SessionOption $opt -
```

CHAPTER 8 Compute

```
Authentication Basic -Credential $c  
  
# Connect to Instance PSSession  
Enter-PSSession $PSSession  
  
# To close connection use: Exit-PSSession
```

You can now remotely manage the Windows instance from your local PowerShell client.

Operating System Lifecycle and Support Policy

When an operating system reaches the end of its support lifecycle, the OS vendor (such as Microsoft) no longer provides security updates for the OS. You should upgrade to the latest version to remain secure.

Here's what you should expect when an OS version reaches the end of its support lifecycle:

- Oracle Cloud Infrastructure no longer provides new images for the OS version. Images that were previously published are deprecated, and are no longer updated.
- Although you can continue to run instances that use deprecated images, Oracle Cloud Infrastructure does not provide any support for operating systems that have reached the end of the support lifecycle.
- If you have an instance that runs an OS version that will be deprecated, and you want to launch new instances with this OS version after the end of support, you can create a custom image of the instance and then use the custom image to launch new instances in the future. For custom Linux images, you must purchase extended support from the OS vendor. For custom Windows images, Extended Security Updates may not be used on Oracle Cloud Infrastructure. Oracle Cloud Infrastructure does not provide any support for custom images that use end-of-support operating systems.

Be aware of these end-of-support dates:

- **Ubuntu 14.04:** Support ended on April 19, 2019.
- **Windows Server 2008 R2:** Support ends on January 14, 2020.

Using NVIDIA GPU Cloud with Oracle Cloud Infrastructure

NVIDIA GPU Cloud (NGC) is a GPU-accelerated cloud platform optimized for deep learning and scientific computing. This topic provides an overview of how to use NGC with Oracle Cloud Infrastructure.

NVIDIA makes available on Oracle Cloud Infrastructure a customized Compute image optimized for the NVIDIA® Tesla Volta™ and Pascal™ GPUs . Running NGC containers on this instance provides optimum performance for deep learning jobs.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Prerequisites

- An Oracle Cloud Infrastructure tenancy. For more information, see [Signing Up for Oracle Cloud Infrastructure](#).
- A cloud network to launch the instance into. For information about setting up cloud networks, see [Using the Console](#) in [VCNs and Subnets](#).
- A key pair, to use for connecting to the instance via SSH. For information about generating a key pair, see [Managing Key Pairs on Linux Instances](#).
- Security group and policy configured for the File Storage service. For more information, see [Managing Groups](#), [Getting Started with Policies](#), and [Details for the File Storage Service](#).
- An NGC API key for authenticating with the NGC service.

To generate your NGC API key

1. Log in to the [NGC website](#).
2. On the [NGC Registry page](#), click **Get API Key**.
3. Click **Generate API Key** and then click **Confirm** to generate the key. If you have an existing API key it will become invalid once you generate a new key.

Launching an instance based on the NGC image

USING THE CONSOLE

1. Open the Console. See [Signing In to the Console](#) for steps to do this.
2. Click **Compute**, choose a compartment you have permission to work in, and then click **Create Instance**.
3. In the **Create Instance** dialog box, specify the instance name, and select the availability domain for the instance.
4. Click **Change Image Source**.
5. On the **Oracle Images** tab, check **NVIDIA GPU Cloud Machine Image**, review and accept the **Agreement for Oracle App "NVIDIA GPU Cloud Machine Image"**, and then click **Select Image**.
6. Select **Virtual Machine** or **Bare Metal Machine** for **Instance Type**.
7. Select the shape you want to use for **Shape**.
8. For **SSH Keys**, click **Choose SSH Key File**, navigate to the location where you saved the public key portion (.pub) of the SSH key file you created, select the file and click **Open**.
9. In the **Configure Networking** section, select the virtual cloud network (VCN) compartment, VCN, subnet compartment, and subnet.
10. Click **Create Instance**.

You should now see the NGC instance with the status of **Provisioning**. Once the status has changed to **Running**, you can connect to the instance. For general information about launching Compute instances, see [Creating an Instance](#).

See the following topics for accessing and working with the instance:

- [Connecting to an Instance](#)
- [Stopping and Starting an Instance](#)
- [Terminating an Instance](#)

When you connect to the instance using SSH you will be prompted for the NGC API key. If you supply the API key at the prompt, the instance will automatically log you into the NGC container registry so that you can run containers from the registry. You can choose not to supply the API key at the prompt and still log in to the instance. You can then log in later to the NGC container registry, see [Logging in to the NGC Container Registry](#) for more information.

USING THE CLI

Oracle Cloud Infrastructure provides a [Command Line Interface \(CLI\)](#) you can use to complete tasks. For more information, see [Quickstart](#) and [Configuration](#). Use the [launch](#) command to create an instance, specifying `image` for **sourceType** and the image OCID

`ocid1.image.oc1..aaaaaaaakn16phck7e3iuii4r4axpwhenw5qtnnsk3tqppajdjzb5nhoma3q` in **InstanceSourceDetails** for **LaunchInstanceDetails**.

Using the File Storage Service for Persistent Data Storage

You can use the [Overview of File Storage](#) for data storage when working with NGC. See the following tasks for creating and working with the File Storage service:

- [Creating File Systems](#)
- [Using the Console](#)
- [Using the API](#)
- [Managing File Systems](#)
- [Using the Command Line Interface \(CLI\)](#)

Using the Block Volume Service for Persistent Data Storage

You can use the Block Volume service for data storage when working with NGC. For more information, see [Overview of Block Volume](#). See the following tasks for creating and working

CHAPTER 8 Compute

with the Block Volume service:

- [Creating a Volume](#)
- [Attaching a Volume](#)
- [Connecting to a Volume](#)

You can also use the CLI to manage block volumes, see the [volume](#) commands.

Examples of Running Containers

You first need to log into the NGC container registry. You can skip this section if you provided your API key when logging into the instance via SSH. If you did not provide your API key when connecting to your instance, then you must perform this step.

To log into the NGC container registry

1. Run the following Docker command:

```
docker login nvcr.io
```

2. When prompted for a username, enter `$oauthtoken`.
3. When prompted for a password enter your NGC API key.

At this point you can run Docker commands and access the NGC container registry from the instance.

Example: MNIST Training Run Using PyTorch Container

This sample demonstrates running the MNIST example under PyTorch. This example downloads the MNIST dataset from the web.

1. Pull and run the PyTorch container with the following Docker commands:

```
docker pull nvcr.io/nvidia/pytorch:17.10  
nvidia-docker run --rm -it nvcr.io/nvidia/pytorch:17.10
```

2. Run the MNIST example with the following commands:

CHAPTER 8 Compute

```
cd /opt/pytorch/examples/mnist  
python main.py
```

Example: MNIST Training Run Using TensorFlow Container

This sample demonstrates running the MNIST example under TensorFlow. This example downloads the MNIST dataset from the web.

1. Pull and run the TensorFlow container with the following Docker commands:

```
nvcr.io/nvidia/tensorflow:17.10  
nvidia-docker run --rm -it nvcr.io/nvidia/tensorflow:17.10
```

2. Run the `MNIST_with_summaries` example with the following commands:

```
cd /opt/tensorflow/tensorflow/examples/tutorials/mnist  
python mnist_with_summaries.py
```

OCI Utilities

Instances created using [Oracle-Provided Images](#) based on Oracle Linux include a pre-installed set of utilities that are designed to make it easier to work with Oracle Linux images. These utilities consist of a service component and related command line tools that can help with managing block volumes (attach, remove, and automatic discovery), secondary VNIC configuration, discovering the public IP address of an instance, and retrieving instance metadata.

The following table summarizes the components that are included in the OCI utilities.

CHAPTER 8 Compute

Name	Description
<code>ocid</code>	The service component of <code>oci-utils</code> . This normally runs as a daemon started via <code>systemd</code> . This service scans for changes in the iSCSI and VNIC device configurations and caches the OCI metadata and public IP address of the instance.
<code>oci-growfs</code>	Expands the root filesystem of the instance to its configured size.
<code>oci-iscsi-config</code>	Used to display and configure iSCSI devices attached to a compute instance. If no command line options are specified, lists devices that need attention.
<code>oci-metadata</code>	Displays metadata for the compute instance. If no command line options are specified, lists all available metadata. Metadata includes the instance OCID, display name, compartment, shape, region, availability domain, creation date, state, image, and any custom metadata that you provide, such as an SSH public key.
<code>oci-network-config</code>	Lists or configures virtual network interface cards (VNICs) attached to the Compute instance. When a secondary VNIC is provisioned in the cloud, it must be explicitly configured on the instance using this script or similar commands.
<code>oci-network-inspector</code>	Displays a detailed report for a given compartment or network.
<code>oci-public-ip</code>	Displays the public IP address of the current system in either human-readable or JSON format.

Installing the OCI Utilities

The OCI utilities are automatically included with instances launched with an Oracle Linux image. They are not currently available on other distributions.

CHAPTER 8 Compute

Much of the `oci-utils` functionality requires that you have the [Oracle Cloud Infrastructure Python SDK](#) and the [Oracle Cloud Infrastructure CLI](#) installed and configured.

You can install the Oracle Cloud Infrastructure CLI using `yum`:

```
yum install python-oci-cli
```

For configuration information, see the [Oracle Cloud Infrastructure Python SDK documentation](#) and the documentation for [configuring](#) the [Oracle Cloud Infrastructure CLI](#).

Updating the OCI Utilities

To update to the latest version of the OCI Utilities:

```
$ sudo yum update oci-utils
```

Using the OCI Utilities

To use the OCI utilities, you first need to start the `ocid` service:

```
sudo systemctl start ocid.service
```

Example output:

```
Redirecting to /bin/systemctl start ocid.service
```

The `ocid` Daemon

DESCRIPTION

The `ocid` daemon is the service component of the `oci-utils`. It monitors for changes in the VNIC and iSCSI configuration of the instance and attempts to automatically attach or detach devices as they appear or disappear - for example, when they are created or deleted using the Oracle Cloud Infrastructure Console, the command line interface (CLI), or the API.

CONFIGURATION

The `ocid` daemon requires root privileges. You can configure root privileges for `ocid` using one of the following methods:

CHAPTER 8 Compute

- Run the `oci setup config` configuration command as root to create SDK configuration files for the host. For more information, see [SDK and Tool Configuration](#).
- Use instance principals by adding the instance to a dynamic group that was granted access to Oracle Cloud Infrastructure services. For more information, see [Managing Dynamic Groups](#).
- Configure `oci-utils` to allow root to use a non-privileged user's Oracle Cloud Infrastructure configuration files. For more information, see the configuration file located in the `/etc/oci-utils.conf.d` directory of the instance.

USAGE

To start the `ocid` daemon using `systemd`:

```
service ocid start
```

To set `ocid` to start automatically during system boot:

```
sudo systemctl enable ocid.service
```

oci-growfs

DESCRIPTION

Expands the root filesystem of the instance to its configured size. This command must be run as root.

USAGE

```
oci-growfs [-y] [-n] [-h]
```

OPTIONS

`-Y`

Answer 'yes' to all prompts.

`-N`

Answer 'no' to all prompts.

CHAPTER 8 Compute

`-H | --HELP`

Displays a summary of the command line options.

EXAMPLE

```
# sudo /usr/libexec/oci-growfs
CHANGE: disk=/dev/sda partition=3: start=17188864 old: size=80486399,end=97675263 new:
size=192526302,end=209715166
Confirm? [y/n]: y
CHANGED: disk=/dev/sda partition=3: start=17188864 old: size=80486399,end=97675263 new:
size=192526302,end=209715166
    meta-data=/dev/sda3          isize=256    agcount=4, agsize=2515200 blks
    =                   sectsz=4096  attr=2,  projid32bit=1
    =                   crc=0        finobt=0  spinodes=0
    data           =                   bsize=4096  blocks=10060800, imaxpct=25
    =                   sunit=0      swidth=0  blks
    naming         =version 2          bsize=4096  ascii-ci=0  ftype=1
    log            =internal           bsize=4096  blocks=4912, version=2
    =                   sectsz=4096  sunit=1   blks, lazy-count=1
    realtime      =none                extsz=4096  blocks=0,  rtextents=0
    data blocks changed from 10060800 to 24065787
```

oci-iscsi-config

DESCRIPTION

Lists and configures iSCSI devices attached to a Compute instance running in Oracle Cloud Infrastructure. When run without any command line options, `oci-iscsi-config` lists devices that need attention.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

CHAPTER 8 Compute

USAGE

```
oci-iscsi-config [-i|--interactive] [-s|--show] [-a | --attach IQN ]  
  [-d IQN | --detach IQN ] [--username username] [--password password]  
  [--help] oci-iscsi-config [-s|--show] [-c | --create-volume size]  
  [--volume-name name] [--destroy-volume OCID ]
```

OPTIONS

`-I | --INTERACTIVE`

Run in interactive mode. This option displays devices that need attention and offers to attach and configure them. Requires root privileges.

`-S | --SHOW`

List all devices. If `ocid` is not running then root privileges are required.

`-A | --ATTACH TARGET`

Attempt to attach the device with the given IQN (a unique ID assigned to a device) or Oracle Cloud Identifier (OCID). When using an IQN, the volume must already be attached (assigned) to the instance in the Console. The Oracle Cloud Infrastructure Python SDK is required for selecting volumes using their OCID. This option can be used multiple times to attach multiple devices at the same time. Requires root privileges.

`-D | --DETACH DEVICE`

Detach the device with the given IQN (a unique ID assigned to a device). If the volume (or any partition of the volume) is mounted, this option will attempt to unmount it first. This option can be used multiple times to detach multiple devices at the same time. Requires root privileges.

`-C | --CREATE-VOLUME SIZE`

Create a new volume of `SIZE` gigabytes and attach it to the current instance. This option requires the Oracle Cloud Infrastructure Python SDK to be installed and configured.

CHAPTER 8 Compute

`--DESTROY-VOLUME OCID`

Destroy the block storage volume with the given OCID. The volume must not be attached to any instances.



Warning

This action is irreversible.

`--VOLUME-NAME NAME`

Set the display name for the volume. This is used with the `--create-volume` option.

`--USERNAME NAME`

Use the specified user name as the CHAP user name when authentication is needed for attaching a device. Not needed when the Oracle Cloud Infrastructure Python SDK is available.

`--PASSWORD PASSWORD`

Use the supplied password as the CHAP password when authentication is needed for attaching a device. This is not needed when the Oracle Cloud Infrastructure Python SDK is available.

`--HELP`

Displays a summary of the command line options.

EXAMPLES

DISPLAYING ISCSI CONFIGURATION

The `oci-iscsi-config` utility works with the `ocid` daemon to monitor device creation and deletion through the command line, console, or SDK and automatically discover those changes. You can use the `--show` option to display a list of all of the devices attached to an instance:

```
# oci-iscsi-config -s
For full functionality of this utility the ocid service must be running
The administrator can start it using this command:
sudo systemctl start ocid.service
```

CHAPTER 8 Compute

```
ocid already running.
Currently attached iSCSI devices:

Target iqn.2015-02.oracle.boot:uefi
Persistent portal: 169.254.0.2:3260
Current portal: 169.254.0.2:3260
State: running
Attached device: sda
Size: 46.6G
Partitions: Device Size Filesystem Mountpoint
sda1 544M vfat /boot/efi
sda2 8G swap [SWAP]
sda3 38G xfs /
```

The following example shows the output of the `--show` option after adding a 50GB block volume using the OCI console:

```
# oci-iscsi-config --show
Currently attached iSCSI devices:

Target iqn.2015-12.com.oracleiaas:abcdefghijklmnopqrstuvwxyz1234567890
Persistent portal: 169.254.2.2:3260
Current portal: 169.254.2.2:3260
State: running
Attached device: sdb
Size: 50G
File system type: Unknown
Mountpoint: Not mounted

Target iqn.2015-02.oracle.boot:uefi
Persistent portal: 169.254.0.2:3260
Current portal: 169.254.0.2:3260
State: running
Attached device: sda
Size: 46.6G
Partitions: Device Size Filesystem Mountpoint
sda1 544M vfat /boot/efi
sda2 8G swap [SWAP]
sda3 38G xfs /
```

CREATING A VOLUME

The following example shows how to create a volume:

CHAPTER 8 Compute

```
# oci-iscsi-config --create-volume 50
For full functionality of this utility the ocid service must be running
The administrator can start it using this command:
sudo systemctl start ocid.service
Creating a new 50 GB volume
Volume abcdefghijklmnopqrstuvwxyz1234567890123456789012345678901234 created
```

DELETING A VOLUME

The following example shows how to destroy a volume:

```
# oci-iscsi-config --destroy-volume
ocid1.volume.oc1.phx.abcdefghijklmnopqrstuvwxyz1234567890123456789012345678901234
```

oci-metadata

DESCRIPTION

Displays metadata for an Oracle Cloud Infrastructure Compute instance.



Note

Instance Metadata

For more information about instance metadata, see [Getting Instance Metadata](#).

USAGE

```
oci-metadata [-h] [-j] [-g key] [--help]
```

OPTIONS

-H | *--HUMAN-READABLE*

Display human readable output (default).

-J | *--JSON*

Display output in JSON.

CHAPTER 8 Compute

`-G | --GET KEY`

Retrieve data only for the specified key.

`--HELP`

Displays a summary of the command line options.

EXAMPLES

GETTING ALL METADATA FOR THE INSTANCE

Running `oci-metadata` with no options returns all metadata for the instance:

```
# oci-metadata
Instance details:
  Display Name: my-example-instance
  Region: phx - us-phoenix-1 (Phoenix, AZ, USA)
  Canonical Region Name: us-phoenix-1
  Availability Domain: cumS:PHX-AD-1
  Fault domain: FAULT-DOMAIN-3
  OCID: ocid1.instance.oc1.phx.exampleuniqueID
  Compartment OCID: ocid1.compartment.oc1..exampleuniqueID
  Instance shape: VM.Standard2.1
  Image ID: ocid1.image.oc1.phx.exampleuniqueID
  Created at: 1569529065596
  state: Running
  agentConfig:
    managementDisabled: False
    monitoringDisabled: False
  Instance Metadata:
    ssh_authorized_keys: ssh-rsa AAAAB3NzaC1yc2EAAAABJQAAAQEApzWsZ0DxxE+jNpsh5C3
ncc18eZ06MCONWnP5D/8eqdGjUqqBU71HvgaIqYHy1PH/B4w/pya56k41FLzdXFVAzCy4NGyE4XZjLcS
SR4jW1As4r3WHY/f61Gc1VicZ39u+bB1XYaHm46zS4WrQEQU4wbz70iwe50nSIhrGpvM5HWYOK0dsV
A7/zzw+yW37NUGa/QeM4/bJvCVg3BVjB6VWdmV7dFwRMeCaVJFQH3wKndvuJib78zoH19sbYm74vzqTY
Si/bVoIz9YnZ4bA3MS0Uqapok/m2M9T27+UA/lz/ILCKXP3+vNcVcjRplanJT/qlzhLiIiBCRo4RsdGx
UIw== rsa-key-20181129
Networking details:
  VNIC OCID: ocid1.vnic.oc1.phx.exampleuniqueID
  VLAN Tag: 2392
  Private IP address: 10.0.0.16
  MAC address: 02:00:17:03:D8:FE
```

CHAPTER 8 Compute

```
Subnet CIDR block: 10.0.0.0/24
Virtual router IP address: 10.0.0.1
```

GETTING ONLY SPECIFIC METADATA

Pass in a key with the `--get` parameter to retrieve only the metadata for that key:

```
# oci-metadata --get state
Instance details:
Instance state: Running
```

oci-network-config

DESCRIPTION

Configures network interfaces for an Oracle Cloud Infrastructure Compute instance.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

USAGE

```
oci-network-config [-h] [-s] [--create-vnic] [--detach-vnic VNIC] [--add-
private-ip]
  [--del-private-ip ip_address] [--private-ip ip_address] [--subnet subnet] [--
vnic-name name]
  [--assign-public-ip] [--vnic OCID] [-a] [-d] [-e ip_address] [-n format]
  [-r] [-X | --exclude item] [-I | --include item] [--quiet]
```

CHAPTER 8 Compute

OPTIONS

`-s | --SHOW`

Show information on all provisioning and interface configuration. This is the default action if no options are given.

`--CREATE-VNIC`

Create a new virtual network interface card (VNIC) and attach it to this instance.

`--DETACH-VNIC OCID`

Detach and delete the VNIC with the given Oracle Cloud Identifier (OCID) or primary IP address. Cannot be the primary VNIC for the instance.

`--ADD-PRIVATE-IP OCID`

Add a secondary private IP to an existing VNIC.

`--DEL-PRIVATE-IP IP_ADDRESS`

Delete the secondary private IP address with the given IP address.

`--PRIVATE-IP IP_ADDRESS`

Assign the given private IP address to the VNIC. Used with the with the `--create-vnic` and `add-private-ip` options,

`--SUBNET SUBNET`

Used with the with the `--create-vnic` option. Connects the new VNIC to the specified subnet.

`--VNIC-NAME NAME`

Used with the with the `--create-vnic` option. Display name for the new VNIC.

`--ASSIGN-PUBLIC-IP`

When used with the `--create-vnic` option, assign a public IP address to the new VNIC.

CHAPTER 8 Compute

`--vnic OCID`

When used with the `--add-private-ip` option, assign the private IP address to the given VNIC

`-A | --AUTO | -C | --CONFIGURE`

Add IP configuration for VNICs that are not configured and delete IP configuration for VNICs that are no longer provisioned.

`-D | --DECONFIGURE`

Deconfigure all VNICs (except the primary). If used with the `-e` option, only the secondary IP addresses are deconfigured.

`-E IP_ADDRESS VNIC_OCID`

Secondary private IP address to configure or deconfigure. Used with `--configure` and `--deconfigure` options.

`-N | -NS FORMAT`

When configuring, place interfaces in namespace identified by the given format. Format can include `$nic` and `$vltag` variables.

`-R | --SSHD`

Start `sshd` in namespace (if `-n` is present).

`-X | --EXCLUDE ITEM`

Persistently exclude the given item from automatic configuration or deconfiguration. Use the `--include` option to include the item again.

`-I | --INCLUDE ITEM`

Include an item that was previously excluded using the `--exclude` option in automatic configuration/deconfiguration.

`-Q | --QUIET`

Do not display information messages.

CHAPTER 8 Compute

`-H | --HELP`

Displays a summary of the command line options.

EXAMPLES

DISPLAYING CURRENT NETWORK CONFIGURATION

Running the `oci-network-config` command with no options returns the network configuration of the current instance:

```
VNIC configuration for instance my-test-instance-20180622-1222

VNIC 1 (primary): my-test-instance-20180622-1222
Hostname: my-test-instance-20180622-1222
OCID: ocid1.vnic.oc1.phx.abcdefg12345678
MAC address: 00:00:17:00:F4:3F
Public IP address: 129.146.110.62
Subnet: Public Subnet cumS:PHX-AD-1 (10.0.0.0/24)

Operating System level network configuration

CONFIG ADDR          SPREFIX          SBITS  VIRTRT          NS          IND IFACE          VLTAG  VLAN
STATE MAC           VNIC
-      10.0.0.3         10.0.0.0         24    10.0.0.1        -           0    ens3            -      -
UP     00:00:17:00:f4:3f  ocid1.vnic.oc1.phx.abcdefg12345678
```

CREATING A NEW VNIC

This example creates a new VNIC named `MY_NEW_VNIC` and attaches it to the instance:

```
# sudo oci-network-config --create-vnic --vnic-name MY_NEW_VNIC
Info: creating VNIC: 10.0.0.4
```

Running `oci-network-config` with the `-s` option shows information for the new VNIC:

```
# sudo oci-network-config -s
VNIC configuration for instance scottb-instance-20180622-1222

VNIC 1 (primary): scottb-instance-20180622-1222
Hostname: scottb-instance-20180622-1222
OCID: ocid1.vnic.oc1.phx.abcdefg12345678
```

CHAPTER 8 Compute

```
MAC address: 00:00:17:00:F4:3F
Public IP address: 129.146.110.62
Subnet: Public Subnet cumS:PHX-AD-1 (10.0.0.0/24)
```

```
VNIC 2: MY_NEW_VNIC
Hostname: scottb-instance-20180622-1222-mynewvnic
OCID: ocid1.vnic.oc1.phx.abcdefg12345678
MAC address: 00:00:17:00:27:A7
Public IP address: None
Subnet: Public Subnet cumS:PHX-AD-1 (10.0.0.0/24)
```

Operating System level network configuration

CONFIG	ADDR	SPREFIX	SBITS	VIRTRT	NS	IND	IFACE	VLTAG	VLAN
STATE	MAC	VNIC							
-	10.0.0.3	10.0.0.0	24	10.0.0.1	-	0	ens3	-	-
UP	00:00:17:00:f4:3f	ocid1.vnic.oc1.phx.abcdefg12345678							
-	10.0.0.4	10.0.0.0	24	10.0.0.1	-	1	ens4	-	-
UP	00:00:17:00:27:a7	ocid1.vnic.oc1.phx.abcdefg12345679							

DETACHING A VNIC

To detach a VNIC from the instance, use the `--detach-vnic` option. Note that the given VNIC cannot be the primary VNIC for the instance:

```
sudo oci-network-config --detach-vnic 00:00:17:00:27:A7
```

oci-network-inspector

DESCRIPTION

Displays a detailed report for a given compartment or network.

USAGE

```
oci-network-inspector [-C OCID] [-N OCID] [--help]
```

OPTIONS

```
-C | --COMPARTMENTOCID
```

Show report for the specified compartment.

CHAPTER 8 Compute

`-N` | `--VCN` *OCID*

Show report for the specified virtual cloud network.

`-H` | `--HELP`

Displays a summary of the command line options.

EXAMPLES

DISPLAYING A DETAILED REPORT FOR A SPECIFIED COMPARTMENT

Running the `oci-network-inspector` command and specifying an OCID with the `-C` parameter returns a detailed network report for that compartment:

```
$ oci-network-inspector -C ocid1.compartment.oc1..aaaaaaaabcdefghijklmnopqrstuvwxyz12345678

Compartment: scottb_sandbox (ocid1.compartment.oc1..aaaaaaaabcdefghijklmnopqrstuvwxyz12345678)

vcn: scottb_vcn
  Security List: Default Security List for scottb_vcn
    Ingress: tcp          0.0.0.0/0:-          ---:22
    Ingress: icmp        0.0.0.0/0:-          code-4:type-3
    Ingress: icmp        10.0.0.0/16:-        code-None:type-3
    Ingress: tcp          0.0.0.0/0:80         ---:80
    Ingress: tcp          0.0.0.0/0:43         ---:43
    Ingress: tcp          0.0.0.0/0:-          ---:-
    Egress : all         ---:-                0.0.0.0/0:-

  Subnet: Public Subnet cumS:PHX-AD-3 Availability domain: cumS:PHX-AD-3
    Cidr_block: 10.0.2.0/24 Domain name: sub999999999999.scottbvcn.oraclevcn.com
    Security List: Default Security List for scottb_vcn
      Ingress: tcp          0.0.0.0/0:-          ---:22
      Ingress: icmp        0.0.0.0/0:-          code-4:type-3
      Ingress: icmp        10.0.0.0/16:-        code-None:type-3
      Ingress: tcp          0.0.0.0/0:80         ---:80
      Ingress: tcp          0.0.0.0/0:43         ---:43
      Ingress: tcp          0.0.0.0/0:-          ---:-
      Egress : all         ---:-                0.0.0.0/0:-

  Subnet: Public Subnet cumS:PHX-AD-2 Availability domain: cumS:PHX-AD-2
    Cidr_block: 10.0.1.0/24 Domain name: sub999999999998.scottbvcn.oraclevcn.com
```


CHAPTER 8 Compute

If the Oracle Cloud Infrastructure SDK is not installed, the `oci-public-ip` utility uses the STUN (Session Traversal Utilities for NAT) protocol to discover IP address. For more information on STUN, see the [STUN Wikipedia article](#).

USAGE

```
oci-public-ip [-h] [-j] [-g] [-s source_IP] [-S STUN_server]  
[-L] [--instance-id OCID] [--help]
```

OPTIONS

`-H` | `--HUMAN-READABLE`

Display human readable output (default).

`-J` | `--JSON`

Display output in JSON.

`-G` | `GET`

Print the IP address only.

`-s` | `--SOURCEIP` *SOURCE_IP*

Specify the source IP address to use.

`-S` | `--STUN-SERVER` *STUN_SERVER*

Specify the STUN server to use.

`-L` | `--LIST-SERVERS`

Print a list of known STUN servers and exit.

`--INSTANCE-ID` *OCID*

Display the public IP address of the given instance instead of the current one. Requires the Oracle Cloud Infrastructure Python SDK to be installed and configured.

CHAPTER 8 Compute

`--HELP`

Displays a summary of the command line options.

EXAMPLES

DISPLAYING CURRENT IP ADDRESS

Running the `oci-public-ip` command with no options returns the IP address of the current instance:

```
# oci-public-ip
Public IP address: 128.0.0.1
```

DISPLAYING THE IP ADDRESS OF ANOTHER INSTANCE

You can pass in the OCID of a running instance with the `--instance-id` option to return the IP address for that instance:

```
# oci-public-ip --instance-id ocid1.instance.oc1.phx.abcdefg12345678
Public IP address: 128.0.0.4
```

LISTING STUN SERVERS

Use the `--list-servers` option to return a list of STUN servers:

```
# oci-public-ip --list-servers
stun.stunprotocol.org
stun.counterpath.net
stun.voxgratia.org
stun.callwithus.com
stun.ekiga.net
stun.ideasip.com
stun.voipbuster.com
stun.voiparound.com
stun.voipstunt.com
```

Compute Shapes

A shape is a template that determines the number of CPUs, amount of memory, and other resources allocated to a newly created instance. The following tables provide basic

CHAPTER 8 Compute

information about the shapes that are available for bare metal machines, virtual machines (VMs), and dedicated virtual machine hosts.



Note

AMD shapes, X5-based and X6-based shapes, and high performance computing (HPC) shapes are not available in Government Cloud realms.

Bare Metal Shapes

Shape	Instance Type	OCP U	Memory (GB)	Local Disk	Network Bandwidth ¹	Max VNICS Total: Linux	Max VNICS Total: Windows
BM.Standard1.36 2	X5-based standard compute	36	256	Block storage only	10 Gbps	36	1
BM.Standard.B1.44	X6-based standard compute	44	512	Block storage only	25 Gbps	44	None
BM.Standard2.52	X7-based standard compute: Intel CPUs	52	768	Block storage only	2 x 25 Gbps	52 total (26 per physical NIC)	27 total (1 on the first physical NIC, 26 on the second)

CHAPTER 8 Compute

Shape	Instance Type	OCP U	Memory (GB)	Local Disk	Network Bandwidth ¹	Max VNICs Total: Linux	Max VNICS Total: Windows
BM.Standard.E2.64	E2-based standard compute: AMD CPUs	64	512	Block storage only	2 x 25 Gbps	75	76 (1 on the first physical NIC, 75 on the second)
BM.DenseIO1.36 ²	X5-based dense I/O compute	36	512	28.8 TB NVMe SSD (9 drives)	10 Gbps	36	1
BM.DenseIO2.52	X7-based dense I/O compute	52	768	51.2 TB NVMe SSD (8 drives)	2 x 25 Gbps	52 total (26 per physical NIC)	27 total (1 on the first physical NIC, 26 on the second)

CHAPTER 8 Compute

Shape	Instance Type	OCP U	Memory (GB)	Local Disk	Network Bandwidth ¹	Max VNICs Total: Linux	Max VNICS Total: Windows
BM.GPU2.2	X7-based GPU compute: 2xP100 NVIDIA GPUs	28	192	Block storage only	2 x 25 Gbps	28	15 (1 on the first physical NIC, 14 on the second)
BM.GPU3.8	X7-based GPU compute: 8xV100 NVIDIA GPUs	52	768	Block storage only	2 x 25 Gbps	52	27 (1 on the first physical NIC, 26 on the second)
BM.HPC2.36	X7-based high frequency compute	36	384	6.7 TB NVMe SSD (1 drive)	1 x 25 Gbps 1 x 100 Gbps RDMA	50	1

1: Network bandwidth is based on expected bandwidth for traffic within a VCN.

2: X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.

VM Shapes

Shape	OCPU	Memory (GB)	Local Disk (TB)	Network Bandwidth 1	Max VNICs Total: Linux	Max VNICs Total: Windows
VM.Standard1.1 2	1	7	Block Storage only	Up to 600 Mbps	2	1
VM.Standard1.2 2	2	14	Block Storage only	Up to 1.2 Gbps	2	1
VM.Standard1.4 2	4	28	Block Storage only	1.2 Gbps	4	1
VM.Standard1.8 2	8	56	Block Storage only	2.4 Gbps	8	1
VM.Standard1.16 2	16	112	Block Storage only	4.8 Gbps	16	1
VM.Standard2.1	1	15	Block Storage only	1 Gbps	2	2
VM.Standard2.2	2	30	Block Storage only	2 Gbps	2	2

CHAPTER 8 Compute

Shape	OCPU	Memory (GB)	Local Disk (TB)	Network Bandwidth 1	Max VNICs Total: Linux	Max VNICs Total: Windows
VM.Standard2.4	4	60	Block Storage only	4.1 Gbps	4	4
VM.Standard2.8	8	120	Block Storage only	8.2 Gbps	8	8
VM.Standard2.16	16	240	Block Storage only	16.4 Gbps	16	16
VM.Standard2.24	24	320	Block Storage only	24.6 Gbps	24	24
VM.Standard.E2.1.Micro	1	1	Block Storage only	480 Mbps	1	-
VM.Standard.E2.1	1	8	Block Storage only	700 Mbps	2	2
VM.Standard.E2.2	2	16	Block Storage only	1.4 Gbps	2	2
VM.Standard.E2.4	4	32	Block Storage only	2.8 Gbps	4	4

CHAPTER 8 Compute

Shape	OCPU	Memory (GB)	Local Disk (TB)	Network Bandwidth 1	Max VNICs Total: Linux	Max VNICs Total: Windows
VM.Standard.E2.8	8	64	Block Storage only	5.6 Gbps	4	4
VM.DenseIO1.4 2	4	60	3.2 TB NVMe SSD	1.2 Gbps	4	1
VM.DenseIO1.8 2	8	120	6.4 TB NVMe SSD	2.4 Gbps	8	1
VM.DenseIO1.16 2	16	240	12.8 TB NVMe SSD	4.8 Gbps	16	1
VM.DenseIO2.8	8	120	6.4 TB NVMe SSD	8.2 Gbps	8	8
VM.DenseIO2.16	16	240	12.8 TB NVMe SSD	16.4 Gbps	16	16
VM.DenseIO2.24	24	320	25.6 TB NVMe SSD	24.6 Gbps	24	24
VM.GPU2.1 (GPU: 1xP100)	12	72	Block Storage only	8 Gbps	12	12

CHAPTER 8 Compute

Shape	OCPU	Memory (GB)	Local Disk (TB)	Network Bandwidth 1	Max VNICs Total: Linux	Max VNICs Total: Windows
VM.GPU3.1 (GPU: 1xV100)	6	90	Block Storage only	4 Gbps	6	6
VM.GPU3.2 (GPU: 2xV100)	12	180	Block Storage only	8 Gbps	12	12
VM.GPU3.4 (GPU: 4xV100)	24	360	Block Storage only	24.6 Gbps	24	24

1: Network bandwidth is based on expected bandwidth for traffic within a VCN.

2: X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.

Dedicated Virtual Machine Host Shapes

Shape	Instance Type	Billed OCPU	Usable OCPU 1	Supported Shapes for Hosted VMs
DVH.Standard2.52	X7-based VM host	52	48	VM.Standard2

1: The difference between billed OCPUs and usable OCPUs is due to OCPUs reserved for hypervisor use.

Installing and Running Oracle Ksplice

Oracle Ksplice enables you to apply important security and other critical kernel updates without a reboot. For more information, see [About Oracle Ksplice](#) and [Ksplice Overview](#).

This topic describes how to install and configure Ksplice. Ksplice is only available for Oracle Linux instances launched on or after February 15, 2017. It is installed on instances launched on or after August 25, 2017, so you just need to run it on these instances to install the available Ksplice patches. For instances launched prior to August 25, 2017, you must install it prior to running it.

Installing Ksplice on instances launched prior to August 25 2017

To install Ksplice you need to connect to your Linux instance by using a Secure Shell (SSH). See [Connecting to an Instance](#) for more information.

1. Use the following SSH command to access the instance.

```
ssh -l opc@<public-ip-address>
```

<public-ip-address> is your instance IP address that you retrieved from the Console, see [Getting the Instance Public IP Address](#).

2. Run the following SSH commands to sudo to the root:

```
sudo bash
```

3. Download the Ksplice installation script with the following SSH command:

```
wget -N https://www.ksplice.com/uptrack/install-uptrack-oc
```

4. Once the script is downloaded, use the following SSH command to install Ksplice:

```
sh install-uptrack-oc
```

Running Ksplice

To run Ksplice you need to connect to your Linux instance by using a Secure Shell (SSH). See [Connecting to an Instance](#) for more information.

1. Use the following SSH command to access the instance.

```
ssh -l opc <public-ip-address>
```

<public-ip-address> is your instance IP address that you retrieved from the Console, see [Getting the Instance Public IP Address](#).

2. Run the following SSH commands to sudo to the root:

```
sudo bash  
cd
```

3. To install available Ksplice patches, run the following SSH command:

```
uptrack-upgrade
```

Automatic Updates

You can configure automatic updates by setting the value of `autoinstall` to `yes` in `/etc/uptrack/uptrack.conf`.



Note

OS Security Updates for Oracle Linux images

Oracle Linux images are updated regularly with the necessary patches, but after you launch an instance using these images, you are responsible for applying the required OS security updates published through the Oracle public Yum server. For more information, see [Installing and Using the Yum Security Plugin](#).

Managing Custom Images

Oracle Cloud Infrastructure uses [images](#) to launch instances. You specify an image to use when you [launch an instance](#).

You can create a custom image of a bare metal instance's boot disk and use it to launch other instances. Instances you launch from your image include the customizations, configuration, and software installed when you created the image.

For details on Windows images, see [Creating Windows Custom Images](#).

Custom images do not include the data from any attached block volumes. For information about backing up volumes, see [Backing Up a Volume](#).



Tip

Oracle Cloud Infrastructure runs on Oracle's high-quality Sun servers. However, any hardware can experience a failure. Follow industry-wide hardware failure best practices to ensure the resilience of your solution. Some best practices include:

- Design your system with redundant compute nodes in different availability domains to support failover capability.
- Create a [custom image](#) of your system drive each time you change the image.
- [Back up](#) your data drives, or sync to spare drives, regularly.

If you experience a hardware failure and have followed these practices, you can terminate the failed instance, launch your custom image to create a new instance, and then apply the backup data.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to create and manage images. If the specified group doesn't need to launch instances or attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.



Tip

When users create a custom image from an instance or launch an instance from a custom image, the instance and image don't have to be in the same compartment. However, users must have access to both compartments.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Limitations and Considerations

- Certain IP addresses are reserved for Oracle Cloud Infrastructure use and may not be used in your address numbering scheme. See [IP Addresses Reserved for Use by Oracle](#) for more information.
- Before you create a custom image of an instance, you must disconnect all iSCSI attachments and remove all iscsid node configurations from the instance. For steps, see [Disconnecting From a Volume](#).
- When you create an image of a running instance, the instance shuts down and remains unavailable for several minutes. The instance restarts when the process completes.
- You cannot create additional custom images of an instance while the instance is engaged in the image creation process. When you start to create a custom image, the system implements a 20-minute timeout, during which you cannot create another image of the same instance. You can, however, create images of different instances at the same time.
- Custom images are available to all users authorized for the compartment in which the image was created.
- Custom images inherit the compatible shapes that are set by default from the base image.
- The maximum size for importing a custom image is 400 GB.
- The maximum size for custom exported images is 400 GB.
- You can create a maximum of 25 custom images per region per root compartment.
- You cannot create an image of an Oracle Database instance.
- If you use a custom image and update the OS kernel on your instance, you must also upload the update to the network drive. See [OS Kernel Updates](#) for more information.
- Editing custom Windows images is not supported due to hardware differences between shapes.

For information about how to deploy any version of any operating system that is supported by the Oracle Cloud Infrastructure hardware, see [Bring Your Own Image \(BYOI\)](#).

X5 and X7 Compatibility for Custom Images

Oracle X5, X6, and X7 servers have different host hardware. As a result, using an X5 or X6 image on an X7 bare metal or virtual machine (VM) instance may not work without additional modifications. Oracle Cloud Infrastructure recommends for X7 hosts that you use the Oracle-provided images for X7. See [Oracle-Provided Image Release Notes](#) for more information about which images support X7. These images have been explicitly created and tested with X7 hardware.

If you attempt to use an existing X5 image on X7 hardware, note the following:

- CentOS 6 and all Windows versions are not cross-compatible.
- Oracle Linux, Ubuntu 16.04, and CentOS 7 are cross-compatible. However, you must update the kernel to the most recent version to install the latest device drivers. To do this, run the following commands from a terminal session:

- **Oracle Linux**

```
yum update
```

- **CentOS 7**

```
yum update
```

- **Ubuntu 16.04**

```
apt-get update  
apt-get dist-upgrade
```

If you attempt to use an X6 image on non-X6 hardware, note the following:

- Oracle Linux 6, all CentOS versions, and all Windows versions are not cross-compatible.
- Oracle Linux 7, Ubuntu 18.04, and Ubuntu 16.04 are cross-compatible. Use the Oracle-provided images for X6.

The primary device drivers that are different between X5, X6, and X7 hosts are:

- Network device drivers
- NVMe drive device drivers

- GPU device drivers

Additional updates may be required depending on how you have customized the image.

Using the Console

To access the Console, you must use a [supported browser](#).

To create a custom image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Find the instance you want to use as the basis for an image.
3. Click the Actions icon (three dots), and then click **Create Custom Image**.
4. Enter a name for the image, and then click **Create Custom Image**. You can [use the API](#) to change the name later, if needed. You cannot use an Oracle-provided image name as a custom image name.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).



Note

Limits

If you see a message indicating that you are at the limit for custom images, you must delete at least one image before you can create another. Or, you can [request a service limit increase](#).

To launch an instance from a custom image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**. Find the custom image you want to use.
2. Click the Actions icon (three dots), and then click **Create Instance**.
3. Provide additional launch options as described in [Creating an Instance](#).

To edit the name or shape of a custom image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. Click the custom image you're interested in.
3. Click **Edit Details**
4. Edit the name, or add and remove compatible shapes for the custom image.
5. Click **Save**.

To manage tags for a custom image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. Click the custom image you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

To delete a custom image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.

2. Find the custom image you want to delete.
3. Click the Actions icon (three dots), and then click **Delete**.
4. Confirm when prompted.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage custom images:

- [DeleteImage](#)
- [GetImage](#)
- [ListImages](#)
- [CreateImage](#)
- [UpdateImage](#)

Creating Windows Custom Images

You can create a Windows custom image of a bare metal or VM instance's boot disk and use it to launch other instances. Instances you launch from your image include the customizations, configuration, and software installed when you created the image. For information about custom images, see [Managing Custom Images](#). For information about the licensing requirements for Windows images, see [Microsoft Licensing on Oracle Cloud Infrastructure](#).

Windows supports two kinds of images: generalized and specialized. Generalized images are images that have been cleaned of instance-specific information. Specialized images are point-in-time snapshots of the boot disk of a running instance, and are useful for creating backups of an instance. Oracle Cloud Infrastructure supports bare metal and VM instances launched from both generalized and specialized custom Windows images.

Generalized images

A generalized image has a generalized OS disk, cleaned of computer-specific information. The images are generalized using Sysprep. Generalized images can be useful in scenarios such as quickly scaling an environment. Generalized images can be configured to preserve the existing `opc` user's account, including the password, at the time the image is created, or configured to recreate the `opc` user account, including generating a new, random password that you retrieve using the API. For background information, see [Sysprep \(Generalize\) a Windows installation](#).

Specialized images

A specialized image has an OS disk that is already fully installed, and is essentially a copy of the original bare metal or VM instance. Specialized images are intended to be used for backups so that you can recover from a failure. Specialized images are useful when you are testing a task and may need to roll back to known good configuration. Specialized images are not recommended for cloning multiple identical bare metal instances or VMs in the same network because of issues with multiple computers having the same computer name and ID. When creating a specialized image, you must remember the `opc` user's password; a new password is not generated when the instance launches, and it cannot be retrieved from the console or API.

Creating a Generalized Image



Warning

Creating a generalized image from an instance will render the instance non-functional, so you should first create a custom image from the instance, and then launch a new instance from the custom image. Steps 1 - 2 describe how to do this. This is the instance that you'll generalize. Alternatively, you can make a backup image of the instance that you can use to launch a replacement instance if needed.



Warning

If you upgrade to PowerShell 5.0/WMF 5.0, you may encounter an issue where **Sysprep** fails which will prevent the image generalization process from completing. If this occurs, you may not be able to log into instances launched from the custom image. See [Unable to log into instance launched from new Windows custom image](#) for more information and how to work around the issue.

1. Create the new image using [To create a custom image](#).
2. Launch an instance from the new image using [To launch an instance from a custom image](#).
3. Connect to the instance using a Remote Desktop client.
4. Go to [Windows Generalized Image Support Files](#) and download to the instance the file matching the Windows version for the instance.
5. Right-click the file, and then click **Run as administrator**.
6. Extract the files to **C:\Windows\Panther**. The following files are extracted into the Panther folder for all Windows Server versions:
 - Generalize.cmd
 - Specialize.cmd
 - unattend.xml
 - Post-Generalize.ps1For Windows Server 2008, the following file is also extracted into the Panther folder:
 - Windows2008-SnapshotUtilities.ps1
7. Optional: If you want to preserve the `opc` user account, edit `C:\Program Files\bmcs\imageType.json` and change the `imageType` setting to `custom`. A new password is not created and the password is not retrievable from the console or API.

If you want to configure the generalized image to recreate the `opc` user account when a new instance is launched from the image, leave the `imageType` setting defaulted to `general`. The new account's password can be retrieved through the API using [GetWindowsInstanceInitialCredentials](#).

8. Right-click **Generalize.cmd**, and then click **Run as administrator**. Keep in mind the following outcomes of running this command:
 - Your connection to the Remote Desktop client may immediately be turned off and you will be logged out of the instance. If this does not occur, you should log out of the instance yourself.
 - Because `sysprep generalize` turns off Remote Desktop, you won't be able to log in to the instance again.
 - Creating a generalized image essentially destroys the instance's functionality.

You should wait for a few minutes before proceeding to the following step to ensure the generalization process has completed.

9. Create the new image using [To create a custom image](#).
10. After you create an image from an instance that has been generalized, we recommend that you terminate the instance. Although it may appear to be running, it won't be fully operable.

Creating a Specialized Image



Important

When creating a specialized image, you must remember the `opc` user's password. It cannot be retrieved from the Console or API.

You create a specialized image the same way you create other custom images. For step-by-step instructions, see [Managing Custom Images](#).

Image Import/Export

Oracle Cloud Infrastructure Compute enables you to share custom images across tenancies and regions using image import/export.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Linux-Based Operating Systems

The following Oracle Cloud Infrastructure operating systems support image import/export:

- Oracle Linux 6, Oracle Linux 7
- CentOS 6, CentOS 7
- Ubuntu 16.04

For more information about these images, see [Oracle-Provided Images](#).

Windows-Based Operating Systems

The following Windows versions support image import/export:

- Windows Server 2008 R2* Standard, Enterprise, Datacenter
- Windows Server 2012 Standard, Datacenter
- Windows Server 2012 R2 Standard, Datacenter
- Windows Server 2016 Standard, Datacenter
- Windows Server 2019 Standard, Datacenter

* Windows Server 2008 R2 reaches end of support on January 14, 2020.

For information about the licensing requirements for Windows images, see [Microsoft Licensing on Oracle Cloud Infrastructure](#).

Verify Your Windows Operating System

When importing custom Windows images, ensure that the version you select matches the Windows image that you imported. Failure to provide the correct version and SKU information could be a violation of your Microsoft Licensing Agreement.

WINDOWS SYSTEM TIME ISSUE ON CUSTOM WINDOWS INSTANCES

If you change the time zone from the default setting on Windows VM instances, when the instance reboots or syncs with the hardware clock, the system time will revert back to the time for the default time zone. However, the time zone setting will stay set to the new time zone, so the system clock will be incorrect. You can fix this by setting the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\TimeZoneInformation` registry key to 1.

Oracle-provided Windows images already have the `RealTimeIsUniversal` registry key set by default, but you must set this for any custom Windows images that you import.

To fix this issue for custom Windows images:

1. Open the Windows Registry Editor and navigate to the `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\TimeZoneInformation` registry key.
2. Create a new `DWORD` key named `RealTimeIsUniversal` and set the value to 1.
3. Reboot the instance.
4. Reset the time and time zone manually.

Bring Your Own Image Scenarios

You can also use image import/export to share custom images from [Bring Your Own Image \(BYOI\)](#) scenarios across tenancies and regions, so you don't need to recreate the image manually in each region. You must go through the steps required to manually create the image in one of the regions, but after this is done, you can export the image, making it available for import in additional tenants and regions. The exported image format is `.oci`,

which is a TAR file that contains a QCOW2 file and Oracle Cloud Infrastructure-specific metadata.

Best practices for replicating an image across regions

You can replicate an image from one region to another region using the Console or API. At a high level:

1. Export the image to an Object Storage bucket in the same region as the image. For steps, see [Exporting an Image](#).
2. Copy the image to an Object Storage bucket in the destination region. For steps, see [Copying Objects](#).
3. Obtain the URL path to the image object. For steps, see [To view object details](#).
4. In the destination region, import the image. Use the URL path as the Object Storage URL. For steps, see [Importing an Image](#).

Best practices for sharing an image across tenancies

You can replicate an image from one tenancy to another tenancy using the Console or API. At a high level:

1. Export the image to an Object Storage bucket in the same region as the image. For steps, see [Exporting an Image](#).
2. Create a pre-authenticated request with read-only access for the image in the destination region. For steps, see [Working with Pre-Authenticated Requests](#).
3. In the destination tenancy, import the image. Use the pre-authenticated request URL as the Object Storage URL. For steps, see [Importing an Image](#).

Object Storage Service URLs

When you import or export custom images using the Console, you might need to specify the Object Storage URL pointing to the location that you want to import the image from or export the image to. Object Storage URLs are structured as follows:

```
https://<host_name>/n/<namespace_name>/b/<bucket_name>/o/<object_name>
```

For example:

```
https://objectstorage.us-phoenix-1.oraclecloud.com/n/MyNamespace/b/MyBucket/o/MyCustomImage.qcow2
```

Pre-Authenticated Requests

When using import/export across tenancies, you need to use an Object Storage pre-authenticated request. See [Working with Pre-Authenticated Requests](#) for steps to create a pre-authenticated request. When you go through these steps, after you click **Create Pre-Authenticated Request**, the **Pre-Authenticated Request Details** dialog box opens. You must make a copy of the pre-authenticated request URL displayed here, because this is the only time this URL is displayed. This is the Object Storage URL that you specify for import/export.



Note

Pre-authenticated requests for a bucket

With image export, if you create the pre-authenticated request for a bucket, you need to append the object name to the generated URL. For example:

```
/o/MyCustomImage.qcow2
```

Exporting an Image

You can use the Console or API to export images, and the exported images are stored in the Oracle Cloud Infrastructure Object Storage service. To perform an image export, you need write access to the Object Storage bucket for the image. For more information, see [Overview of Object Storage](#) and [Let users write objects to Object Storage buckets](#).

To export an image using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. Find the custom image you want to export, click the Actions icon (three dots), and then click **Export Custom Image**.
3. In the **Export Image** dialog box, specify the Object Storage location to export the image to. You have two options here: You can select a compartment and bucket, and then enter a name for the exported image, or you can enter the Object Storage URL.
4. Click **Export Image**.

After you click **Export Image**, the image status changes to **EXPORTING**. You can still launch instances while the image is exporting, but you can't delete the image until the export has finished. To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

When the export is complete, the image status changes to **AVAILABLE**. If the image status changes to **AVAILABLE**, but you don't see the exported image in the Object Storage location you specified, this means that the export failed, and you will need to go through the steps again to export the image.

Importing an Image

You can use the Console or API to import exported images from Object Storage. To import an image, you need read access to the Object Storage object containing the image. For more information, see [Let users download objects from Object Storage buckets](#).

To import an image using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. Click **Import Image**.
3. Select the compartment that you want to import the image to.

4. Enter a name for the image.
5. Select the **Operating System**:
 - For Linux images, select **Linux**.
 - For Windows images, select **Windows**. Select the **Operating System Version**, and then certify that the selected operating system complies with Microsoft licensing agreements.
6. Specify the **Object Storage URL** where the image is stored. When importing across tenancies, you must specify a pre-authenticated request URL.
7. Select the **Image Type**.
8. Select the **Launch Mode**:
 - For custom images where the image format is `.oci`, Oracle Cloud Infrastructure selects the applicable launch mode based on the launch mode for the source image.
 - For custom images exported from Oracle Cloud Infrastructure where the image type is QCOW2, select **Native Mode**.
 - To import other custom images select **Paravirtualized Mode** or **Emulated Mode**. For more information, see [Bring Your Own Image \(BYOI\)](#).
9. **Tags**: Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
10. Click **Import Image**.

After you click **Import Image**, you'll see the imported image in the **Custom Images** list for the compartment, with a status of **IMPORTING**. To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

When the import completes successfully, the status changes to **AVAILABLE**. If the status does not change, or no entry appears in the **Custom Images** list, the import failed. If the

import failed, ensure you have read access to the Object Storage object, and that the object contains a supported image.

Editing Image Details

You can edit the details of custom images, such as the image name and compatible shapes for the image. For more information, see [To edit the name or shape of a custom image](#) in [Managing Custom Images](#).

Managing Tags for an Image

You can apply tags to your resources, such as images, to help you organize them according to your business needs. You can apply tags at the time you import an image, or you can update the image later with the desired tags.

To manage tags for an image

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. Click the image you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following API operations for custom image import/export:

- [ExportImage](#): Exports a custom image to Object Storage.
- [CreateImage](#): To import an exported image, specify [ImageSourceDetails](#) in the request body.
- [AddImageShapeCompatibilityEntry](#): Adds a shape to the compatible shapes list for the image.
- [RemoveImageShapeCompatibilityEntry](#): Removes a shape from the compatible shapes list for the image.

X5 and X7 Compatibility for Image Import/Export

Oracle X5, X6, and X7 servers have different host hardware. As a result, using an X5 or X6 image on an X7 bare metal or virtual machine (VM) instance may not work without additional modifications. Oracle Cloud Infrastructure recommends for X7 hosts that you use the Oracle-provided images for X7. See [Oracle-Provided Image Release Notes](#) for more information about which images support X7. These images have been explicitly created and tested with X7 hardware.

If you attempt to use an existing X5 image on X7 hardware, note the following:

- CentOS 6 and all Windows versions are not cross-compatible.
- Oracle Linux, Ubuntu 16.04, and CentOS 7 are cross-compatible. However, you must update the kernel to the most recent version to install the latest device drivers. To do this, run the following commands from a terminal session:

- **Oracle Linux**

```
yum update
```

- **CentOS 7**

```
yum update
```

- **Ubuntu 16.04**

```
apt-get update  
apt-get dist-upgrade
```

If you attempt to use an X6 image on non-X6 hardware, note the following:

- Oracle Linux 6, all CentOS versions, and all Windows versions are not cross-compatible.
- Oracle Linux 7, Ubuntu 18.04, and Ubuntu 16.04 are cross-compatible. Use the Oracle-provided images for X6.

The primary device drivers that are different between X5, X6, and X7 hosts are:

- Network device drivers
- NVMe drive device drivers
- GPU device drivers

Additional updates may be required depending on how you have customized the image.

Bring Your Own Image (BYOI)

The Bring Your Own Image (BYOI) feature enables you to bring your own versions of operating systems to the cloud as long as the underlying hardware supports it. The services do not depend on the OS you run.

The BYOI feature:

- Enables lift-and-shift cloud migration projects.
- Supports both old and new operating systems.
- Encourages experimentation.
- Increases infrastructure flexibility.



Note

Licensing Requirements

You must comply with all licensing requirements when you upload and start instances based on OS images that you supply.

Bringing Your Own Image

A critical part of any lift-and-shift cloud migration project is the migration of on-premises virtual machines (VMs) to the cloud. You can import your on-premises virtualized root volumes to Oracle Cloud Infrastructure using the custom image import feature, and then launch Compute instances using those images.

You can import Windows and Linux-based custom images.

- **Windows images**

These Windows versions support custom image import:

- Windows Server 2008 R2* Standard, Enterprise, Datacenter
- Windows Server 2012 Standard, Datacenter
- Windows Server 2012 R2 Standard, Datacenter
- Windows Server 2016 Standard, Datacenter
- Windows Server 2019 Standard, Datacenter

* Windows Server 2008 R2 reaches [end of support](#) on January 14, 2020.

For steps to import a Windows image, see [Importing Custom Windows Images](#).

For information about the licensing requirements for Windows images, see [Microsoft Licensing on Oracle Cloud Infrastructure](#).

- **Linux images**

These Linux distributions support custom image import:

Linux Distribution	Supported Versions	Preferred Launch Mode
RHEL	7 or later	Paravirtualized
	4.5, 5.5, 5.6, 5.9, 5.11, 6.5, 6.9	Emulated
CentOS	7 or later	Paravirtualized
	4.0, 4.8, 5.11, 6.9	Emulated

Linux Distribution	Supported Versions	Preferred Launch Mode
Oracle Linux	7 or later	Paravirtualized
	4.5, 4.8, 5.8, 5.11, 6.2, 6.5	Emulated
Ubuntu	13.04 or later	Paravirtualized
	12.04	Emulated
FreeBSD	12 or later	Paravirtualized
	8, 9, 10, 11	Emulated
Debian	8 or later	Paravirtualized
	5.0.10, 6.0, 7	Emulated
SUSE	12.2 or later	Paravirtualized
	11, 12.1	Emulated

You might also have success importing other distributions of Linux.

For steps to import a Linux-based image, see [Importing Custom Linux Images](#).

Bringing Your Own Hypervisor

- **Bring your own KVM:** You can bring your own operating system images or older operating systems such as Ubuntu 6.x, RHEL 3.x, CentOS 5.4 using KVM on bare metal instances. For more information, see [Getting Started: Oracle Linux KVM Image for Oracle Cloud Infrastructure](#) and [Installing and Configuring KVM on Bare Metal Instances with Multi-VNIC](#).
- **Bring your own Hyper-V:** You can bring your own operating system images or older operating systems such as Windows Server 2003, Windows Server 2008, and older Linux -based operating systems using Hyper-V on bare metal instances. For a full list of supported Hyper-V guests, see [Supported Windows guest operating systems for Hyper-](#)

[V on Windows Server](#) and [Supported Linux and FreeBSD virtual machines for Hyper-V on Windows](#). See [Deploying Hyper-V on Oracle Cloud Infrastructure](#) for more information.

Importing Custom Windows Images

The Compute service enables you to import Windows images that were created outside of Oracle Cloud Infrastructure. For example, you can import images running on your on-premises physical or virtual machines (VMs), or VMs running in Oracle Cloud Infrastructure Classic. You can then launch your imported images on Compute virtual machines.



Note

Imported images are not supported on AMD shapes.

For information about the licensing requirements for Windows images, see [Microsoft Licensing on Oracle Cloud Infrastructure](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Supported Operating Systems

These Windows versions support custom image import:

- Windows Server 2008 R2* Standard, Enterprise, Datacenter
- Windows Server 2012 Standard, Datacenter
- Windows Server 2012 R2 Standard, Datacenter

CHAPTER 8 Compute

- Windows Server 2016 Standard, Datacenter
- Windows Server 2019 Standard, Datacenter

* Windows Server 2008 R2 reaches [end of support](#) on January 14, 2020.



Note

- Oracle Cloud Infrastructure has tested the operating systems listed previously and will support customers in ensuring that instances launched from these images and built according to the guidelines in this topic are accessible using RDP.
- For OS editions not listed previously, Oracle Cloud Infrastructure will provide commercially reasonable support to customers in an effort to get instances that are launched from these images accessible via RDP.
- Support from Oracle Cloud Infrastructure in launching an instance from a custom OS does not ensure that the operating system vendor also supports the instance.
- Oracle Cloud Infrastructure licenses and charges the Windows licensing fee for all instances launched using an imported Windows OS image. This applies whether or not those instances are registered with Oracle Cloud Infrastructure's Microsoft Key Management service.

Windows Source Image Requirements

Custom images must meet the following requirements:

- The maximum image size is 400 GB.
- The image must be set up for BIOS boot.

- Only one disk is supported, and it must be the boot drive with a valid master boot record (MBR) and boot loader. You can migrate additional data volumes after you import the image's boot volume.
- The boot process must not require additional data volumes to be present for a successful boot.
- The disk image cannot be encrypted.
- The disk image must be a VMDK or QCOW2 file.
 - Create the image file by cloning the source volume, not by creating a snapshot.
 - VMDK files must be either the "single growable" (monolithicSparse) type or the "stream optimized" (streamOptimized) type, both of which consist of a single VMDK file. All other VMDK formats, such as those that use multiple files, split volumes, or contain snapshots, are not supported.
- The network interface must use DHCP to discover the network settings. When you import a custom image, existing network interfaces are not recreated. Any existing network interfaces are replaced with a single NIC after the import process is complete. You can attach additional VNICs after you launch the imported instance.
- The network configuration must not hardcode the MAC address for the network interface.

Preparing Windows VMs for Import

Before you can import a custom Windows image, you must prepare the image to ensure that instances launched from the image can boot correctly and that network connections will work.

You can perform the tasks described in this section on the running source system. If you have concerns about modifying the live source system, you can export the image as-is, import it into Oracle Cloud Infrastructure, and then launch an instance based on the custom image. You can then [connect to the instance using the VNC console](#) and perform the preparation steps.



Important

The system drive where Windows is installed will be imported to Oracle Cloud Infrastructure. All partitions on the drive will follow through the imported image. Any other drives will not be imported and you must re-create them on the instance after import. You will then need to manually move the data on the non-system drives.

To prepare a Windows VM for import:

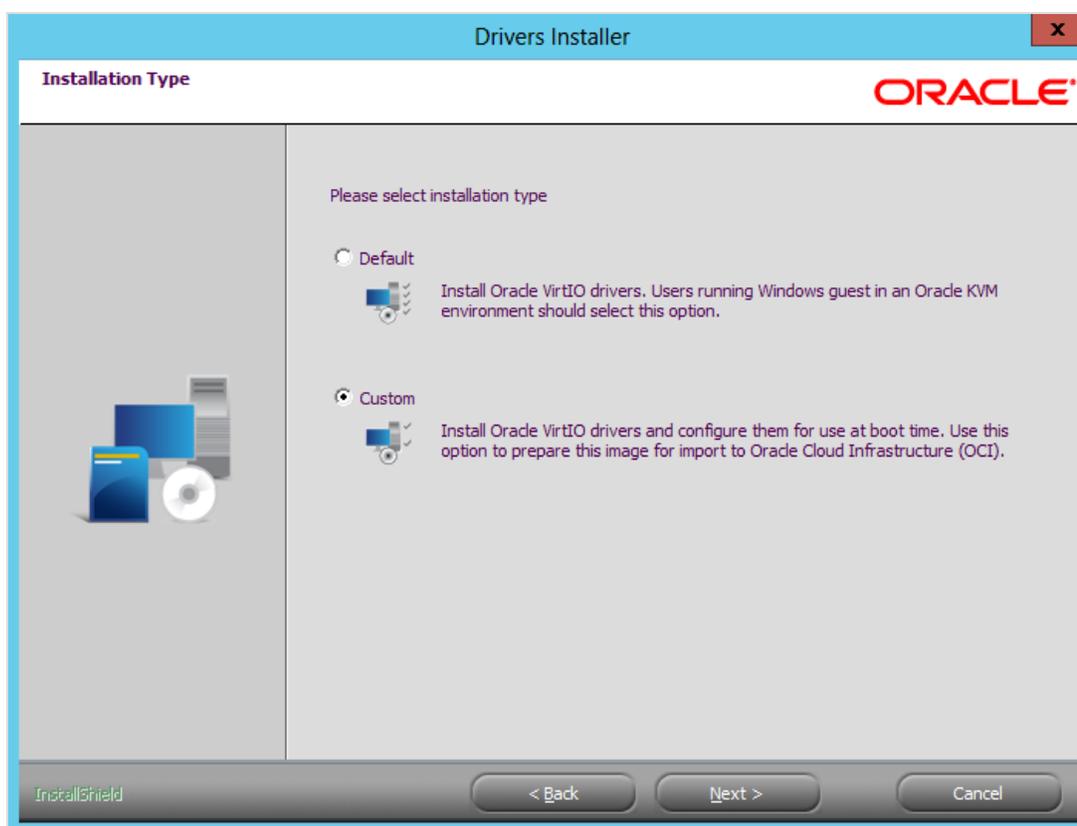
1. Follow your organization's security guidelines to ensure that the Windows system is secured. This can include, but is not limited to the following tasks:
 - Install the latest security updates for the operating system and installed applications.
 - Enable the firewall, and configure it so that you only enable the rules which are needed.
 - Disable unnecessary privileged accounts.
 - Use strong passwords for all accounts.
2. Configure Remote Desktop Protocol (RDP) access to the image:
 - a. [Enable Remote Desktop connections](#) to the image.
 - b. [Modify the Windows Firewall inbound port rule](#) to allow RDP access for both Private and Public network location types. When you import the image, the Windows Network Location Awareness service will identify the network connection as a Public network type.
3. Determine whether the current Windows license type is a volume license by running the following command in PowerShell:

```
Get-CimInstance -ClassName SoftwareLicensingProduct | where {$_.PartialProductKey} | select ProductKeyChannel
```

If the license is not a volume license, after you import the image, you will update the license type.

4. If you plan to launch the imported image on more than one VM instance, [create a generalized image](#) of the boot disk. A generalized image is cleaned of computer-specific information, such as unique identifiers. When you create instances from a generalized image, the unique identifiers are regenerated. This prevents two instances that are created from the same image from colliding on the same identifiers.
5. Create a backup of the root volume.
6. If the VM has remotely attached storage, such as NFS or block volumes, configure any services that rely on this storage to start manually. Remotely attached storage is not available the first time that an imported instance boots on Oracle Cloud Infrastructure.
7. Ensure that all network interfaces use DHCP, and that the MAC address and IP addresses are not hardcoded. See your system documentation for steps to perform network configuration for your system.
8. Download the Oracle Windows VirtIO drivers:
 - a. Log in to the [Oracle Software Delivery Cloud site](#).
 - b. In the **All Categories** list, select **Release**.
 - c. Type **Oracle Linux 7.6** in the search box and click **Search**.
 - d. Add **REL: Oracle Linux 7.6.x** to your cart, and then click **Checkout**.
 - e. In the **Platforms/Languages** list, select **x86 64 bit**. Click **Continue**.
 - f. Accept the license agreement and then click **Continue**.
 - g. Select the check box next to **Oracle VirtIO Drivers Version for Microsoft Windows 1.1.4**. Clear the other check boxes.
 - h. Click **Download** and then follow the prompts.

9. Install the Oracle VirtIO drivers for Windows:
 - a. Follow the prompts in the installation wizard. On the **Installation Type** page, select **Custom**, as shown in the following screenshot.



- b. Reboot the VM.
10. Stop the VM.
11. Clone the stopped VM as a VMDK or QCOW2 file, and then export the image from your virtualization environment. See the tools documentation for your virtualization environment for steps.

Importing a Windows-Based VM

After you prepare a Windows image for import, follow these steps to import the image:

1. Upload the VMDK or QCOW2 file to Oracle Cloud Infrastructure:
 - a. [Upload the file to an Object Storage bucket](#). You can upload the file using the Console or using the [command line interface \(CLI\)](#). If you use the CLI, use the following command:

```
oci os object put -bn <destination_bucket_name> --file <path_to_the_VMDK_or_QCOW2_file>
```
 - b. Copy the URL of the file that you uploaded: On the **Bucket Details** page, click the Actions icon (three dots) next to the file, and then click **Details**. Copy the **URL Path (URI)**.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
3. Click **Import Image**.
4. In the **Create in Compartment** list, select the compartment that you want to import the image to.
5. Enter a name for the image.
6. For the **Operating System**, select **Windows**.
7. In the **Operating System Version** list, select the version of Windows.
8. Confirm that you chose the operating system version that complies with your Microsoft licensing agreement, and then select the compliance check box.



Important

Failure to provide the correct version and SKU information could be a violation of your Microsoft Licensing Agreement.

9. In the **Object Storage URL** field, paste the URL of the file that you uploaded.

10. For the **Image Type**, select the file type of the image, either **VMDK** or **QCOW2**.
11. In the **Launch Mode** area, select **Paravirtualized Mode**.
12. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
13. Click **Import Image**.

The imported image appears in the **Custom Images** list for the compartment, with a status of **IMPORTING**. When the import completes successfully, the status changes to **AVAILABLE**.

If the status doesn't change, or no entry appears in the **Custom Images** list, the import failed. Ensure that you have read access to the Object Storage object, and that the object contains a supported image.
14. Complete the post-import tasks.

Post-Import Tasks for Windows Images

After you import a custom Windows-based image, do the following:

1. [Create an instance based on the custom image](#). For the image source, select **Custom Images**, and then select the image that you imported.
2. [Enable Remote Desktop Protocol \(RDP\) access](#) to the Compute instance.
3. [Connect to the instance using RDP](#).
4. If the instance requires any remotely attached storage, such as [block volumes](#) or [file storage](#), create and attach it.
5. [Create and attach any required secondary VNICs](#).
6. Test that all applications are working as expected.
7. Reset any services that were set to start manually.

CHAPTER 8 Compute

8. Register the instance with the Oracle-provided Key Management Service (KMS) server:
 - a. On the instance, open PowerShell as Administrator.
 - b. To set the KMS endpoint, run the following command:

```
slmgr /skms 169.254.169.253:1688
```

- c. If the Windows license type that you noted while preparing the image isn't a volume license, you must update the license type. Run the following command:

```
slmgr /ipk <setup key>
```

<setup key> is the KMS client setup key that corresponds to the version of Windows that you imported:

Windows Version	KMS Client Setup Key
Windows Server 2008 R2 Standard	YC6KT-GKW9T-YTKYR-T4X34-R7VHC
Windows Server 2008 R2 Enterprise	489J6-VHDMP-X63PK-3K798-CPX3Y
Windows Server 2008 R2 Datacenter	74YFP-3QFB3-KQT8W-PMXWJ-7M648
Windows Server 2012 Standard	XC9B7-NBPP2-83J2H-RHMBY-92BT4
Windows Server 2012 Datacenter	48HP8-DN98B-MYWDG-T2DCC-8W83P
Windows Server 2012 R2 Standard	D2N9P-3P6X9-2R39C-7RTCD-MDVJX
Windows Server 2012 R2 Datacenter	W3GGN-FT8W3-Y4M27-J84CP-Q3VJ9
Windows Server 2016 Standard	WC2BQ-8NRM3-FDDYY-2BFGV-KHKQY
Windows Server 2016 Datacenter	CB7KF-BWN84-R7R2Y-793K2-8XDDG
Windows Server 2019 Standard	N69G4-B89J2-4G8F4-WWYCC-J464C
Windows Server 2019 Datacenter	WMDGN-G9PQG-XVXX-R3X43-63DFG

- d. To activate Windows, run the following command:

```
slmgr /ato
```

- e. To verify the license status, run the following command:

```
Get-CimInstance -ClassName SoftwareLicensingProduct | where {$_.PartialProductKey} |  
select Description, LicenseStatus
```

If the `LicenseStatus` is 1, the instance is properly licensed. It might take up to 48 hours for the license status to update.

Importing Custom Linux Images

The Compute service enables you to import Linux-based images that were created outside of Oracle Cloud Infrastructure. For example, you can import images running on your on-premises physical or virtual machines (VMs), or VMs running in Oracle Cloud Infrastructure Classic. You can then launch your imported images on Compute virtual machines.



Note

Imported images are not supported on AMD shapes.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Supported Operating Systems

You can launch imported Linux VMs in either paravirtualized mode or emulated mode.

CHAPTER 8 Compute

Paravirtualized mode offers better performance than emulated mode. We recommend that you use paravirtualized mode if your OS supports it. Linux-based operating systems running the kernel version 3.4 or later support paravirtualized drivers. You can verify your system's kernel version using the [uname](#) command.

To verify the kernel version using the uname command

Run the following command:

```
uname -a
```

The output should look similar to this sample:

```
Linux ip_bash 4.14.35-1818.2.1.el7uek.x86_64 #2 SMP Mon Aug 27 21:16:31 PDT 2018 x86_64 x86_64 x86_64
GNU/Linux
```

The kernel version is the number at the first part of output string. In the sample output shown previously, the version is 4.14.35.

These Linux distributions support custom image import:

Linux Distribution	Supported Versions	Preferred Launch Mode
RHEL	7 or later	Paravirtualized
	4.5, 5.5, 5.6, 5.9, 5.11, 6.5, 6.9	Emulated
CentOS	7 or later	Paravirtualized
	4.0, 4.8, 5.11, 6.9	Emulated
Oracle Linux	7 or later	Paravirtualized
	4.5, 4.8, 5.8, 5.11, 6.2, 6.5	Emulated
Ubuntu	13.04 or later	Paravirtualized
	12.04	Emulated

Linux Distribution	Supported Versions	Preferred Launch Mode
FreeBSD	12 or later	Paravirtualized
	8, 9, 10, 11	Emulated
Debian	8 or later	Paravirtualized
	5.0.10, 6.0, 7	Emulated
SUSE	12.2 or later	Paravirtualized
	11, 12.1	Emulated

**Note**

- Oracle Cloud Infrastructure has tested the operating systems listed in the previous table and will support customers in ensuring that instances launched from these images and built according to the guidelines in this topic are accessible using SSH.
- For OS versions not listed in the previous table, Oracle Cloud Infrastructure will provide commercially reasonable support to customers in an effort to get instances that are launched from these images accessible via SSH.
- Support from Oracle Cloud Infrastructure in launching an instance from a custom OS does not ensure that the operating system vendor also supports the instance. Customers running [Oracle Linux](#) on Oracle Cloud Infrastructure automatically have access to Oracle Linux Premier Support.



Tip

If your image supports paravirtualized drivers, you can convert your existing emulated mode instances into paravirtualized instances. Create a custom image of your instance, export it to Object Storage, and then reimport it using paravirtualized mode.

Linux Source Image Requirements

Custom images must meet the following requirements:

- The maximum image size is 400 GB.
- The image must be set up for BIOS boot.
- Only one disk is supported, and it must be the boot drive with a valid master boot record (MBR) and boot loader. You can migrate additional data volumes after you import the image's boot volume.
- The boot process must not require additional data volumes to be present for a successful boot.
- The boot loader should use LVM or a UUID to locate the boot volume.
- The disk image cannot be encrypted.
- The disk image must be a VMDK or QCOW2 file.
 - Create the image file by cloning the source volume, not by creating a snapshot.
 - VMDK files must be either the "single growable" (monolithicSparse) type or the "stream optimized" (streamOptimized) type, both of which consist of a single VMDK file. All other VMDK formats, such as those that use multiple files, split volumes, or contain snapshots, are not supported.
- The network interface must use DHCP to discover the network settings. When you import a custom image, existing network interfaces are not recreated. Any existing network interfaces are replaced with a single NIC after the import process is complete.

You can attach additional VNICs after you launch the imported instance.

- The network configuration must not hardcode the MAC address for the network interface.

We recommend that you enable certificate-based SSH, however this is optional. If you want your image to automatically use SSH keys supplied from the **User Data** field when you launch an instance, you can install [cloud-init](#) when preparing the image. See [Creating an Instance](#) for more information about the **User Data** field.

Preparing Linux VMs for Import

Before you import a custom Linux image, you must prepare the image to ensure that instances launched from the image can boot correctly and that network connections will work. Do the following:

1. Optionally, [configure your Linux image to support serial console connections](#). A console connection can help you remotely troubleshoot malfunctioning instances, such as an imported image that does not complete a successful boot.
2. Create a backup of the root volume.
3. If the VM has remotely attached storage, such as NFS or block volumes, configure any services that rely on this storage to start manually. Remotely attached storage is not available the first time that an imported instance boots on Oracle Cloud Infrastructure.
4. Ensure that all network interfaces use DHCP, and that the MAC address and IP addresses are not hardcoded. See your system documentation for steps to perform network configuration for your system.
5. Stop the VM.
6. Clone the stopped VM as a VMDK or QCOW2 file, and then export the image from your virtualization environment. See the tools documentation for your virtualization environment for steps.

Importing a Linux-Based VM

After you prepare a Linux image for import, follow these steps to import the image:

1. Upload the VMDK or QCOW2 file to Oracle Cloud Infrastructure:
 - a. [Upload the file to an Object Storage bucket](#). You can upload the file using the Console or using the [command line interface \(CLI\)](#). If you use the CLI, use the following command:

```
oci os object put -bn <destination_bucket_name> --file <path_to_the_VMDK_or_QCOW2_file>
```

- b. Copy the URL of the file that you uploaded: On the **Bucket Details** page, click the Actions icon (three dots) next to the file, and then click **Details**. Copy the **URL Path (URI)**.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
 3. Click **Import Image**.
 4. In the **Create in Compartment** list, select the compartment that you want to import the image to.
 5. Enter a name for the image.
 6. For the **Operating System**, select **Linux**.
 7. In the **Object Storage URL** field, paste the URL of the file that you uploaded.
 8. For the **Image Type**, select the file type of the image, either **VMDK** or **QCOW2**.
 9. Depending on your image's version of Linux, in the **Launch Mode** area, select **Paravirtualized Mode** or **Emulated Mode**. If your [image supports paravirtualized drivers](#), we recommend that you select paravirtualized mode.
 10. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

11. Click **Import Image**.

The imported image appears in the **Custom Images** list for the compartment, with a status of **IMPORTING**. When the import completes successfully, the status changes to **AVAILABLE**.

If the status doesn't change, or no entry appears in the **Custom Images** list, the import failed. Ensure that you have read access to the Object Storage object, and that the object contains a supported image.

12. Complete the post-import tasks.

Post-Import Tasks for Linux Images

After you import a custom Linux-based image, do the following:

1. [Create an instance based on the custom image](#). For the image source, select **Custom Images**, and then select the image that you imported.
2. [Connect to the instance using SSH](#).
3. If the instance requires any remotely attached storage, such as [block volumes](#) or [file storage](#), create and attach it.
4. [Create and attach any required secondary VNICs](#).
5. Test that all applications are working as expected.
6. Reset any services that were set to start manually.
7. If you enabled serial console access to the image, test it by [creating a serial console connection to the instance](#).

See the [current issues and workarounds](#) for known issues with imported custom images.

Enabling Serial Console Access for Imported Linux Images

You can configure your custom Linux image to support connections using the serial console feature in the Compute service.

For more information about serial console connections, and steps to troubleshoot if your image has network connectivity issues after it is launched, see [Instance Console Connections](#).

CHAPTER 8 Compute

The serial console connection in Oracle Cloud Infrastructure uses the first serial port, `ttyS0`, on the VM. The boot loader and the operating system should be configured to use `ttyS0` as a console terminal for both input and output.

CONFIGURING THE BOOT LOADER

The steps to configure the boot loader to use `ttyS0` as a console terminal for both input and output depend on the GRUB version. Run the following command on the operating system to determine the GRUB version:

```
grub install --version
```

If the version number returned is 2.x, use the steps for GRUB 2. For earlier versions, use the steps for GRUB.

To configure GRUB2

1. Run the following command to modify the GRUB configuration file:

```
sudo vi /etc/default/grub
```

2. Confirm that the configuration file contains the following:

```
GRUB_SERIAL_COMMAND="serial --unit=0 --speed=115200"
GRUB_TERMINAL="serial console"
```

3. Append the following to the end of the `GRUB_CMDLINE_LINUX` line:

```
console=tty1 console=ttyS0,115200
```

If `GRUB_CMDLINE_LINUX` does not exist, create this line, using `GRUB_CMDLINE_OUTPUT` as a template.

4. Regenerate the GRUB2 configuration using the following command:

```
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

If you have a beta version of GRUB 2, use this command instead:

```
sudo grub-mkconfig -o /boot/grub/grub.cfg
```

To configure GRUB

1. Run the following command to modify the GRUB configuration file:

```
sudo vi /boot/grub/grub.conf
```

2. Add following after the line containing `timeout`:

```
serial --unit=0 --speed=115200
terminal --timeout=5 serial console
```

3. Append the following to each `kernel` line:

```
console=tty1 console=ttyS0,115200
```

CONFIGURING THE OPERATING SYSTEM

The operating system may already be configured to use `ttyS0` as a console terminal for both input and output. To verify, run the following command:

```
sudo vi /etc/securetty
```

Check the file for `ttyS0`. If you don't see it, append `ttyS0` to the end of the file.

VALIDATING SERIAL CONSOLE ACCESS

After completing the steps to enable serial console access to the image, you should validate that serial console access is working by testing the image with serial console in your virtualization environment. Consult the documentation for your virtualization environment for steps to do this. Verify that the boot output displays in the serial console output and that there is interactive input after the image has booted.

TROUBLESHOOTING THE SERIAL CONSOLE

If no output is displayed on the serial console, verify in the configuration for your virtualization environment that the serial console device is attached to the first serial port.

If the serial console displays output, but there is no interactive input available, check that there is a terminal process listening on the `ttyS0` port. To do this, run the following command:

```
ps aux | grep ttyS0
```

CHAPTER 8 Compute

This command should output a terminal process that is listening on the ttyS0 port. For example, if your system is using `getty`, you will see the following output:

```
/sbin/getty ttyS0
```

If you don't see this output, it is likely that a login process is not configured for the serial console connection. To resolve this, enable the init settings, so that a terminal process is listening on the ttyS0 at startup.

For example, if your system is using `getty`, add the following command to the init settings to run on system startup:

```
getty -L 9600 ttyS0 vt102
```

The steps to do this will vary depending on the operating system, so consult the documentation for the image's operating system.

OS Kernel Updates



Note

This topic applies only to Linux instances that were launched before February 15, 2017. Linux instances launched on or after February 15, 2017 boot directly from the image and do not require further action for kernel updates.

Oracle Cloud Infrastructure boots each instance from a network drive. This configuration requires additional actions when you update the OS kernel.

Oracle Cloud Infrastructure uses Unified Extensible Firmware Interface (UEFI) firmware and a Preboot eXecution Environment (PXE) interface on the host server to load iPXE from a Trivial File Transfer Protocol (TFTP) server. The iPXE implementation runs a script to boot Oracle Linux. During the boot process, the system downloads the kernel, the `initrd` file, and the kernel boot parameters from the network. The instance does not use the host's GRUB boot loader.

Normally, the `yum update kernel-uek` command edits the GRUB configuration file, either `grub.cfg` or `grub.conf`, to configure the next boot. Since bare metal instances do not use the GRUB boot loader, changes to the GRUB configuration file are not implemented. When you update the kernel on your instance, you also must upload the update to the network to ensure a successful boot process. The following approaches address this need:

- Instances launched from an Oracle-provided image include an Oracle yum plug-in that seamlessly handles the upload when you run the `yum update kernel-uek` command.
- If you use a custom image based on an Oracle-provided image, the included yum plug-in will continue to work, barring extraordinary changes.
- If you install your own package manager, you must either write your own plug-in or upload the kernel, `initrd`, and kernel boot parameters manually.

Oracle Yum Plug-in

On instances launched with an [Oracle-provided image](#), you can find the Oracle yum plug-in at:

```
/usr/share/yum-plugins/kernel-update-handler.py
```

The plug-in configuration is at:

```
/etc/yum/pluginconf.d
```

The plug-in looks for two variables in the `/etc/sysconfig/kernel` file, `UPDATEDEFAULT` and `DEFAULTKERNEL`. It picks up the updates only when the first variable is set to "yes" and the `DEFAULTKERNEL` value matches the kernel being updated. For example:

```
# UPDATEDEFAULT specifies if new-kernel-pkg should make
# new kernels the default
UPDATEDEFAULT=yes

# DEFAULTKERNEL specifies the default kernel package type
DEFAULTKERNEL=kernel-uek
```

Oracle-provided images incorporate the Unbreakable Enterprise Kernel (UEK). If you want to switch to a non-UEK kernel, you must update the `DEFAULTKERNEL` value to "kernel" before you run `yum update kernel`.

Manual Updates



Tip

Oracle recommends using the Oracle yum plug-in to update the kernel.

If you manually upload the updates, there are four relevant URLs:

```
http://169.254.0.3/kernel
```

```
http://169.254.0.3/initrd
```

```
http://169.254.0.3/cmdline
```

```
http://169.254.0.3/activate
```

The first three URLs are for uploading files (HTTP request type PUT). The fourth URL is for activating the uploaded files (HTTP request type POST). The system discards the uploaded files if they are not activated before the host restarts.

The kernel and initrd are simple file uploads. The cmdline upload must contain the kernel boot parameters found in the `grub.cfg` or `grub.conf` file, depending on the Linux version. The following example is an entry from the `/boot/efi/EFI/redhat/grub.cfg` file in Red Hat Linux 7. The highlighted text represents the parameters to upload.

```
kernel /boot/vmlinuz-4.1.12-37.5.1.el6uek.x86_64 ro root=UUID=8079e287-53d7-4b3d-b708-c519cf6829c8 rd_NO_LUKS KEYBOARDTYPE=pc KEYTABLE=us netroot=iscsi:@169.254.0.2::3260:iface1:eth0::iqn.2015-02.oracle.boot:uefi rd_NO_MD SYSFONT=latarcyrheb-sun16 ifname=eth0:90:e2:ba:a2:e3:80 crashkernel=auto iscsi_initiator=iqn.2015-02.rd_NO_LVM ip=eth0:dhcp rd_NO_DM LANG=en_US.UTF-8 console=tty0 console=ttyS0,9600 iommu=on
```

The following command returns what is being uploaded to the cmdline file.

```
cat /tmp/cmdline
```

A typical response resembles the following.

```
ro root=UUID=8079e287-53d7-4b3d-b708-c519cf6829c8 rd_NO_LUKS KEYBOARDTYPE=pc KEYTABLE=us netroot=iscsi:@169.254.0.2::3260:iface1:eth0::iqn.2015-02.oracle.boot:uefi rd_NO_MD SYSFONT=latarcyrheb-
```

CHAPTER 8 Compute

```
sun16 ifname=eth0:90:e2:ba:a2:e3:80 crashkernel=auto iscsi_initiator=iqn.2015-02. rd_NO_LVM ip=eth0:dhcp
rd_NO_DM LANG=en_US.UTF-8 console=tty0 console=ttyS0,9600 iommu=on
```

The following commands update the `cmdline` and `initrd` files, and then activate the changes.

```
CKSUM=`md5sum /tmp/cmdline | cut -d ' ' -f 1`
```

```
sudo curl -X PUT --data-binary @/tmp/cmdline -H "Content-MD5: $CKSUM" http://169.254.0.3/cmdline
```

```
CKSUM=`md5sum /boot/initramfs-3.8.13-118.8.1.el7uek.x86_64.img | cut -d ' ' -f 1`
```

```
sudo curl -X PUT --data-binary @/boot/initramfs-3.8.13-118.8.1.el7uek.x86_64.img -H "Content-MD5:
$CKSUM" http://169.254.0.3/initrd
```

```
sudo curl -X POST http://169.254.0.3/activate
```

Managing Key Pairs on Linux Instances

Instances launched using Oracle Linux, CentOS, or Ubuntu images use an SSH key pair instead of a password to authenticate a remote user (see [Security Credentials](#)). A key pair consists of a private key and public key. You keep the private key on your computer and provide the public key every time you launch an instance.

When you connect to an instance using SSH, you provide the path to the key pair file in the SSH command. You can have as many key pairs as you want, or you can keep it simple and use one key pair for all or several of your instances.

To create key pairs, you can use a third-party tool such as OpenSSH on UNIX-style systems (including Linux, Solaris, BSD, and OS X) or PuTTY Key Generator on Windows.

Prerequisites

If you're using a UNIX-style system, you probably already have the `ssh-keygen` utility installed. To determine if it's installed, type `ssh-keygen` on the command line. If it's not installed, you can download OpenSSH for UNIX from <http://www.openssh.com/portable.html> and install it.

If you're using a Windows operating system you will need PuTTY and the PuTTY Key Generator. Download PuTTY and PuTTYgen from <http://www.putty.org> and install them.

Creating an SSH Key Pair on the Command Line

1. Open a shell or terminal for entering the commands.
2. At the prompt, enter `ssh-keygen` and provide a name and passphrase when prompted. The keys will be created with the default values: RSA keys of 2048 bits.

Alternatively, you can type a complete `ssh-keygen` command, for example:

```
ssh-keygen -t rsa -N "" -b 2048 -C "<key_name>" -f <path/root_name>
```

The command arguments are shown in the following table:

Argument	Description
<code>-t rsa</code>	Use the RSA algorithm.
<code>-N "<passphrase>"</code>	A passphrase to protect the use of the key (like a password). If you don't want to set a passphrase, don't enter anything between the quotes. A passphrase is not required. You can specify one as a security measure to protect the private key from unauthorized use.
<code>-b 2048</code>	Generate a 2048-bit key. You don't have to set this if 2048 is acceptable, as 2048 is the default. A minimum of 2048 bits is recommended for SSH-2 RSA.
<code>-C "<key_name>"</code>	A name to identify the key.
<code>-f <path/root_name></code>	The location where the key pair will be saved and the root name for the files.

Creating an SSH Key Pair Using PuTTY Key Generator

1. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
 - In the **Key** menu, confirm that the default value of **SSH-2 RSA key** is selected.
 - For the **Type of key to generate**, accept the default key type of **RSA**.
 - Set the **Number of bits in a generated key** to 2048 if it is not already set.
3. Click **Generate**.
4. Move your mouse around the blank area in the PuTTY window to generate random data in the key.

When the key is generated, it appears under **Public key for pasting into OpenSSH authorized_keys file**.
5. A **Key comment** is generated for you, including the date and time stamp. You can keep the default comment or replace it with your own more descriptive comment.
6. Leave the **Key passphrase** field blank.
7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.

The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

You can name the key anything you want, but use the `ppk` file extension. For example, `mykey.ppk`.
8. Select *all* of the generated key that appears under **Public key for pasting into OpenSSH authorized_keys file**, copy it using **Ctrl + C**, paste it into a text file, and then save the file in the same location as the private key.

(Do not use **Save public key** because it does not save the key in the OpenSSH format.)

You can name the key anything you want, but for consistency, use the same name as the private key and a file extension of `pub`. For example, `mykey.pub`.

9. Write down the names and location of your public and private key files. You will need the public key when launching an instance. You will need the private key to access the instance via SSH.

Now that you have a key pair, you're ready to launch instances as described in [Creating an Instance](#).

Creating an Instance

You can create an instance using the Console or API. When you create an instance, it is automatically attached to a virtual network interface card (VNIC) in the cloud network's subnet and given a private IP address from the subnet's CIDR. You can either let the IP address be automatically assigned, or specify a particular address of your choice. The private IP address lets instances within the cloud network communicate with each other. They can instead use fully qualified domain names (FQDNs) if you've set up the cloud network for DNS (see [DNS in Your Virtual Cloud Network](#)).

If the subnet is public, you can optionally assign the instance a public IP address. A public IP address is required to communicate with the instance over the Internet, and to establish a Secure Shell (SSH) or RDP connection to the instance from outside the cloud network. For more information, see [Access to the Internet](#).



Tip

If this is your first time creating an instance, consider following the [Getting Started Tutorial](#) for a guided workflow through the steps required to create an instance.



Note

Partner images and pre-built Oracle enterprise images are not available in Government Cloud realms.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.



Tip

When you create an instance, several other resources are involved, such as an image, a cloud network, and a subnet. Those other resources can be in the same compartment with the instance or in other compartments. You must have the required level of access to each of the compartments involved in order to launch the instance. This is also true when you attach a volume to an instance; they don't have to be in the same compartment, but if they're not, you need the required level of access to each of the compartments.

For administrators: The simplest policy to enable users to create instances is listed in [Let users launch Compute instances](#). It gives the specified group general access to managing instances and images, along with the required level of access to attach existing block volumes to the instances. If the group needs to *create* block volumes, they'll need the ability to *manage* block volumes (see [Let volume admins manage block volumes, backups, and volume groups](#)).

Partner Image Catalog

If the group needs to create instances based on partner images, they'll need the *manage* permission for app-catalog-listing to create subscriptions to images from the Partner Image catalog. See [Let users list and subscribe to images from the Partner Image catalog](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Recommended Networking Launch Types

When you launch a virtual machine (VM) instance, by default, Oracle Cloud Infrastructure chooses a recommended networking type for the VNIC based on the instance shape and OS image. The networking interface handles functions such as disk input/output and network communication. The following options are available:

- **Paravirtualized networking:** For general purpose workloads such as enterprise applications, microservices, and small databases. Paravirtualized networking also provides increased flexibility to use the same image across different hardware platforms.
- **Hardware-assisted (SR-IOV) networking:** Single root input/output virtualization. For low-latency workloads such as video streaming, real-time applications, and large or clustered databases. Hardware-assisted (SR-IOV) networking uses the VFIO driver framework.

The following table lists the default and supported networking types for [VM shapes](#).

Shape Type	Default Networking Type	Supported Networking Types
VM.Standard1	SR-IOV	Paravirtualized, SR-IOV
VM.Standard2	Paravirtualized	Paravirtualized, SR-IOV
VM.Standard.E2	Paravirtualized	Paravirtualized only
VM.DenseIO1	SR-IOV	Paravirtualized, SR-IOV
VM.DenseIO2	Paravirtualized	Paravirtualized, SR-IOV
VM.GPU2	SR-IOV	Paravirtualized, SR-IOV
VM.GPU3	SR-IOV	Paravirtualized, SR-IOV

To use paravirtualized networking, you must also use an image that supports paravirtualized networking. Paravirtualized networking is supported on these [Oracle-provided images](#):

- **Oracle Linux 7, Oracle Linux 6:** Images published in March 2019 or later.
- **CentOS 7, CentOS 6:** Images published in July 2019 or later.
- **Ubuntu 18.04, Ubuntu 16.04:** Images published in March 2019 or later.
- **Windows Server 2016:** Images published in August 2019 or later.

SR-IOV networking is supported on all Oracle-provided images.

You can create an instance that uses a specific networking type instead of the default. However, depending on compatibility between the shape and image that you choose, the instance might not launch properly. You can test whether it succeeded by [connecting to the instance](#). If the connection fails, the networking type is not supported. Relaunch the instance using a supported networking type.

Using the Console

To create a Linux instance

Prerequisites

To create a Linux instance, you'll need:

- A virtual cloud network (VCN) to launch the instance in. For information about setting up cloud networks, see [Overview of Networking](#).
- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the instance via SSH. The following sample public key is abbreviated for readability:

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA...lo/gKMLVM2xzclxJr/Hc26biw3TXWGEakrK1OQ== rsa-key-20160304
```

For information about generating a key pair, see [Managing Key Pairs on Linux Instances](#).

To create a Linux instance using the Console:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Then, choose a **Compartment** you have permission to work in.
2. Click **Create Instance**.
3. On the **Create Compute Instance** page, you specify the resources to use for your instance. By default, your instance launches in the current compartment, and the resources you choose also come from the current compartment.

On the **Create Compute Instance** page, you specify the resources to use for the instance. By default, the instance launches in the current compartment, and the resources that you choose also come from the current compartment.

Specify the following:

- **Name your instance:** The name for the instance. You can add or change the name later. The name doesn't need to be unique, because an Oracle Cloud Identifier (OCID) uniquely identifies the instance.
- **Choose an operating system or image source:** The source of the image to use for booting the instance. When you click **Change Image Source**, the **Browse All Images** page opens with the operating system or image source options. The following options are available:
 - **Platform Images:** Pre-built images for Oracle Cloud Infrastructure. See [Oracle-Provided Images](#) for a list of these images.
 - **Oracle Images:** Pre-built Oracle enterprise images and solutions enabled for Oracle Cloud Infrastructure.
 - **Partner Images:** Trusted third-party images published by Oracle partners. Click the down arrow in the row for an image to view and change the image build, or to view additional details about the image. For more information, see [Overview of Marketplace](#) and [Working with Listings](#).
 - **Custom Images:** Custom images created or imported into your Oracle Cloud Infrastructure environment. See [Managing Custom Images](#) for more information.

- **Boot Volumes:** Boot volumes available for creating a new instance in your Oracle Cloud Infrastructure environment. See [Boot Volumes](#) for more information.
- **Image OCID:** Create an instance using a specific version of an image by providing the image OCID. See [Oracle-Provided Image Release Notes](#) to determine the image OCID for Oracle-provided images.
- **Availability Domain:** The availability domain in which you want to run the instance.
- **Instance Type:** Select **Virtual Machine** or **Bare Metal Machine**.
- **Instance Shape:** A template that determines the number of CPUs, amount of memory, and other resources allocated to a newly created instance.
When you click **Change Shape**, the **Browse All Shapes** dialog box opens with a list of the virtual machine (VM) or bare metal shapes that are available for the instance type that you selected. Incompatible shapes don't appear in the list. See [Compute Shapes](#) for more information about bare metal and VM shapes. Choose a shape and then click **Select Shape**.
- **Configure networking:** The network details for the instance. In this section, you configure the following:
 - **Virtual cloud network compartment:** The compartment containing the network in which to create the instance.
 - **Virtual cloud network:** The network in which to create the instance.
 - **Subnet compartment:** The compartment containing a subnet within the cloud network to attach the instance to.
 - **Subnet:** A subnet within the cloud network to attach the instance to. The subnets are either public or private. Private means the instances in that subnet can't have public IP addresses. For more information, see [Access to the Internet](#). Subnets can also be either AD-specific or regional (regional ones have "*regional*" after the name). We recommend using regional subnets. For more information, see [About Regional Subnets](#).

- **Use network security groups to control traffic:** Select this check box to add the instance's primary VNIC to at least one network security group (NSG) of your choice. NSGs have security rules that apply only to the VNICs in that NSG. For more information, see [Network Security Groups](#).
- If the subnet is public, you can optionally assign the instance a public IP address. A public IP address makes the instance accessible from the internet. Select the **Assign a public IP address** option. For more information, see [Access to the Internet](#).
- **Boot volume:** Size and encryption options for the instance's boot volume.
 - To specify a custom size for the boot volume, select the **Custom boot volume size (in GB)** check box. Then, enter a custom size from 50 GB to 32 TB. The specified size must be larger than the default boot volume size for the selected image. See [Custom Boot Volume Sizes](#) for more information.
 - For VM instances, you can optionally select the **Use in-transit encryption** check box. See [Block Volume Encryption](#) for more information. If you are using your own Key Management encryption key for the boot volume, then this key is also used for in-transit encryption. Otherwise, the Oracle-provided encryption key is used.
 - Boot volumes are encrypted by default, but you can optionally encrypt the data in this volume using your own Key Management encryption key. To use Key Management for your encryption needs, select the **Choose a key from Key Management to encrypt this volume** check box. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#). If you enable this option, this key is used for both data at rest encryption and in-transit encryption.
 - The Block Volume elastic performance feature enables you to change the volume performance for boot volumes, but this can only be modified after

the instance has been launched. When the instance is created, the boot volume is configured with the default volume performance set to **Balanced**. See [Changing the Performance of a Volume](#) for how to modify this setting. See [Block Volume Elastic Performance](#) for more information about this feature.

- **Add SSH key:** The public key portion of the key pair that you want to use for SSH access to the instance. You can drag and drop single key files into the box. To provide multiple keys, click **Choose files**, then press and hold down the Command key (on Mac) or the CTRL key (on Windows) while selecting files.
- **Show Advanced Options:** Advanced networking and management options. On the **Management** tab, configure the following:
 - **Choose a compartment for your instance:** The compartment in which you want to launch the instance.
 - **Choose a fault domain:** The fault domain to use for the instance. If you do not specify the fault domain, the system selects one for you. The fault domain cannot be changed after you create the instance. If you want to use a different fault domain, you must terminate the instance and launch a new instance in the preferred fault domain. For more information, see [Fault Domains](#) and [Best Practices for Your Compute Instance](#).
 - **User data:** Data to be used by cloud-init to run custom scripts or provide custom cloud-init configuration. Click **Choose File** to select the script file, or paste the script into the text box. The file or script does not need to be base64-encoded, because the Console performs this encoding when the information is submitted. For information about how to take advantage of user data, see the [cloud-init documentation](#).
 - **Enable monitoring:** Select this check box to collect metrics for this instance. When enabled, the OracleCloudAgent software on the instance emits metrics for this instance to the Monitoring service using the oci_computeagent metric namespace.

This option is available for supported images only. Legacy versions of supported images may also require installation of the OracleCloudAgent software. For more information, see [Enabling Monitoring for Compute Instances](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

On the **Networking** tab, configure the following:

- **Private IP Address:** Optional. An available private IP address of your choice from the subnet's CIDR. If you don't specify a value, the private IP address is automatically assigned.
- **Hostname:** Optional. A hostname to be used for DNS within the cloud network. Available only if the VCN and subnet both have DNS labels. For more information, see [DNS in Your Virtual Cloud Network](#).
- **Launch Options:** Optional. The networking launch type. Available only for VMs. For more information, see [Recommended Networking Launch Types](#).

On the **Image** tab, you can optionally change the image build. By default, the latest build of the image is used to create the instance. You can select an older build of the image that is compatible with the shape you selected. Only compatible image builds are displayed in the list. You must select a shape before you can change the image build.

On the **Host** tab, you can optionally choose to launch the instance on a dedicated virtual machine host. This enables you to run the instance in isolation, so that it is not running on shared infrastructure. To do this, select the **Launch the virtual machine on a dedicated host** check box, and then select a dedicated virtual machine host from the drop-down list. Before you can place an instance on a dedicated virtual machine host, you must create a dedicated virtual machine host

in the same availability domain and fault domain as the instance. You can only place an instance on a dedicated virtual machine host at the time that you create the instance. For more information, see [Dedicated Virtual Machine Hosts](#).

4. Click **Create**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

After the instance is provisioned, details about it appear in the instance list. To view additional details, including IP addresses, click the instance name.

When the instance is fully provisioned and running, you can connect to it using SSH as described in [Connecting to an Instance](#).

You also can attach a volume to the instance, provided the volume is in the same availability domain. For background information about volumes, see [Overview of Block Volume](#).

To create a Windows instance

Prerequisites

To create a Windows instance, you'll need:

- A virtual cloud network (VCN) to launch the instance in. For information about setting up VCNs, see [Overview of Networking](#).
- A VCN security rule that enables Remote Desktop Protocol (RDP) access so you can connect to your instance. Specifically, you need a stateful ingress rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port. For more information, see [Security Rules](#). You can implement this security rule in either a network security group that you will add this Windows instance to, or a security list that is used by the instance's subnet.

To enable RDP access

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the cloud network you're interested in.
3. To add the rule to a network security group that the instance belongs to:
 - a. Under **Resources**, click **Network Security Groups**. Then click the network security group that you're interested in.
 - b. Click **Add Rules**.
 - c. Enter the following values for the rule:
 - **Stateless:** Leave the check box unselected
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** RDP (TCP/3389)
 - **Source Port Range:** All
 - **Destination Port Range:** 3389
 - d. When done, click **Add**.
4. Or, to add the rule to a security list that is used by the instance's subnet:
 - a. Under **Resources**, click **Security Lists**. Then click the security list you're interested in.
 - b. Click **Add Ingress Rules**.

- c. Enter the following values for the rule:
 - **Stateless:** Leave the check box unselected
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** RDP (TCP/3389)
 - **Source Port Range:** All
 - **Destination Port Range:** 3389
- d. When done, click **Add Ingress Rules**.

To create a Windows instance using the Console:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Then, choose a **Compartment** you have permission to work in.
2. Click **Create Instance**.
3. On the **Create Compute Instance** page, you specify the resources to use for the instance. By default, the instance launches in the current compartment, and the resources that you choose also come from the current compartment.

Specify the following:

- **Name your instance:** The name for the instance. You can add or change the name later. The name doesn't need to be unique, because an Oracle Cloud Identifier (OCID) uniquely identifies the instance.



Important

Use only these ASCII characters in the instance name: uppercase letters (A-Z), lowercase letters (a-z), numbers (0-9), and hyphens (-). See this [known issue](#) for more information.

- **Choose an operating system or image source:** The source of the image to use for booting the instance. When you click **Change Image Source**, the **Browse All Images** page opens with the operating system or image source options. The following options are available:
 - **Platform Images:** Pre-built images for Oracle Cloud Infrastructure. See [Oracle-Provided Images](#) for a list of these images.
 - **Oracle Images:** Pre-built Oracle enterprise images and solutions enabled for Oracle Cloud Infrastructure.
 - **Partner Images:** Trusted third-party images published by Oracle partners. Click the down arrow in the row for an image to view and change the image build, or to view additional details about the image. For more information, see [Overview of Marketplace](#) and [Working with Listings](#).
 - **Custom Images:** Custom images created or imported into your Oracle Cloud Infrastructure environment. See [Managing Custom Images](#) for more information.
 - **Boot Volumes:** Boot volumes available for creating a new instance in your Oracle Cloud Infrastructure environment. See [Boot Volumes](#) for more information.
 - **Image OCID:** Create an instance using a specific version of an image by providing the image OCID. See [Oracle-Provided Image Release Notes](#) to determine the image OCID for Oracle-provided images.
- **Availability Domain:** The availability domain in which you want to run the instance.
- **Instance Type:** Select **Virtual Machine** or **Bare Metal Machine**.
- **Instance Shape:** A template that determines the number of CPUs, amount of memory, and other resources allocated to a newly created instance. When you click **Change Shape**, the **Browse All Shapes** dialog box opens with a list of virtual machine (VM) or bare metal shapes that are available for the instance type that you selected. Incompatible shapes don't appear in the list. See

[Compute Shapes](#) for more information about bare metal and VM shapes. Choose a shape and then click **Select Shape**.

- **Configure networking:** The network details for the instance. In this section, you configure the following:
 - **Virtual cloud network compartment:** The compartment containing the network in which to create the instance.
 - **Virtual cloud network:** The network in which to create the instance.
 - **Subnet compartment:** The compartment containing a subnet within the cloud network to attach the instance to.
 - **Subnet:** A subnet within the cloud network to attach the instance to. The subnets are either public or private. Private means the instances in that subnet can't have public IP addresses. For more information, see [Access to the Internet](#). Subnets can also be either AD-specific or regional (regional ones have "*regional*" after the name). We recommend using regional subnets. For more information, see [About Regional Subnets](#).
 - **Use network security groups to control traffic:** Select this check box to add the instance's primary VNIC to at least one network security group (NSG) of your choice. NSGs have security rules that apply only to the VNICs in that NSG. For more information, see [Network Security Groups](#).
 - If the subnet is public, you can optionally assign the instance a public IP address. A public IP address makes the instance accessible from the internet. Select the **Assign a public IP address** option. For more information, see [Access to the Internet](#).
- **Boot volume:** Size and encryption options for the instance's boot volume.
 - To specify a custom size for the boot volume, select the **Custom boot volume size (in GB)** check box. Then, enter a custom size from 50 GB (256 GB for Oracle-provided Windows images) to 32 TB. The specified size must be larger than the selected image's default boot volume size. See [Custom Boot Volume Sizes](#) for more information.

- For VM instances, you can optionally select the **Use in-transit encryption** check box. See [Block Volume Encryption](#) for more information. If you are using your own Key Management encryption key for the boot volume, then this key is also used for in-transit encryption. Otherwise, the Oracle-provided encryption key is used.
- Boot volumes are encrypted by default, but you can optionally encrypt the data in this volume using your own Key Management encryption key. To use Key Management for your encryption needs, select the **Choose a key from Key Management to encrypt this volume** check box. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#).
- The Block Volume elastic performance feature enables you to change the volume performance for boot volumes, but this can only be modified after the instance has been launched. When the instance is created, the boot volume is configured with the default volume performance set to **Balanced**. See [Changing the Performance of a Volume](#) for how to modify this setting. See [Block Volume Elastic Performance](#) for more information about this feature.
- **Add SSH Key:** An SSH key pair is only required for Linux instances. For Windows instances, you connect to the instance using a password. An initial password will be provided when you finish launching the instance.
- **Show Advanced Options:** Advanced networking and management options.
On the **Management** tab, configure the following:
 - **Choose a compartment for your instance:** The compartment in which you want to launch the instance.
 - **Choose a fault domain:** The fault domain to use for the instance. If you do not specify the fault domain, the system selects one for you. The fault domain cannot be changed after you create the instance. If you want to use

a different fault domain, you must terminate the instance and launch a new instance in the preferred fault domain. For more information, see [Fault Domains](#) and [Best Practices for Your Compute Instance](#).

- **User data:** Data to be used by cloudbase-init to run custom scripts or provide custom cloudbase-init configuration. Click **Choose File** to select the script file, or paste the script into the text box. The file or script does not need to be base64-encoded, because the Console performs this encoding when the information is submitted. For information about how to take advantage of user data, see the [cloudbase-init documentation](#).



Warning

Do not include anything in the script that could trigger a reboot, because this could impact the instance launch and cause it to fail. Any actions requiring a reboot should only be performed once the instance state is **RUNNING**.

- **Enable monitoring:** Select this check box to collect metrics for this instance. When enabled, the OracleCloudAgent software on the instance emits metrics for this instance to the Monitoring service using the oci_computeagent metric namespace.
This option is available for supported images only. Legacy versions of supported images may also require installation of the OracleCloudAgent software. For more information, see [Enabling Monitoring for Compute Instances](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag

namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

On the **Networking** tab, configure the following:

- **Private IP address:** Optional. An available private IP address of your choice from the subnet's CIDR. If you don't specify a value, the private IP address is automatically assigned.
- **Hostname:** Optional. A hostname to be used for DNS within the cloud network. Available only if the VCN and subnet both have DNS labels. For more information, see [DNS in Your Virtual Cloud Network](#).
- **Launch Options:** Optional. The networking launch type. Available only for VMs. For more information, see [Recommended Networking Launch Types](#).

On the **Image** tab, you can optionally change the image build. By default, the latest build of the image is used to create the instance. You can select an older build of the image that is compatible with the shape you selected. Only compatible image builds are displayed in the list. You must select a shape before you can change the image build.

On the **Host** tab, you can optionally choose to launch the instance on a dedicated virtual machine host. This enables you to run the instance in isolation, so that it is not running on shared infrastructure. To do this, select the **Launch the virtual machine on a dedicated host** check box, and then select a dedicated virtual machine host from the list. Before you can place an instance on a dedicated virtual machine host, you must create a dedicated virtual machine host in the same availability domain and fault domain as the instance. You can only place an instance on a dedicated virtual machine host at the time you create the instance. For more information, see [Dedicated Virtual Machine Hosts](#).

4. Click **Create**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

After the instance is provisioned, details about it appear in the instance list. To view additional details, including IP addresses and the initial Windows password, click the instance name.

When the instance is fully provisioned and running, you can connect to it using Remote Desktop as described in [Connecting to an Instance](#).

You also can attach a volume to the instance, provided the volume is in the same availability domain. For background information about volumes, see [Overview of Block Volume](#).

Managing Tags for an Instance

You can apply tags to your resources, such as instances, to help you organize them according to your business needs. You can apply tags when you create an instance, or you can update the instance later with the tags that you want.

To manage tags for an instance

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance that you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage instances:

- [ListInstances](#)
- [LaunchInstance](#)
- [GetInstance](#)
- [UpdateInstance](#)
- [TerminateInstance](#)
- [GetWindowsInstanceInitialCredentials](#)

Oracle Cloud Infrastructure enables you to launch instances from images published by Oracle partners from the Partner Image catalog. Use these APIs to work with the Partner Image catalog listings:

- [AppCatalogListing](#)
- [AppCatalogListingResourceVersion](#)
- [AppCatalogListingResourceVersionAgreements](#)
- [AppCatalogListingResourceVersionSummary](#)
- [AppCatalogListingSummary](#)
- [AppCatalogSubscription](#)
- [AppCatalogSubscription](#)

Managing Compute Instances

You can simplify the management of your Compute instances using resources such as instance configurations and instance pools.

An instance configuration is a template that defines the settings to use when creating Compute instances.

An instance pool is a group of instances within the same region that are created based off of an instance configuration.

Instance Configurations

An instance configuration defines the settings to use when creating Compute instances, including details such as the base image, shape, and metadata. You can also specify the associated resources for the instance, such as block volume attachments and network configuration.

For steps to create an instance configuration, see [Creating an Instance Configuration](#).

To modify an existing instance configuration, create a new instance configuration with the desired settings.

For steps to delete an instance configuration, see [Deleting an Instance Configuration](#).

Use these API operations to work with instance configurations:

- [CreateInstanceConfiguration](#)
- [DeleteInstanceConfiguration](#)
- [GetInstanceConfiguration](#)
- [UpdateInstanceConfiguration](#)
- [ListInstanceConfigurations](#)

Instance Pools

Instance pools give you the ability to provision and create multiple Compute instances based off of the same instance configuration, within the same region. They also enable integration with other services, such as the Load Balancing service and IAM service, making it easier to manage groups of instances.

You create an instance pool using an existing instance configuration. For steps, see [Creating an Instance Pool](#).

You can automatically adjust the number of instances in an instance pool based on performance metrics such as CPU utilization. To do this, you enable autoscaling for the instance pool. For background information and steps, see [Autoscaling](#).

After you have created an instance pool, you can update the size and attach or detach load balancers from the Console. To update additional settings, you need to use the CLI, API, or SDKs.

If you need to update the instance configuration, create a new instance configuration and then update the instance pool to use the new instance configuration. For more information, see [Updating an Instance Pool](#).

A cluster network is a special kind of instance pool that is designed for massive, high-performance computing jobs. For more information, see [Managing Cluster Networks](#).

For steps to delete an instance pool, see [Deleting an Instance Pool](#).



Warning

When you delete an instance pool all of its resources will be permanently deleted, including associated instances, attached boot volumes, and block volumes.

Instance Pool Lifecycle States

The following list describes the different lifecycle states for instance pools.

- **Provisioning:** When you create an instance pool, this is the first state the instance pool is in. Instances for the instance pool are being configured based on the specified instance configuration.
- **Starting:** The instances are being launched. At this point, the only action you can take is to terminate the instance pool.
- **Running:** The instances are created and running.
- **Stopping:** The instances are in the process of being shut down.
- **Stopped:** The instances are shut down.
- **Scaling:** Once an instance pool has been created, if you update the instance pool size, it will go into this state while creating (for increases in size) or terminating (for decreases

in size) instances. At this point, the only action you can take is to terminate the instance pool.

- **Terminating:** The instances and associated resources are being terminated.
- **Terminated:** The instance pool, all its instances and associated resources are terminated.

When working with instance configurations and instance pools, keep the following in mind:

- You can't delete an instance configuration if it is associated with at least one instance pool.
- You can use the same instance configuration for multiple instance pools. However, an instance pool can only have one instance configuration associated with it.
- If the instance pool has been in the scaling or provisioning state for an extended period of time, it may be because the number of instances requested has exceeded your tenancy's service limits for that shape and availability domain. Check your tenancy's [_service limits](#) for Compute. If you want to request a limit increase, you need to [contact support](#). For more information, see [Service Limits](#). If this occurs, you need to terminate the instance pool and re-create it.
- If you modify the instance configuration for an instance pool, existing instances that are part of that pool will not change. Any new instances created after the instance configuration change will use the new instance configuration. New instances will not be created unless you have increased the size of the instance pool or terminate existing instances.
- If you decrease the size of an instance pool, the oldest instances will be terminated first.

Use these API operations to manage instance pools:

- [CreateInstancePool](#)
- [GetInstancePool](#)
- [ResetInstancePool](#)
- [SoftresetInstancePool](#)

- [StartInstancePool](#)
- [StopInstancePool](#)
- [TerminateInstancePool](#)
- [UpdateInstancePool](#)
- [ListInstancePools](#)
- [ListInstancePoolInstances](#)
- [AttachLoadBalancer](#)
- [DetachLoadBalancer](#)
- [GetInstancePoolLoadBalancerAttachment](#)

Creating an Instance Configuration

Instance configurations allow you to define the settings to use when creating Compute instances.

You use an instance configuration when you want to create one or more instances in an instance pool. For background information about instance pools, see [Managing Compute Instances](#).

You can also use an instance configuration to launch individual instances that are not part of a pool. To do this, use the SDKs, command line interface (CLI), or API.

In the Console, you create an instance configuration using an existing Compute instance as a template. If you want to create an instance configuration by specifying a list of configuration settings, use the SDKs, CLI, or API.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).

Tagging Resources

You can add tags to your resources to help you organize them according to your business needs. You can add tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Using the Console

To create an instance configuration

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance that you want to use as a template to create the instance configuration.
3. Click **Actions**, and then click **Create Instance Configuration**.
4. Select the compartment you want to create the instance configuration in.
5. Specify a name for the instance configuration. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API).
6. **Show Tagging Options:** Optionally, you can add tags. If you have permissions to create a resource, you also have permissions to add free-form tags to that resource. To add a defined tag, you must have permissions to use the tag namespace. For more

information about tagging, see [Resource Tags](#). If you are not sure if you should add tags, skip this option (you can add tags later) or ask your administrator.

7. Click **Create Instance Configuration**.

To manage tags for an instance configuration

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Configurations**.
2. Click the instance configuration that you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Add tags** to add new ones.

For more information, see [Resource Tags](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [CreateInstanceConfiguration](#) operation to create an instance configuration.

Creating an Instance Pool

Instance pools give you the ability to provision and create multiple Compute instances based off the same configuration, within the same region. For more information about instance pools and instance configurations, see [Managing Compute Instances](#).

Optionally, you can associate a load balancer with an instance pool. If you do this, when you add an instance to the instance pool, the instance is automatically added to the load balancer's backend set. After the instance reaches a healthy state (the instance is listening on the configured port number), incoming traffic is automatically routed to the new instance. For background information about the Load Balancing service, see [Overview of Load Balancing](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).



Important

See this [known issue](#) for information about the policy statements that are required if the instance configuration or load balancer associated with the instance pool includes defined tags.

Tagging Resources

You can add tags to your resources to help you organize them according to your business needs. You can add tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Distributing Instances Across Fault Domains for High Availability

By default, the instances in a pool are distributed across all [fault domains](#) in a best-effort manner based on capacity. If capacity isn't available in one fault domain, the instances are placed in other fault domains to allow the instance pool to launch successfully.

In a high availability scenario, you can require that the instances in a pool are evenly distributed across each of the fault domains that you specify. When sufficient capacity isn't available in one of the fault domains, the instance pool will not launch or scale successfully, and a work request for the instance pool will return an "out of capacity" error. To fix the capacity error, either wait for capacity to become available, or use the [UpdateInstancePool](#) operation to update the placement configuration (the availability domain and fault domain) for the instance pool.

Prerequisites

Before you can create an instance pool, you need:

- An instance configuration. An instance configuration is a template that defines the settings to use when creating instances. For more information, see [Creating an Instance Configuration](#).



Note

You cannot create an instance pool from an instance configuration where the image source is a boot volume.

- If you want to associate the instance pool with a load balancer, you need a load balancer and backend set. For steps to create a load balancer, see [Managing a Load Balancer](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click **Create Instance Pool**.
3. In the **Create in Compartment** list, select the compartment that you want to create the instance pool in.
4. Enter a name for the instance pool. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API).
5. Specify the targeted **Number of Instances** for the instance pool.
6. To select the instance configuration that you want to use, enter the following:
 - **Instance Configuration Compartment:** The compartment that the instance configuration is in.
 - **Instance Configuration:** The instance configuration to use.
7. If you want to associate a load balancer with the instance pool, select the **Attach a Load Balancer** check box. Then, enter the following:
 - **Load Balancer Compartment:** The compartment that the load balancer is in.
 - **Load Balancer:** The load balancer to associate with the instance pool.
 - **Backend Set:** The name of the backend set on the load balancer to add instances to.
 - **Port:** The server port on the instances to which the load balancer must direct traffic. This value applies to all instances that use this load balancer attachment.
 - **VNIC:** The VNIC to use when adding the instance to the backend set. Instances that belong to a backend set are also called backend servers. The private IP address is used. This value applies to all instances that use this load balancer attachment.

For background information about load balancers, see [Overview of Load Balancing](#).

8. Select the location where you want to place the instances:
 - a. In the **Availability Domain** list, select the availability domain to launch the instances in.
 - b. If you want the instances in the pool to be placed evenly in one or more fault domains, select the **Distribute instances evenly across selected fault domains** check box. Then, select the fault domains to place the instances in. For more information, see [Distributing Instances Across Fault Domains for High Availability](#).
 - c. In the **Primary VNIC** section, configure the network details for the instances:
 - **Virtual cloud network compartment:** The compartment containing the network to create the instances in.
 - **Virtual cloud network:** The virtual cloud network (VCN) to create the instances in.
 - **Subnet compartment:** The compartment containing a subnet within the cloud network to attach the instances to.
 - **Subnet:** A subnet within the cloud network to attach the instances to. The subnets are either public or private. Private means the instances in that subnet can't have public IP addresses. For more information, see [Access to the Internet](#). Subnets can also be either AD-specific or regional (regional ones have "*regional*" after the name). We recommend using regional subnets. For more information, see [About Regional Subnets](#).
 - d. If secondary VNICs are defined by the instance configuration, a **Secondary VNIC** section appears. Select the compartments, and then select the secondary VCN and subnet for the instance pool.
9. If you want the instance pool to create instances in more than one availability domain, click **+ Additional Selection** and select a different availability domain for the instance pool. Then, repeat the previous step to specify the network information for the second availability domain.
10. **Show Tagging Options:** Optionally, you can add tags. If you have permissions to create a resource, you also have permissions to add free-form tags to that resource. To

add a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should add tags, skip this option (you can add tags later) or ask your administrator.

11. Click **Create Instance Pool**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to create and manage instance pools:

- [CreateInstancePool](#)
- [AttachLoadBalancer](#)
- [DetachLoadBalancer](#)
- [GetInstancePoolLoadBalancerAttachment](#)

Updating an Instance Pool

You can update the number of instances for an instance pool.

Optionally, you can associate a load balancer with an instance pool. If you do this, when you add an instance to the instance pool, the instance is automatically added to the load balancer's backend set. After the instance reaches a healthy state (the instance is listening on the configured port number), incoming traffic is automatically routed to the new instance. For background information about the Load Balancing service, see [Overview of Load Balancing](#).

To update other settings for an instance pool, use the command line interface (CLI), SDKs, or REST APIs.

You can automatically adjust the number of instances in an instance pool based on performance metrics such as CPU utilization. To do this, you enable autoscaling for the instance pool. For background information and steps, see [Autoscaling](#).

CHAPTER 8 Compute

See [Managing Compute Instances](#) for more information about instance pools and instance configurations.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).



Important

See this [known issue](#) for information about the policy statements that are required if the instance configuration or load balancer associated with the instance pool includes defined tags.

Tagging Resources

You can add tags to your resources to help you organize them according to your business needs. You can add tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Using the Console

To update the instance pool size

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click

Instance Pools.

2. Click the name of the instance pool that you're interested in.
3. Click **Edit**.
4. Specify the updated number of instances for the instance pool, and then click **Save Changes**.

When you update the instance pool size, it triggers a scaling event. Keep the following in mind:

- If the instance pool lifecycle state is **RUNNING**, the instance pool will create new instances or terminate existing instances at that time, to match the new size of the instance pool. Instances are terminated in the order that they were created, first-in, first-out.
- If the instance pool lifecycle state is **STOPPED**, for an increase in size, new instances will be configured for the instance pool, but won't be launched. For a decrease in size, the instances will be terminated.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).



Important

If the instance pool has been in the scaling or provisioning state for an extended period of time, it may be because the number of instances requested has exceeded your tenancy's service limits for that shape and availability domain. Check your tenancy's [service limits](#) for Compute.

To attach a load balancer to an instance pool

You must have a load balancer and backend set to associate with the instance pool. For steps

to create a load balancer, see [Managing a Load Balancer](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the name of the instance pool that you're interested in.
3. In the **Resources** section, click **Load Balancers**.
4. Click **Attach a Load Balancer**.
5. Enter the following:
 - **Load Balancer Compartment:** The compartment that the load balancer is in.
 - **Load Balancer:** The load balancer to associate with the instance pool.
 - **Backend Set:** The name of the backend set on the load balancer to add instances to.
 - **Port:** The server port on the instances to which the load balancer must direct traffic. This value applies to all instances that use this load balancer attachment.
 - **VNIC:** The VNIC to use when adding the instance to the backend set. Instances that belong to a backend set are also called backend servers. The private IP address is used. This value applies to all instances that use this load balancer attachment.
6. Click **Attach**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

To detach a load balancer from an instance pool

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the name of the instance pool that you're interested in.
3. In the **Resources** section, click **Load Balancers**.
4. Click the Actions icon (three dots) for the load balancer.

5. Click **Detach**, and then click **Detach** to confirm.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

To manage tags for an instance pool

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the instance pool that you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Add tags** to add new ones.

For more information, see [Resource Tags](#).

Using the API

To update other instance pool configuration settings, use the CLI, SDKs, or REST APIs. For information about using the CLI, see [Command Line Interface \(CLI\)](#). For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For instance pool configuration settings, such as the instance configuration, display name, tags, or availability domain selections, use the [UpdateInstancePool](#) operation.

To manage the load balancers that are associated with an instance pool, use the [AttachLoadBalancer](#) and [DetachLoadBalancer](#) operations.

To update the configuration used by the instance pool when creating instances you can either:

- Create a new instance configuration with the desired settings. You can do this using the Console. For steps, see [Creating an Instance Configuration](#). To do this using the API, use the [CreateInstanceConfiguration](#) operation.
- Update the existing instance configuration for the instance pool. You can only update the display name and tags of existing instance configurations. For any other updates, create a new instance configuration with the settings you want to use. To update the display

name or tags, use the [UpdateInstanceConfiguration](#) operation. You cannot use the Console to update instance configuration settings.

Stopping and Starting the Instances in an Instance Pool

You can stop and start the instances in an instance pool as needed to update software or resolve error conditions.

Stopping or Restarting an Instance From Within the Instance

In addition to using the API and Console, you can stop and restart instances using the commands available in the operating system when you are logged in to the instance. Stopping an instance using the instance's OS does not stop billing for that instance. If you stop the instances in an instance pool this way, be sure to also stop the instance pool from the Console or API.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Resource Billing for Stopped Instances

For both VM and bare metal instances, billing depends on the [shape](#) that you use to create the instance:

- **Standard shapes:** Stopping an instance pool pauses billing. However, stopped instances continue to count toward your service limits.
- **Dense I/O shapes:** Billing continues for stopped instance pools because of the attached NVMe storage, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance pool](#).
- **GPU shapes:** Billing continues for stopped instance pools, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance pool](#).
- **HPC shapes:** Billing continues for stopped instance pools because of the attached NVMe storage, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance pool](#).

Stopping an instance using the instance's OS does not stop billing for that instance. If you stop the instances in an instance pool this way, be sure to also stop the instance pool from the Console or API.

For more information about how instances running Microsoft Windows Server are billed when they are stopped, see [How am I charged for Windows Server on Oracle Cloud Infrastructure?](#)

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the instance pool that contains the instances that you want to stop or start.

3. Click one of the following actions:

- **Start:** Starts all instances in the instance pool.
- **Stop:** Shuts down all instances in the instance pool.
- **Reboot:** Gracefully reboots all instances in the instance pool by sending a shutdown command to the operating system, and then powers the instances back on.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage the lifecycle state of the instances in an instance pool, use these operations:

- [StartInstancePool](#)
- [StopInstancePool](#)
- [ResetInstancePool](#)
- [SoftresetInstancePool](#)

Deleting an Instance Configuration

You can permanently delete instance configurations that you no longer need.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Configurations**.
2. Click the instance configuration that you want to delete.
3. Click **Delete**, and then click **Delete Instance Configuration**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [DeleteInstanceConfiguration](#) operation to delete an instance configuration.

Deleting an Instance Pool

You can permanently delete instance pools that you no longer need.



Warning

When you delete an instance pool all of its resources will be permanently deleted, including associated instances, attached boot volumes, and block volumes.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to instance pools and instance configurations, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the instance pool that you want to delete.
3. Click **Actions**, and then click **Terminate**.
4. Enter the name of the instance pool and click **Terminate**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [TerminateInstancePool](#) operation to delete an instance pool.

Autoscaling

Autoscaling enables you to automatically adjust the number of Compute instances in an instance pool based on performance metrics such as CPU utilization. This helps you provide consistent performance for your end users during periods of high demand, and helps you reduce your costs during periods of low demand.

You select a performance metric to monitor, and set thresholds that the performance metric must reach to trigger an autoscaling event. When system usage meets a threshold,

autoscaling dynamically allocates resources in near-real time. As load increases, instances are automatically provisioned: the instance pool *scales out*. As load decreases, instances are automatically removed: the instance pool *scales in*.

When an instance pool scales in, instances are terminated in this order: the number of instances is balanced across availability domains, and then balanced across fault domains. Finally, within a fault domain, the oldest instance is terminated first.

Autoscaling relies on performance metrics that are collected by the Monitoring service. These performance metrics are aggregated into one-minute time periods and then averaged across the instance pool. When three consecutive values (that is, the average metrics for three consecutive minutes) meet the threshold, an autoscaling event is triggered.

A cooldown period between autoscaling events lets the system stabilize at the updated level. The cooldown period starts when the instance pool reaches a steady state. Autoscaling continues to evaluate performance metrics during the cooldown period. When the cooldown period ends, autoscaling adjusts the instance pool's size again if needed.

For background information about the Monitoring service, see [Monitoring Overview](#). For information about the metrics emitted by Compute instances, see [Compute Instance Metrics](#).



Note

Autoscaling is not available in Oracle Cloud Infrastructure Government Cloudrealms.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to autoscaling configurations, see [Let users manage Compute autoscaling configurations](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Prerequisites

- You have an instance pool. Optionally, you can attach a load balancer to the instance pool. For steps to create an instance pool and attach a load balancer, see [Creating an Instance Pool](#).
- Monitoring is enabled on the instances in the instance pool. For steps to enable monitoring, see [Enabling Monitoring for Compute Instances](#).
- The instance pool supports the maximum number of instances that you want to scale to. This limit is determined by your tenancy's service limits. For more information, see [Service Limits](#).

Using the Console

To create an autoscaling configuration

Each instance pool can have one autoscaling configuration.

You can create an autoscaling configuration from the **Instance Pools** page or the **Autoscaling Configurations** page.

To create an autoscaling configuration from the Instance Pools page

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the instance pool that you want to apply an autoscaling configuration to.
3. Click **Actions**, and then click **Create Autoscaling Configuration**.
4. Select the **Compartment** that you want to create the autoscaling configuration in.
5. Enter a name for the autoscaling configuration.
6. Select the **Instance Pool** to apply the autoscaling configuration to. The current instance pool is selected by default.
7. In the **Cooldown in Seconds** box, enter the minimum amount of time to wait between scaling events. The cooldown period gives the system time to stabilize before rescaling. The minimum value is 300 seconds.
8. In the **Autoscaling Policy** area, define the criteria that trigger autoscaling actions and the actions to take:
 - a. Enter a name for the autoscaling policy.
 - b. Select the **Performance Metric** that triggers an increase or decrease in the number of instances in the instance pool. For example, CPU utilization or memory utilization.

- c. In the **Scaling Limits** area, specify the number of instances in the instance pool:
 - **Minimum Number of Instances:** The minimum number of instances that the instance pool is allowed to decrease to.
 - **Maximum Number of Instances:** The maximum number of instances that the instance pool is allowed to increase to. The maximum number of instances depends on the limits for your tenancy. For more information, see [Service Limits](#).
 - **Initial Number of Instances:** The number of instances to launch in the instance pool immediately after autoscaling is enabled. After autoscaling retrieves performance metrics, the number of instances is automatically adjusted from this initial number to a number that is based on the limits that you set.
- d. In the **Scaling Rule** area, specify the thresholds that the performance metric must reach to trigger a scaling event:
 - Select a **Scale-Out Operator** and **Threshold Percentage** at which to increase the number of instances. Then enter the **Number of Instances to Add**. For example, when CPU utilization is greater than 90%, add 10 instances to the instance pool.
 - Select a **Scale-In Operator** and **Threshold Percentage** at which to decrease the number of instances. Then enter the **Number of Instances to Remove**. For example, when CPU utilization is less than 50%, remove 5 instances from the instance pool.
9. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

10. Click **Create**.

Autoscaling runs. The cooldown period starts when the instance pool's status changes from **SCALING** to **RUNNING**.

To create an autoscaling configuration from the Autoscaling Configuration page

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
2. Click **Create Autoscaling Configuration**.
3. Select the **Compartment** that you want to create the autoscaling configuration in.
4. Enter a name for the autoscaling configuration.
5. Select the **Instance Pool** to apply the autoscaling configuration to.
6. In the **Cooldown in Seconds** box, enter the minimum amount of time to wait between scaling events. The cooldown period gives the system time to stabilize before rescaling. The minimum value is 300 seconds.
7. In the **Autoscaling Policy** area, define the criteria that trigger autoscaling actions and the actions to take:
 - a. Enter a name for the autoscaling policy.
 - b. Select the **Performance Metric** that triggers an increase or decrease in the number of instances in the instance pool. For example, CPU utilization or memory utilization.

- c. In the **Scaling Limits** area, specify the number of instances in the instance pool:
 - **Minimum Number of Instances:** The minimum number of instances that the instance pool is allowed to decrease to.
 - **Maximum Number of Instances:** The maximum number of instances that the instance pool is allowed to increase to. The maximum number of instances depends on the limits for your tenancy. For more information, see [Service Limits](#).
 - **Initial Number of Instances:** The number of instances to launch in the instance pool immediately after autoscaling is enabled. After autoscaling retrieves performance metrics, the number of instances is automatically adjusted from this initial number to a number that is based on the limits that you set.
- d. In the **Scaling Rule** area, specify the thresholds that the performance metric must reach to trigger a scaling event:
 - Select a **Scale-Out Operator** and **Threshold Percentage** at which to increase the number of instances. Then enter the **Number of Instances to Add**. For example, when CPU utilization is greater than 90%, add 10 instances to the instance pool.
 - Select a **Scale-In Operator** and **Threshold Percentage** at which to decrease the number of instances. Then enter the **Number of Instances to Remove**. For example, when CPU utilization is less than 50%, remove 5 instances from the instance pool.
8. In the **Autoscaling Policy** area, define the criteria that trigger autoscaling actions and the actions to take:
 - a. Enter a name for the autoscaling policy.
 - b. Select the **Performance Metric** that triggers an increase or decrease in the number of instances in the instance pool. For example, CPU utilization or memory utilization.

- c. In the **Scaling Limits** area, specify the number of instances in the instance pool:
 - **Minimum Number of Instances:** The minimum number of instances that the instance pool is allowed to decrease to.
 - **Maximum Number of Instances:** The maximum number of instances that the instance pool is allowed to increase to. The maximum number of instances depends on the limits for your tenancy. For more information, see [Service Limits](#).
 - **Initial Number of Instances:** The number of instances to launch in the instance pool immediately after autoscaling is enabled. After autoscaling retrieves performance metrics, the number of instances is automatically adjusted from this initial number to a number that is based on the limits that you set.
- d. In the **Scaling Rule** area, specify the thresholds that the performance metric must reach to trigger a scaling event:
 - Select a **Scale-Out Operator** and **Threshold Percentage** at which to increase the number of instances. Then enter the **Number of Instances to Add**. For example, when CPU utilization is greater than 90%, add 10 instances to the instance pool.
 - Select a **Scale-In Operator** and **Threshold Percentage** at which to decrease the number of instances. Then enter the **Number of Instances to Remove**. For example, when CPU utilization is less than 50%, remove 5 instances from the instance pool.
9. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

10. Click **Create**.

Autoscaling runs. The cooldown period starts when the instance pool's status changes from **SCALING** to **RUNNING**.

To edit an autoscaling configuration

You can change these characteristics of an autoscaling configuration:

- Name
 - Cooldown period between autoscaling actions
1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
 2. Click the autoscaling configuration that you're interested in.
 3. Click **Edit**.
 4. Update the **Name** or **Cooldown in Seconds**, and then click **Save Changes**.

To edit an autoscaling policy

You can change these characteristics of an autoscaling policy:

- Name
 - Which performance metric triggers an autoscaling action
 - Scale-out and scale-in operators and thresholds
 - The number of instances to add or remove
1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
 2. Click the autoscaling configuration that you're interested in.

3. In the **Autoscaling Policies** area, click **Edit**.
4. Make your changes, and then click **Save**.

To disable an autoscaling configuration

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
2. Click the autoscaling configuration that you're interested in.
3. Click **Disable**, and then confirm when prompted.

To delete an autoscaling configuration

When you delete an autoscaling configuration, the instance pool remains in its most recent state.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
2. Click the autoscaling configuration that you're interested in.
3. Click **Delete**, and then confirm when prompted.

To manage tags for an autoscaling configuration

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
2. Click the autoscaling configuration that you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [Autoscaling API](#) to manage autoscaling configurations and policies.

Managing Cluster Networks

A cluster network is a pool of high performance computing (HPC) instances that are connected with a high-bandwidth, ultra low-latency network. Each node in the cluster is a bare metal machine located in close physical proximity to the other nodes. A remote direct memory access (RDMA) network between nodes provides latency as low as single-digit microseconds, comparable to on-premises HPC clusters.

Cluster networks are designed for highly demanding parallel computing workloads. For example:

- Computational fluid dynamics simulations for automotive or aerospace modeling
- Financial modeling and risk analysis
- Biomedical simulations
- Trajectory analysis and design for space exploration
- Artificial intelligence and big data workloads

Cluster networks are built on top of the [instance pools](#) feature. Most operations in the instance pool are managed directly by the cluster network, though you can monitor and add tags to the underlying instance pool.

For more information about how to access and store the data that you want to process in your cluster networks, see [FastConnect Overview](#), [Overview of File Storage](#), [Overview of Object Storage](#), and [Overview of Block Volume](#).



Note

Cluster networks are not available in Government Cloud realms.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Supported Regions and Availability Domains for Cluster Networks

Cluster networks are supported in the following regions:

- Germany Central (Frankfurt)
- Japan East (Tokyo)
- UK South (London)
- US East (Ashburn)

The availability domain that you create the cluster network in must have cluster network-capable hardware. Typically, to be able to create the multiple HPC instances that are contained in a cluster network, you must [request a service limit increase](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to cluster networks, see [Let users manage Compute instance configurations, instance pools, and cluster networks](#).



Important

See this [known issue](#) for information about the policy statements that are required if the instance configuration or load balancer associated with the cluster network includes defined tags.

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Prerequisites

Create an instance configuration for the instance pool that is managed by the cluster network. To do this:

- a. [Create an instance](#) with the following settings:
 - **Choose an operating system or image source:** Click **Change Image Source**, and then click **Oracle Images**. Select the Oracle HPC cluster networking image.

- **Instance type:** Select **Bare Metal Machine**.
 - **Instance Shape:** Select the **BM.HPC2.36** shape. For more information about this shape, see [Compute Shapes](#).
- b. [Create an instance configuration](#) using the instance that you created in the previous step as a template.
- Optionally, you can delete the instance after you create the instance configuration.

Using the Console

To create a cluster network

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Cluster Networks**.
2. Click **Create Cluster Network**.
3. Select the compartment that you want to create the cluster network in.
4. Specify a name for the cluster network. It doesn't have to be unique, and you can change it later.
5. Select the availability domain to run the cluster network in. Only the availability domains with cluster network-capable hardware can be selected.
6. In the **Configure networking** section, specify the network that you want to use to administer the cluster network. This network is separate from the closed RDMA network between nodes within the cluster. Enter the following information:
 - **Select a virtual cloud network:** The virtual cloud network (VCN) for the cluster network.
 - **Select a subnet:** The subnet for the cluster network.
7. In the **Configure instance pool** section, enter the following:
 - **Instance pool name:** A name for the instance pool that is managed by the cluster network.
 - **Number of instances:** The number of instances in the pool.

- **Select an instance configuration:** Select the instance configuration to use when creating the instances in the cluster network's instance pool, as described in the [prerequisites](#).
8. **Show Advanced Tagging Options:** Optionally, you can add tags. If you have permissions to create a resource, you also have permissions to add free-form tags to that resource. To add a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should add tags, skip this option (you can add tags later) or ask your administrator.
 9. Click **Create Cluster Network**.

To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

If the required number of instances is not available or if some instances fail to launch, the cluster network is not created. Wait a few minutes, and then try launching the cluster network again.

To edit the name of a cluster network

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Cluster Networks**.
2. Click the cluster network that you're interested in.
3. Click **Edit Name**.
4. Enter a new name, and then click **Save Changes**.

To manage tags for a cluster network

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Cluster Networks**.
2. Click the cluster network that you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Add Tags** to add new ones.

For more information, see [Resource Tags](#).

To delete a cluster network



Warning

When you delete a cluster network, all of its resources are permanently deleted, including associated instances, attached boot volumes, and block volumes.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Cluster Networks**.
2. Click the cluster network that you're interested in.
3. Click **Terminate**, and then confirm when prompted.
To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to work with cluster networks:

- [CreateClusterNetwork](#)
- [GetClusterNetwork](#)
- [ListClusterNetworks](#)
- [ListClusterNetworkInstances](#)
- [UpdateClusterNetwork](#)

- [ChangeClusterNetworkCompartment](#)
- [TerminateClusterNetwork](#)

Dedicated Virtual Machine Hosts

The Oracle Cloud Infrastructure Compute service's dedicated virtual machine host feature provides you with the ability to run your Compute virtual machine (VM) instances on dedicated servers that are a single tenant and not shared with other customers. This enables scenarios where you have compliance and regulatory requirements for isolation that prevent you from using shared infrastructure.

Support and Limitations

When you create a dedicated virtual machine host, you select a shape for the host, see [Dedicated Virtual Machine Host Shapes](#) for the available shapes and shape details for dedicated virtual machine host. Note that there is a difference between the number listed for billed OCPUs compared to available OCPUs, this is because four OCPUs are reserved for virtual machine management.

You are billed for the dedicated virtual machine host as soon as you create it, but you are not billed for any of the individual VM instances you place on it. You will still be billed for image licensing costs if they apply to the image you are using for the VM instances.

For instances launched on a dedicated virtual machine host, all of the VM.Standard2 shapes are supported, for details about these shapes, see [VM Shapes](#). Most of the Compute service features for VM instances are supported for instances running on dedicated virtual machine hosts, however the following features are not supported:

- Instance configurations
- Instance pools
- Autoscaling

Reboot migration is also not supported for dedicated virtual machine hosts, in this scenario, you need to manually migrate the instance. See [Moving an Instance with Manual Migration](#) for this process.

You can mix VM instances with different shapes on the same dedicated virtual machine host. This may impact the maximum number of instances you can place on the dedicated virtual machine host, for more information see [Optimizing Capacity on your Dedicated Virtual Machine Host](#).

Managing Dedicated Virtual Machine Hosts

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The simplest policy to enable users to work with dedicated virtual machine hosts is listed in [Let users manage Compute dedicated virtual machine hosts](#). It gives the specified group general access to launching instances on and managing dedicated virtual machine hosts.

See [Let users launch Compute instances on dedicated virtual machine hosts](#) for an example of a policy that allows users to launch instances on dedicated virtual machine hosts without giving them full administrator access to dedicated virtual machine hosts.

Creating a Dedicated Virtual Machine Host

You need to create a dedicated virtual machine host before you can place any instances on it. When creating the dedicated virtual machine host, you select an availability domain and fault domain to launch it in. All the VM instances you place on the host will subsequently be created in this availability domain and fault domain. You also select a compartment when you create the dedicated virtual machine host, but you can move it to a new compartment later without impacting any of the instances placed on it. You can also create the instances in a different compartment than the dedicated virtual machine, or move them to difference compartments after they have been launched.

To create a dedicated virtual machine host using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Dedicated Virtual Machine Hosts**.
2. Click **Create Dedicated Virtual Machine Host**.
3. Fill in the required volume information:
 - **Compartment:** The compartment for the dedicated virtual machine host.
 - **Name:** A user-friendly name or description.
 - **Availability Domain:** The availability domain for the dedicated virtual machine host.
 - **Fault Domain:** The fault domain for the dedicated virtual machine host.
 - **Shape:** The shape to use for the dedicated virtual machine host.
4. Click **Create Dedicated Virtual Machine Host**.

To create a dedicated virtual machine host using the CLI

Open a command prompt and run:

```
oci compute dedicated-vm-host create --dedicated-vm-host-shape DVH.Standard2.52 --wait-for-state ACTIVE
--display-name <display_name> --availability-domain <availability_domain> --compartment-id <compartment_
ID>
```

It can take up to 15 minutes for the dedicated virtual machine host to be fully created. It must be in the `ACTIVE` state before you can launch an instance on it.

To query the current state of a dedicated virtual machine host using the CLI, run the following command:

```
oci compute dedicated-vm-host get --dedicated-vm-host-id <dedicatedVMhost_ID>
```

Deleting a Dedicated Virtual Machine Host

To delete a dedicated virtual machine host using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Dedicated Virtual Machine Hosts**.
2. In the list of dedicated virtual machine hosts, find the dedicated virtual machine hosts you want to delete.
3. Click the name of the dedicated virtual machine host to display the details page.
4. Click **Delete**, and then respond to the confirmation prompt.

If you try to delete a dedicated virtual machine host that still has running instances hosted on it, the delete operation will fail, you need to ensure that all of the instances hosted on it have been terminated. To check if there are any instances still running on the dedicated virtual machine host, go to the **Details** page for the dedicated virtual machine host, and click **Hosted Instances** in the **Resources** section. You need to perform this step for each compartment in your tenancy that may have instances running on your dedicated virtual machine host. To change the compartment for the **Host Instances** list, select a different compartment from the **Table Scope** drop down.

To delete a dedicated virtual machine host using the CLI

Open a command prompt and run:

```
oci compute dedicated-vm-host delete --dedicated-vm-host-id <dedicated_VM_host_ID>
```

Before you can delete a dedicated machine host, all of the instances running on it must be terminated.

To list the instances running on a dedicated virtual machine host using the CLI, run the following command:

```
oci compute dedicated-vm-host list --compartment-id <compartment_ID> --dedicated-vm-host-id <dedicatedVMhost_ID>
```

You need to run this command for every compartment in your tenancy that may have instances that you placed on the dedicated virtual machine host that you want to delete.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations for working with dedicated virtual machine hosts:

- [CreateDedicatedVmHost](#)
- [DeleteDedicatedVmHost](#)
- [ListDedicatedVmHosts](#)
- [ListDedicatedVmHostShapes](#)
- [ListDedicatedVmHostInstances](#)
- [ListDedicatedVmHostInstanceShapes](#)
- [UpdateDedicatedVmHost](#)
- [ChangeDedicatedVmHostCompartment](#)

Instances on Dedicated Virtual Machine Hosts

Placing an Instance on a Dedicated Virtual Machine Host

You place an instance on a dedicated virtual machine host at the time that you create the instance. The steps are the same as creating a regular instance, you just need specify that you want to create the instance on a dedicated virtual machine host when you create the instance. See [Creating an Instance](#) for the steps to create an instance. Once you get to the **Advanced Options** section of the form, using the following steps to place the instance in a dedicated virtual machine host.

To place an instance on a dedicated virtual machine host using the Console

1. Perform the initial steps to create an instance based on an image and shape type that support placement on a dedicated virtual machine, up to the **Configure networking** section.
2. Click **Show Advanced Options**, and then click the **Host** tab.
3. Check **Launch the virtual machine on a dedicated host**.
4. Select the dedicated virtual machine host you want to place the instance on from the list.
5. Click **Create**.



Note

Only dedicated virtual machine hosts with sufficient capacity to launch an instance based on the shape you have specified will be displayed in the **Dedicated Virtual Machine Host** drop down list. If you have a dedicated virtual machine host and it does not appear in the list, you need to do one of the following to place the instance on a dedicated virtual machine host:

- Terminate instances you no longer need on the dedicated virtual machine host to free up capacity.
- Choose another smaller shape for the instance you are trying to place on the dedicated virtual machine host.
- Create a new dedicated virtual machine host to place the instance on.

For more information, see [Optimizing Capacity on your Dedicated Virtual Machine Host](#).

If you're using the CLI or REST API to create the instance, you just need to pass the dedicated virtual machine host OCID in the optional parameter `dedicatedVmHostId` when you use the [LaunchInstance](#) operation. If you try to launch an instance with a shape requiring more capacity than what is available on the dedicated virtual machine host you are trying to place it on, the launch operation will fail. To avoid this you can use the [ListDedicatedVmHosts](#) operation and pass the shape you want to use when launching the instance in the `InstanceShapeNameQueryParam` parameter. This will return all the dedicated virtual machine hosts that you can place the instance in.

CHAPTER 8 Compute

The following example demonstrates how to call this operation in the CLI to return all the dedicated virtual machine hosts with sufficient capacity for you to place an instance launched using the VM.Standard2.16 shape:

```
compute dedicated-vm-host list --compartment-id <compartment_ID> --instance-shape-name VM.Standard2.16
```

Auditing your Dedicated Virtual Machine Host

To fully meet requirements for some compliance scenarios you may be required to validate that your instances are running on a dedicated virtual machine host and not using shared infrastructure. The Oracle Cloud Infrastructure Audit service provides you with the functionality to do this. Use the steps described in the [Viewing Audit Log Events](#) to access the log events for the dedicated virtual machine host.

The steps described in the **To search log events** section walk you through how to retrieve the log events with the data you need to verify that your instances are running on a dedicated virtual machine host. For this procedure:

- Ensure that you select the dedicated virtual machine host's compartment and not the compartment for the instances hosted on it.
- Use the dedicated virtual machine host's OCID as the search keyword.

Once you have retrieved the log events for the dedicated virtual machine host, view the log event lower-level details, and check the contents of the `responsePayload` property. This property should contain the OCIDs for the instances running on the dedicated virtual machine host.

Optimizing Capacity on your Dedicated Virtual Machine Host

When you place an instance on a dedicated virtual machine host using the Console, only dedicated virtual machine hosts with sufficient capacity to launch an instance based on the shape you have specified will be displayed in the **Dedicated Virtual Machine Host** drop down list. If you don't see your dedicated virtual machine host in the list, to understand why, it may help to understand how instances are launched in this scenario.

When you place an instance on a dedicated virtual machine host, Oracle Cloud Infrastructure launches them in a manner to optimize performance. For example, a dedicated virtual

machine host created based on the DVH.Standard2.52 shape has two sockets with 24 cores configured per socket. Instances are placed so that each instance will only use resources local to a single physical socket. In scenarios where you are creating and terminating instances with a mix of shapes, this can result in inefficient distribution of resources, meaning that not all OCPUs on a dedicated virtual machine host are available to be used. In this scenario, it may appear that a dedicated virtual machine has enough OCPUs to launch an additional instance on it, but the instance will fail to launch because of their distribution.

In this example, if you are launching instances using a shape with 16 OCPUs on a dedicated virtual machine host, you can only launch a maximum of two instances using that shape, you cannot launch a third instance with 16 OCPUs, even though the remaining number of OCPUs showing for the dedicated virtual machine host is 16. You can launch additional instances using shapes with a smaller number of OCPUs.

When designing your cloud footprint, we recommend that you plan to always launch the largest instance first.

Connecting to an Instance

You can connect to a running instance by using a Secure Shell (SSH) or Remote Desktop connection. Most UNIX-style systems include an SSH client by default. To connect to a Linux instance from a Windows system, you can download a free SSH client called PuTTY from <http://www.putty.org>.

Required IAM Policy

To connect to a running instance with SSH, you don't need an IAM policy to grant you access. However, to SSH you need the public IP address of the instance (see [Prerequisites](#) below). If there's a policy that lets you launch an instance, that policy probably also lets you get the instance's IP address. The simplest policy that does both is listed in [Let users launch Compute instances](#).

For administrators: Here's a more restrictive policy that lets the specified group get the IP address of existing instances and use power actions on the instances (e.g., stop, start, etc.),

CHAPTER 8 Compute

but not launch or terminate instances. The policy assumes the instances and the cloud network are together in a single compartment (XYZ):

```
Allow group InstanceUsers to read virtual-network-family in compartment XYZ
```

```
Allow group InstanceUsers to use instance-family in compartment XYZ
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Prerequisites

You'll need the following information to connect to the instance:

- For Linux instances: The full path to the key pair that you used when you launched the instance. For information about generating key pairs, see [Managing Key Pairs on Linux Instances](#).
- The default user name for the instance. If you used an Oracle-provided Linux, CentOS, or Windows image to launch the instance, the user name is `opc`. If you used the Ubuntu image to launch the instance, the user name is `ubuntu`.
- The public IP address of the instance. You can get the address from the Instance Details page in the Console. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Then, select your instance. Alternatively, you can use the Core Services API [ListVnicAttachments](#) and [GetVnic](#) operations.
- For Windows instances: If you're connecting to the instance for the first time, you will need the initial password for the instance. You can get the password from the Instance Details page in the Console.

Connecting to a Linux Instance

Log in to your instance using SSH.

Connecting to Your Linux Instance from a Unix-style System

1. Use the following command to set the file permissions so that only you can read the file:

```
$ chmod 400 <private_key>
```

<private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.

2. Use the following SSH command to access the instance.

```
$ ssh -i <private_key> <username>@<public-ip-address>
```

<private_key> is the full path and name of the file that contains the private key associated with the instance you want to access.

<username> is the default name for the instance. For Oracle Linux and CentOS images, the default user name is `opc`. For the Ubuntu image, the default name is `ubuntu`.

<public-ip-address> is your instance IP address that you retrieved from the Console.

Connecting to Your Linux Instance from a Windows System

1. Open `putty.exe`.
2. In the **Category** pane, expand **Window**, and then select **Translation**.
3. In the **Remote character set** drop-down list, select **UTF-8**. The default locale setting on Linux-based instances is UTF-8, and this configures PuTTY to use the same locale.
4. In the **Category** pane, select **Session** and enter the following:

- **Host Name (or IP address):**

<username>@*<public-ip-address>*

<username> is the default name for the instance. For Oracle Linux and CentOS images, the default user name is `opc`. For the Ubuntu image, the default name is `ubuntu`.

<public-ip-address> is your instance public IP address that you retrieved from the Console

- **Port:** 22
 - **Connection type:** SSH
5. In the **Category** pane, expand **Connection**, expand **SSH**, and then click **Auth**.
 6. Click **Browse**, and then select your private key.
 7. Click **Open** to start the session.
If this is your first time connecting to the instance, you might see a message that the server's host key is not cached in the registry. Click **Yes** to continue the connection.

Connecting to a Windows Instance

You can connect to a Windows instance by using a Remote Desktop connection. Most Windows systems include a Remote Desktop client by default.

To enable Remote Desktop Protocol (RDP) access to the Windows instance, you need to add a stateful ingress [security rule](#) for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port. You can implement this security rule in either a [network security group](#) that the Windows instance belongs to, or a [security list](#) that is used by the instance's subnet.

To enable RDP access

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the cloud network you're interested in.
3. To add the rule to a network security group that the instance belongs to:
 - a. Under **Resources**, click **Network Security Groups**. Then click the network security group that you're interested in.
 - b. Click **Add Rules**.
 - c. Enter the following values for the rule:

- **Stateless:** Leave the check box unselected
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** RDP (TCP/3389)
 - **Source Port Range:** All
 - **Destination Port Range:** 3389
- d. When done, click **Add**.
4. Or, to add the rule to a security list that is used by the instance's subnet:
- a. Under **Resources**, click **Security Lists**. Then click the security list you're interested in.
 - b. Click **Add Ingress Rules**.
 - c. Enter the following values for the rule:
 - **Stateless:** Leave the check box unselected
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0
 - **IP Protocol:** RDP (TCP/3389)
 - **Source Port Range:** All
 - **Destination Port Range:** 3389
 - d. When done, click **Add Ingress Rules**.

Connecting to Your Windows Instance from a Remote Desktop Client

1. Open the Remote Desktop client.
2. In the **Computer** field, enter the public IP address of the instance you want to connect to. Your public IP is the instance address you get from the Console.

3. The **User name** is **opc**. Depending on the Remote Desktop client you are using, you might have to connect to the instance before you can enter this credential.
4. Click **Connect** to start the session.
5. Accept the certificate if you are prompted to do so.
6. If you are connecting to the instance for the first time, enter the initial password that was provided to you by Oracle Cloud Infrastructure when you launched the instance. You will be prompted to change the password as soon as you log in. Your new password must be at least 12 characters long and must comply with [Microsoft's password policy](#). Otherwise, enter the password that you created. If you are using a custom image, you might need to know the password for the instance that the image was created from. For details about Windows custom images, see [Creating Windows Custom Images](#).
7. Press **Enter**.

Instance Console Connections

The Oracle Cloud Infrastructure Compute service provides console connections that enable you to remotely troubleshoot malfunctioning instances, such as:

- An imported or customized image that does not complete a successful boot.
- A previously working instance that stops responding.

There are two types of instance console connections:

- Serial console connections
- VNC console connections

Creating the Instance Console Connection

Before you can connect to the serial console or VNC console, you need to create the instance console connection.

To create the console connection for an instance

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the list of instances, find the instance you want to access the serial console for, and then click the instance name.
3. In the **Resources** section, click **Console Connections**.
4. Click **Create Console Connection**.
5. Specify the public key (.pub) portion for the SSH key. You can browse to a public key file on your computer or paste your public key into the text box. Then, click **Create Console Connection**.

When the console connection has been created and is available, the status changes to **ACTIVE**.

Connecting to the Serial Console

After you have created the console connection for the instance, you can then connect to the serial console by using a Secure Shell (SSH) connection. When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must reauthenticate to connect again.



Note

Serial console connections for VM instances launched before September 2017

Serial console connections only work for VM instances launched in September 2017 or later.



Note

Serial console connections for bare metal instances launched before November 2017

Serial console connections only work for Bare Metal instances launched in November 2017 or later.

Connecting from Mac OS X and Linux Operating Systems

You connect to the serial console by using an SSH client. Mac OS X and most Linux distributions by default include the SSH client OpenSSH.

To connect to the serial console for an instance using OpenSSH on Mac OS X or Linux

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
2. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
3. Click the Actions icon (three dots), and then click **Connect with SSH**.
4. Select **LINUX/MAC OS** for **PLATFORM**.
5. Click **Copy** to copy the string to the clipboard.
6. Paste the connection string copied from the previous step to a terminal window on a Mac OS X or Linux system, and press **Enter** to connect to the console.

If you are not using the default SSH key or ssh-agent, you can modify the serial console connection string to include the identity file flag, `-i` to specify the SSH key to use. You must specify this for both the SSH connection and the SSH ProxyCommand, as shown in the following line:

```
ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W %h:%p -p 443...
```

7. Press **Enter** again to activate the console.

Connecting from Windows Operating Systems

Windows does not include an SSH client by default, so you need to install one. You can use [PuTTY](#), or there are options that include a version of OpenSSH such as:

- [Git for Windows](#)
- [Windows Subsystem for Linux](#)

The steps to connect to the serial console from the PuTTY client are different from the steps for OpenSSH.

To connect to the serial console for an instance on Microsoft Windows

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
2. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
3. Click the Actions icon (three dots), and then click **Connect with SSH**.
4. If you are using PuTTY, select **WINDOWS** for **PLATFORM**.
If you are using OpenSSH, select **LINUX/MAC OS** for **PLATFORM**.
5. Click **Copy** to copy the string to the clipboard.
6. Paste the connection string copied from the previous step to PuTTY or your OpenSSH client and press **Enter** to connect to the console.
7. Press **Enter** again to activate the console.

Connecting to the VNC Console



Warning

The VNC console connection uses SSH port forwarding to create a secure connection from your local system to the VNC server attached to your instance's console. While this is a secure way to use VNC over the internet, owners of multiuser systems should be aware that opening a port on the local system makes it available to all of the users on that system until a VNC client connects. For this reason, we don't recommend using this product on a multiuser system unless you take proper actions to secure the port or you isolate the VNC client by running it in a virtual environment, such as [Oracle VM VirtualBox](#).

After you create the console connection for the instance, you need to set up a secure tunnel to the VNC server on the instance, and then you can connect with a VNC client.



Note

VNC console connections for VM instances launched before October 13, 2017

VNC console connections only work for VM instances launched on October 13, 2017 or later.



Note

VNC console connections for bare metal instances launched before February 21, 2019

VNC console connections only work for bare metal instances launched on February 21, 2019, or later, using one of the following shapes:

- BM.GPU2.2
- BM.HPC2.36
- BM.Standard2.52
- BM.DenseIO2.52
- BM.Standard.B1.44

To set up a secure tunnel to the VNC server on the instance using OpenSSH on Mac OS X or Linux

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
2. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
3. Click the Actions icon (three dots), and then click **Connect with VNC**.
4. Select **LINUX/MAC OS** for **PLATFORM**.
5. Click **Copy** to copy the string to the clipboard.
6. Paste the connection string copied from the previous step to a terminal window on a Mac OS X or Linux system, and press **Enter** to set up the secure connection.
7. After the connection is established, open your VNC client and specify `localhost` as the host to connect to and `5900` as the port to use.



Note

Mac OS X Screen Sharing.app Not Compatible with VNC Console Connections

The Mac OS X built-in VNC client, Screen Sharing.app does not work with VNC console connections in Oracle Cloud Infrastructure. Use another VNC client, such as [Real VNC Viewer](#) or [Chicken](#).

To set up a secure tunnel to the VNC server on the instance using PowerShell on Windows

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Click the instance you want to connect to.
2. On the **Instances Details** page, in the **Resources** section, click **Console Connections**.
3. Click the Actions icon (three dots), and then click **Connect with VNC**.
4. Select **WINDOWS** for **PLATFORM**.
5. Click **Copy** to copy the string to the clipboard.
6. Paste the connection string copied from the previous step to [Windows Powershell](#) and press **Enter** to set up the secure connection.
7. After the connection is established, open your VNC client and specify `localhost` as the host to connect to and `5900` as the port to use.



Note

Secure Connection Warning

When you connect, you may see a warning from the VNC client that the connection is not encrypted. Since you are connecting through SSH, the connection is secure, so this is not an issue.

Troubleshooting Instances from Instance Console Connections

After you are connected with an instance console connection, you can perform various tasks, such as:

- Edit system configuration files.
- Add or reset the SSH keys for the **opc** user.

Both of these tasks require you to boot into a bash shell, in maintenance mode.



Note

The following tasks describe steps specific to instances running Oracle Linux 7, connecting from OpenSSH. Other OS versions and SSH clients may require different steps.

To boot into maintenance mode

1. Reboot the instance from the Console
 - In the Console, on the **Instances Details page**, click **Reboot**.
2. When the reboot process starts, switch back to the terminal window, and you see

CHAPTER 8 Compute

Console messages start to appear in the window. As soon as you see the GRUB boot menu appear, use the up/down arrow key to stop the automatic boot process, enabling you to use the boot menu.

3. In the boot menu, highlight the top item in the menu, and type `e` to edit the boot entry.
4. In edit mode, use the down arrow key to scroll down through the entries until you reach the line that starts with either `linuxefi` for instances running Oracle Linux 7.x, or `kernel` for instances running Oracle Linux 6.x.
5. At the end of that line, add the following:

```
init=/bin/bash
```

6. Reboot the instance from the terminal window by entering the keyboard shortcut **CTRL+X**.

When the instance has rebooted, you'll see the Bash shell command line prompt, and you can proceed with either of the following procedures.

To edit the system configuration files

1. From the Bash shell, run the following command to load the SELinux policies to preserve the context of the files you are modifying:

```
/usr/sbin/load_policy -i
```

2. Run the following command to remount the root partition with read/write permissions:

```
/bin/mount -o remount, rw /
```

3. Edit the configuration files as needed to try to recover the instance.
4. After you have finished editing the configuration files, to start the instance from the existing shell, run the following command:

```
exec /usr/lib/systemd/systemd
```

Alternatively, to reboot the instance, run the following command:

```
/usr/sbin/reboot -f
```

To add or reset the SSH key for the `opc` user

1. From the Bash shell, run the following command to load the SELinux policies to preserve the context of the files you are modifying:

```
/usr/sbin/load_policy -i
```

2. Run the following command to remount the root partition with read/write permissions:

```
/bin/mount -o remount, rw /
```

3. From the Bash shell, run the following command to change to the SSH key directory for the **opc** user:

```
cd ~opc/.ssh
```

4. Rename the existing authorized keys file with the following command:

```
mv authorized_keys authorized_keys.old
```

5. Replace the contents of the public key file with the new public key file with the following command:

```
echo '<contents of .pub key file>' >> authorized_keys
```

6. Restart the instance by running the following command:

```
/usr/sbin/reboot -f
```

Exiting the Instance Console Connection

To exit the serial console connection

When using SSH, the `~` character at the beginning of a new line is used as an escape character.

- To exit the serial console, enter:

```
~.
```

- To suspend the SSH session, enter:

```
~^z
```

The ^ character represents the **CTRL** key

- To see all the SSH escape commands, enter:

```
~?
```

To exit the VNC console connection

1. Close the VNC client.
2. In the Terminal or Powershell window, type `CTRL C`

When you are finished using the console connection, delete the connection for the instance.

To delete the console connection for an instance

1. In the Console, on the **Instances Details** page, in the **Resources** section, click **Console Connections**.
2. Click the Actions icon (three dots), click **Delete**, and then click **OK** to confirm.

Adding Users on an Instance

If you created your instance using an Oracle-provided Linux or CentOS image, you can use SSH to access your instance from a remote host as the `opc` user. If you created your instance using the Ubuntu image, you can use SSH to access your instance from a remote host as the `ubuntu` user. After logging in, you can add users on your instance.

If you created your instance using an Oracle-provided Windows image, you can create new users after you log on to the instance through a Remote Desktop client.

Creating Additional SSH-Enabled Users on Linux Instances

If you do not want to share your SSH key, you can create additional SSH-enabled users:

- Generate SSH key pairs for the users offline.
- Add the new users.
- Append a public key to the `~/.ssh/authorized_keys` file for each new user.



Tip

If you re-create an instance from an Oracle-provided image, users and SSH public keys that you added or edited manually (that is, users that weren't defined in the machine image) must be added again.

If you need to edit the `~/.ssh/authorized_keys` file of a user on your instance, start a second SSH session before you make any changes to the file and ensure that it remains connected while you edit the file. If the `~/.ssh/authorized_keys` file becomes corrupted or you inadvertently make changes that lock you out of the instance, you can use the backup SSH session to fix or revert the changes. Before closing the backup SSH session, test all changes you made by logging in with the new or updated SSH key.

The new users then can SSH to the instance using the appropriate private keys.

To create an additional SSH-enabled user:

1. Generate an SSH key pair for the new user. See [Managing Key Pairs on Linux Instances](#).
2. Copy the public key value to a text file for use later in this procedure.

CHAPTER 8 Compute

3. Log in to your instance. See [Connecting to an Instance](#).

4. Become the root user:

```
sudo su
```

5. Create the new user:

```
useradd <new_user>
```

6. Create a `.ssh` directory in the new user's home directory:

```
mkdir /home/<new_user>/.ssh
```

7. Copy the SSH public key that you saved to a text file into the `/home/new_user/.ssh/authorized_keys` file:

```
echo <public_key> > /home/<new_user>/.ssh/authorized_keys
```

8. Change the owner and group of the `/home/username/.ssh` directory to the new user:

```
chown -R <new_user>:<group> /home/<new_user>/.ssh
```

9. To enable `sudo` privileges for the new user, run the `visudo` command and edit the `/etc/sudoers` file as follows:

a. In `/etc/sudoers`, look for:

```
%<username> ALL=(ALL) NOPASSWD: ALL
```

b. Add the following line immediately after the preceding line:

```
%<group> ALL=(ALL) NOPASSWD: ALL
```

You can now log in as the new user.

Creating Additional Users on a Windows Instance

To create a new user on a Windows Instance:

1. Log in to your instance using a Remote Desktop client.
2. On the **Start** menu, click **Control Panel**.
3. Click **User Accounts**, and then click **User Accounts** again.
4. Click **Manage User Accounts**.
5. Click **Manage Another Account**.
6. Click **Add User Account**.
7. Enter a **User name** and **Password**.
8. Confirm the password, and then create a **Password hint**.
9. Click **Next**.
10. Verify the account, and then click **Finish**.

You can now log in as the new user.

Displaying the Console for an Instance

You can capture and display the serial console data for an instance. The data includes configuration messages that occur when the instance boots, such as kernel and BIOS messages, and is useful for checking the status of the instance or diagnosing problems. Note that the *raw* console data, including multi-byte characters, is captured.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to manage console history data. If the specified group doesn't need to launch instances or attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.

Using the API

For information about using the API and signing requests, see [REST APIs](#).

Use these API operations to manage the serial console logs:

- [CaptureConsoleHistory](#)
- [DeleteConsoleHistory](#)
- [GetConsoleHistory](#)
- [GetConsoleHistoryContent](#)
- [ListConsoleHistories](#)

Getting Instance Metadata

The metadata for an instance includes information such as the instance's OCID, display name, hostname, region, availability domain, fault domain, compartment, shape, image, creation date, state, tags, and any custom metadata that you provide, such as an SSH public key. The instance metadata also includes the [region identifier](#) for an instance, such as `us-phoenix-1`, in the `canonicalRegionName` field.

You can find some of this information in the Console on the **Instance Details** page, or you can get all of it by logging in to the instance and using the metadata service. The service runs on every instance and is an HTTP endpoint listening on `169.254.169.254`.

Required IAM Policy

No IAM policy is required if you're logged in to the instance and using cURL to get the metadata.

For administrators: Users can also get instance metadata through the Compute API (for example, with [GetInstance](#)). The policy in [Let users launch Compute instances](#) covers that ability. If the specified group doesn't need to launch instances or attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Accessing Instance Metadata on Oracle-Provided Images

You can get instance metadata for Oracle-provided images by using cURL on Linux instances or an internet browser on Windows instances.

Instance metadata is available at the following URLs:

- All of the instance information:

```
http://169.254.169.254/opc/v1/instance/
```

- Only the custom metadata:

```
http://169.254.169.254/opc/v1/instance/metadata/
```

- The custom metadata only for a specified key name:

```
http://169.254.169.254/opc/v1/instance/metadata/<key-name>
```

<key-name> is `ssh_authorized_keys`, `user_data`, or any custom key name that you provided when you launched the instance. (For information about using the Core Services API to provide `user_data` to cloud-init, see [LaunchInstanceDetails](#).)

- Information about the virtual network interface cards (VNICs) that are attached to the instance:

```
http://169.254.169.254/opc/v1/vnics/
```

CHAPTER 8 Compute

Here's an example response that shows of all of the information for an instance:

```
{
  "availabilityDomain": "cumS:PHX-AD-1",
  "faultDomain": "FAULT-DOMAIN-2",
  "compartmentId": "ocid.compartment.oc1..exampleuniqueID",
  "displayName": "my-example-instance",
  "hostname": "my-hostname",
  "id": "ocid1.instance.oc1.phx.exampleuniqueID",
  "image": "ocid1.image.oc1.phx.exampleuniqueID",
  "metadata": {
    "ssh_authorized_keys": "ssh-rsa
AAAAB3NzaC1yc2EAAAABJQAAAQEApzWsZ0DxxE+jNpsh5C3ncc18eZ06MCONWnP5D/8eqdGjUqqBU7lHvgaIqYHy1PH/B4w/pya56k4
1FLzdXFVAzCy4NGyE4XZjLcSSR4jWlAs4r3WHY/f61Gc1VicZ39u+bB1XYaHm46zS4WrQEQuU4wbz70iwe50nSIhrGpvM5HWYOK0ds
VA7/zzw+yW37NUGa/QeM4/bJvCVg3BVjB6VWdmV7dFwRMeCaVJFQH3wKndvuJib78zoH19sbYm74vzqTYSi/bVoIz9YnZ4bA3MS0Uqa
poK/m2M9T27+UA/lz/ILCKXP3+vNcVcjRplanJT/qlzhLiIiBCRo4RsdGxUIw== rsa-key-20181129"
  },
  "region": "phx",
  "canonicalRegionName": "us-phoenix-1",
  "shape": "VM.Standard2.1",
  "state": "Running",
  "timeCreated": 1569358494495,
  "agentConfig": {
    "monitoringDisabled": false
  },
  "definedTags": {
    "Operations": {
      "CostCenter": "42"
    }
  },
  "freeformTags": {
    "Department": "Finance"
  }
}
```

For more information about the data that is returned, see [Instance](#).

Here's an example response that shows the VNICs that are attached to an instance:

```
[ {
  "vnicId" : "ocid1.vnic.oc1.phx.exampleuniqueID",
  "privateIp" : "10.0.3.6",
  "vlanTag" : 11,
}
```

CHAPTER 8 Compute

```
"macAddr" : "02:00:17:00:12:D3",
"virtualRouterIp" : "10.0.3.1",
"subnetCidrBlock" : "10.0.3.0/24",
"nicIndex" : 0
}, {
  "vnicId" : "ocid1.vnic.oc1.phx.exampleuniqueID",
  "privateIp" : "10.0.4.3",
  "vlanTag" : 12,
  "macAddr" : "02:00:17:00:13:13",
  "virtualRouterIp" : "10.0.4.1",
  "subnetCidrBlock" : "10.0.4.0/24",
  "nicIndex" : 0
} ]
```

To use cURL to get Linux instance metadata

1. [Connect to a Linux instance](#) using SSH.
2. Use cURL to issue a GET request to the instance metadata URL that you're interested in. For example:

```
curl -L http://169.254.169.254/opc/v1/instance/
```

To use an internet browser to get Windows instance metadata

1. [Connect to a Windows instance](#) by using a Remote Desktop connection.
2. Open an internet browser and then navigate to the instance metadata URL that you're interested in.

Updating Instance Metadata

The Oracle Cloud Infrastructure Compute service lets you add and update custom metadata for an instance using the [Command Line Interface \(CLI\)](#) or [REST APIs](#).

When you create an instance using the [LaunchInstance](#) operation you can specify custom metadata for the instance in the [LaunchInstanceDetails](#) datatype's `metadata` or `extendedMetadata` attributes. To update an instance's metadata, use the [UpdateInstance](#)

operation, specifying the custom metadata in the [UpdateInstanceDetails](#) datatype's `metadata` or `extendedMetadata` attributes. The `metadata` attribute supports key/value string pairs while the `extendedMetadata` attribute supports nested JSON objects.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to rename an instance. If the specified group doesn't need to launch instances or attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the API

When you use the `UpdateInstance` operation, the instance's metadata will be the combination of the values specified in the [UpdateInstanceDetails](#) datatype's `metadata` or `extendedMetadata` attributes. Any set of key/value pairs specified for these attributes in the `UpdateInstance` operation will replace the existing values for these attributes, so you need

to include all the metadata values for the instance in each call, not just the ones you want to add. If you leave the attribute empty when calling `UpdateInstance`, the existing metadata values in that attribute will be used. You cannot specify a value for the same metadata key twice, this will cause the `UpdateInstance` operation to fail due to there being duplicate keys.

To understand this, consider the example scenario where you created an instance using the `LaunchInstance` operation and specified the following key/value pair for the `metadata` attribute:

```
"myCustomMetadataKey" : "myCustomMetadataValue"
```

If you then call the `UpdateInstance` operation, and add new metadata by specifying additional key/value pairs in the `extendedMetadata` attribute, but you leave the `metadata` attribute empty, do not include the `myCustomMetadataKey` key/value in the `extendedMetadata` attribute, as this will cause the operation to fail since that key already exists. If you do specify values for the metadata attribute, you need to include the `myCustomMetadataKey` key/value to maintain it in the instance's metadata. In this case, you can specify it in either of the attributes.

There are two reserved keys, `user_data` and `ssh_authorized_keys`, that can only be set for an instance at launch time, they cannot be updated later. If you use the metadata attribute to add or update metadata to an instance, you need to ensure that you include the values specified at launch time for both these keys, otherwise the `UpdateInstance` operation will fail.

Best Practices for Updating an Instance's Metadata

When using the `UpdateInstance` operation, Oracle recommends the following:

- Use the [GetInstance](#) operation to retrieve the existing custom metadata for the instance to ensure that you include the values you want to maintain in the appropriate attributes when you call `UpdateInstance`. The metadata values are returned in the `metadata` and `extendedMetadata` attributes for the [Instance](#). For a code example demonstrating this, see the [UpdateInstanceExample](#) in the [SDK for Java](#).
- Unless you are updating custom metadata that was added using the `metadata` attribute, use the `extendedMetadata` attribute to add custom metadata. Otherwise you need to include the launch time values for the `user_data` and `ssh_authorized_keys` reserved

keys. If you use the `metadata` attribute to add values and you leave out the values for these reserved keys or specify different values for them, the `UpdateInstance` call will fail.

Renaming an Instance

You can rename an instance without changing its Oracle Cloud Identifier (OCID).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to rename an instance. If the specified group doesn't need to launch instances or attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [UpdateInstance](#) operation to change the name of an instance.

Moving a Compute Instance to a New Host

This topic covers how to relocate a virtual machine or a bare metal instance by using [reboot migration](#) or a [manual process](#).



Note

Dedicated virtual machine hosts do not support reboot migration. To relocate these instances use the process described in [Moving an Instance with Manual Migration](#).

Reboot Migration

For instances with a date in the **Reboot Maintenance** field (available in the Console, CLI, and SDKs), you can reboot your instance to move it to new infrastructure. After you reboot the instance, the **Reboot Maintenance** field is cleared. This change indicates that the instance was moved successfully.

Prerequisites for Reboot Migration

1. Prepare the instance for reboot migration:
 - Ensure that any remote block volumes defined in `/etc/fstab` use the [recommended options](#).
 - Ensure that any File Storage service (NFS) mounts use the `nofail` option.

- If you use the [Oracle-provided script](#) to configure secondary VNICs, ensure it runs automatically at startup.

Moving an Instance with Reboot Migration

After you complete the prerequisites:

1. Stop any running applications.
2. Reboot the instance.
3. Confirm that the **Reboot Maintenance** field no longer has a date.
4. Start and test any applications on the instance.

Moving an Instance with Manual Migration

For instances without a date in the **Reboot Maintenance** field (available in the Console, CLI, and SDKs), you must move the instance manually. This method requires that you terminate the instance, and then launch a new instance from the retained boot volume. Instances that have additional VNICs, secondary IP addresses, remote attached block volumes, or that belong to a backend set of a load balancer require additional steps.

Limitations and Warnings for Manual Migration

Be aware of the following limitations and warnings when performing a manual migration:

- Any public IP addresses assigned to your instance from a [reserved public pool](#) are retained. Any that were not assigned from a reserved public IP address pool will change. Private IP addresses do not change.
- MAC addresses, CPUIDs, and other unique hardware identifiers do change during the move. If any applications running on the instance use these identifiers for licensing or other purposes, be sure to take note of this information before moving the instance to help you manage the change.

Prerequisites for Manual Migration

1. Before moving the instance, document all critical details:
 - The instance's region, availability domain, and fault domain.
 - The instance's display name.
 - All private IP addresses, names, and subnets. Note that the instance can have multiple VNICs, and each VNIC can have multiple secondary IP addresses.
 - All private DNS names. The instance can have multiple VNICs, and each VNIC can have multiple secondary IP addresses. Each private IP address can have a DNS name.
 - Any [public IP addresses](#) assigned from a reserved public pool. Note that the instance can have multiple VNICs, and each VNIC can have multiple secondary private IP addresses. Each VNIC and secondary private IP address can have an attached public IP address.
 - Any remote block volumes attached to the instance.
 - Any tags on the instance or attached resources.
2. Prepare the instance for manual migration:
 - Ensure that any remote block volumes defined in `/etc/fstab` use the [recommended options](#).
 - Ensure that any File Storage service (NFS) mounts use the `nofail` option.
 - If you have statically defined any network interfaces belonging to secondary VNICs using their MAC addresses, such as those defined in `/etc/sysconfig/network-scripts/ifcfg*`, those interfaces will not start due to the change in the MAC address. Remove the static mapping.
 - If you use the [Oracle-provided script](#) to configure secondary VNICs, ensure it runs automatically at startup.

Moving an Instance Manually

After you complete the prerequisites:

1. Stop any running applications.
2. Ensure that those applications will not start automatically.



Warning

When the relocated instance starts for the first time, remote block volumes, secondary VNICs, or any resource that relies on them, will not be attached. The absence of these resources can cause application issues.

3. If your instance has local NVMe storage (dense instances), you must back up this data:
 - a. [Create](#) and [attach](#) one or more remote block volumes to the instance.
 - b. Copy the data from the NVMe devices to the remote block volumes.
4. Unmount any remote block volumes or File Storage service (NFS) mounts.
5. Back up all remote block volumes. See [Overview of Block Volume Backups](#) for more information.
6. Create a backup of the root volume.



Important

Do not generalize or specialize Windows instances.

7. Terminate the instance:

Using the Console

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.

- b. In the list of instances, find the instance you want to terminate.
- c. Click the highlighted name of the instance to display the instance details.
- d. Click **Terminate**, and then respond to the confirmation prompt. Ensure that `Permanently delete the attached Boot Volume` is unchecked to preserve the boot volume associated with the instance.
Terminated instances temporarily remain in the list of instances with the status **Terminated**.

Using the API

To terminate the instance, use the [TerminateInstance](#) operation and pass the `preserveBootVolume` parameter set to `true` in the request.

Using the CLI

To terminate the instance, use the [terminate](#) operation and set the `preserve-boot-volume` option to `true`.

8. Create a new instance using the boot volume from the terminated instance.
9. In the launch instance flow, specify the private IP address that was attached to the primary VNIC. If the public IP address was assigned from a reserved IP address pool, be sure to assign the same IP address.
10. When the instance state changes to `RUNNING`, **Stop** the instance.
11. Recreate any secondary VNICs and secondary IP addresses.
12. Attach any remote block volumes.



Note

This step includes any volumes used to back up local NVMe devices. Copy the data onto the NVMe storage on the new instance, and then detach the volumes.

13. Start the instance.
14. Start and test any applications on the instance.
15. Configure the applications to start automatically, as required.
16. Recreate the required tags.
17. (Optional) After you confirm that the instance and applications are healthy, you can delete the volume backups.

Migrating an Instance from a Local to Remote Boot Volume

Oracle Cloud Infrastructure introduced [remote boot volumes](#) November 15, 2017. Remote boot volumes include significant improvements over local boot volumes including lower boot times, encryption of data at rest and in-transit, zero downtime snapshots, improved availability, and more. Instances launched before November 15, 2017 use local boot volumes. Local boot volumes are being deprecated, we recommend that you migrate any instances currently using a local boot volume to a remote boot volume. This topic describes that process.

Limitations and Warnings for Manual Migration

Be aware of the following limitations and warnings when performing a manual migration:

- Any public IP addresses assigned to your instance from a [reserved public pool](#) are retained. Any that were not assigned from a reserved public IP address pool will

change. Private IP addresses do not change.

- MAC addresses, CPUIDs, and other unique hardware identifiers do change during the move. If any applications running on the instance use these identifiers for licensing or other purposes, be sure to take note of this information before moving the instance to help you manage the change.

Prerequisites for Manual Migration

1. Before moving the instance, document all critical details:
 - The instance's region, availability domain, and fault domain.
 - The instance's display name.
 - All private IP addresses, names, and subnets. Note that the instance can have multiple VNICs, and each VNIC can have multiple secondary IP addresses.
 - All private DNS names. The instance can have multiple VNICs, and each VNIC can have multiple secondary IP addresses. Each private IP address can have a DNS name.
 - Any [public IP addresses](#) assigned from a reserved public pool. Note that the instance can have multiple VNICs, and each VNIC can have multiple secondary private IP addresses. Each VNIC and secondary private IP address can have an attached public IP address.
 - Any remote block volumes attached to the instance.
 - Any tags on the instance or attached resources.
2. Prepare the instance for manual migration:
 - Ensure that any remote block volumes defined in `/etc/fstab` use the [recommended options](#).
 - Ensure that any File Storage service (NFS) mounts use the `nofail` option.
 - If you have statically defined any network interfaces belonging to secondary VNICs using their MAC addresses, such as those defined in `/etc/sysconfig/network-scripts/ifcfg*`, those interfaces will not start due to the change in the MAC address. Remove the static mapping.

- If you use the [Oracle-provided script](#) to configure secondary VNICs, ensure it runs automatically at startup.

Migrating an Instance Manually

After you complete the prerequisites:

1. Stop any running applications.
2. Ensure that those applications will not start automatically.



Warning

When the relocated instance starts for the first time, remote block volumes, secondary VNICs, or any resource that relies on them, will not be attached. The absence of these resources can cause application issues.

3. If your instance has local NVMe storage (dense instances), you must back up this data:
 - a. [Create](#) and [attach](#) one or more remote block volumes to the instance.
 - b. Copy the data from the NVMe devices to the remote block volumes.
4. Unmount any remote block volumes or File Storage service (NFS) mounts.
5. Back up all remote block volumes. See [Overview of Block Volume Backups](#) for more information.
6. Create a custom image of the instance using the steps described in the [Using the Console](#) or [Using the API](#) sections of [Managing Custom Images](#).



Important

Do not generalize or specialize Windows instances.

7. Terminate the instance:

Using the Console

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
- b. In the list of instances, find the instance you want to terminate.
- c. Click the highlighted name of the instance to display the instance details.
- d. Click **Terminate**, and then respond to the confirmation prompt. Ensure that the **Permanently delete the attached Boot Volume** check box is cleared to preserve the boot volume associated with the instance.
Terminated instances temporarily remain in the list of instances with the status **Terminated**.

Using the API

To terminate the instance, use the [TerminateInstance](#) operation and pass the `preserveBootVolume` parameter set to `true` in the request.

Using the CLI

To terminate the instance, use the [terminate](#) operation and set the `preserve-boot-volume` option to `true`.

8. Create a new instance using the custom image from the terminated instance, see the steps described in the [Using the Console](#) or the [Using the API](#) sections of [Managing Custom Images](#).
9. In the launch instance flow, specify the private IP address that was attached to the primary VNIC. If the public IP address was assigned from a reserved IP address pool, be sure to assign the same IP address.
10. When the instance state changes to `RUNNING`, **Stop** the instance.

11. Recreate any secondary VNICs and secondary IP addresses.
12. Attach any remote block volumes, see [Attaching a Volume](#) for more information.



Note

This step includes any volumes used to back up local NVMe devices. Copy the data onto the NVMe storage on the new instance, and then detach the volumes.

13. Start the instance.
14. Start and test any applications on the instance.
15. Configure the applications to start automatically, as required.
16. Recreate the required tags.
17. (Optional) After you confirm that the instance and applications are healthy, you can delete the volume backups.

Moving Compute Resources to a Different Compartment

You can move Compute resources such as instances, instance pools, and custom images from one compartment to another.

When you move a Compute resource to a new compartment, associated resources such as boot volumes and VNICs are not moved.

After you move the resource to the new compartment, inherent policies apply immediately and affect access to the resource through the Console. For more information, see [Managing Compartments](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CHAPTER 8 Compute

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The following policies allow users to move Compute resources to a different compartment:

```
Allow group ComputeCompartmentMovers to manage instance-family in tenancy
Allow group ComputeCompartmentMovers to manage compute-management-family in tenancy
Allow group ComputeCompartmentMovers to manage auto-scaling-configurations in tenancy
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Using the Console

To move an instance to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the **List Scope** section, select a compartment.
3. Click the instance that you're interested in.
4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.
To track the progress of the operation, you can monitor the associated work request. For more information, see [Using the Console to View Work Requests](#).
7. If there are alarms monitoring the instance, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.
8. Optionally, move the resources that are attached to the instance to the new compartment.

To move an instance configuration to a different compartment



Note

Most of the properties for an existing instance configuration, including the compartment, cannot be modified after you create the instance configuration. Although you can move an instance configuration to a different compartment, you will not be able to use the instance configuration to manage instance pools in the new compartment. If you want to update an instance configuration to point to a different compartment, you should instead create a new instance configuration in the target compartment. For steps, see [Creating an Instance Configuration](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Configurations**.
2. In the **List Scope** section, select a compartment.
3. Click the instance configuration that you're interested in.
4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

To move an instance pool to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. In the **List Scope** section, select a compartment.
3. Click the instance pool that you're interested in.

4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.
7. Optionally, update the instance pool with an instance configuration that points to the new compartment. Do the following:
 - a. Create a new instance configuration in the new compartment. You can do this using the Console or the API. For steps, see [Creating an Instance Configuration](#).
 - b. Update the instance pool with the new instance configuration. You can do this using the API. For steps, see [Updating an Instance Pool](#).
8. Optionally, move the instances and other resources that are associated with the instance pool to the new compartment.

To move an autoscaling configuration to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Autoscaling Configurations**.
2. In the **List Scope** section, select a compartment.
3. Click the autoscaling configuration that you're interested in.
4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

To move a custom image to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**.
2. In the **List Scope** section, select a compartment.
3. Click the custom image that you're interested in.

4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

To move a cluster network to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Cluster Networks**.
2. In the **List Scope** section, select a compartment.
3. Click the cluster network that you're interested in.
4. Click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.
7. Optionally, move the instances and other resources that are associated with the cluster network to the new compartment.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to move Compute resources to different compartments:

- [ChangeInstanceCompartment](#)
- [ChangeInstanceConfigurationCompartment](#)
- [ChangeInstancePoolCompartment](#)
- [ChangeAutoScalingConfigurationCompartment](#)
- [ChangeImageCompartment](#)
- [ChangeClusterNetworkCompartment](#)

Stopping and Starting an Instance

You can stop and start an instance as needed to update software or resolve error conditions.

For steps to manage the lifecycle state of instances in an instance pool, see [Stopping and Starting the Instances in an Instance Pool](#).

Stopping or Restarting an Instance From Within the Instance

In addition to using the API and Console, you can stop and restart instances using the commands available in the operating system when you are logged in to the instance. Stopping an instance using the instance's OS does not stop billing for that instance. If you stop an instance this way, be sure to also stop it from the Console or API.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to stop or start an existing instance. If the specified group doesn't need to launch instances or

attach volumes, you could simplify that policy to include only `manage instance-family`, and remove the statements involving `volume-family` and `virtual-network-family`.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Rebooting Your Virtual Machine (VM) Instance During Planned Maintenance

When an underlying Oracle Cloud Infrastructure component needs to undergo maintenance, you are notified before the impact to your VM instances. If your VM instances are scheduled for a maintenance reboot, you can proactively reboot, or stop and start your instances using the Console, API, or CLI at any time before the scheduled reboot. This let you control how and when your applications experience downtime. Customer-managed VM maintenance is supported on standard and GPU instance shapes, including Oracle-provided platform images and custom images imported from outside of Oracle Cloud Infrastructure.

To identify the VM instances that you can proactively reboot, using the Console, check the **Maintenance Reboot** field for the instance. If the instance has a maintenance reboot scheduled and can be proactively rebooted, this field displays the date and start time for the reboot. To check this using the API, use the `timeMaintenanceRebootDue` field for the [Instance](#). For VM instances with a boot volume, additional iSCSI block volume attachments, and a single VNIC, you can proceed to reboot, or stop and start the instance. If you have non-iSCSI (paravirtualized or emulated) block volume attachments or secondary VNICs, you need to first detach these resources before you reboot your instance.

When you reboot, or stop and start the instance, it is migrated to a different physical VM host. Once the **Maintenance Reboot** field is blank, the instance is no longer impacted by the maintenance event. If you choose not to reboot before the scheduled time, then Oracle Cloud Infrastructure will reboot and migrate your instances within a 24-hour period after the scheduled time.

To make it easier to locate and perform these actions on your VM instances, you can use Search with a predefined query to find all instances that have a maintenance reboot scheduled.

To find all the instances scheduled for a maintenance reboot

1. In the Console, append "/a/query" to the end of your base Console URL. For example, <https://console.us-ashburn-1.oraclecloud.com/a/query>.
2. Click **Select Sample Query**, and then click **Query for all instances which have an upcoming scheduled maintenance**.

Resource Billing for Stopped Instances

For both VM and bare metal instances, billing depends on the [shape](#) that you use to create the instance:

- **Standard shapes:** Stopping an instance pauses billing. However, stopped instances continue to count toward your service limits.
- **Dense I/O shapes:** Billing continues for stopped instances because of the attached NVMe storage, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance](#).
- **GPU shapes:** Billing continues for stopped instances, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance](#).
- **HPC shapes:** Billing continues for stopped instances because of the attached NVMe storage, and related resources continue to count toward your service limits. To halt billing and remove related resources from your service limits, you must [terminate the instance](#).

Stopping an instance using the instance's OS does not stop billing for that instance. If you stop an instance this way, be sure to also stop it from the Console or API.

For more information about how instances running Microsoft Windows Server are billed when they are stopped, see [How am I charged for Windows Server on Oracle Cloud Infrastructure?](#).

Hardware Reclamation for Stopped Bare Metal Instances

When a bare metal instance remains in the stopped state for longer than 48 hours, the instance is taken offline and the physical hardware is reclaimed. The next time that you restart the instance, it starts on different physical hardware. There are no changes to the block volumes, boot volumes, and instance metadata, including the ephemeral and public IP addresses.

However, the following properties do change when a bare metal instance restarts on different physical hardware: the MAC addresses and the host serial number. You might also notice changes in the BIOS firmware version, BIOS settings, and CPU microcode. If you want to keep the same physical hardware, do not stop the instance using the Console or the API, SDKs, or CLI. Instead, shut down the instance [using the instance's OS](#). When you want to restart the instance, use the Console or the API, SDKs, or CLI.

This behavior applies to Linux instances that use the following [shapes](#):

- BM.Standard1.36
- BM.Standard.B1.44
- BM.Standard2.52
- BM.Standard.E2.64

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance that you want to stop or start.
3. Click one of the following actions:
 - **Start:** Restarts a stopped instance. After the instance is restarted, the **Stop** action is enabled.
 - **Stop:** Shuts down the instance. After the instance is powered off, the **Start** action is enabled.
 - **Reboot:** Shuts down the instance, and then restarts it.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [InstanceAction](#) operation to restart an instance.

The following actions are only available using the API:

- SOFTSTOP
- SOFTRESET

Terminating an Instance

You can permanently terminate (delete) instances that you no longer need. Any attached VNICs and volumes are automatically detached when the instance terminates. Eventually, the instance's public and private IP addresses are released and become available for other instances. By default, the instance's boot volume is deleted when you terminate the instance, however you can preserve the boot volume associated with the instance, so that you can attach it to a different instance as a data volume, or use it to launch a new instance.



Warning

If your instance has NVMe storage, terminating it securely erases the NVMe drives and the data that was on those drives becomes completely unrecoverable. Make sure you back up important data before terminating an instance. For more information, see [Protecting Data on NVMe Devices](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users launch Compute instances](#) includes the ability to terminate an instance (with or without an attached block volume).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For reference material about writing policies for instances, cloud networks, or other Core Services API resources, see [Details for the Core Services](#).

Using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. In the list of instances, find the instance you want to terminate.
3. Click the name of the instance to display the instance details.
4. Click **Terminate**, and then respond to the confirmation prompt.
If you want to preserve the boot volume associated with the instance, clear the **Permanently delete the attached Boot Volume** check box.
Terminated instances temporarily remain in the list of instances with the status **Terminated**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [TerminateInstance](#) operation to terminate an instance.

Enabling Monitoring for Compute Instances

This topic describes how to enable monitoring for Compute instances that use supported images.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Except for Compute instances, all resources that support the Monitoring service emit metrics by default. A Compute instance using a supported image emits metrics when it has the required instance configuration and OracleCloudAgent software.

Supported Images

Compute instances with [Oracle-provided images](#) support the configuration setting to enable monitoring. When this setting is enabled for an instance and the OracleCloudAgent software is installed on the instance, the instance emits metrics.



Note

Because legacy images require installation of the OracleCloudAgent software, we recommend that you select the latest image, which already has the software installed. If you need a legacy image, then select an image dated on or after November 16, 2018 (except Ubuntu, which must be dated after February 28, 2019).

Prerequisites

- Service gateways or public IP addresses: The Compute instances must have either [service gateways](#) or public IP addresses to send metrics to the Monitoring service.
- IAM policies: To create and update Compute instances, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information, see [Creating an Instance](#).
- SSH: To connect to a Compute instance, you must have an SSH key. For more information, see [Connecting to an Instance](#).

Process Overview: Enabling Monitoring for a New Compute Instance

Following is the process for configuring a new Compute instance to emit metrics.

Task 1: Create a monitoring-enabled instance

Steps depend on the date of the image used to create the instance.

Latest version of supported image



Note

Like the latest version, some recent versions of supported images also have the OracleCloudAgent software installed. Compare to the date listed in [Supported Images](#).

While defining properties for the new instance, set the property that enables the instance for monitoring. For instructions, see [Using the Console](#) or [Using the API](#).

Legacy version of supported image

A legacy version of a supported image is one provided before the date listed in [Supported Images](#).

While defining properties for the new instance, set the property that enables the instance for monitoring and do one of the following:

- In the Console: While defining properties for the new instance, provide a script to install the OracleCloudAgent software onto the instance during the instance creation process.
- Complete the instance creation process and then [manually install the OracleCloudAgent software](#).

Task 2: (Optional) Create a service gateway

If your instances do not have public IP addresses, set up a service gateway on the virtual cloud network (VCN). The service gateway allows the instances to send metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
- When setting up routing for the subnet that contains your instances, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Setting Up a Service Gateway in the Console](#).

Find Out if Monitoring Has Your Metrics

To find out if Monitoring has your metrics

This task involves either querying instance metrics or viewing instance configuration.

- Query metrics to determine if Monitoring is receiving metrics emitted by the instance. For instructions, see [Using the Console](#) or [Using the API](#).

Not seeing metrics for your instance?

If you don't see any metric charts, your Compute instance might not be emitting metrics. See the following possible causes and resolutions.

Possible cause	How to check	Resolution
Monitoring is disabled on the instance.	Review the instance configuration.	Enable monitoring.
No OracleCloudAgent software exists on the instance (occurs with older images).	Connect to the instance and look for the software.	Install the software.

Possible cause	How to check	Resolution
The instance cannot access the Monitoring service because its VCN does not use the Internet.	Review the instance's IP address. If it's not public, then a service gateway is needed.	Set up a service gateway.
The instance does not use a supported image .	Review Supported Images .	Create an instance with a supported image .
New instance in a new compartment: The IAM policies required for the instance to publish metrics to Monitoring are not yet initialized. More information: IAM policies are automatically created for new instances and are immediately available, unless the instances are in a new compartment. For a new instance in a new compartment, the policies can take up to 20 minutes to initialize, which delays the emission of metrics.	(not applicable)	Check back after 10 or 20 minutes.

- View the instance configuration to determine if monitoring is enabled. Monitoring-enabled instances may require installation of the OracleCloudAgent software before metrics are emitted. For instructions, see [Using the Console](#) or [Using the API](#).

Process Overview: Enabling Monitoring for an Existing Compute Instance

Following is the process for configuring an existing Compute instance to emit metrics.

Task 1: Enable monitoring

Update the instance configuration to enable monitoring. For instructions, see [Using the Console](#) or [Using the API](#).

Task 2: Install the OracleCloudAgent software

Choose the operating system corresponding to the instance.

Linux

This section covers CentOS, Linux, and Ubuntu images.

1. [Enable monitoring on the instance.](#)
2. Connect to the instance.
For step-by-step instructions, see [Connecting to an Instance](#).
3. Run the script corresponding to the image used by the instance.

CentOS 6.x, Oracle Linux 6.x

```
#!/bin/sh

cd ~
curl -O https://objectstorage.us-phoenix-1.oraclecloud.com/n/oci-i3/b/agents/o/pool%2F4a97146b-2c7f-4a6c-9d10-7bca4a6b27b3%2Foracle-cloud-agent-0.0.13-196.el6.x86_64.rpm -v
```

CentOS 7.x, Oracle Linux 7.x

```
#!/bin/sh
```

CHAPTER 8 Compute

```
cd ~
curl -O https://objectstorage.us-phoenix-1.oraclecloud.com/n/oci-i3/b/agents/o/pool%2Fc2d4e19d-46f2-4331-bd28-aa8b109eec67%2Foracle-cloud-agent-0.0.13-196.e17.x86_64.rpm -v
```

Ubuntu 16.04, 18.04



Note

Installation of the OracleCloudAgent software on instances using Ubuntu images requires [Snapcraft](#). To install Snapcraft, run the following commands, in sequence:

```
sudo apt update
```

```
sudo apt install snapd
```

```
sudo snap install oracle-cloud-agent --classic
```

4. Enter the relevant command to run the OracleCloudAgent software on the instance.

Example 1: CentOS or Oracle Linux image

```
sudo yum install -y <instance-agent-filename>
```

Example 2: Ubuntu image

(Same instructions as noted in the previous step: This command installs and runs the software.)

```
sudo snap install oracle-cloud-agent --classic
```

Metrics are now emitted by the instance.

Windows

1. [Enable monitoring on the instance.](#)
2. Download the OracleCloudAgent software from the following URL.
https://objectstorage.us-phoenix-1.oraclecloud.com/p/RDDCRrNT05WB19mI52QJTOXCq1whXFZHXCQP74b4ttg/n/image/en/b/windows_instance_agents/o/OracleCloudAgentSetup.msi
3. Connect to the instance.
For step-by-step instructions, see [Connecting to an Instance](#).
4. Copy the downloaded OracleCloudAgent software to the instance.
5. As a user with administrative privileges, enter the relevant command to run the OracleCloudAgent software on the instance.

```
msiexec /qb /i <instance-agent-filename>
```

Metrics are now emitted by the instance.

Task 3: (Optional) Create a service gateway

If your instances do not have public IP addresses, set up a service gateway on the virtual cloud network (VCN). The service gateway allows the instances to send metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
- When setting up routing for the subnet that contains your instances, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Setting Up a Service Gateway in the Console](#).

Updating the OracleCloudAgent Software

We recommend always running the latest version of the OracleCloudAgent software.

If the installed software can access the internet, then no action is needed. The software periodically checks for newer versions and automatically installs the latest version when different from the current version.

If the installed software does not have access to the internet, then a manual update is required. For example, a Compute instance that has no internet gateway or service gateway cannot access the internet. In this situation, the software cannot complete its checks for newer versions.

To find the version of the installed OracleCloudAgent software

Do one of the following:

- For Oracle Linux/CentOS, run the following command:

```
sudo yum info oracle-cloud-agent
```

- For Ubuntu, run the following command:

```
snap info oracle-cloud-agent
```

- For Windows, do one of the following.
 - In Control Panel, select **Programs and Features** and then find the version number provided for "Oracle Cloud Agent."
 - In PowerShell, run the following command:

```
Get-WmiObject -Class Win32_Product | Where-Object { $_.Name -eq "Oracle Cloud Agent" }
```

Example output:

```
IdentifyingNumber : {exampleuniqueidentifer}  
Name               : Oracle Cloud Agent  
Vendor            : Oracle Corporation
```

CHAPTER 8 Compute

```
Version      : 0.0.10.0
Caption     : Oracle Cloud Agent
```

To manually update the OracleCloudAgent software on a Compute instance

Do one of the following:

- Temporarily allow the instance to access the internet.
- Redo [the installation steps](#), using the latest version.

OracleCloudAgent software versions

Linux versions

Linux versions include CentOS 6.x, Oracle Linux 6.x, CentOS 7.x, Oracle Linux 7.x, Ubuntu 16.04, and Ubuntu 18.04.

Linux version	Date	Changes
0.0.13	November 4, 2019	Fix a bug in handling monitoring service internal server errors
0.0.11	September 13, 2019	Fix retry strategy for sending metrics and refresh security tokens
0.0.10	July 15, 2019	Fix for correct handling of forced termination of the oracle-cloud-agent-updater

Windows versions

Windows version	Date	Changes
0.0.11.0	November 5, 2019	Fixed: Fix a bug in handling monitoring service internal server errors
0.0.10.0	September 13, 2019	Fixed: <ul style="list-style-type: none">• Fix retry strategy for sending metrics and refresh security tokens• Fix for correct handling of forced termination of the oracle-cloud-agent-update
0.0.9.0	June 6, 2019	Fixed: Bug fix where agent restarts when telemetry or auth service returns 5xx

Using the Console

Use the Console to create a Compute instance with monitoring enabled and, for new instances using legacy images, to run the script that installs the required OracleCloudAgent software.

To create a monitoring-enabled instance

Steps depend on the date of the image used to create the instance.

Latest version of supported image



Note

Like the latest version, some recent versions of



supported images also have the OracleCloudAgent software installed. Compare to the date listed in [Supported Images](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Choose a **Compartment** you have permission to work in, and then click **Create Instance**.
2. In the **Create Instance** dialog box, select the latest version of a [supported image](#). For more information about launching instances, see [Creating an Instance](#).
3. Select **Enable monitoring**.
4. Update other configuration as needed and then click **Create Instance**.
The newly created monitoring-enabled instance emits metrics to the Monitoring service.

Legacy version of supported image

A legacy version of a supported image is one provided before the date listed in [Supported Images](#). Legacy images require you to install the OracleCloudAgent software.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**. Choose a **Compartment** you have permission to work in, and then click **Create Instance**.
2. In the **Create Instance** dialog box, select the legacy version of the [supported image](#). For more information about launching instances, see [Creating an Instance](#).
3. Select **Enable monitoring**.
4. Click **Show Advanced Options**.
5. Select **Paste cloud-init script**.
6. Copy and paste the script corresponding to the image used by the instance.

CentOS 6.x, Oracle Linux 6.x

```
#!/bin/sh
curl -O https://objectstorage.us-phoenix-1.oraclecloud.com/n/oci-i3/b/agents/o/pool%2F4a97146b-2c7f-4a6c-9d10-7bca4a6b27b3%2Foracle-cloud-agent-0.0.13-196.el6.x86_64.rpm
yum install -y ~/oracle-cloud-agent-0.0.13-196.el6.x86_64.rpm -v
```

CentOS 7.x, Oracle Linux 7.x

```
#!/bin/sh
curl -O https://objectstorage.us-phoenix-1.oraclecloud.com/n/oci-i3/b/agents/o/pool%2Fc2d4e19d-46f2-4331-bd28-aa8b109eec67%2Foracle-cloud-agent-0.0.13-196.el7.x86_64.rpm -v
yum install -y ~/oracle-cloud-agent-0.0.13-196.el7.x86_64.rpm -v
```

Windows Server 2012 R2, 2016



Note

For legacy versions of Windows images, make sure cloudbase-init is supported. See <https://docs.cloud.oracle.com/iaas/releasenotes/changes/595afbb7-de0c-4934-8074-5b1ed6be1b56/>.

```
#ps1_sysnative
cd \Users\opc\Desktop
Start-BitsTransfer -Source "https://objectstorage.us-phoenix-1.oraclecloud.com/p/RDDCrNT05WB19mI52QJTOXCq1whXFZHXCQP74b4ttg/n/imagegen/b/windows_instance_agents/o/OracleCloudAgentSetup.msi" -Destination "c:\Users\opc\Desktop\OracleCloudAgentSetup.msi"
msiexec /i "c:\Users\opc\Desktop\OracleCloudAgentSetup.msi" /quiet /L*v "c:\Users\opc\Desktop\OracleCloudAgentSetup.log"
```

Windows Server 2008 R2

Download the OracleCloudAgent software from the following URL and manually install it on the instance.

https://objectstorage.us-phoenix-1.oraclecloud.com/p/RDDCRrNTto5WB19mI52QJTOXCq1whXFZHXCQP74b4ttg/n/imagegallery/b/windows_instance_agents/o/OracleCloudAgentSetup.msi

Ubuntu 16.04, 18.04



Note

Installation of the OracleCloudAgent software on instances using Ubuntu images requires [Snapcraft](#). To install Snapcraft, run the following commands, in sequence:

```
sudo apt update
```

```
sudo apt install snapd
```

```
sudo snap install oracle-cloud-agent --classic
```

7. Update other configuration as needed and then click **Create Instance**. The newly created monitoring-enabled instance emits metrics to the Monitoring service.

To find out if monitoring is enabled or if Monitoring is receiving metrics

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance you're interested in.

3. On the instance detail page, under **Resources**, click **Metrics**.

If you see metric charts with data, then the Monitoring service is receiving metrics from this instance. For a list of metrics related to Compute instances, see [Compute Instance Metrics](#).

If you see a message that monitoring is not enabled, or that the OracleCloudAgent software needs to be installed, then complete those tasks.

To enable monitoring on an existing instance

1. Go to the Metrics page for the instance:
 - a. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
 - b. Click the instance you're interested in.
 - c. On the instance detail page, under **Resources**, click **Metrics**.
If monitoring is not enabled (and the instance uses a supported image), then a button is available for enabling monitoring.

2. Click **Enable monitoring**.

If you see metric charts with data, then the Monitoring service is receiving metrics from this instance. For a list of metrics related to Compute instances, see [Compute Instance Metrics](#).

If you see a message that the OracleCloudAgent software needs to be installed, then see [Task 2: Install the OracleCloudAgent software](#).

Using the API

Use the API to enable monitoring on a new or existing instance. After monitoring is enabled, you can [install the OracleCloudAgent software](#).

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To create a monitoring-enabled instance

Steps depend on the date of the image used to create the instance.

Latest version of supported image



Note

Like the latest version, some recent versions of supported images also have the OracleCloudAgent software installed. Compare to the date listed in [Supported Images](#).

Use the [LaunchInstance](#) API operation, specifying the latest version of a [supported image](#), and include the following parameter setting.

```
{
  "agentConfig":
  {
    "isMonitoringDisabled":false
  }
}
```

Legacy version of supported image

A legacy version of a supported image is one provided before the date listed in [Supported Images](#). Legacy images require you to install the OracleCloudAgent software.

1. Use the [LaunchInstance](#) API operation, specifying the legacy version of the [supported image](#), and include the following parameter setting.

```
{
  "agentConfig":
```

CHAPTER 8 Compute

```
{
  "isMonitoringDisabled":false
}
```

2. Install the OracleCloudAgent software onto the newly created instance. See [Task 2: Install the OracleCloudAgent software](#).

To find out if monitoring is enabled or if Monitoring is receiving metrics

To query metrics, use the [SummarizeMetricsData](#) API operation. Returned metrics indicate that the Monitoring service received metrics from the instance.

To determine instance agent configuration (`isMonitoringDisabled` value), use the [GetInstance](#) or [ListInstances](#) operation.

To enable monitoring on an existing instance

Use the [UpdateInstance](#) API operation and include the following parameter setting.

```
{
  "agentConfig":
  {
    "isMonitoringDisabled":false
  }
}
```

Compute Metrics

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

There are multiple Monitoring service metric namespaces related to Compute resources:

- **oci_computeagent:** Metrics related to the OracleCloudAgent software on Compute instances. See [Compute Instance Metrics](#).
- **oci_compute_infrastructure_health:** Metrics related to the health of bare metal hardware. See [Infrastructure Health Metrics](#).

Compute Instance Metrics

You can monitor the health, capacity, and performance of your Compute instances by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_computeagent` (the OracleCloudAgent software on Compute instances).

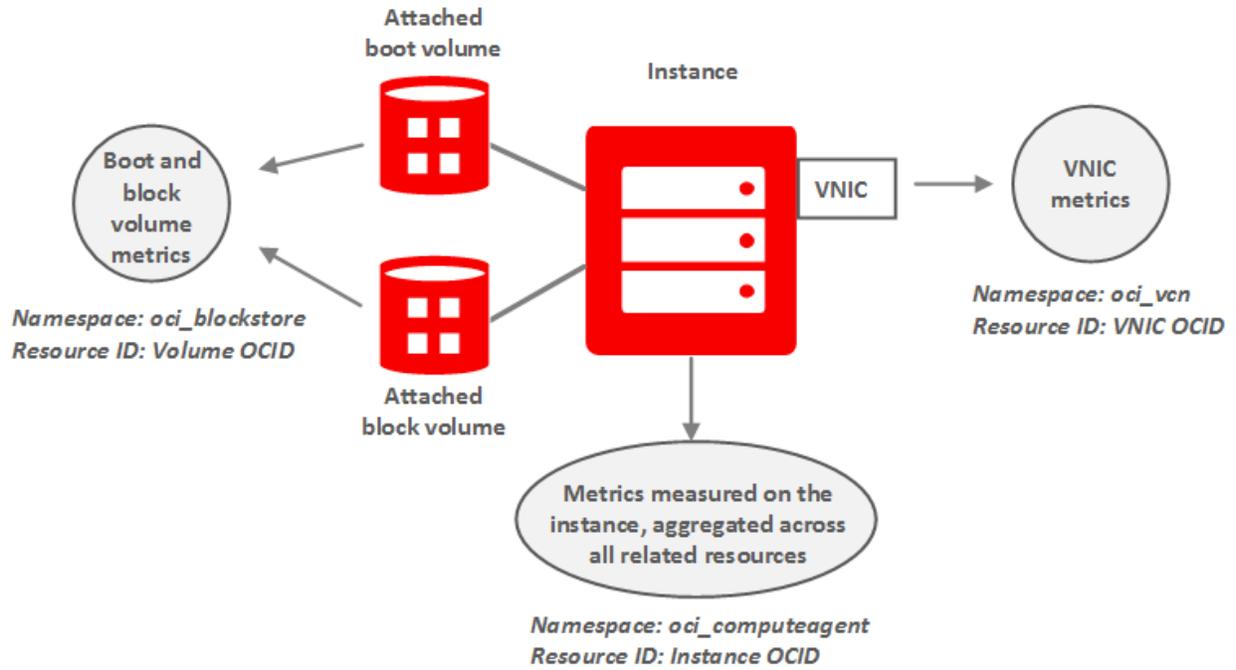
You can view these metrics for individual Compute instances, and for all the instances in an instance pool.

Resources: [Monitoring-enabled](#) Compute instances.

Overview of Metrics for an Instance and Related Resources

This section gives an overall picture of the different types of metrics available for an instance and its storage and network devices. See the following diagram and table for a summary.

CHAPTER 8 Compute



Metric Namespace	Resource ID	Where Measured	Available Metrics
oci_computeagent	Instance OCID	On the instance. The metrics in this namespace are aggregated across all the related resources on the instance. For example, <code>DiskBytesRead</code> is aggregated across all the instance's attached storage volumes, and <code>NetworkBytesIn</code> is aggregated across all the instance's attached VNICs.	See Available Metrics: oci_computeagent .
oci_blockstore	Boot or block volume OCID	By the Block Volume service. The metrics are for an individual volume (either boot volume or block volume).	See Block Volume Metrics .
oci_vcn	VNIC OCID	By the Networking service. The metrics are for an individual VNIC.	See VNIC Metrics .

Prerequisites

- **IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- **Metrics exist in Monitoring:** The resources that you want to monitor must emit metrics to the Monitoring service.

- Compute instances: To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).

Available Metrics: `oci_computeagent`

The Compute instance metrics help you measure activity level and throughput of Compute instances. The metrics listed in the following table are available for any [monitoring-enabled](#) Compute instance. You must enable monitoring on the instances to get these metrics.

The metrics in this namespace are aggregated across all the related resources on the instance. For example, `DiskBytesRead` is aggregated across all the instance's attached storage volumes, and `NetworkBytesIn` is aggregated across all the instance's attached VNICs.

You also can use the Monitoring service to create [custom queries](#).

Each metric includes the following dimensions:

AVAILABILITYDOMAIN

The availability domain where the instance resides.

FAULTDOMAIN

The fault domain where the instance resides.

IMAGEID

The OCID of the image for the instance.

INSTANCEPOOLID

The [instance pool](#) that the instance belongs to.

REGION

The region where the instance resides.

RESOURCEDISPLAYNAME

The friendly name of the instance.

CHAPTER 8 Compute

RESOURCEID

The OCID of the instance.

SHAPE

The shape of the instance.

CHAPTER 8 Compute

Metric	Metric Display Name	Unit	Description	Dimensions
CpuUtilization	CPU Utilization	percent	Activity level from CPU. Expressed as a percentage of total time. For instance pools, the value is averaged across all instances in the pool.	availabilityDomain faultDomain imageId instancePoolId region resourceDisplayName resourceId shape
DiskBytesRead ^{1, 3}	Disk Read Bytes	bytes	Read throughput. Expressed as bytes read per interval.	
DiskBytesWritten ^{1, 3}	Disk Write Bytes	bytes	Write throughput. Expressed as bytes written per interval.	

CHAPTER 8 Compute

Metric	Metric Display Name	Unit	Description	Dimensions
DiskIopsRead ^{1, 3} ₋	Disk Read I/O	operations	Activity level from I/O reads. Expressed as reads per interval.	
DiskIopsWritten ^{1, 3} ₋	Disk Write I/O	operations	Activity level from I/O writes. Expressed as writes per interval.	

CHAPTER 8 Compute

Metric	Metric Display Name	Unit	Description	Dimensions
MemoryUtilization ¹ <u>1</u>	Memory Utilization	percent	Space currently in use. Measured by pages. Expressed as a percentage of used pages. For instance pools, the value is averaged across all instances in the pool.	
NetworksBytesIn ^{1,2} <u>2</u>	Network Receive Bytes	bytes	Network receipt throughput. Expressed as bytes received.	

Metric	Metric Display Name	Unit	Description	Dimensions
NetworksBytesOut 1, 2	Network Transmit Bytes	bytes	Network transmission throughput. Expressed as bytes transmitted.	

¹This metric is a cumulative counter that shows monotonically increasing behavior for each session of the OracleCloudAgent software, resetting when the operating system is restarted.

²The Networking service provides additional metrics (in the `oci_vcn` metric namespace) for each [VNIC](#) on the instance. For more information, see [Network Metrics](#).

³The Block Volume service provides additional metrics (in the `oci_blockstore` metric namespace) for each volume attached to the instance. For more information, see [Block Volume Metrics](#).

Using the Console

To view default metric charts for a single Compute instance

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance that you're interested in.
3. Under **Resources**, click **Metrics**. If this is a bare metal instance, in the **Metric Namespace** list, select **oci_computeagent**.
The Metrics page displays a default set of charts for the current instance.

Not seeing any metric charts for the instance?

If you don't see any metric charts, your Compute instance might not be emitting metrics. See the following possible causes and resolutions.

Possible cause	How to check	Resolution
Monitoring is disabled on the instance.	Review the instance configuration.	Enable monitoring.
No OracleCloudAgent software exists on the instance (occurs with older images).	Connect to the instance and look for the software.	Install the software.
The instance cannot access the Monitoring service because its VCN does not use the Internet.	Review the instance's IP address. If it's not public, then a service gateway is needed.	Set up a service gateway.

Possible cause	How to check	Resolution
The instance does not use a supported image .	Review Supported Images .	Create an instance with a supported image .
New instance in a new compartment: The IAM policies required for the instance to publish metrics to Monitoring are not yet initialized. More information: IAM policies are automatically created for new instances and are immediately available, unless the instances are in a new compartment. For a new instance in a new compartment, the policies can take up to 20 minutes to initialize, which delays the emission of metrics.	(not applicable)	Check back after 10 or 20 minutes.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).
For information about notifications for alarms, see [Notifications Overview](#).

To view default metric charts for resources related to a Compute instance

- **For an attached block volume:** While viewing the instance's details, click **Attached Block Volumes**, and then click the volume you're interested in. Click **Metrics** to see the volume's charts. For more information about the emitted metrics, see [Block Volume Metrics](#).
- **For the attached boot volume:** While viewing the instance's details, click **Boot Volume**, and then click the volume you're interested in. Click **Metrics** to see the

volume's charts. For more information about the emitted metrics, see [Block Volume Metrics](#).

- **For an attached VNIC:** While viewing the instance's details, click **Attached VNICs**, and then click the VNIC you're interested in. Click **Metrics** to see the charts for the VNIC. For more information about the emitted metrics, see [Networking Metrics](#).

To view default metric charts for all Compute instances in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. Select a compartment.
3. For **Metric Namespace**, select **oci_computeagent**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

To view default metric charts for the instances in an instance pool

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instance Pools**.
2. Click the instance pool that you're interested in.
3. Under **Resources**, click **Metrics**.

The Metrics page displays a default set of charts for the current instance pool.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

Infrastructure Health Metrics

You can monitor the health, capacity, and performance of your Compute bare metal instances by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_compute_infrastructure_health`.

Resources: Bare metal Compute instances.

Overview of Metrics: `oci_compute_infrastructure_health`

The infrastructure health metrics help you monitor the health of the infrastructure for your bare metal instances, including hardware components such as the CPU, motherboard, DIMM, and NVMe drives. You can use the metrics to identify hardware issues, and [proactively take action](#) to minimize the impact on your applications.

REQUIRED IAM POLICY

To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user

authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

Available Metrics: `oci_compute_infrastructure_health`

The metric listed in the following table is automatically available for each bare metal instance that you create. You do not need to enable monitoring on the instance to get this metric.

You also can use the Monitoring service to create [custom queries](#).

The metric includes the following dimensions:

FAULTCLASS

The type of [hardware issue](#):

- `CPU`: A fault has been detected in one or more CPUs.
- `MEM-BOOT`: A fault in the memory subsystem was detected during instance launch or a recent reboot.
- `MEM-RUNTIME`: A fault in the memory subsystem was detected.
- `MGMT-CONTROLLER`: A fault in the instance management controller has been detected.
- `PCI`: A fault in the PCI subsystem has been detected.

RESOURCEDISPLAYNAME

The friendly name of the instance.

RESOURCEID

The OCID of the instance.

Metric	Metric Display Name	Unit	Description	Dimensions
<code>health_status</code>	Infrastructure Health Status	Issues	Number of issues. Any non-zero value indicates a health defect.	<code>faultClass</code> <code>resourceDisplayName</code> <code>resourceId</code>

Using the Console

To view infrastructure health metrics for a single Compute instance

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance that you're interested in.
3. Under **Resources**, click **Metrics**.
4. In the **Metric Namespace** list, select **oci_compute_infrastructure_health**.
The Metrics page displays a default set of charts for the current instance.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).
For information about notifications for alarms, see [Notifications Overview](#).

To view infrastructure health metrics for all Compute instances in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. Select a compartment.
3. For **Metric Namespace**, select **oci_compute_infrastructure_health**.
The **Service Metrics** page dynamically updates to show charts for each metric that is emitted by the selected metric namespace.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).
For information about notifications for alarms, see [Notifications Overview](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

Compute Performance

The content in the sections below apply to **Category 7** and **Section 3.a** of the [Oracle PaaS and IaaS Public Cloud Services Pillar documentation](#).

Oracle Cloud Infrastructure provides a variety of instance configurations in both bare metal and virtual machine (VM) shapes. Each shape varies on multiple dimensions including memory, CPU cores, network bandwidth, and the option of local NVMe SSD storage found in DenseIO shapes.

Oracle Cloud Infrastructure provides a service-level agreement (SLA) for NVMe performance. Measuring performance is complex and open to variability.

A NVMe drive also has non-uniform drive performance over the period of drive usage. A NVMe drive performs differently when tested brand new compared to when tested in a steady-state after some duration of usage. New drives have not incurred many write/erase cycles and the inline garbage collection has not had a significant impact on IOPS performance. To achieve the goal of reproducibility and reduced variability, our testing focuses on the steady-state duration of the NVMe drive's operation.

Testing Methodology



Warning

Before running any tests, protect your data by making a backup of your data and operating system environment to prevent any data loss. The tests described in this document will overwrite the data on the disk, and cause data corruption.

Summary: To capture the IOPS measure, first provision a shape such as the new BM.DenseIO2.52, and then use the [Gartner Cloud Harmony test suite](#) to run tests on an instance running the latest supported Oracle Linux image for each NVMe drive target.

Instructions:

1. Launch an instance based on the latest supported Oracle Linux image and select a shape such as the new BM.DenseIO2.52. For launch instructions, see [Creating an Instance](#).
2. Run the Gartner Cloud Harmony test suite tests on the instance for each NVMe drive target. The following is an example of a command that will work for all shapes and drives on the shape:

```
sudo ./run.sh `ls /dev/nvme[0-9]n1 | sed -e 's/\/\//--target=\/\/'`
--nopurge -noprecondition --fio_direct=1 --fio_size=10g --test=iops
--skip_blocksize=512b --skip_blocksize=8k --skip_blocksize=16k
--skip_blocksize=32k --skip_blocksize=64k --skip_blocksize=128k
--skip_blocksize=1m
```

The SLA for NVMe drive performance is measured against 4k block sizes with 100% random write workload on DenseIO shapes where the drive is in a steady-state of operation.

Performance Benchmarks

The following table lists the minimum IOPS for the specified shape to meet the SLA, given the testing methodology with 4k block sizes for 100% random write tests using the tests described in the previous section.

Shape	Minimum Supported IOPS
VM.DenseIO1.4	200k
VM.DenseIO1.8	250k
VM.DenseIO1.16	400k
BM.DenseIO1.36	2.5MM
VM.DenseIO2.8	250k

CHAPTER 8 Compute

Shape	Minimum Supported IOPS
VM.DenseIO2.16	400k
VM.DenseIO2.24	800k
BM.DenseIO2.52	3.0MM

While the NVMe drives are capable of higher IOPS, Oracle Cloud Infrastructure currently guarantee this minimum level of IOPS performance.

Frequently Asked Questions

Q: I suspect a slowdown in my NVMe drive performance. Is there a SLA violation?

A: We test hosts on a regular basis to ensure that our low-level software updates do not regress performance. In the event you have reproduced the testing methodology and your drive's performance does not meet the terms in the SLA please contact your Oracle sales team.

Q: Why does the testing methodology not represent a diversity of IO workloads such as random reads and writes to reflect real world IO?

A: We focused on reproducibility and we believe the tests provide a significant indicator of overall drive performance.

Q: Will Oracle Cloud Infrastructure change the tests in this document?

A: We will make changes to provide greater customer value through better guarantees and improved reproducibility.

Compute Health Monitoring for Bare Metal Instances

Compute health monitoring for bare metal instances is a feature that provides notifications about hardware issues with your bare metal instances. With the health monitoring feature, you can monitor the health of the hardware for your bare metal instances, including components such as the CPU, motherboard, DIMM, and NVMe drives. You can use the notifications to identify problems, letting you proactively redeploy your instances to improve availability.

Health monitoring notifications are emailed to the tenant administrator within one business day of the error occurring. This warning helps you to take action before any potential hardware failure and redeploy your instances to healthy hardware to minimize the impact on your applications.

You can also use the [infrastructure health metrics](#) available in the Monitoring service to create alarms and [notifications](#) based on hardware issues.

Error Messages and Troubleshooting

This section contains information about the most common health monitoring error messages and provides troubleshooting suggestions for you to try for your bare metal instance.

A fault has been detected in one or more CPUs

Fault class: CPU

Details: This error indicates that a processor or one or more cores have failed in your instance. Your instance might be inaccessible or there might be fewer available cores than expected.

Troubleshooting steps:

- If the instance is inaccessible, you must replace it using the steps in [Moving a Compute Instance to a New Host](#).
- If your instance is available, check for the expected number of cores:

CHAPTER 8 Compute

- On Linux-based systems, run the following command:

```
nproc --all
```

- On Windows-based systems, open Resource Monitor.

Compare the core count to the expected values documented in [Compute Shapes](#). If the number of cores is less than expected and this reduction impacts your application, we recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

A fault in the memory subsystem was detected during instance launch or a recent reboot

Fault class: MEM-BOOT

Details: This error indicates that one or more failed DIMMs were detected in your instance while the instance was being launched or rebooted. Any failed DIMMs have been disabled.

Troubleshooting steps: The total amount of memory in the instance will be lower than expected. If this impacts your application, we recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

To check for the amount of memory in the instance:

- On Linux-based systems, run the following command:

```
awk '$3=="kB"{$2=$2/1024**2;$3="GB";} 1' /proc/meminfo | column -t | grep MemTotal
```

- On Windows-based systems, open Resource Monitor.

The expected values are documented in [Compute Shapes](#).

A fault in the memory subsystem was detected

Fault class: MEM-RUNTIME

Details: This error indicates that one or more non-critical errors were detected on a DIMM in your instance. The instance might have unexpectedly rebooted in the last 72 hours.

Troubleshooting steps:

- If the instance has unexpectedly rebooted in the last 72 hours, one or more DIMMs might have been disabled. To check for the total amount of memory in the instance:
 - On Linux-based systems, run the following command:

```
awk '$3=="kB"{$2=$2/1024**2;$3="GB";} 1' /proc/meminfo | column -t | grep MemTotal
```

- On Windows-based systems, open Resource Monitor.

If the total memory in the instance is lower than expected, then one or more DIMMs have failed. If this impacts your application, we recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

- If the instance has not unexpectedly rebooted, it is at increased risk of doing so. During the next reboot, one or more DIMMs might be disabled. We recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

A fault in the instance management controller has been detected

Fault class: MGMT-CONTROLLER

Details: This error indicates that a device used to manage your instance might have failed. You might not be able to use the Console, CLI, SDKs, or APIs to stop, start, or reboot your instance. This functionality will still be available from within the instance using the standard operating system commands. You also might not be able to create a console connection to your instance. You will still be able to terminate your instance.

Troubleshooting steps: If this loss of control impacts your application, we recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

A fault in the PCI subsystem has been detected

Fault class: PCI

Details: This error indicates that one or more of the PCI devices in your instance have failed or are not operating at peak performance.

Troubleshooting steps:

- If you cannot [connect to the instance](#) over the network, the NIC might have failed. Use the Console or CLI to stop the instance and then start the instance. For steps, see [Stopping and Starting an Instance](#).

If you're still unable to connect to the instance over the network, you might be able to connect to it using a console connection. Follow the steps in [Connecting to the Serial Console](#) or [Connecting to the VNC Console](#) to establish a console connection and then reboot the instance. If the instance remains inaccessible, you must replace it using the steps in [Moving a Compute Instance to a New Host](#).

- An NVMe device may have failed.

On Linux-based systems, run the command `sudo lsblk` to get a list of the attached NVMe devices.

On Windows-based systems, open Disk Manager. Check the count of NVMe devices against the expected number of devices in [Compute Shapes](#).

If you determine that an NVMe device is missing from the list of devices for your instance, we recommend that you replace the instance using the steps in [Moving a Compute Instance to a New Host](#).

Microsoft Licensing on Oracle Cloud Infrastructure

This topic provides information about the licensing requirements to use Microsoft products on Oracle Cloud Infrastructure.

Moving Microsoft Licenses to Oracle Cloud Infrastructure: Microsoft License Mobility

Microsoft Volume Licensing customers can move eligible Microsoft server application licenses purchased under a Volume Licensing agreement to Oracle Cloud Infrastructure. To do this, you must enroll in the License Mobility through Microsoft Software Assurance benefit. This

benefit is included with an active Software Assurance contract. You don't need to purchase additional Microsoft software licenses, and there are no associated mobility fees.

For more information about this Microsoft benefit, see [Microsoft License Mobility through Software Assurance](#).

Eligibility Requirements

To enroll in Microsoft License Mobility through Software Assurance, you must be a Microsoft Volume License customer with eligible server application products. The following are key requirements:

- Windows Server operating systems, desktop client operating systems, and desktop applications such as Microsoft Office are not eligible under License Mobility through Software Assurance.
- Active Software Assurance coverage is required on eligible licenses migrated to Oracle Cloud Infrastructure.
- All licenses used to run and access your licensed software, such as server licenses, processor licenses, Client Access Licenses (CALs), External Connector (EC) licenses, and server management licenses, require active Software Assurance coverage. Your rights to run licensed software and manage instances on Oracle Cloud Infrastructure expire with the expiration of the Software Assurance coverage on those licenses.
- Eligible Volume Licensing programs include the Microsoft Enterprise Agreement, Microsoft Enterprise Subscription Agreement, and Microsoft Open Value agreement, where Software Assurance is included, and other Volume Licensing programs where Software Assurance is an option, such as the Microsoft Open License agreement and the Microsoft Select Plus agreement.
- You may move Microsoft licenses from on-premises or another cloud services provider only after more than 90 days have passed since the last license move.
- Eligible Microsoft licenses on Oracle Cloud Infrastructure must be maintained for a minimum period of 90 days in a specific Oracle Cloud Infrastructure region. After the 90-day period, you may move the licensed software to a shared host in another Oracle Cloud Infrastructure region.

- Any Microsoft Server licenses permitted on Oracle Cloud Infrastructure must be eligible according to the latest [Microsoft Product Terms](#). It is your responsibility to verify that the licenses you bring to Oracle Cloud Infrastructure are eligible according to the latest Microsoft Product Terms.

Enrolling in License Mobility through Software Assurance

All customers using License Mobility through Software Assurance must complete a license verification process. Microsoft will verify that you have eligible licenses with active Software Assurance and send confirmation when the verification process is complete.

You can deploy your application server software before completing the verification process, but you must submit the license verification form within 10 days of deployment.

You are responsible for managing true ups and renewals as required under your Volume Licensing agreement.

You must submit a new form each time that you deploy additional licenses, when you renew your agreement, and when you deploy any previously unverified products.

To enroll in License Mobility through Software Assurance:

1. Verify that you are a Microsoft Volume Licensing customer with eligible application server licenses that are covered by active Software Assurance.
2. Download the license verification form:
 - a. Go to the [Microsoft Product Licensing search page](#).
 - b. In the **Document Type** area, select **License Verification**.
 - c. Filter the results by language, region, and business sector. Note that the verification form is not available in the WW (World Wide) region.
 - d. Download the **LicenseMobilityVerif** document.

3. Complete the license verification form. To specify Oracle as the Authorized Mobility Partner, provide the following information:
 - **Authorized Mobility Partner Name:** Oracle America, Inc.
 - **Authorized Mobility Partner Website URL:** <http://www.oracle.com/>
 - **Authorized Mobility Partner Email Address:** microsoftlm_us_grp@oracle.com

For instructions to complete the form, see the [Microsoft License Mobility Verification Guide \(PDF\)](#).

4. Submit the completed verification form to both Microsoft and Oracle:
 - **Microsoft:** Submit the form through your Microsoft reseller or directly to the email address in the form.
 - **Oracle:** Send the form to microsoftlm_us_grp@oracle.com.

Microsoft and Oracle verify that the product licenses for the workloads you deploy to Oracle Cloud Infrastructure are eligible according to the terms of your License Mobility through Software Assurance benefit. Microsoft will communicate your verification status to you and to Oracle as an Authorized Mobility Partner.

Using Microsoft Windows on Oracle Cloud Infrastructure: FAQ

Oracle Cloud Infrastructure is licensed to provide Microsoft software offerings, including being a Microsoft Authorized License Mobility Partner.

For the latest Microsoft licensing requirements, refer to the [Microsoft Product Terms](#).

If you can't find the answer to your question here, or you need more assistance running Microsoft products on Oracle Cloud Infrastructure, contact [Oracle Support](#).

General Questions

What OS editions of Microsoft Windows Server are supported?

Oracle-provided images

These Windows versions are available for [Oracle-provided images](#):

- Windows Server 2008 R2* Enterprise
- Windows Server 2012 R2 Standard, Datacenter
- Windows Server 2016 Standard, Datacenter

Bring Your Own Image (BYOI)

These Windows versions support custom image import:

- Windows Server 2008 R2* Standard, Enterprise, Datacenter
- Windows Server 2012 Standard, Datacenter
- Windows Server 2012 R2 Standard, Datacenter
- Windows Server 2016 Standard, Datacenter
- Windows Server 2019 Standard, Datacenter

* Windows Server 2008 R2 reaches [end of support](#) on January 14, 2020.

If you don't need to migrate your Windows OS licenses, you can use the [Bring Your Own Image](#) process to migrate your Windows image to Oracle Cloud Infrastructure.

Is Windows Server 2019 available as an Oracle-provided image?

No, Windows Server 2019 is currently not available as an [Oracle-provided image](#).

If you're interested in Windows Server 2019, contact [Oracle Support](#).

Is Windows Server 2019 available as a Bring Your Own Image (BYOI) image?

Yes, you can import your own Windows Server 2019 image. For source image requirements and steps to import an image, see [Importing Custom Windows Images](#).

What VM and bare metal options are available for Windows Server operating systems?

The following table shows support for Microsoft Windows Server operating systems on Oracle Cloud Infrastructure.

Use Case	Bare Metal Machines	Virtual Machines (VMs)	License
Use an Oracle-provided Windows Server operating system image for Windows Server 2016, Windows Server 2012 R2, or Windows Server 2008 R2.	Supported	Supported	Volume license issued by Oracle Cloud Infrastructure
Bring your own virtual machine image . You can import your own custom virtual machine Windows Server OS image.	Not supported	Supported	Volume license issued by Oracle Cloud Infrastructure
Bring your own Windows Server ISO image. You can import your own custom Windows Server ISO image. You must use iPXE boot.	Supported	Not supported	Customer-owned license
Bring your own hypervisor. You can use a Windows Server 2016 Datacenter hypervisor host provided by Oracle Cloud Infrastructure and import your own VM images.	Supported	Not supported	Volume license issued by Oracle Cloud Infrastructure

Does Oracle Cloud Infrastructure support Bring Your Own Image (BYOI) for Windows Server?

Yes, you are permitted to import your own generalized custom image of Windows Server.

When you create an instance with an imported image on a VM or a shared bare metal machine, Oracle Cloud Infrastructure licenses the instance. For more information about imported images, see [Creating Windows Custom Images](#).

If you want to use your own license, BYOI is supported only for bare metal machines on a dedicated host.

How am I charged for Windows Server on Oracle Cloud Infrastructure?

The cost of a Microsoft Windows Server license is an additional cost, on top of the underlying Compute instance price. You pay separately for the Compute instance and the Windows Server license. For more information about Microsoft Windows Server pricing, see [Compute Pricing](#).

Billing continues for the Windows Server license when the [instance is stopped](#). To halt billing for the Windows Server license, you must [terminate \(delete\) the instance](#).

How does Windows Server get updated with the latest patches?

You must [update your VCN's security list](#) to enable egress traffic for port 80 (HTTP) and port 443 (HTTPS) to install patches from Microsoft. Oracle Cloud Infrastructure enables automatic updates for Microsoft Windows Server and uses the default settings for applying Windows Server patches.

Can I take a snapshot image after customizing a running Windows Server instance?

Yes, there are several options available on both bare metal and virtual machines:

- [Create a custom image](#): Creates a custom image that you can use to launch other instances. Instances that you launch from your image include the customizations, configuration, and software installed when you created the image.
- [Clone a boot volume](#): Makes a copy of an existing boot volume without needing to go through the backup and restore process. A boot volume clone is a point-in-time direct disk-to-disk deep copy of the source boot volume, so all the data that is in the source boot volume when the clone is created is copied to the boot volume clone.
- [Back up a block volume](#): Makes a point-in-time backup of data on a block volume. You can restore a backup to a new volume either immediately after a backup or at a later time that you choose.
- [Back up a boot volume](#): Makes a backup of a boot volume. Boot volume backup capabilities are the same as block volume backup capabilities and are in-region only. Windows boot volume backups cannot be copied across regions.

Can I export a custom Windows Server image?

Yes, exporting custom Windows Server operating system images is supported. For steps, see [Image Import/Export](#).

Licensing - Windows Server

What is BYOL?

BYOL stands for "bring your own license." BYOL enables you use software licenses that you already own to deploy software on Oracle Cloud Infrastructure, without any additional licensing fees. This process uses the [License Mobility through Microsoft Software Assurance](#) benefit provided by Microsoft. You must have active Software Assurance with Microsoft to bring your licenses to Oracle Cloud Infrastructure.

What is Microsoft License Mobility?

License Mobility through Software Assurance is a Microsoft benefit that permits you to move your eligible Microsoft licenses to cloud services providers such as Oracle Cloud Infrastructure. Oracle is an Authorized Mobility Partner for License Mobility.

With License Mobility through Software Assurance, you can deploy eligible application servers on bare metal hosts or virtual shared hardware in Oracle Cloud Infrastructure. An example of an application eligible for License Mobility through Software Assurance is Microsoft SQL. Windows Server operating systems are not eligible.

You may move Microsoft licenses from on-premises or another Authorized Mobility Partner only after more than 90 days have passed since the last license move.

For more information about this Microsoft benefit, see [Microsoft License Mobility through Software Assurance](#). For steps to move your Microsoft licenses to Oracle Cloud Infrastructure, see [Moving Microsoft Licenses to Oracle Cloud Infrastructure: Microsoft License Mobility](#).

Is Oracle a Microsoft Authorized Mobility Partner?

Yes, Oracle is an Authorized Mobility Partner for the Microsoft License Mobility through Software Assurance benefit.

Can I bring my own license for Microsoft Windows Server to Oracle Cloud Infrastructure?

Yes. You can bring your own license (BYOL) for Microsoft Windows Server, subject to the [Microsoft Product Terms](#). You are responsible for managing your own licenses to maintain compliance with Microsoft licensing terms.

The following table shows the BYOL requirements for Microsoft licenses on Oracle Cloud Infrastructure.

CHAPTER 8 Compute

Microsoft License	Bare Metal Machines and Dedicated Virtual Machine Hosts	Virtual Machines (Multi-Tenant Shared Host)
Windows Server	Supported with restrictions. You can BYOL on a bare metal dedicated host. You must import your own custom ISO image and use the iPXE boot process. BYOL is not supported for Oracle-provided images .	Not supported. Shared hosts must use Oracle-provided images that include the Microsoft license.
SQL Server Subject to the Microsoft Product Terms	Supported. You must have License Mobility through Software Assurance.	Supported. You must have License Mobility through Software Assurance and use an Oracle-provided image .
MSDN	Supported. Non-production use only.	Not supported.
Microsoft Office	Not supported.	Not supported.

Microsoft License	Bare Metal Machines and Dedicated Virtual Machine Hosts	Virtual Machines (Multi-Tenant Shared Host)
Windows 7, Windows 8, and Windows 10	Not supported.	Not supported.
Other Microsoft applications	Supported. Subject to the Microsoft Product Terms.	Supported. You must have License Mobility through Software Assurance and use an Oracle-provided image .

Application licenses require [License Mobility through Software Assurance](#) when running on Oracle Cloud Infrastructure VM instances.

Questions about your licensing rights should be directed to Microsoft or your Microsoft reseller.

Can I use virtual machines and bring my own license for Microsoft Windows Server to Oracle Cloud Infrastructure?

You cannot migrate your Windows Server OS licenses when using Oracle Cloud Infrastructure virtual machines.

However, you can [bring your own hypervisor](#) (KVM) to run a Windows Server VM with your own Windows Server OS license.

The following restrictions apply:

- You can use VMs with their own license only if you use bring your own hypervisor on a dedicated bare metal host.

- BYOL of Microsoft Windows Server is not supported for VMs running on a shared host. Oracle Cloud Infrastructure-provided VMs offer Windows Server.
 - You can use a bare metal instance under bring your own hypervisor.
 - You must install and manage a hypervisor (KVM or Hyper-V) and launch your own VMs. This will ensure isolation, because all Oracle VMs are running on a dedicated bare metal server. The VMs can run Windows Server if they are licensed through MSDN (development use only).
- BYOL on a dedicated host is only permitted with License Mobility through Active Software Assurance. [Follow the license mobility process](#) to move your SQL Server license to Oracle Cloud Infrastructure.

Licensing - Other Microsoft Software

What other Microsoft applications can I bring to Oracle Cloud Infrastructure?

Any Microsoft Server licenses permitted on Oracle Cloud Infrastructure must be eligible according to the latest [Microsoft Product Terms](#). It is your responsibility to verify that the licenses you bring to Oracle Cloud Infrastructure are eligible according to the latest Microsoft Product Terms. All products that are currently eligible for License Mobility and covered by Software Assurance are eligible for BYOL.

Can I bring my own SQL Server license to Oracle Cloud Infrastructure?

Yes, you can bring your own SQL Server license using License Mobility through Active Software Assurance. The following restrictions apply:

- When you move your Microsoft SQL license using the license mobility process, the Microsoft Windows Server license is not included. Microsoft Windows Server licenses are not permitted to be moved under License Mobility. Windows Server operating systems must use the license issued by Oracle Cloud Infrastructure.

- Perpetual licenses can be moved from on-premises or other cloud providers only after more than 90 days have passed since the last license move.
- End-of-support versions are not supported on shared host virtual machines on Oracle Cloud Infrastructure.

[Follow the license mobility process](#) to move your SQL Server license to Oracle Cloud Infrastructure.

Can I use my MSDN license on Microsoft Windows Server on Oracle Cloud Infrastructure?

Yes, you can use your MSDN license on Oracle Cloud Infrastructure if you use the Oracle Cloud Infrastructure bare metal offering.

You can use virtual machines with your own MSDN license only if you [bring your own hypervisor](#) (KVM). You cannot use your MSDN license when using virtual machines on Oracle Cloud Infrastructure.

Can I buy a MSDN subscription from Oracle Cloud Infrastructure?

No, Oracle does not sell MSDN subscriptions. Contact Microsoft or your Microsoft reseller.

Can I use a MSDN license for a production environment?

No, MSDN subscription licenses are for development, testing, or demonstration purposes only.

How can I remote access to a Windows Server instance on Oracle Cloud Infrastructure?

Follow the steps to [connect to a Windows instance](#). Windows operating systems permit remote access for a maximum of two users using Remote Desktop Services (RDS) for Administration purposes.

RDS Client Access Licenses (CALs) are required for each user or device using Remote Desktop.

Does Oracle Cloud Infrastructure offer additional Remote Desktop Services licenses for applications running on Windows VMs?

No, Oracle Cloud Infrastructure does not offer Microsoft RDS (Remote Desktop Server) Subscriber Access Licenses (SALs). You can bring your own license (BYOL) and use your RDS Client Access Licenses (CALs) on Oracle Cloud Infrastructure bare metal or virtual machines only if you have active Software Assurance coverage and move those licenses using the [license mobility process](#).

Can I bring my own RDS CALs if I want more than two users to access my Windows Server instance?

Yes, you can use your Remote Desktop Services (RDS) Client Access Licenses (CALs) on Oracle Cloud Infrastructure if you use the Oracle Cloud Infrastructure bare metal offering. In addition, you can use virtual machines with their own MSDN license if you [bring your own hypervisor](#) (KVM).

You can use your RDS CAL licenses on Oracle Cloud Infrastructure virtual machines only if you have active Software Assurance coverage and move your CALs using the [license mobility process](#).

Other Windows Server Questions

Are there user data capabilities when launching Windows Server images?

Yes, [Oracle-provided Windows images](#) include cloudbase-init installed by default. You can use cloudbase-init to run PowerShell scripts, batch scripts, or other user data content on instance launch. Cloudbase-init is the equivalent of cloud-init on Linux-based images.

Can I use Windows Remote Management on Oracle Cloud Infrastructure?

Yes, Microsoft Windows Remote Management (WinRM) is enabled by default on [Oracle-provided Windows images](#). WinRM enables you to remotely manage the operating system.

What is Microsoft end of support?

Microsoft establishes the support lifecycle policy for its products. When a product reaches the end of its support lifecycle, Microsoft no longer provides security updates for the product. You should upgrade to the latest version to remain secure.

What happens when Windows Server 2008 R2 reaches end of support?

Windows Server 2008 R2 reaches the [end of its support lifecycle](#) on January 14, 2020. After this date, you can import your own Windows 2008 R2 image and run your existing instances, but are at a higher risk of security issues, incompatibility, or failures. Extended Security Updates may not be used on Oracle Cloud Infrastructure. Oracle does not provide any operating system support for end-of-support operating systems.

Oracle Cloud Infrastructure does not provide platform images after the end of support date. However, you can import your own image and launch it on a shared host VM.

There are no restrictions to running end-of-support operating systems on bare metal machines on a dedicated host. You may bring your own image (BYOI) of a Windows Server 2008 R2 image, but you must import a custom OS image and run the image on a dedicated host.

Updating the Linux iSCSI Service to Restart Automatically

Oracle Cloud Infrastructure supports iSCSI attached remote boot and block volumes to Compute instances. These iSCSI attached volumes are managed by the Linux iSCSI initiator service, `iscsid`. In scenarios where this service is stopped for any reason, such as the service crashes or a system administrator inadvertently stops the service, it's important that this service is automatically restarted immediately.

The following platform images distributed by Oracle Cloud Infrastructure are configured so that the `iscsid` service restarts automatically:

- Oracle Linux 7 images released February 26, 2019 and later. See the release notes for [Oracle-Linux-7.6-Gen2-GPU-2019.02.20-0](#) and [Oracle-Linux-7.6-2019.02.20-0](#).
- Oracle Linux 6 images released February 26, 2019 and later. See the release notes [Oracle-Linux-6.10-2019.02.22-0](#).
- CentOS 7 images released February 25, 2019 and later. See the release notes for [CentOS-7-2019.02.23-0](#).

Instances created from earlier versions of CentOS 7.x and Oracle Linux platform images, or any versions of Oracle Cloud Infrastructure CentOS 6.x and Ubuntu platform images do not have this configuration. You should update these existing instances and custom images created from these images so that the `iscsid` service restarts automatically. You should also check this configuration on your imported paravirtualized custom images and any instances launched from these images and update the configuration as needed.

This topic describes how to update the `iscsid` service on an instance so that it will restart automatically.



Note

Configuring an instance to automatically restart the `iscsid` service does not require a reboot and will increase the stability of your infrastructure.

Oracle Linux 7

Run the following command to update the `iscsid` service on your Oracle 7 Linux instances:

```
sudo yum update -y iscsi-initiator-utils
```

After running this command, the version of the `iscsid` service should be 6.2.0.874 or newer.

Run the following command to check the version:

CHAPTER 8 Compute

```
yum info iscsi-initiator-utils
```

This update does not require a system reboot and will not make any changes to your instances beyond configuring `iscsid` to restart automatically.

Oracle Linux 6

Run the following command to update the `iscsid` service on your Oracle 6 Linux instances:

```
sudo yum update -y iscsi-initiator-utils
```

After running this command, the version of the `iscsid` service should be 6.2.0.873 or newer.

Run the following command to check the version:

```
yum info iscsi-initiator-utils
```

This update does not require a system reboot and will not make any changes to your instances beyond configuring `iscsid` to restart automatically.

CentOS 7.x



Important

Do not directly edit the `systemd` **`iscsid.service`** file. You should instead create an override to ensure that the `restart` option isn't overwritten the next time the `iscsid` service is updated.

On your CentOS 7 instances run the following command to create an override file:

```
sudo systemctl edit iscsid.service
```

Paste and save the following into the file:

```
[Service]
Restart=always
```

Run the following commands to reload `systemd` and restart the `iscsid` service:

CHAPTER 8 Compute

```
sudo systemctl daemon-reload
sudo systemctl restart iscsid
```

CentOS 6.x

On your CentOS 6 instances run the following command to install the [monit](#) package:

```
sudo yum install monit
```

Create the **/etc/monit/conf.d/iscsid.conf** file and include the following commands:

```
check process iscid with pidfile /run/iscsid.pid
start program = "/etc/init.d/open-iscsi start" with timeout 60 seconds
stop program = "/etc/init.d/open-iscsi stop"
```

Run the following command to start the `monit` service:

```
/etc/init.d/monit start
```

Ubuntu 18



Important

Do not directly edit the `systemd iscsid.service` file, instead create an override to ensure that the `restart` option isn't overwritten the next time the `iscsid` service is updated.

On your Ubuntu 18 instances run the following command to create an override file:

```
sudo systemctl edit iscsid.service
```

Paste and save the following into the file:

```
[Service]
Restart=
Restart=always
```

Run the following commands to reload `systemd` and restart the `iscsid` service:

CHAPTER 8 Compute

```
sudo systemctl daemon-reload
sudo systemctl restart iscsid
```

Ubuntu 16



Important

Do not directly edit the `systemd iscsid.service` file, instead create an override to ensure that the `restart` option isn't overwritten the next time the `iscsid` service is updated.

On your Ubuntu 16 instances run the following command to create an override file:

```
sudo systemctl edit iscsid.service
```

Paste and save the following into the file:

```
[Service]
Restart=
Restart=always
```

Run the following commands to reload `systemd` and restart the `iscsid` service:

```
sudo systemctl daemon-reload
sudo systemctl restart iscsid
```

Ubuntu 14

On your Ubuntu 14 instances run the following command to install the [monit](#) package:

```
sudo apt-get install monit
```

Create the `/etc/monit/conf.d/iscsid.conf` file and include the following commands:

```
check process iscsid with pidfile /run/iscsid.pid
start program = "/etc/init.d/open-iscsi start" with timeout 60 seconds
stop program = "/etc/init.d/open-iscsi stop"
```

Run the following command to start the `monit` service:

```
/etc/init.d/monit start
```

Testing the `iscsid` Service Update

Perform these steps to verify that the `iscsid` service has been updated successfully, and that it restarts automatically.



Warning

Do not perform these steps on a production instance. If the `iscsid` service fails to restart the instance may become unresponsive.

1. Run the following command to confirm that the `iscsid` service is running:

```
ps -ef | grep iscsid
```

2. Run the following command to stop the `iscsid` service:

```
sudo pkill -9 iscsid
```

3. Wait 60 seconds and then run the following command to verify that the `iscsid` service has restarted:

```
ps -ef | grep iscsid
```

CHAPTER 9 Container Engine for Kubernetes

This chapter explains how to define and create Kubernetes clusters to enable the deployment, scaling, and management of containerized applications.

Overview of Container Engine for Kubernetes

Oracle Cloud Infrastructure Container Engine for Kubernetes is a fully-managed, scalable, and highly available service that you can use to deploy your containerized applications to the cloud. Use Container Engine for Kubernetes (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud-native applications. You specify the compute resources that your applications require, and Container Engine for Kubernetes provisions them on Oracle Cloud Infrastructure in an existing OCI tenancy.

Container Engine for Kubernetes uses Kubernetes - the open-source system for automating deployment, scaling, and management of containerized applications across clusters of hosts. Kubernetes groups the containers that make up an application into logical units (called pods) for easy management and discovery. Container Engine for Kubernetes uses versions of Kubernetes certified as conformant by the [Cloud Native Computing Foundation \(CNCF\)](#).

You can access Container Engine for Kubernetes to define and create Kubernetes clusters using the Console and the REST API. You can access the clusters you create using the Kubernetes command line (kubectl), the Kubernetes Dashboard, and the Kubernetes API.

Container Engine for Kubernetes is integrated with Oracle Cloud Infrastructure Identity and Access Management (IAM), which provides easy authentication with native Oracle Cloud Infrastructure identity functionality.

For an introductory tutorial, see [Creating a Cluster with Oracle Cloud Infrastructure Container Engine for Kubernetes](#).



Note

Container Engine for Kubernetes is not available in Oracle Cloud Infrastructure Government Cloud realms.

Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud

network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Note that to perform certain operations on clusters created by Container Engine for Kubernetes, you might require additional permissions granted via a Kubernetes RBAC role or clusterrole. See [About Access Control and Container Engine for Kubernetes](#).

Container Engine for Kubernetes Capabilities and Limits

In each region that is enabled for your tenancy, you can create three clusters (Monthly Universal Credits) or one cluster (Pay-as-You-Go or Promo) by default. Each cluster you create can have a maximum of 1000 nodes. See [Service Limits](#).

Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For more details about policies for Container Engine for Kubernetes, see:

- [Policy Configuration for Cluster Creation and Deployment](#)
- [Details for Container Engine for Kubernetes](#)

Preparing for Container Engine for Kubernetes

Before you can use Container Engine for Kubernetes to create a Kubernetes cluster:

- You must have access to an Oracle Cloud Infrastructure tenancy. The tenancy must be subscribed to one or more of the regions in which Container Engine for Kubernetes is available (see [Availability by Region Name and Region Code](#)).
- Your tenancy must have sufficient quota on different types of resource (see [Service Limits](#)). More specifically:
 - Compute instance quota: To create a Kubernetes cluster, at least one compute instance (node) must be available in the tenancy. However, you'll probably want more than this minimum. For example, to create a highly available cluster in a region with three availability domains (ADs), at least three compute instances must be available (one in each availability domain).
 - Block volume quota: If you intend to create Kubernetes persistent volumes, sufficient block volume quota must be available in each availability domain to meet the persistent volume claim. Persistent volume claims must request a minimum of 50 gigabytes. See [Creating a Persistent Volume Claim](#).
 - Load balancer quota: If you intend to create a load balancer to distribute traffic between the nodes running a service in a Kubernetes cluster, sufficient load balancer quota must be available in the region. See [Creating Load Balancers to Distribute Traffic Between Cluster Nodes](#).
- Within your tenancy, there must already be a compartment to contain the necessary network resources (such as a VCN, subnets, internet gateway, route table, security lists). If such a compartment does not exist already, you will have to create it. Note that the network resources can reside in the root compartment. However, if you expect multiple teams to create clusters, best practice is to create a separate compartment for each team.
- Within the compartment, network resources (such as a VCN, subnets, internet gateway, route table, security lists) must be appropriately configured in each region in which you want to create and deploy clusters. For example, to create a highly available cluster in a region with three availability domains, the VCN must include:

- For worker nodes: a regional subnet (recommended), or three AD-specific subnets (one in each of the availability domains).
- For load balancers: optionally (but usually) an additional regional subnet (recommended), or an additional two AD-specific subnets (each in a different availability domain).

Best practice is to use regional subnets to make failover across availability domains simpler to implement.

When creating a new cluster, you can have Container Engine for Kubernetes automatically create and configure new network resources for the new cluster, or you can specify existing network resources. If you specify existing network resources, you or somebody else must have already configured those resources appropriately. See [Network Resource Configuration for Cluster Creation and Deployment](#).

- Within the root compartment of your tenancy, a policy statement (`Allow service OKE to manage all-resources in tenancy`) must be defined to give Container Engine for Kubernetes access to resources in the tenancy. See [Create Required Policy for Container Engine for Kubernetes](#)
- To create and/or manage clusters, you must belong to one of the following:
 - The tenancy's Administrators group
 - A group to which a policy grants the appropriate Container Engine for Kubernetes permissions. If you are creating or modifying clusters using the Console, or want Container Engine for Kubernetes to automatically create and configure new network resources for a new cluster, policies must also grant the group the following permissions:
 - VCN_READ and VCN_CREATE
 - SUBNET_READ and SUBNET_CREATE
 - COMPARTMENT_INSPECT
 - INTERNET_GATEWAY_CREATE
 - NAT_GATEWAY_CREATE

- ROUTE_TABLE_UPDATE
- SECURITY_LIST_CREATE

If you want to create a service gateway to enable applications deployed on a cluster to pull images from Oracle Cloud Infrastructure Registry (or to use other Oracle Cloud Infrastructure resources) without exposing data to the public internet, a policy must also grant the group the SERVICE_GATEWAY_CREATE permission.

See [Create Required Policy for Groups](#).

- To perform operations on a cluster:
 - You must have installed and configured the Kubernetes command line tool kubectl (see the [kubectl documentation](#)).
 - You must have downloaded your own copy of the cluster's kubeconfig configuration file (see [Downloading a kubeconfig File to Enable Cluster Access](#)). Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded.
 - You must have appropriate permissions to access the cluster (see [About Access Control and Container Engine for Kubernetes](#)).

Availability by Region Name and Region Code

Container Engine for Kubernetes is available in the following regions. Note that you have to use the region code in some commands. In some cases, you might have to use shortened versions of availability domain names.

CHAPTER 9 Container Engine for Kubernetes

Region Name	Region Code	Shortened Availability Domain Names
US East (Ashburn)	iad	<ul style="list-style-type: none"> US-ASHBURN-AD-1 US-ASHBURN-AD-2 US-ASHBURN-AD-3
Germany Central (Frankfurt)	fra	<ul style="list-style-type: none"> EU-FRANKFURT-1-AD-1 EU-FRANKFURT-1-AD-2 EU-FRANKFURT-1-AD-3
UK South (London)	lhr	<ul style="list-style-type: none"> UK-LONDON-1-AD-1 UK-LONDON-1-AD-2 UK-LONDON-1-AD-3
India West (Mumbai)	bom	<ul style="list-style-type: none"> AP-MUMBAI-1-AD-1
US West (Phoenix)	phx	<ul style="list-style-type: none"> PHX-AD-1 PHX-AD-2 PHX-AD-3
Brazil East (Sao Paulo)	gru	<ul style="list-style-type: none"> SA-SAOPAULO-1-AD-1
South Korea Central (Seoul)	icn	<ul style="list-style-type: none"> AP-SEOUL-1-AD-1
Australia East (Sydney)	syd	<ul style="list-style-type: none"> AP-SYDNEY-1-AD-1
Japan East (Tokyo)	nrt	<ul style="list-style-type: none"> AP-TOKYO-1-AD-1
Canada Southeast (Toronto)	yyz	<ul style="list-style-type: none"> CA-TORONTO-1-AD-1
Switzerland North (Zurich)	zrh	<ul style="list-style-type: none"> EU-ZURICH-1-AD-1

Network Resource Configuration for Cluster Creation and Deployment

Before you can use Container Engine for Kubernetes to create and deploy clusters in the regions in a tenancy:

- The tenancy's root compartment must include a policy to allow Container Engine for Kubernetes to perform operations in the tenancy. See [Create Required Policy for Container Engine for Kubernetes](#).
- Within the tenancy, there must already be a compartment to contain the necessary network resources (such as a VCN, subnets, internet gateway, route table, security lists). If such a compartment does not exist already, you will have to create it. Note that the network resources can reside in the root compartment. However, if you expect multiple teams to create clusters, best practice is to create a separate compartment for each team.
- Within the compartment, network resources (such as a VCN, subnets, internet gateway, route table, security lists) must be appropriately configured in each region in which you want to create and deploy clusters. When creating a new cluster, you can have Container Engine for Kubernetes automatically create and configure new network resources for a new 'quick cluster'. Alternatively, you can explicitly specify the existing network resources to use for a 'custom cluster'. If you specify existing network resources, you or somebody else must have already configured those resources appropriately, as described in this topic.

This topic describes the necessary configuration for each network resource. To see details of a typical configuration, see [Example Network Resource Configurations](#).

For an introductory tutorial, see [Creating a Cluster with Oracle Cloud Infrastructure Container Engine for Kubernetes](#).

Root Compartment Configuration

You have to define a policy for the tenancy's root compartment to enable Container Engine for Kubernetes to perform operations on the tenancy. See [Create Required Policy for Container Engine for Kubernetes](#).

VCN Configuration

The VCN in which you want to create and deploy clusters must be configured as follows:

- The VCN must have a CIDR block defined that is large enough for the number of subnets you specify for the clusters you create. For example, to create a highly available cluster in a region with three availability domains will typically require two regional subnets (recommended) or five AD-specific subnets to support the necessary number of worker nodes and load balancers. However, you can create clusters with fewer subnets. A /16 CIDR block would be large enough for almost all use cases (10.0.0.0/16 for example). The CIDR block you specify for the VCN must not overlap with the CIDR block you specify for pods and for the Kubernetes services (see [CIDR Blocks and Container Engine for Kubernetes](#)).
- The VCN must have an appropriate number of subnets defined. For example, to create a highly available cluster in a region with three availability domains, the VCN must include:
 - For worker nodes: a regional subnet (recommended), or three AD-specific subnets (one in each of the availability domains).
 - For load balancers: optionally (but usually) an additional regional subnet (recommended), or an additional two AD-specific subnets (each in a different availability domain).

However, you can create clusters with fewer worker nodes, and fewer or no load balancers, and therefore require fewer subnets. Best practice is to use regional subnets to make failover across availability domains simpler to implement. See [Subnet Configuration](#).

- The VCN must have security lists defined for worker node subnets and load balancer subnets (if specified). See [Security List Configuration](#).

In addition:

- Oracle recommends DNS Resolution is selected for the VCN.
- If you expect applications deployed on a cluster to require worker nodes to initiate connections to the internet, the VCN must have an internet gateway (if the worker

nodes are in public subnets) or a NAT gateway (if the worker nodes are in private subnets). See [Internet Gateway Configuration](#) and [NAT Gateway Configuration](#).

- If you expect applications deployed on a cluster to pull images from Oracle Cloud Infrastructure Registry (or to use other Oracle Cloud Infrastructure resources) and you don't want to expose the data to the public internet, you can define a service gateway in the VCN. See [Service Gateway Configuration](#).
- If the VCN has a NAT gateway, an internet gateway, or a service gateway, it must have a route table with appropriate rules defined. See [Route Table Configuration](#).

See [VCNs and Subnets](#) and [Example Network Resource Configurations](#).

Internet Gateway Configuration

If you intend to deploy worker nodes in public subnets, and you expect deployed applications to require the worker nodes to initiate connections to the internet, the VCN must have an internet gateway. The internet gateway must be specified as the target for the destination CIDR block 0.0.0.0/0 in a route rule in a worker node route table.

See [VCNs and Subnets](#) and [Example Network Resource Configurations](#).

NAT Gateway Configuration

If you intend to deploy worker nodes in private subnets, and you expect deployed applications to require the worker nodes to initiate connections to the internet, the VCN must have a NAT gateway. The NAT gateway must be specified as the target for the destination CIDR block 0.0.0.0/0 in a route rule in a worker node route table.

See [NAT Gateway](#) and [Example Network Resource Configurations](#).

Service Gateway Configuration

If you expect applications deployed on a cluster to pull images from Oracle Cloud Infrastructure Registry (or to use other Oracle Cloud Infrastructure resources) and you want to protect data from the public internet, you can set up a service gateway. Setting up a service gateway enables worker nodes to access other resources in the same region without exposing data to the public internet.

When setting up the service gateway, create it in the same VCN and compartment as the worker nodes, and select the **All <region> Services in Oracle Services Network** option.

Having created the service gateway, it must be specified as the target for **All <region> Services in Oracle Services Network** in a route rule in the worker node route table.

Note that if you expect deployed applications to require access to public endpoints or services not supported by a service gateway (for example, to download updates or patches), configure additional network resources (such as a NAT gateway) to access the internet.

See [Access to Oracle Services: Service Gateway](#) and [Example Network Resource Configurations](#).

Route Table Configuration

If you intend to deploy worker nodes in public subnets, and you expect deployed applications to require the worker nodes to initiate connections to the internet, create an internet gateway. Then create a worker node route table with a route rule that specifies the internet gateway as the target for the destination CIDR block 0.0.0.0/0.

If you intend to deploy worker nodes in private subnets, and you expect deployed applications to require the worker nodes to initiate connections to the internet, create a NAT gateway. Then create a worker node route table with a route rule that specifies the NAT gateway as the target for the destination CIDR block 0.0.0.0/0.

If you expect applications deployed on a cluster to pull images from Oracle Cloud Infrastructure Registry (or to use other Oracle Cloud Infrastructure resources) and you don't want to expose the data to the public internet, create a service gateway. Then create a worker node route table with a route rule that specifies the service gateway as the target for **All <region> Services in Oracle Services Network**.

See [Internet Gateway](#), [NAT Gateway](#), [Access to Oracle Services: Service Gateway](#), and [Example Network Resource Configurations](#).

DHCP Options Configuration

The VCN in which you want to create and deploy clusters must have DHCP Options configured. The default value for **DNS Type** of Internet and VCN Resolver is acceptable.

See [DHCP Options](#) and [Example Network Resource Configurations](#).

Security List Configuration

The VCN in which you want to create and deploy clusters must have security lists defined for worker node subnets and load balancer subnets (if specified). The security lists for worker node subnets and load balancer subnets must be different. The security list for load balancer subnets must be unique and for their exclusive use.

Worker nodes are created with public or private IP addresses, according to whether you specify public or private subnets when defining the node pools in a cluster. Container Engine for Kubernetes must be able to access worker nodes.

See [Security Lists](#) and [Example Network Resource Configurations](#).

PUBLIC WORKER NODE SUBNET SECURITY LIST CONFIGURATION

When configuring a security list for public worker node subnets, the security list must have:

- stateless ingress and egress rules that allow all traffic between different worker node subnets
- stateless ingress and egress rules that allow all traffic between worker node subnets and load balancer subnets (if specified)
- ingress rules to allow Container Engine for Kubernetes to access worker nodes on port 22 from the following source CIDR blocks:
 - 130.35.0.0/16
 - 134.70.0.0/17
 - 138.1.0.0/16
 - 140.91.0.0/17
 - 147.154.0.0/16
 - 192.29.0.0/16

Optionally, you can include ingress rules for public worker node subnets to:

CHAPTER 9 Container Engine for Kubernetes

- explicitly allow SSH access to worker nodes on port 22 (see [Connecting to Worker Nodes in Public Subnets Using SSH](#))
- allow inbound traffic to the worker nodes on the default NodePort range of 30000 to 32767 (see the [Kubernetes documentation](#))

Optionally, you can include an egress rule that allows all outbound traffic to the internet.

PRIVATE WORKER NODE SUBNET SECURITY LIST CONFIGURATION

When configuring a security list for private worker node subnets, the security list must have:

- stateless ingress and egress rules that allow all traffic between the different worker node subnets
- stateless ingress and egress rules that allow all traffic between worker node subnets and load balancer subnets

Optionally, you can include ingress rules for private worker node subnets to:

- explicitly allow SSH access to worker nodes on port 22 from within the VCN CIDR block (see [Connecting to Worker Nodes in Private Subnets Using SSH](#))
- allow inbound traffic to the worker nodes on the default NodePort range of 30000 to 32767 from within the VCN CIDR block (see the [Kubernetes documentation](#))

Optionally, you can include an egress rule that allows all outbound traffic to the internet.

Subnet Configuration

The characteristics of the cluster you want to create, and the number of availability domains in the region in which you want to deploy the cluster, will determine the number of subnets to configure. For example, to create a highly available cluster in a region with three availability domains will require:

- For worker nodes: a regional subnet (recommended), or three AD-specific subnets (one in each of the availability domains).

- For load balancers: optionally (but usually) an additional regional subnet (recommended), or an additional two AD-specific subnets (each in a different availability domain).

Best practice is to use regional subnets to make failover across availability domains simpler to implement.

The VCN in which you want to create and deploy clusters must have at least one subnet defined in which to deploy worker nodes. Worker node subnets can be either public, or private for additional security (as specified by the **Subnet access** property). The number of worker node subnets to create depends on the region in which you are creating the cluster:

- If you are creating a cluster in a region with multiple availability domains, you can define a single regional subnet (recommended), or multiple AD-specific subnets (one in each of the availability domains).
- If you are creating a cluster in a region with a single availability domain, you can define a single regional subnet (recommended), or a single AD-specific subnet.

You have the option to define and use load balancers in clusters you create. If you want to define and use load balancers, the VCN in which you want to create and deploy clusters must have at least one subnet defined to host the load balancers. Load balancer subnets can be public or private (as specified by the **Subnet access** property). However, load balancers are optional, so you might not define load balancer subnets at all. The number of load balancer subnets to define depends on the region in which you are creating the cluster:

- If you are creating a cluster in a region with three availability domains, you can define:
 - zero or one load balancer regional subnet (recommended)
 - zero or two load balancer AD-specific subnets. If you define two load balancer AD-specific subnets, they must be in different availability domains.
- If you are creating a cluster in a region with a single availability domain, you can define zero or one load balancer subnet:
 - zero or one load balancer regional subnet (recommended)
 - zero or one load balancer AD-specific subnet.

In addition, all subnets must have the following properties set as shown:

- **Route Table:** The name of a route table, if one has been created, that has a route rule specifying an internet gateway (for public worker node subnets) or NAT gateway (for private worker node subnets) as the target for the destination CIDR block 0.0.0.0/0, and/or a route rule specifying a service gateway as the target for **All <region> Services in Oracle Services Network**.
- **DHCP options:** Default.

The CIDR blocks you specify for worker node and load balancer subnets must not overlap with CIDR blocks you specify for pods running in the cluster (see [CIDR Blocks and Container Engine for Kubernetes](#)).

Worker node subnets must have different security lists to load balancer subnets.

See [VCNs and Subnets](#) and [Example Network Resource Configurations](#).

Example Network Resource Configurations

Before you can use Container Engine for Kubernetes to create and deploy clusters in the regions in a tenancy:

- The tenancy's root compartment must include a policy to allow Container Engine for Kubernetes to perform operations in the tenancy. See [Create Required Policy for Container Engine for Kubernetes](#).
- Within the tenancy, there must already be a compartment to contain the necessary network resources (such as a VCN, subnets, internet gateway, route table, security lists). If such a compartment does not exist already, you will have to create it. Note that the network resources can reside in the root compartment. However, if you expect multiple teams to create clusters, best practice is to create a separate compartment for each team.
- Within the compartment, network resources (such as a VCN, subnets, internet gateway, route table, security lists) must be appropriately configured in each region in which you want to create and deploy clusters. When creating a new cluster, you can have Container Engine for Kubernetes automatically create and configure new network

resources for a new 'quick cluster'. Alternatively, you can explicitly specify the existing network resources to use for a 'custom cluster'. If you specify existing network resources, you or somebody else must have already configured those resources appropriately. See [Network Resource Configuration for Cluster Creation and Deployment](#).

This topic gives examples of how you might configure network resources for highly available 'custom cluster' creation and deployment in a region with three availability domains:

- for public clusters, where you want worker nodes hosted in public AD-specific subnets that can be accessed directly from the internet (see [Example 1: Example Network Resource Configuration for a Highly Available Public Cluster in a Region with Three Availability Domains, Using AD-Specific Subnets](#))
- for private clusters, where you want worker nodes hosted in private AD-specific subnets that can only be accessed from within the VCN (see [Example 2: Example Network Resource Configuration for a Highly Available Private Cluster in a Region with Three Availability Domains, Using AD-Specific Subnets](#))
- for public clusters, where you want worker nodes hosted in a public regional subnet that can be accessed directly from the internet (see [Example 3: Example Network Resource Configuration for a Highly Available Public Cluster in a Region with Three Availability Domains, Using a Regional Subnet](#))

Note that all the examples in this topic include a service gateway to enable worker nodes to access other Oracle Cloud Infrastructure resources in the same region (such as Oracle Cloud Infrastructure Registry) without exposing data to the public internet. However, you might be expecting applications deployed on the cluster to require access to public endpoints or services not supported by a service gateway. For example, to download updates or patches. If so, configure additional network resources (such as a NAT gateway) to access the internet.

For an introductory tutorial, see [Creating a Cluster with Oracle Cloud Infrastructure Container Engine for Kubernetes](#).

Example 1: Example Network Resource Configuration for a Highly Available Public Cluster in a Region with Three Availability Domains, Using AD-Specific Subnets

This example assumes you want worker nodes hosted in three public AD-specific subnets that can be accessed directly from the internet.

EXAMPLE NETWORK RESOURCE CONFIGURATION

Resource	Example
VCN	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: acme-dev-vcn • CIDR Block: 10.0.0.0/16 • DNS Resolution: Selected
Internet Gateway	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: gateway-0
Service Gateway	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: service-gateway-0 • Services: All <region> Services in Oracle Services Network

Resource	Example
Route Table	<p>Two route tables created manually, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: routetable-0, with a route rule defined as follows: <ul style="list-style-type: none"> ◦ Destination CIDR block: 0.0.0.0/0 ◦ Target Type: Internet Gateway ◦ Target Internet Gateway: gateway-0 • Name: routetable-1, with a route rule defined as follows: <ul style="list-style-type: none"> ◦ Destination: All <region> Services in Oracle Services Network ◦ Target Type: Service Gateway ◦ Target: service-gateway-0
DHCP Options	<p>Created automatically and defined as follows:</p> <ul style="list-style-type: none"> • DNS Type set to Internet and VCN Resolver

Resource	Example
Security Lists	<p>Two created (in addition to the default security list) manually, named, and defined as follows:</p> <ul style="list-style-type: none">• Security List Name: workers• Security List Name: loadbalancers <p>For details of the ingress rules and egress rules defined for the workers security list and the loadbalancers security list, see Example Security List Configurations for a Highly Available Public Cluster Using AD-Specific Subnets.</p>

Resource	Example
Subnets	<p>Three worker node AD-specific subnets created manually, named, and defined as follows:</p> <ul style="list-style-type: none">• Name: workers-1 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD1◦ CIDR Block: 10.0.10.0/24◦ Route Table: routetable-1◦ Subnet access: Public◦ DNS Resolution: Selected◦ DHCP Options: Default◦ Security List: workers• Name: workers-2 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD2◦ CIDR Block: 10.0.11.0/24◦ Route Table: routetable-1◦ Subnet access: Public◦ DNS Resolution: Selected◦ DHCP Options: Default◦ Security List: workers• Name: workers-3 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD3◦ CIDR Block: 10.0.12.0/24◦ Route Table: routetable-1◦ Subnet access: Public◦ DNS Resolution: Selected

Resource	Example
	<ul style="list-style-type: none"> ◦ DHCP Options: Default ◦ Security List: workers <p>Two load balancer AD-specific subnets created, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: loadbalancers-1 with the following properties: <ul style="list-style-type: none"> ◦ Availability Domain: AD1 ◦ CIDR Block: 10.0.20.0/24 ◦ Route Table: routetable-0 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: loadbalancers • Name: loadbalancers-2 with the following properties: <ul style="list-style-type: none"> ◦ Availability Domain: AD2 ◦ CIDR Block: 10.0.21.0/24 ◦ Route Table: routetable-0 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: loadbalancers

EXAMPLE SECURITY LIST CONFIGURATIONS FOR A HIGHLY AVAILABLE PUBLIC CLUSTER USING AD-SPECIFIC SUBNETS

In the example VCN, two security lists have been created (in addition to the default security list) to control access to and from the public worker node AD-specific subnets and the load

CHAPTER 9 Container Engine for Kubernetes

balancer AD-specific subnets. The two security lists are named 'workers' and 'loadbalancers' respectively.

The workers security list has the following ingress and egress rules for the public worker node AD-specific subnets:

Example Ingress Rules in a Security List for Public Worker Node AD-Specific Subnets:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateless	10.0.11.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
3	Stateless	10.0.12.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
4	Stateful	0.0.0.0/0	ICMP	n/a	n/a	3, 4	<p>Allows: ICMP traffic for: 3, 4 Destination Unreachable: Fragmentation Needed and Don't Fragment was Set</p> <p>Description: This rule enables worker nodes to receive Path MTU Discovery fragmentation messages.</p>
5	Stateful	130.35.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
6	Stateful	134.70.0.0/17	TCP	All	22	n/a	Allows: TCP traffic for ports: 22 SSH Remote Login Protocol Description: This rule enables Container Engine for Kubernetes to access worker nodes.
7	Stateful	138.1.0.0/16	TCP	All	22	n/a	Allows: TCP traffic for ports: 22 SSH Remote Login Protocol Description: This rule enables Container Engine for Kubernetes to access worker nodes.

CHAPTER 9 Container Engine for Kubernetes

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
8	Stateful	140.91.0.0/17	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>
9	Stateful	147.154.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>

CHAPTER 9 Container Engine for Kubernetes

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
10	Stateful	192.29.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>
11	Stateful	0.0.0.0/0	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This <i>optional</i> rule enables inbound SSH traffic from the internet on port 22 to access worker nodes.</p>

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
12	Stateful	0.0.0.0/0	TCP	All	30000 - 32767	n/a	Allows: TCP traffic for ports: 30000 - 32767 Description: This <i>optional</i> rule enables inbound traffic to the worker nodes on the default NodePort range of 30000-32767 (see the Kubernetes documentation).

Example Egress Rules in a Security List for Public Worker Node AD-Specific Subnets:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateless	10.0.11.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
3	Stateless	10.0.12.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
4	Stateful	All <region> Services in Oracle Services Network	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables outbound access to the regional service gateway.

CHAPTER 9 Container Engine for Kubernetes

The loadbalancers security list has the following ingress and egress rules for load balancer AD-specific subnets:

Example Ingress Rules in a Security List for Load Balancer AD-Specific Subnets:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for all ports: all Description: This rule enables incoming public traffic to service load balancers.

Example Egress Rules in a Security List for Load Balancer AD-Specific Subnets:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for ports: all Description: This rule enables responses from a web application through the service load balancers.

Example 2: Example Network Resource Configuration for a Highly Available Private Cluster in a Region with Three Availability Domains, Using AD-Specific Subnets

This example assumes you want worker nodes hosted in three private AD-specific subnets that can only be accessed from within the VCN.

EXAMPLE NETWORK RESOURCE CONFIGURATION

Resource	Example
VCN	Created manually, and defined as follows: <ul style="list-style-type: none">• Name: acme-dev-vcn• CIDR Block: 10.0.0.0/16• DNS Resolution: Selected
Internet Gateway	Created manually, and defined as follows: <ul style="list-style-type: none">• Name: gateway-0
NAT Gateway	Created manually, and defined as follows: <ul style="list-style-type: none">• Name: nat-gateway-0
Service Gateway	Created manually, and defined as follows: <ul style="list-style-type: none">• Name: service-gateway-0• Services: All <region> Services in Oracle Services Network

Resource	Example
Route Table	<p>Two route tables created manually, named, and defined as follows:</p> <ul style="list-style-type: none">• Name: routetable-0, with a route rule defined as follows:<ul style="list-style-type: none">◦ Destination CIDR block: 0.0.0.0/0◦ Target Type: Internet Gateway◦ Target Internet Gateway: gateway-0• Name: routetable-1, with two route rules defined as follows:<ul style="list-style-type: none">◦ Rule 1:<ul style="list-style-type: none">▪ Destination CIDR block: 0.0.0.0/0▪ Target Type: NAT Gateway▪ Target NAT Gateway: nat-gateway-0◦ Rule 2:<ul style="list-style-type: none">▪ Destination: All <region> Services in Oracle Services Network▪ Target Type: Service Gateway▪ Target: service-gateway-0
DHCP Options	<p>Created automatically and defined as follows:</p> <ul style="list-style-type: none">• DNS Type set to Internet and VCN Resolver

Resource	Example
Security Lists	<p>Two created (in addition to the default security list) manually, named, and defined as follows:</p> <ul style="list-style-type: none">• Security List Name: workers• Security List Name: loadbalancers <p>For details of the ingress rules and egress rules defined for these security lists, see Example Security List Configurations for a Highly Available Private Cluster Using AD-Specific Subnets.</p>

Resource	Example
Subnets	<p>Three worker node AD-specific subnets created manually, named, and defined as follows:</p> <ul style="list-style-type: none">• Name: workers-1 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD1◦ CIDR Block: 10.0.10.0/24◦ Route Table: routetable-1◦ Subnet access: Private◦ DNS Resolution: Selected◦ DHCP Options: Default◦ Security List: workers• Name: workers-2 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD2◦ CIDR Block: 10.0.11.0/24◦ Route Table: routetable-1◦ Subnet access: Private◦ DNS Resolution: Selected◦ DHCP Options: Default◦ Security List: workers• Name: workers-3 with the following properties:<ul style="list-style-type: none">◦ Availability Domain: AD3◦ CIDR Block: 10.0.12.0/24◦ Route Table: routetable-1◦ Subnet access: Private◦ DNS Resolution: Selected

Resource	Example
	<ul style="list-style-type: none"> ◦ DHCP Options: Default ◦ Security List: workers <p>Two load balancer AD-specific subnets created, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: loadbalancers-1 with the following properties: <ul style="list-style-type: none"> ◦ Availability Domain: AD1 ◦ CIDR Block: 10.0.20.0/24 ◦ Route Table: routetable-0 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: loadbalancers • Name: loadbalancers-2 with the following properties: <ul style="list-style-type: none"> ◦ Availability Domain: AD2 ◦ CIDR Block: 10.0.21.0/24 ◦ Route Table: routetable-0 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: loadbalancers

EXAMPLE SECURITY LIST CONFIGURATIONS FOR A HIGHLY AVAILABLE PRIVATE CLUSTER USING AD-SPECIFIC SUBNETS

In the example VCN, two security lists have been created (in addition to the default security list) to control access to and from the private worker node AD-specific subnets and the load

CHAPTER 9 Container Engine for Kubernetes

balancer AD-specific subnets. The two security lists are named 'workers' and 'loadbalancers' respectively.

The workers security list has the following ingress and egress rules for private worker node subnets:

Example Ingress Rules in a Security List for Private Worker Node AD-Specific Subnets:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateless	10.0.11.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
3	Stateless	10.0.12.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
4	Stateful	10.0.0.0/16	TCP	All	22	n/a	Allows: TCP traffic for ports: 22 SSH Remote Login Protocol Description: This <i>optional</i> rule enables inbound SSH traffic from the VCN on port 22 to access worker nodes.

Example Egress Rules in a Security List for Private Worker Node AD-Specific Subnets:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateless	10.0.11.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
3	Stateless	10.0.12.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
4	Stateful	All <region> Services in Oracle Services Network	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables outbound access to the regional service gateway.

CHAPTER 9 Container Engine for Kubernetes

The loadbalancers security list has the following ingress and egress rules for load balancer AD-specific subnets:

Example Ingress Rules in a Security List for a Load Balancer AD-Specific Subnet:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for all ports: all Description: This rule enables incoming public traffic to service load balancers.

Example Egress Rules in a Security List for a Load Balancer AD-Specific Subnet:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for ports: all Description: This rule enables responses from a web application through the service load balancers.

Example 3: Example Network Resource Configuration for a Highly Available Public Cluster in a Region with Three Availability Domains, Using a Regional Subnet

This example assumes you want worker nodes hosted in a public regional subnet that can be accessed directly from the internet.

EXAMPLE NETWORK RESOURCE CONFIGURATION

Resource	Example
VCN	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: acme-dev-vcn • CIDR Block: 10.0.0.0/16 • DNS Resolution: Selected
Internet Gateway	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: gateway-0
Service Gateway	Created manually, and defined as follows: <ul style="list-style-type: none"> • Name: service-gateway-0 • Services: All <region> Services in Oracle Services Network

Resource	Example
Route Table	<p>Two route tables created manually, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: routetable-0, with a route rule defined as follows: <ul style="list-style-type: none"> ◦ Destination CIDR block: 0.0.0.0/0 ◦ Target Type: Internet Gateway ◦ Target Internet Gateway: gateway-0 • Name: routetable-1, with a route rule defined as follows: <ul style="list-style-type: none"> ◦ Destination: All <region> Services in Oracle Services Network ◦ Target Type: Service Gateway ◦ Target: service-gateway-0
DHCP Options	<p>Created automatically and defined as follows:</p> <ul style="list-style-type: none"> • DNS Type set to Internet and VCN Resolver

Resource	Example
Security Lists	<p>Two created (in addition to the default security list) manually, named, and defined as follows:</p> <ul style="list-style-type: none"> • Security List Name: workers • Security List Name: loadbalancers <p>For details of the ingress rules and egress rules defined for the workers security list and the loadbalancers security list, see Example Security List Configurations for a Highly Available Public Cluster Using AD-Specific Subnets.</p>
Subnets	<p>One worker node regional subnet created manually, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: workers-rs with the following properties: <ul style="list-style-type: none"> ◦ CIDR Block: 10.0.10.0/24 ◦ Route Table: routetable-1 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: workers <p>One load balancer regional subnet created, named, and defined as follows:</p> <ul style="list-style-type: none"> • Name: loadbalancers-rs with the following properties: <ul style="list-style-type: none"> ◦ CIDR Block: 10.0.20.0/24 ◦ Route Table: routetable-0 ◦ Subnet access: Public ◦ DNS Resolution: Selected ◦ DHCP Options: Default ◦ Security List: loadbalancers

EXAMPLE SECURITY LIST CONFIGURATIONS FOR A HIGHLY AVAILABLE PUBLIC CLUSTER USING REGIONAL SUBNETS

In the example VCN, two security lists have been created (in addition to the default security list) to control access to and from a public worker node regional subnet and a load balancer regional subnet. The two security lists are named 'workers' and 'loadbalancers' respectively.

CHAPTER 9 Container Engine for Kubernetes

The workers security list has the following ingress and egress rules for the public worker node regional subnet:

Example Ingress Rules in a Security List for a Public Worker Node Regional Subnet:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateful	0.0.0.0/0	ICMP	n/a	n/a	3, 4	Allows: ICMP traffic for: 3, 4 Destination Unreachable: Fragmentation Needed and Don't Fragment was Set Description: This rule enables worker nodes to receive Path MTU Discovery fragmentation messages.

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
3	Stateful	130.35.0.0/16	TCP	All	22	n/a	Allows: TCP traffic for ports: 22 SSH Remote Login Protocol Description: This rule enables Container Engine for Kubernetes to access worker nodes.
4	Stateful	134.70.0.0/17	TCP	All	22	n/a	Allows: TCP traffic for ports: 22 SSH Remote Login Protocol Description: This rule enables Container Engine for Kubernetes to access worker nodes.

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
5	Stateful	138.1.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>
6	Stateful	140.91.0.0/17	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>

CHAPTER 9 Container Engine for Kubernetes

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
7	Stateful	147.154.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>
8	Stateful	192.29.0.0/16	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This rule enables Container Engine for Kubernetes to access worker nodes.</p>

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
9	Stateful	0.0.0.0/0	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: This <i>optional</i> rule enables inbound SSH traffic from the internet on port 22 to access worker nodes.</p>
10	Stateful	0.0.0.0/0	TCP	All	30000 - 32767	n/a	<p>Allows: TCP traffic for ports: 30000 - 32767</p> <p>Description: This <i>optional</i> rule enables inbound traffic to the worker nodes on the default NodePort range of 30000-32767 (see the Kubernetes documentation).</p>

Example Egress Rules in a Security List for a Public Worker Node Regional Subnet:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	10.0.10.0/24	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables intra-VCN traffic.
2	Stateful	All <region> Services in Oracle Services Network	All	n/a	n/a	n/a	Allows: All traffic for all ports Description: This rule enables outbound access to the regional service gateway.

CHAPTER 9 Container Engine for Kubernetes

The loadbalancers security list has the following ingress and egress rules for load balancer regional subnets:

Example Ingress Rules in a Security List for a Load Balancer Regional Subnet:

#	Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for all ports: all Description: This rule enables incoming public traffic to service load balancers.

Example Egress Rules in a Security List for a Load Balancer Regional Subnet:

#	Type	Dest. CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
1	Stateless	0.0.0.0/0	TCP	All	All	n/a	Allows: TCP traffic for ports: all Description: This rule enables responses from a web application through the service load balancers.

CIDR Blocks and Container Engine for Kubernetes

When configuring the VCN and the worker node and load balancer subnets for use with Container Engine for Kubernetes, you specify CIDR blocks to indicate the network addresses that can be allocated to the resources. See [Network Resource Configuration for Cluster Creation and Deployment](#).

When creating a cluster with Container Engine for Kubernetes, you specify:

- CIDR blocks for the Kubernetes services
- CIDR blocks that can be allocated to pods running in the cluster (see [Creating a Kubernetes Cluster](#))

Note the following:

- The CIDR block you specify for the VCN must not overlap with the CIDR block you specify for the Kubernetes services.
- The CIDR blocks you specify for pods running in the cluster must not overlap with CIDR blocks you specify for worker node and load balancer subnets.

Policy Configuration for Cluster Creation and Deployment

Before you can use Container Engine for Kubernetes to create and deploy clusters in the regions in a tenancy, the tenancy's root compartment must include a policy to allow Container Engine for Kubernetes to perform operations in the tenancy. See [Create Required Policy for Container Engine for Kubernetes](#).

When a tenancy is created, an Administrators group is automatically created for the tenancy. Users that are members of the Administrators group can perform any operation on resources in the tenancy. If all the users that will be working with Container Engine for Kubernetes are already members of the Administrators group, there's no need to create additional policies (aside from the policy to allow Container Engine for Kubernetes to perform operations in the tenancy). However, if you want to enable users that are not members of the Administrators group to use Container Engine for Kubernetes, you must create policies to enable the groups to which those users do belong to perform operations on resources in the tenancy or in

individual compartments. Some policies are required, some are optional. See [Create Required Policy for Groups](#) and [Create One or More Additional Policies for Groups](#).

Note that in addition to the above policies managed by IAM, you can also use the Kubernetes RBAC Authorizer to enforce additional fine-grained access control for users on specific clusters via Kubernetes RBAC roles and clusterroles. See [About Access Control and Container Engine for Kubernetes](#).

Create Required Policy for Container Engine for Kubernetes

To create and manage clusters in your tenancy, Container Engine for Kubernetes must have access to all resources in the tenancy. To give Container Engine for Kubernetes the necessary access, create a policy for the service as follows:

1. In the Console, open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**. A list of the policies in the compartment you're viewing is displayed.
2. Select the tenancy's root compartment from the list on the left.
3. Click **Create Policy**.
4. Enter the following:
 - **Name:** A unique name for the policy (for example, `oke-service`). The name must be unique across all policies in your tenancy. You cannot change this later. Avoid entering confidential information.
 - **Description:** A friendly description. You can change this later if you want to. Avoid entering confidential information.
 - **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in YYYY-MM-DD format. For more information, see [Policy Language Version](#).
 - **Statement:** The following policy statement:

```
Allow service OKE to manage all-resources in tenancy
```

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

Create Required Policy for Groups

To create, update, and delete clusters and node pools, users that are not members of the Administrators group must have permissions to work with cluster-related resources. To give users the necessary access, you must create a policy with a number of required policy statements for the groups to which those users do belong:

1. In the Console, open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**. A list of the policies in the compartment you're viewing is displayed.
2. Select the tenancy's root compartment or an individual compartment containing cluster-related resources from the list on the left.
3. Click **Create Policy**.
4. Enter the following:
 - **Name:** A name for the policy (for example, `acme-dev-team-oke-required-policy`) that is unique within the compartment. If you are creating the policy in the tenancy's root compartment, the name must be unique across all policies in your tenancy. You cannot change this later. Avoid entering confidential information.
 - **Description:** A friendly description. You can change this later if you want to. Avoid entering confidential information.
 - **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were

current on a specific date, select **Use Version Date** and enter that date in YYYY-MM-DD format. For more information, see [Policy Language Version](#).

- **Statement:** The following required policy statements to enable users to use Container Engine for Kubernetes to create, update, and delete clusters and node pools:
 - Allow group <group-name> to manage instance-family in <location>
 - Allow group <group-name> to use subnets in <location>
 - Allow group <group-name> to read virtual-network-family in <location>
 - Allow group <group-name> to use vnics in <location>
 - Allow group <group-name> to inspect compartments in <location>

The following required policy statement to enable users to perform any operation on cluster-related resources (this 'catch-all' policy effectively makes all users administrators insofar as cluster-related resources are concerned):

- Allow group <group-name> to manage cluster-family in <location>

In the above policy statements, replace <location> with either `tenancy` (if you are creating the policy in the tenancy's root compartment) or `compartment <compartment-name>` (if you are creating the policy in an individual compartment).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

Create One or More Additional Policies for Groups

To enable users that are not members of the Administrators group to use Container Engine for Kubernetes, create additional policies to enable the groups to which those users do belong to perform operations on cluster-related resources as follows:

1. In the Console, open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**. A list of the policies in the compartment you're viewing is displayed.
2. Select the tenancy's root compartment or an individual compartment containing cluster-related resources from the list on the left.
3. Click **Create Policy**.
4. Enter the following:
 - **Name:** A name for the policy (for example, `acme-dev-team-oke-additional-policy`) that is unique within the compartment. If you are creating the policy in the tenancy's root compartment, the name must be unique across all policies in your tenancy. You cannot change this later. Avoid entering confidential information.
 - **Description:** A friendly description. You can change this later if you want to. Avoid entering confidential information.
 - **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in YYYY-MM-DD format. For more information, see [Policy Language Version](#).
 - **Statement:** A suitable policy statement to allow existing groups to perform operations on cluster-related resources. In the example policy statements below, replace `<location>` with either `tenancy` (if you are creating the policy in the tenancy's root compartment) or `compartment <compartment-name>` (if you are creating the policy in an individual compartment):
 - To enable users in the `acme-dev-team` group to automatically create and configure associated new network resources when creating new 'quick

clusters', policies must also grant the group:

- the VCN_READ and VCN_CREATE permissions. Enter a policy statement like `Allow group acme-dev-team to manage vcns in <location>`
 - the SUBNET_READ and SUBNET_CREATE permissions. Enter a policy statement like `Allow group acme-dev-team to manage subnets in <location>`
 - the INTERNET_GATEWAY_CREATE permission. Enter a policy statement like `Allow group acme-dev-team to manage internet-gateways in <location>`
 - the NAT_GATEWAY_CREATE permission. Enter a policy statement like `Allow group acme-dev-team to manage nat-gateways in <location>`
 - the ROUTE_TABLE_UPDATE permission. Enter a policy statement like `Allow group acme-dev-team to manage route-tables in <location>`
 - the SECURITY_LIST_CREATE permission. Enter a policy statement like `Allow group acme-dev-team to manage security-lists in <location>`
- To enable users in the `acme-dev-team-cluster-viewers` group to simply list the clusters, enter a policy statement like `Allow group acme-dev-team-cluster-viewers to inspect clusters in <location>`.
 - To enable users in the `acme-dev-team-pool-admins` group to list, create, update, and delete node pools, enter a policy statement like `Allow group acme-dev-team-pool-admins to use cluster-node-pools in <location>`.
 - To enable users in the `acme-dev-team-auditors` group to see details of operations performed on clusters, enter a policy statement like `Allow`

```
group acme-dev-team-auditors to read cluster-work-requests in
<location>.
```

- To enable users in the acme-dev-team-sgw group to create a service gateway to enable worker nodes to access other resources in the same region without exposing data to the public internet, enter a policy statement like `Allow group acme-dev-team-sgw to manage service-gateways in <location>`.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

About Kubernetes Clusters and Nodes

A Kubernetes cluster is a group of nodes. The nodes are the machines running applications. Each node can be a physical machine or a virtual machine. The node's capacity (its number of CPUs and amount of memory) is defined when the node is created. A cluster can be organized into namespaces, to divide the cluster's resources between multiple uses. A cluster comprises:

- one or more master nodes (for high availability, typically there will be a number of master nodes)
- one or more worker nodes (sometimes known as minions)

The master nodes in a cluster run a number of processes:

- kube-apiserver to support API operations via the Kubernetes command line tool (kubectl) and the REST API, and includes admissions controllers required for advanced Kubernetes operations

CHAPTER 9 Container Engine for Kubernetes

- kube-controller-manager to manage different Kubernetes components (for example, replication controller, endpoints controller, namespace controller, and serviceaccounts controller)
- kube-scheduler to control where in the cluster to run jobs
- etcd to store the cluster's configuration data

Each worker node runs two Kubernetes processes:

- kubelet to communicate with the master nodes
- kube-proxy to handle networking

Each worker node also runs the Docker runtime.

The Kubernetes processes running on the master nodes are collectively referred to as the Kubernetes Control Plane. Together, the Control Plane processes monitor and record the state of the cluster and distribute requested operations between the nodes in the cluster.

When an application running on a worker node comprises multiple containers, Kubernetes groups the containers into a single logical unit called a pod for easy management and discovery. The containers in the pod share the same networking namespace and the same storage space, and can be managed as a single object by the Kubernetes Control Plane. A number of pods providing the same functionality can be grouped into a single logical set known as a service.

A Kubernetes manifest file comprises instructions in a yaml or json file that specify how to deploy an application to the node or nodes in a Kubernetes cluster. The instructions include information about the Kubernetes deployment, the Kubernetes service, and other Kubernetes objects to be created on the cluster. The manifest is commonly also referred to as a pod spec, or as a deployment.yaml file (although other filenames are allowed). The parameters to include in a Kubernetes manifest file are described in the [Kubernetes documentation](#).

A node pool is a subset of machines within a cluster that all have the same configuration. Node pools enable you to create pools of machines within a cluster that have different configurations. For example, you might create one pool of nodes in a cluster as virtual machines, and another pool of nodes as bare metal machines. A cluster must have a minimum of one node pool, but a node pool need not contain any worker nodes.

Creating a Kubernetes Cluster

You can use Container Engine for Kubernetes to create new Kubernetes clusters. To create a cluster, you must either belong to the tenancy's Administrators group, or belong to a group to which a policy grants the CLUSTER_MANAGE permission. In addition, a policy in the root compartment must grant Container Engine for Kubernetes access to all resources in the tenancy. See [Policy Configuration for Cluster Creation and Deployment](#).

You first specify basic details for the new cluster (the cluster name, and the Kubernetes version to install on master nodes). You can then create the cluster in one of two ways:

- Using default settings to create a 'quick cluster' with new network resources as required. This approach is the fastest way to create a new cluster. If you accept all the default values, you can create a new cluster in just a few clicks. New network resources for the 'quick cluster' are created automatically, including one regional subnet for worker nodes, and another regional subnet for load balancers. The regional subnet for load balancers will be public, but you can specify whether the regional subnet for worker nodes will be public or private. Note that if you specify a private regional subnet for worker nodes in the 'quick cluster', a NAT gateway is also created (in addition to an internet gateway). To create a 'quick cluster', you must belong to a group to which a policy grants the necessary permissions to create the new network resources (see [Create One or More Additional Policies for Groups](#)).
- Using custom settings to create a 'custom cluster'. This approach gives you the most control over the new cluster. You can explicitly define the new cluster's properties. And you can explicitly specify which existing network resources to use, including the existing public or private subnets in which to create worker nodes and load balancers. The subnets can be regional subnets (recommended) or AD-specific subnets. Note that although you will usually define node pools immediately when defining a new 'custom cluster', you don't have to. You can create a 'custom cluster' with no node pools, and add node pools later.

Regardless of how you create a cluster, Container Engine for Kubernetes gives names to worker nodes in the following format:

```
oke-c<part-of-cluster-OCID>-n<part-of-node-pool-OCID>-s<part-of-subnet-OCID>-<slot>
```

CHAPTER 9 Container Engine for Kubernetes

where:

- `oke` is the standard prefix for all worker nodes created by Container Engine for Kubernetes
- `c<part-of-cluster-OCID>` is a portion of the cluster's OCID, prefixed with the letter `c`
- `n<part-of-node-pool-OCID>` is a portion of the node pool's OCID, prefixed with the letter `n`
- `s<part-of-subnet-OCID>` is a portion of the subnet's OCID, prefixed with the letter `s`
- `<slot>` is an ordinal number of the node in the subnet (for example, 0, 1)

For example, if you specified a cluster is to have two nodes in a node pool, the two nodes might be named:

- `oke-cywiqripuyg-nsgagklgnst-st2qczvnmba-0`
- `oke-cywiqripuyg-nsgagklgnst-st2qczvnmba-1`

Do not change the auto-generated names that Container Engine for Kubernetes gives to worker nodes.

To ensure high availability, Container Engine for Kubernetes:

- creates the Kubernetes Control Plane on multiple Oracle-managed master nodes (distributing the master nodes across different availability domains in a region, where supported)
- creates worker nodes in each of the fault domains in an availability domain (distributing the worker nodes as evenly as possible across the fault domains, subject to any other infrastructure restrictions)

Using the Console to create a 'Quick Cluster' with Default Settings

To create a 'quick cluster' with default settings and new network resources using Container Engine for Kubernetes:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in, and in which you want to create both the new cluster and the associated network resources.
3. On the **Cluster List** page, click **Create Cluster**.
4. Either just accept the default configuration details for the new cluster, or specify alternatives as follows:
 - **Name:** The name of the new cluster. Either accept the default name or enter a name of your choice. Avoid entering confidential information.
 - **Kubernetes Version:** The version of Kubernetes to run on the master nodes and worker nodes of the cluster. Either accept the default version or select a version of your choice. Amongst other things, the Kubernetes version you select determines the default set of admission controllers that are turned on in the created cluster (the set follows the recommendation given in the [Kubernetes documentation](#) for that version).
5. Select **Quick Create** to create a new cluster with default settings, along with new network resources for the new cluster.

The **Create Virtual Cloud Network** panel shows the network resources that will be created for you by default.
6. Select either **Private** or **Public** to specify whether to create a private or a public regional subnet to host worker nodes (note that a public regional subnet is always created to host load balancers in a 'quick cluster', regardless of your selection here):
 - **Private:** Select to create a private regional subnet to host worker nodes (along with the public regional subnet to host load balancers).
 - **Public:** Select to create a public regional subnet to host worker nodes (along with the public regional subnet to host load balancers).

The **Create Node Pool** panel shows the fixed properties of the first node pool in the cluster that will be created for you:

- the name of the node pool (always pool1)
- the compartment in which the node pool will be created (always the same as the one in which the new network resources will reside)
- the version of Kubernetes that will run on each worker node in the node pool (always the same as the version specified for the master nodes)
- the image to use on each node in the node pool

The **Create Node Pool** panel also contains some node pool properties that you can change, but which have been given sensible defaults.

7. Either just accept all the default configuration details and skip ahead to the next step to create the cluster immediately, or specify alternatives as follows:
 - a. Either accept the default configuration details for the node pool, or specify alternatives in the **Create Node Pool** panel as follows:
 - **Shape:** The shape to use for each node in the node pool. The shape determines the number of CPUs and the amount of memory allocated to each node. The list shows only those shapes available in your tenancy that are supported by Container Engine for Kubernetes.
 - **Number of Nodes:** The number of worker nodes to create in the node pool, placed in the regional subnet created for the 'quick cluster'. The nodes are distributed as evenly as possible across the availability domains in a region (or in the case of a region with a single availability domain, across the fault domains in that availability domain).
 - **Public SSH Key:** (Optional) The public key portion of the key pair you want to use for SSH access to each node in the node pool. The public key is installed on all worker nodes in the cluster. Note that if you don't specify a public SSH key, Container Engine for Kubernetes will provide one. However, since you won't have the corresponding private key, you will not have SSH access to the worker nodes. Note that if you specify that you want the worker nodes in the 'quick cluster' to be hosted in a private regional subnet,

you cannot use SSH to access them directly (see [Connecting to Worker Nodes in Private Subnets Using SSH](#)).

- **Kubernetes Labels:** One or more labels (in addition to a default label) to add to worker nodes in the node pool to enable the targeting of workloads at specific node pools.
- b. Either accept the defaults for the remaining cluster details, or specify alternatives in the **Additional Add Ons** panel as follows:
- **Kubernetes Dashboard Enabled:** Select if you want to use the Kubernetes Dashboard to deploy and troubleshoot containerized applications, and to manage Kubernetes resources. See [Starting the Kubernetes Dashboard](#).
 - **Tiller (Helm) Enabled:** Select if you want Tiller (the server portion of Helm) to run in the Kubernetes cluster. With Tiller running in the cluster, you can use Helm to manage Kubernetes resources.
- c. (Optional) Select **View Detail Page After This Cluster Is Requested** to return to the **Cluster Details** tab (rather than the **Cluster List** page) in the Console at the end of the cluster creation process.
8. Click **Create** to create the new network resources and the new cluster.

Container Engine for Kubernetes starts creating:

- the network resources (such as the VCN, internet gateway, NAT gateway, route tables, security lists, a regional subnet for worker nodes and another regional subnet for load balancers), with auto-generated names in the format `oke-
<resource-type>-quick-<cluster-name>-<creation-date>`
- the cluster, with the name you specified
- the node pool, named `pool1`
- worker nodes, with auto-generated names in the format `oke-c<part-of-
cluster-OCID>-n<part-of-node-pool-OCID>-s<part-of-subnet-OCID>-
<slot>`

Do not change the resource names that Container Engine for Kubernetes has auto-generated. Note that if the cluster is not created successfully for some reason (for example, if you have insufficient permissions or if you've exceeded the cluster limit for the tenancy), any network resources created during the cluster creation process are not deleted automatically. You will have to manually delete any such unused network resources.

9. Click **Close** to return to the Console.

If you selected **View Detail Page After This Cluster Is Requested**, you return to the **Cluster Details** tab in the console. If you didn't select **View Detail Page After This Cluster Is Requested**, you return to the **Cluster List** page.

Initially, the new cluster appears in the Console with a status of Creating. When the cluster has been created, it has a status of Active.

Container Engine for Kubernetes also creates a Kubernetes kubeconfig configuration file that you use to access the cluster using kubectl and the Kubernetes Dashboard.

Using the Console to create a 'Custom Cluster' with Explicitly Defined Settings

To create a 'custom cluster' with explicitly defined settings and existing network resources using Container Engine for Kubernetes:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in, and in which you want to create the new cluster.
3. On the **Cluster List** page, click **Create Cluster**.
4. Specify configuration details for the new cluster:
 - **Name:** The name of the new cluster. Either accept the default name or enter a name of your choice. Avoid entering confidential information.

- **Kubernetes Version:** The version of Kubernetes to run on the master nodes of the cluster. Either accept the default version or select a version of your choice. Amongst other things, the Kubernetes version you select determines the default set of admission controllers that are turned on in the created cluster (the set follows the recommendation given in the [Kubernetes documentation](#) for that version).
5. Select **Custom Create** to create a new cluster by explicitly defining the new cluster's properties and which existing network resources to use.
 6. Specify the existing network resources to use for the new cluster in the **Network Selection** panel:
 - **Network Compartment:** The compartment in which the existing network resources reside.
 - **VCN:** The existing virtual cloud network that has been configured for cluster creation and deployment. See [VCN Configuration](#).
 - **Kubernetes Service LB Subnets:** Optionally, the existing subnets that have been configured to host load balancers. Load balancer subnets must be different from worker node subnets, can be public or private, and can be regional (recommended) or AD-specific. You don't have to specify any load balancer subnets. However, if you do specify load balancer subnets, the number of load balancer subnets to specify depends on the region in which you are creating the cluster and whether the subnets are regional or AD-specific.

If you are creating a cluster in a region with three availability domains, you can specify:

 - Zero or one load balancer regional subnet (recommended).
 - Zero or two load balancer AD-specific subnets. If you specify two AD-specific subnets, the two subnets must be in different availability domains.

If you are creating a cluster in a region with a single availability domain, you can specify:

- Zero or one load balancer regional subnet (recommended).
- Zero or one load balancer AD-specific subnet.

See [Subnet Configuration](#).

- **Kubernetes Service CIDR Block:** The available group of network addresses that can be exposed as Kubernetes services (ClusterIPs), expressed as a single, contiguous IPv4 CIDR block. For example, 10.96.0.0/16. The CIDR block you specify must not overlap with the CIDR block for the VCN. See [CIDR Blocks and Container Engine for Kubernetes](#).
 - **Pods CIDR Block:** The available group of network addresses that can be allocated to pods running in the cluster, expressed as a single, contiguous IPv4 CIDR block. For example, 10.244.0.0/16. The CIDR block you specify must not overlap with the CIDR blocks for subnets in the VCN, and can be outside the VCN CIDR block. See [CIDR Blocks and Container Engine for Kubernetes](#).
7. Specify whether to encrypt Kubernetes secrets at rest in the etcd key-value store for the cluster using the Key Management service. If you do want to encrypt Kubernetes secrets in the etcd key-value store, select the **Encrypt Using Customer-Managed Keys** option, and select:
- **Choose a Vault in <compartment-name>:** The vault that contains the master encryption key, from the list of vaults in the specified compartment. By default, <compartment-name> is the compartment in which you are creating the cluster, but you can select a different compartment by clicking **Change Compartment**.
 - **Choose a Key in <compartment-name>:** The name of the master encryption key, from the list of keys in the specified compartment. By default, <compartment-name> is the compartment in which you are creating the cluster, but you can select a different compartment by clicking **Change Compartment**. Note that you cannot change the master encryption key after the cluster has been created.

Note that if you do want to use encryption, a suitable master encryption key, dynamic group, and policy must already exist before you can create the cluster. For more information, see [Encrypting Kubernetes Secrets At Rest in Etcd](#).

8. Specify remaining details for the cluster in the **Additional Add Ons** panel:
 - **Kubernetes Dashboard Enabled:** Select if you want to use the Kubernetes Dashboard to deploy and troubleshoot containerized applications, and to manage Kubernetes resources. See [Starting the Kubernetes Dashboard](#).
 - **Tiller (Helm) Enabled:** Select if you want Tiller (the server portion of Helm) to run in the Kubernetes cluster. With Tiller running in the cluster, you can use Helm to manage Kubernetes resources.
9. Click **Continue**.
10. (Optional) Specify configuration details for the first node pool in the cluster in the **Node Pool** panel:
 - **Name:** A name of your choice for the new node pool. Avoid entering confidential information.
 - **Version:** The version of Kubernetes to run on each worker node in the node pool. By default, the version of Kubernetes specified for the master nodes is selected. The Kubernetes version on worker nodes must be either the same version as that on the master nodes, or an earlier version that is still compatible. See [Kubernetes Versions and Container Engine for Kubernetes](#).
 - **Image:** The image to use on each node in the node pool. An image is a template of a virtual hard drive that determines the operating system and other software for the node.
 - **Shape:** The shape to use for each node in the node pool. The shape determines the number of CPUs and the amount of memory allocated to each node. The list shows only those shapes available in your tenancy that are supported by Container Engine for Kubernetes.
 - **Number of Nodes:** The number of worker nodes to create in the node pool, placed in the availability domains you select, and in the regional subnet (recommended) or AD-specific subnet you specify for each availability domain.

- **Availability Domain 1:**

- **Availability Domain:** An availability domain in which to place worker nodes.
- **Subnet:** A regional subnet (recommended) or AD-specific subnet configured to host worker nodes. If you specified load balancer subnets, the worker node subnets must be different. The subnets you specify can be public or private, and can be regional (recommended) or AD-specific. See [Subnet Configuration](#).

Optionally click **Add Availability Domain** to select additional domains and subnets in which to place worker nodes.

When they are created, the worker nodes are distributed as evenly as possible across the availability domains you select (or in the case of a single availability domain, across the fault domains in that availability domain).

- **Public SSH Key:** (Optional) The public key portion of the key pair you want to use for SSH access to each node in the node pool. The public key is installed on all worker nodes in the cluster. Note that if you don't specify a public SSH key, Container Engine for Kubernetes will provide one. However, since you won't have the corresponding private key, you will not have SSH access to the worker nodes. Note that you cannot use SSH to access directly any worker nodes in private subnets (see [Connecting to Worker Nodes in Private Subnets Using SSH](#)).
 - **Kubernetes Labels:** One or more labels (in addition to a default label) to add to worker nodes in the node pool to enable the targeting of workloads at specific node pools.
11. (Optional) Click **Add node pool** and specify configuration details for a second and subsequent node pools in the cluster.

If you define multiple node pools in a cluster, you can host all of them on a single AD-specific subnet. However, it's best practice to host different node pools for a cluster on a regional subnet (recommended) or on different AD-specific subnets (one in each availability domain in the region).
 12. Click **Review** to confirm the resources that will be used and created.

13. (Optional) Select **View Detail Page After This Cluster Is Requested** to return to the **Cluster Details** tab (rather than the **Cluster List** page) in the Console at the end of the cluster creation process.
14. Click **Create** to create the new cluster.
Container Engine for Kubernetes starts creating the cluster with the name you specified. If you specified details for one or more node pools, Container Engine for Kubernetes creates:
 - node pools with the names you specified
 - worker nodes with auto-generated names in the format `oke-c<part-of-cluster-OCID>-n<part-of-node-pool-OCID>-s<part-of-subnet-OCID>-<slot>`Do not change the auto-generated names of worker nodes.
If you selected **View Detail Page After This Cluster Is Requested**, you return to the **Cluster Details** tab in the console. If you didn't select **View Detail Page After This Cluster Is Requested**, you return to the **Cluster List** page.
15. Initially, the new cluster appears in the Console with a status of Creating. When the cluster has been created, it has a status of Active.
Container Engine for Kubernetes also creates a Kubernetes kubeconfig configuration file that you use to access the cluster using kubectl and the Kubernetes Dashboard.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [CreateCluster](#) operation to create a cluster.

Downloading a kubeconfig File to Enable Cluster Access

When you create a cluster, you need to download a Kubernetes configuration file (commonly known as a 'kubeconfig' file) for the cluster. The kubeconfig file (by default named `config`

and stored in the `$HOME/.kube` directory) provides the necessary details to access the cluster using `kubectl` and the Kubernetes Dashboard.

You have to follow a number of steps to download the `kubeconfig` file. Having completed the steps, you can start using `kubectl` and the Kubernetes Dashboard to manage the cluster.

To download the `kubeconfig` file:

Step 1: Generate an API signing key pair

If you already have an API signing key pair, go straight to the next step. If not:

1. Use OpenSSL commands to generate the key pair in the required PEM format. If you're using Windows, you'll need to install Git Bash for Windows and run the commands with that tool. See [How to Generate an API Signing Key](#).
2. Copy the contents of the public key to the clipboard (you'll need to paste the value into the Console later).

Step 2: Upload the public key of the API signing key pair

1. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
2. Click **Add Public Key**.
3. Paste the public key's value into the window and click **Add**.
The key is uploaded and its fingerprint is displayed (for example, `d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:13`).

Step 3: Install and configure the Oracle Cloud Infrastructure CLI

1. Install the Oracle Cloud Infrastructure CLI version 2.6.4 (or later). See [Quickstart](#).
2. Configure the Oracle Cloud Infrastructure CLI. See [Configuration](#).

Step 4: Download the kubeconfig file

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster you want to access using `kubectl` and the Kubernetes Dashboard. The **Cluster** page shows details of the cluster.
4. Click the **Access Kubeconfig** button to display the **How to Access Kubeconfig** dialog box.
5. Create a directory to contain the kubeconfig file. By default, the expected directory name is `$HOME/.kube`.

For example, on Linux, enter the following command (or copy and paste it from the **How to Access Kubeconfig** dialog box):

```
$ mkdir -p $HOME/.kube
```

6. Run the Oracle Cloud Infrastructure CLI command to download the kubeconfig file and save it in a location accessible to `kubectl` and the Kubernetes Dashboard. For example, on Linux, enter the following command (or copy and paste it from the **How to Access Kubeconfig** dialog box):

```
$ oci ce cluster create-kubeconfig --cluster-id ocid1.cluster.oc1.phx.aaaaaaaaae... --file $HOME/.kube/config --region us-phoenix-1 --token-version 2.0.0
```

where `ocid1.cluster.oc1.phx.aaaaaaaaae...` is the OCID of the current cluster. For convenience, the command in the **How to Access Kubeconfig** dialog box already includes the cluster's OCID.

Note that if a kubeconfig file already exists in the location you specify, details about the cluster will be added as a new context to the existing kubeconfig file. The `current-context`: element in the kubeconfig file will be set to point to the newly-added context.

7. Set the value of the `KUBECONFIG` environment variable to point to the name and location of the kubeconfig file. For example, on Linux, enter the following command (or copy and paste it from the **How to Access Kubeconfig** dialog box):

```
$ export KUBECONFIG=$HOME/.kube/config
```

Step 5: Verify that kubectl is available

1. Verify that kubectl is available by entering the following command:

```
$ kubectl version
```

The response shows:

- the version of kubectl installed and running locally
- the version of Kubernetes (strictly speaking, the version of the kube-apiserver) running on the cluster's master node

Note that the kubectl version must be within one minor version (older or newer) of the Kubernetes version running on the master node. If kubectl is more than one minor version older or newer, install an appropriate version of kubectl. See [Kubernetes version and version skew support policy](#) in the Kubernetes documentation.

2. Verify that kubectl can connect to the cluster by entering the following command:

```
$ kubectl get nodes
```

Information about the nodes in the cluster is shown.

You can now use kubectl and the Kubernetes Dashboard to perform operations on the cluster.

Notes about kubeconfig Files

Note the following about kubeconfig files:

- A single kubeconfig file can include the details for multiple clusters, as multiple contexts. The cluster on which operations will be performed is specified by the `current-context`: element in the kubeconfig file.
- A kubeconfig file includes an Oracle Cloud Infrastructure CLI command that dynamically generates an authentication token and inserts it when you run a kubectl command. The

Oracle Cloud Infrastructure CLI must be available on your shell's executable path (for example, `$PATH` on Linux).

- The tokens generated by the Oracle Cloud Infrastructure CLI command in the kubeconfig file are short-lived, cluster-scoped, and specific to individual users. As a result, you cannot share kubeconfig files between users to access Kubernetes clusters.
- The Oracle Cloud Infrastructure CLI command in the kubeconfig file uses your current CLI profile when generating an authentication token. If you have defined multiple profiles in different tenancies in the CLI configuration file (for example, in `~/.oci/config`), specify which profile to use when generating the authentication token as follows. In both cases, `<profile-name>` is the name of the profile defined in the CLI configuration file:
 - Add `--profile` to the `args`: section of the kubeconfig file as follows:

```
user:
  exec:
    apiVersion: client.authentication.k8s.io/v1beta1
    args:
      - ce
      - cluster
      - generate-token
      - --cluster-id
      - <cluster ocid>
      - --profile
      - <profile-name>
    command: oci
    env: []
```

- Set the `OCI_CLI_PROFILE` environment variable to the name of the profile defined in the CLI configuration file before running `kubectl` commands. For example:

```
$ export OCI_CLI_PROFILE=<profile-name>
$ kubectl get nodes
```

Upgrading kubeconfig files from Version 1.0.0 to Version 2.0.0

Container Engine for Kubernetes currently supports two versions of kubeconfig file:

- version 1.0.0
- version 2.0.0

Enhancements in kubeconfig version 2.0.0 files provide security improvements for your Kubernetes environment, including short-lived cluster-scoped tokens with automated refreshing, and support for instance principals to access Kubernetes clusters. Additionally, authentication tokens are generated on-demand for each cluster, so kubeconfig version 2.0.0 files cannot be shared between users to access Kubernetes clusters (unlike kubeconfig version 1.0.0 files).

Note that support for kubeconfig version 1.0.0 files will be discontinued on November 15, 2019. Prior to that date, you must upgrade any kubeconfig version 1.0.0 files to version 2.0.0. Follow the instructions below to determine the current version of kubeconfig files, and how to upgrade them to version 2.0.0.

Determine the kubeconfig file version

To determine the version of a cluster's kubeconfig file:

1. In a terminal window, enter the following command to see the format of the kubeconfig file currently pointed at by the KUBECONFIG environment variable:

```
$ kubectl config view
```

2. If the kubeconfig file is version 1.0.0, you see a response in the following format:

```
users:  
- name: <username>  
  user:  
    token: <token-value>
```

If you see a response in the above format, you have to upgrade the kubeconfig file. See [Upgrade a kubeconfig version 1.0.0 file to version 2.0.0](#).

3. If the kubeconfig file is version 2.0.0, you see a response in the following format:

```
user:  
  exec:  
    apiVersion: client.authentication.k8s.io/v1beta1  
    args:  
    - ce
```

```
- cluster
- generate-token
- --cluster-id
- <cluster ocid>
command: oci
env: []
```

If you see a response in the above format, no further action is required.

Upgrade a kubeconfig version 1.0.0 file to version 2.0.0

To upgrade a kubeconfig version 1.0.0 file:

1. Confirm the Oracle Cloud Infrastructure CLI version 2.6.4 (or later) is installed by entering:

```
oci -version
```

2. If the Oracle Cloud Infrastructure CLI version is earlier than version 2.6.4, upgrade the CLI to a later version. See [Upgrading the CLI](#).
3. Follow the instructions to download the kubeconfig file (see [Step 4: Download the kubeconfig file](#)). Running the `oci ce cluster create-kubeconfig` command shown in the **How to Access Kubeconfig** dialog box upgrades the existing kubeconfig version 1.0.0 file. If you change the name or location of the kubeconfig file, set the `KUBECONFIG` environment variable to point to the new name and location of the file.
4. Confirm the kubeconfig file is now version 2.0.0:

- a. In a terminal window, enter:

```
$ kubectl config view
```

- b. Confirm that that the response is in the following format:

```
user:
  exec:
    apiVersion: client.authentication.k8s.io/v1beta1
    args:
      - ce
      - cluster
      - generate-token
      - --cluster-id
```

```
- <cluster ocid>  
command: oci  
env: []
```

Modifying a Kubernetes Cluster

You can use Container Engine for Kubernetes to modify the node pool details of existing Kubernetes clusters.

You can change:

- the name of the node pool
- the number of node pools in a cluster by adding new node pools, or deleting existing node pools
- the number of worker nodes in a node pool, and the availability domains and subnets in which to place them
- the version of Kubernetes to run on new worker nodes

However, note that you cannot change:

- the name of the cluster
- the shape of existing worker nodes
- the operating system running on existing worker nodes
- the version of Kubernetes running on existing worker nodes
- the master encryption key (if specified when the cluster was created)

Also note that you must not change the auto-generated names of resources that Container Engine for Kubernetes has created (such as the names of worker nodes).

Using the Console

To modify an existing Kubernetes cluster:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster you want to modify.
4. Use the buttons across the top of the **Cluster** page as follows:
 - If you want to download the kubeconfig configuration file for the cluster, click the **Access Kubeconfig** button (see [Downloading a kubeconfig File to Enable Cluster Access](#)).
 - If you want to add a new node pool to the cluster, click the **Add Node Pool** button and enter details for the new node pool.
 - If you want to delete the cluster along with its master nodes and worker nodes, click the **Delete Cluster** button.
 - If a newer version of Kubernetes is available than the one running on the master nodes in the cluster, the **Upgrade Available** button is enabled. If you want to upgrade the master nodes to a newer version, click **Upgrade Available** (see [Upgrading the Version of Kubernetes Running on Master Nodes](#)).
5. Use the **Cluster Details** tab to see information about the cluster, including:
 - The status of the cluster, and of the node pools in the cluster.
 - The cluster's OCID.
 - The Kubernetes version running on the master nodes in the cluster.
 - The address of the Kubernetes endpoint.
6. Use the **Node Pools** tab to:
 - View a summary of each node pool and the worker nodes within it.
 - Change the name of a node pool by selecting **Edit** from the **Actions** menu.

- Change the number and placement of worker nodes in a node pool by selecting **Scale** from the **Actions** menu and specifying a different number of worker nodes, different availability domains, and different regional subnets (recommended) or AD-specific subnets.
 - View and edit configuration details of specific worker nodes by selecting **Show Node Details** and clicking the name of the worker node.
 - Delete a node pool by selecting **Delete Node Pool** from the **Actions** menu.
7. Use the **Getting Started** tab to:
- View and copy the commands to start the Kubernetes Dashboard to view the deployed application running on nodes in the cluster (see [Starting the Kubernetes Dashboard](#)).
 - View and copy the commands to download and deploy a sample nginx application using the Kubernetes command line tool kubectl from the instructions in a manifest file (see [Deploying a Sample Nginx App on a Cluster Using kubectl](#)).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [Update Cluster](#) and the [UpdateNodePool](#) operations to modify an existing Kubernetes cluster.

Deleting a Kubernetes Cluster

You can delete a cluster along with its master nodes, worker nodes, and node pools.

Note the following:

- When you delete a cluster, no other resources created during the cluster creation process or associated with the cluster (such as VCNs, internet gateways, NAT gateways, route tables, security lists, load balancers, and block volumes) are deleted

automatically. If you want to delete these resources, you have to do so manually.

- Container Engine for Kubernetes creates the worker nodes (compute instances) in a cluster with auto-generated names in the format `oke-c<part-of-cluster-OCID>-n<part-of-node-pool-OCID>-s<part-of-subnet-OCID>-<slot>`. Do not change the auto-generated names of worker nodes. If you do change the auto-generated name of a worker node and then delete the cluster, the renamed worker node is not deleted. You would have to delete the renamed worker node manually.

Using the Console

To delete a Kubernetes cluster using Container Engine for Kubernetes:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the Delete icon beside the name of the cluster to delete, and confirm that you want to delete it.

You can also delete a cluster using the **Delete Cluster** button on the **Cluster** page.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [DeleteCluster](#) operation to delete a cluster.

Monitoring Clusters

Having created a cluster, you can monitor the status of the cluster itself, and the nodes and node pools within it.

Using the Console

To monitor a Kubernetes cluster:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.

The **Status** column on the **Cluster List** page shows a summary status for each individual cluster and its master nodes. Clusters can have one of the following statuses:

Cluster Status	Explanation	Possible Reason
Creating	Cluster is in the process of being created.	Application is being deployed.
Active	Cluster is running normally.	Master nodes are running normally.
Failed	Cluster is not running due to an unrecoverable error.	Possible reasons: <ul style="list-style-type: none"> • a problem setting up load balancers • an error installing cluster add-ons (Tiller, Kubernetes dashboard) • conflicts in networking ranges
Deleting	Cluster is in the process of being deleted. Application no longer required, so resources in the process of being released.	Application no longer required, so resources in the process of being released.

Cluster Status	Explanation	Possible Reason
Deleted	Cluster has been deleted. Application no longer required, so resources have been released.	Application no longer required, so resources have been released.
Updating	Version of Kubernetes on the master nodes is in the process of being upgraded.	A newly supported version of Kubernetes has become available.

Note that the cluster's summary status is not necessarily directly related to the status of node pools and nodes within the cluster.

3. On the **Cluster List** page, click the name of the cluster for which you want to see detailed status.

The **Cluster Details** tab shows the summary status for the cluster and its master nodes.

4. Use the **Node Pools** tab to see the status of individual nodes within each node pool. Nodes can have one of the following statuses:

Node Status	Explanation	Possible Reason
Creating	Node is being created.	Compute instance in the process of being created.
Active	Node is running normally.	Node is running normally.
Updating	Node is in the process of being updated.	Container Engine for Kubernetes is performing an operation on the node.

Node Status	Explanation	Possible Reason
Deleting	Node is in the process of being deleted.	Application no longer required, so resources in the process of being released.
Deleted	Node has been deleted.	Application no longer required, so resources have been released.
Inactive	Node still exists, but is not running.	Compute resource has a status of Stopped, Stopping, or Down For Maintenance.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [GetCluster](#) and [GetNodePool](#) operations to monitor the status of Kubernetes clusters.

Monitoring Operations of Container Engine for Kubernetes

You can monitor operations performed by Container Engine for Kubernetes as follows:

- You can monitor and manage operations performed on a particular cluster by Container Engine for Kubernetes using the Work Requests tab of the cluster's Summary page.
- You can view all operations performed by Container Engine for Kubernetes as log events using the Oracle Cloud Infrastructure Audit service.

Using the Console

To monitor and manage operations performed on a particular cluster:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster for which you want to monitor and manage operations.

The **Cluster** page shows information about the cluster.

To view all operations performed by Container Engine for Kubernetes as log events:

1. In the Console, open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Policies**.
2. Choose a **Compartment** you have permission to work in.
3. Search and filter to show the operations performed by Container Engine for Kubernetes. See [Viewing Audit Log Events](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [GetWorkRequest](#), [DeleteWorkRequest](#), [ListWorkRequestErrors](#), [ListWorkRequestLogs](#), and [ListWorkRequests](#) operations to monitor and manage operations performed by Container Engine for Kubernetes.

Accessing a Cluster Using kubectl

You can use the Kubernetes command line tool kubectl to perform operations on a cluster you've created with Container Engine for Kubernetes. Before you can use kubectl to access a cluster, you need to specify the cluster on which to perform operations by downloading the

cluster's kubeconfig file. Note that a Oracle Cloud Infrastructure CLI command in the kubeconfig file generates authentication tokens that are short-lived, cluster-scoped, and specific to individual users. As a result, you cannot share kubeconfig files between users to access Kubernetes clusters.

To access a cluster using kubectl:

1. If you haven't already done so, install kubectl (see the [kubectl documentation](#)).
2. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
3. In a terminal window, enter `kubectl` followed by the command for the operation you want to perform on the cluster. For a list of available commands and options, see the [kubectl documentation](#).

Note that you must have the appropriate permissions to run the command you enter. See [About Access Control and Container Engine for Kubernetes](#).

Starting the Kubernetes Dashboard

Kubernetes Dashboard is a web-based user interface that you can use as an alternative to the Kubernetes kubectl command line tool to:

- deploy containerized applications to a Kubernetes cluster
- troubleshoot your containerized applications

You use the Kubernetes Dashboard to get an overview of applications running on a cluster, as well as to create or modify individual Kubernetes resources. The Kubernetes Dashboard also reports the status of Kubernetes resources in the cluster, and any errors that have occurred. Note that to use the Kubernetes Dashboard, it must have been enabled when the cluster was initially created.

CHAPTER 9 Container Engine for Kubernetes

In contrast to the Kubernetes Dashboard, Container Engine for Kubernetes enables you to create and delete Kubernetes clusters and node pools, and to manage the associated compute, network, and storage resources.

Before you can use the Kubernetes Dashboard to access a cluster, you need to specify the cluster on which to perform operations by downloading the cluster's kubeconfig file. Note that an Oracle Cloud Infrastructure CLI command in the kubeconfig file generates authentication tokens that are short-lived, cluster-scoped, and specific to individual users. As a result, you cannot share kubeconfig files between users to access Kubernetes clusters.

To start the Kubernetes Dashboard:

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. In a text editor, create a file (for example, called oke-admin-service-account.yaml) with the following content:

```
apiVersion: v1
kind: ServiceAccount
metadata:
  name: oke-admin
  namespace: kube-system
---
apiVersion: rbac.authorization.k8s.io/v1beta1
kind: ClusterRoleBinding
metadata:
  name: oke-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: oke-admin
  namespace: kube-system
```

The file defines an administrator service account and a clusterrolebinding, both called oke-admin.

3. Create the service account and the clusterrolebinding in the cluster by entering:

```
$ kubectl apply -f <filename>
```

CHAPTER 9 Container Engine for Kubernetes

where `<filename>` is the name of the file you created earlier. For example:

```
$ kubectl apply -f oke-admin-service-account.yaml
```

The output from the above command confirms the creation of the service account and the clusterrolebinding:

```
serviceaccount "oke-admin" created
clusterrolebinding.rbac.authorization.k8s.io "oke-admin" created
```

You can now use the oke-admin service account to view and control the cluster, and to connect to the Kubernetes dashboard.

4. Obtain an authentication token for the oke-admin service account by entering:

```
$ kubectl -n kube-system describe secret $(kubectl -n kube-system get secret | grep oke-admin | awk '{print $1}')
```

The output from the above command includes an authentication token (a long alphanumeric string) as the value of the `token:` element, as shown below:

```
Name:          oke-admin-token-gwbp2
Namespace:    kube-system
Labels:       <none>
Annotations:  kubernetes.io/service-account.name: oke-admin
              kubernetes.io/service-account.uid: 3a7fcd8e-e123-11e9-81ca-0a580aed8570
Type:         kubernetes.io/service-account-token
Data
====
ca.crt:       1289 bytes
namespace:    11 bytes
token:        eyJh_____pxlQ
```

In the example above, `eyJh_____pxlQ` (abbreviated for readability) is the authentication token.

5. Copy the value of the `token:` element from the output. You will use this token to connect to the dashboard.
6. In a terminal window, enter `kubectl proxy` to start the Kubernetes Dashboard.

7. Open a browser and go to `http://localhost:8001/api/v1/namespaces/kube-system/services/https:kubernetes-dashboard:/proxy/#!/login` to display the Kubernetes Dashboard.
8. In the Kubernetes Dashboard, select **Token** and paste the value of the `token:` element you copied earlier into the **Token** field.
9. In the Kubernetes Dashboard, click **Sign In**, and then click **Overview** to see the applications deployed on the cluster.

Deploying a Sample Nginx App on a Cluster Using kubectl

Having created a Kubernetes cluster using Container Engine for Kubernetes, you'll typically want to try it out by deploying an application on the nodes in the cluster. For convenience, the **Cluster** page includes a **Getting Started** tab that makes it easy to view and copy the commands to:

- download the kubeconfig configuration file for the cluster
- download and deploy a sample Nginx application using the Kubernetes command line tool kubectl from the instructions in a manifest file
- start the Kubernetes Dashboard to view the deployed application running on nodes in the cluster

To deploy the sample nginx application:

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. In a terminal window, deploy the sample Nginx application by entering `kubectl create -f https://k8s.io/docs/tasks/run-application/deployment.yaml`



Tip

If the command fails to connect to <https://k8s.io/docs/tasks/run-application/deployment.yaml>, go to the url in a browser and download the manifest file `deployment.yaml` to a local directory. Repeat the `kubectl create` command and specify the local location of the `deployment.yaml` file.

3. Use the Kubernetes Dashboard or `kubectl` to confirm that the sample application has deployed successfully. For example:
 - a. Enter `kubectl proxy` to start the Kubernetes Dashboard.
 - b. Open a browser and go to `http://localhost:8001/api/v1/namespaces/kubernetes/services/https:kubernetes-dashboard:/proxy/` to display the Kubernetes Dashboard.
 - c. Click **Overview** to see the applications deployed on the cluster.

You can see the Nginx sample application has been deployed as two pods, on two nodes in the cluster.

Pulling Images from Registry during Deployment

During the deployment of an application to a Kubernetes cluster, you'll typically want one or more images to be pulled from a Docker registry. In the application's manifest file you specify the images to pull, the registry to pull them from, and the credentials to use when pulling the images. The manifest file is commonly also referred to as a pod spec, or as a `deployment.yaml` file (although other filenames are allowed).

CHAPTER 9 Container Engine for Kubernetes

If you want the application to pull images that reside in Oracle Cloud Infrastructure Registry, you have to perform two steps:

- You have to use `kubectl` to create a Docker registry secret. The secret contains the Oracle Cloud Infrastructure credentials to use when pulling the image. When creating secrets, Oracle strongly recommends you use the latest version of `kubectl` (see the [kubectl documentation](#)).
- You have to specify the image to pull from Oracle Cloud Infrastructure Registry, including the repository location and the Docker registry secret to use, in the application's manifest file.

To create a Docker registry secret:

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the `KUBECONFIG` environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. In a terminal window, enter:

```
$ kubectl create secret docker-registry <secret-name> --docker-server=<region-code>.ocir.io --  
docker-username='<tenancy-namespace>/<oci-username>' --docker-password='<oci-auth-token>' --  
docker-email='<email-address>'
```

where:

- `<secret-name>` is a name of your choice, that you will use in the manifest file to refer to the secret. For example, `ocirsecret`
- `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `iad`. See [Availability by Region Name and Region Code](#) for the list of region codes.
- `ocir.io` is the Oracle Cloud Infrastructure Registry name.
- `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy containing the repository from which the application is to pull the image (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some

older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`).

- `<oci-username>` is the username to use when pulling the image. The username must have access to the tenancy specified by `<tenancy-name>`. For example, `jdoue@acme.com`. If your tenancy is federated with Oracle Identity Cloud Service, use the format `oracleidentitycloudservice/<username>`
- `<oci-auth-token>` is the auth token of the user specified by `<oci-username>`. For example, `k]j64r{1sJSSF-;)K8`
- `<email-address>` is an email address. An email address is required, but it doesn't matter what you specify. For example, `jdoue@acme.com`

Note the use of single quotes around strings containing special characters.

For example, combining the previous examples, you might enter:

```
$ kubectl create secret docker-registry ocirsecret --docker-server=phx.ocir.io --docker-username='ansh81vrulzp/jdoue@acme.com' --docker-password='k]j64r{1sJSSF-;)K8' --docker-email='jdoue@acme.com'
```

Having created the Docker secret, you can now refer to it in the application manifest file.

To specify the image to pull from Oracle Cloud Infrastructure Registry, along with the Docker secret to use, during deployment of an application to a cluster:

1. Open the application's manifest file in a text editor.
2. Add the following sections to the manifest file:
 - a. Add a `containers` section that specifies the name and location of the container you want to pull from Oracle Cloud Infrastructure Registry, along with other deployment details.
 - b. Add an `imagePullSecrets` section to the manifest file that specifies the name of the Docker secret you created to access the Oracle Cloud Infrastructure Registry.

Here's an example of what the manifest might look like when you've added the `containers` and `imagePullSecrets` sections:

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx-image
spec:
  containers:
    - name: nginx
      image: phx.ocir.io/ansh81vrulzp/project01/nginx-lb:latest
      imagePullPolicy: Always
      ports:
        - name: nginx
          containerPort: 8080
          protocol: TCP
  imagePullSecrets:
    - name: ocirsecret
```

3. Save and close the manifest file.

Encrypting Kubernetes Secrets At Rest in Etcd

The master nodes in a Kubernetes cluster store sensitive configuration data (such as authentication tokens, passwords, and SSH keys) as Kubernetes secret objects in etcd. Etcd is an open source distributed key-value store that Kubernetes uses for cluster coordination and state management. In the Kubernetes clusters created by Container Engine for Kubernetes, etcd writes and reads data to and from block storage volumes in the Oracle Cloud Infrastructure Block Volume service. Although the data in block storage volumes is encrypted, Kubernetes secrets at rest in etcd itself are not encrypted by default.

For additional security, when you create a new cluster you can specify that Kubernetes secrets at rest in etcd are to be encrypted using the Oracle Cloud Infrastructure Key Management service (see [Overview of Key Management](#)). Before you can create a cluster where Kubernetes secrets are encrypted in the etcd key-value store, you have to:

- know the name and OCID of a suitable master encryption key in Key Management
- create a dynamic group that includes all clusters in the compartment in which you are going to create the new cluster
- create a policy authorizing the dynamic group to use the master encryption key

Having created the cluster and specified that you want Kubernetes secrets at rest in the etcd key-value store to be encrypted, you can optionally restrict the use of the master encryption key by modifying the dynamic group to include just that cluster.

Note the following:

- You can only select the option to encrypt the Kubernetes secrets in the cluster's etcd key-value store when creating a new 'custom cluster'. You cannot encrypt Kubernetes secrets in the etcd key-value stores of existing 'custom clusters', or in the etcd key-value stores of 'quick clusters'.
- You can only select the option to encrypt Kubernetes secrets in the cluster's etcd key-value store if you specify Kubernetes version 1.13.x or later as the version of Kubernetes to run on the master nodes of the cluster.
- After you've specified a master encryption key for a new cluster and created the cluster, do not subsequently delete the master encryption key in the Key Management service. As soon as you schedule a key for deletion in Key Management, the Kubernetes secrets stored for the cluster in etcd become inaccessible. If you have already scheduled the key for deletion, it might still be in the Pending Deletion state. If that is the case, cancel the scheduled key deletion (see [To cancel the deletion of a key](#)) to restore access to the Kubernetes secrets. If you allow the scheduled key deletion operation to complete and the master encryption key to be deleted, the Kubernetes secrets stored for the cluster in etcd are permanently inaccessible. As a result, cluster upgrades will fail. In this situation, you have no choice but to delete and recreate the cluster.

Using the Console

To create a new 'custom cluster' where Kubernetes secrets are encrypted in the cluster's etcd key-value store:

1. Log in to the Console.
2. If you know the OCID of the master encryption key to use to encrypt Kubernetes secrets, go straight to the next step. Otherwise:
 - If a suitable master encryption key already exists in Key Management but you're not sure of its OCID, follow the instructions in [To view key details](#) and make a note of the master encryption key's OCID.
 - If a suitable master encryption key does not already exist in Key Management, follow the instructions in [To create a new key](#) to create one. Having created a new master encryption key, make a note of its OCID.
3. Create a new dynamic group containing all the clusters in the compartment in which you intend to create the new cluster:
 - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Dynamic Groups**.
 - b. Follow the instructions in [To create a dynamic group](#), and give the dynamic group a name (for example, `acme-oke-kms-dyn-grp`).
 - c. Enter a rule that includes all clusters in the compartment in the format:

```
ALL {resource.type = 'cluster', resource.compartment.id = '<compartment-ocid>'}
```

where `<compartment-ocid>` is the OCID of the compartment in which you intend to create the new cluster.

For example:

```
ALL {resource.type = 'cluster', resource.compartment.id =  
'ocid1.compartment.oc1..aaaaaaa23_____smwa'}
```

- d. Click **Create Dynamic Group**.

Having created a dynamic group that includes all clusters in the compartment, you can now create a policy to give the dynamic group access to the master encryption key in Key Management.

4. Create a new policy to enable use of the master encryption key:
 - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
 - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `acme-oke-kms-dyn-grp-policy`).
 - c. Enter a policy statement to give the dynamic group access to the master encryption key, in the format:

```
Allow dynamic-group <dynamic-group-name> to use keys in compartment <compartment-name>
where target.key.id = '<key-OCID>'
```

where:

- `<dynamic-group-name>` is the name of the dynamic group you created earlier.
- `<compartment-name>` is the name of the compartment containing the master encryption key.
- `<key-OCID>` is the OCID of the master encryption key in Key Management.

For example:

```
Allow dynamic-group <acme-oke-kms-dyn-grp> to use keys in compartment acme-kms-key-
compartment where target.key.id = 'ocid1.key.oc1.iad.annrl_____trfg'
```

- d. Click **Create** to create the new policy.
5. Follow the instructions to create a new 'custom cluster' in [Using the Console to create a 'Custom Cluster' with Explicitly Defined Settings](#), select the **Encrypt Using Customer-Managed Keys** option, and select:
 - **Choose a Vault in <compartment-name>**: The vault that contains the master encryption key, from the list of vaults in the specified compartment. By default, `<compartment-name>` is the compartment in which you are creating the cluster, but you can select a different compartment by clicking **Change Compartment**.

- **Choose a Key in <compartment-name>:** The name of the master encryption key, from the list of keys in the specified compartment. By default, <compartment-name> is the compartment in which you are creating the cluster, but you can select a different compartment by clicking **Change Compartment**. Note that you cannot change the master encryption key after the cluster has been created.
6. (Optional) Having created the cluster, for additional security:
 - a. Make a note of the OCID of the new cluster you just created.
 - b. Restrict the use of the master encryption key by modifying the dynamic group rule you created earlier to explicitly specify the OCID of the new cluster, rather than all clusters in the compartment. For example:

```
resource.id = 'ocid1.cluster.oc1.iad.aaaaaaaaaf_____yg5q'
```

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [CreateCluster](#) operation to create a cluster.

Connecting to Worker Nodes Using SSH

If you provided a public SSH key when creating the node pool in a cluster, the public key is installed on all worker nodes in the cluster. On UNIX and UNIX-like platforms (including Solaris and Linux), you can then connect through SSH to the worker nodes using the ssh utility (an SSH client) to perform administrative tasks.

Note the following instructions assume the UNIX machine you use to connect to the worker node:

- Has the ssh utility installed.
- Has access to the SSH private key file paired with the SSH public key that was specified when the cluster was created.

How to connect to worker nodes using SSH depends on whether you specified public or private subnets for the worker nodes when defining the node pools in the cluster.

Connecting to Worker Nodes in Public Subnets Using SSH

Before you can connect to a worker node in a public subnet using SSH, you must define an ingress rule in the subnet's security list to allow SSH access. The ingress rule must allow access to port 22 on worker nodes from source 0.0.0.0/0 and any source port, as follows:

Type	Source CIDR	IP Protocol	Source Port Range	Dest. Port Range	Type and Code	Allows: and Description:
Stateful	0.0.0.0/0	TCP	All	22	n/a	<p>Allows: TCP traffic for ports: 22 SSH Remote Login Protocol</p> <p>Description: Enables SSH access.</p>

To connect to a worker node in a public subnet through SSH from a UNIX machine using the ssh utility:

1. Find out the IP address of the worker node to which you want to connect. You can do this in a number of ways:
 - Using kubectl. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#). Then in a terminal window, enter `kubectl get nodes` to see the public IP addresses of worker nodes in node pools in the cluster.

- Using the Console. In the Console, display the **Cluster List** page and then select the cluster to which the worker node belongs. Click **Node Pools** to see the public IP addresses of worker nodes in every node pool in the cluster.
 - Using the REST API. Use the [ListNodePools](#) operation to see the public IP addresses of worker nodes in a node pool.
2. In the terminal window, enter `ssh opc@<node_ip_address>` to connect to the worker node, where `<node_ip_address>` is the IP address of the worker node that you made a note of earlier. For example, you might enter `ssh opc@192.0.2.254`.

Note that if the SSH private key is not stored in the file or in the path that the ssh utility expects (for example, the ssh utility might expect the private key to be stored in `~/.ssh/id_rsa`), you must explicitly specify the private key filename and location in one of two ways:

- Use the `-i` option to specify the filename and location of the private key. For example, `ssh -i ~/.ssh/my_keys/my_host_key_filename opc@192.0.2.254`
- Add the private key filename and location to an SSH configuration file, either the client configuration file (`~/.ssh/config`) if it exists, or the system-wide client configuration file (`/etc/ssh/ssh_config`). For example, you might add the following:

```
Host 192.0.2.254 IdentityFile ~/.ssh/my_keys/my_host_key_filename
```

For more about the ssh utility's configuration file, enter `man ssh_config`

Note also that permissions on the private key file must allow you read/write/execute access, but prevent other users from accessing the file. For example, to set appropriate permissions, you might enter `chmod 600 ~/.ssh/my_keys/my_host_key_filename`. If permissions are not set correctly and the private key file is accessible to other users, the ssh utility will simply ignore the private key file.

Connecting to Worker Nodes in Private Subnets Using SSH

Worker nodes in private subnets have private IP addresses only (they do not have public IP addresses). They can only be accessed by other resources inside the VCN. Oracle recommends using bastion hosts to control external access (such as SSH) to worker nodes in private subnets. A bastion host is in a public subnet, has a public IP address, and is accessible

from the internet. For more information about bastion hosts, see the white paper [Bastion Hosts: Protected Access for Virtual Cloud Networks](#).

About Access Control and Container Engine for Kubernetes

To perform operations on a Kubernetes cluster, you must have appropriate permissions to access the cluster.

For most operations on Kubernetes clusters created and managed by Container Engine for Kubernetes, Oracle Cloud Infrastructure Identity and Access Management (IAM) provides access control. A user's permissions to access clusters comes from the groups to which they belong. The permissions for a group are defined by policies. Policies define what actions members of a group can perform, and in which compartments. Users can then access clusters and perform operations based on the policies set for the groups they are members of.

IAM provides control over:

- whether a user can create or delete clusters
- whether a user can add, remove, or modify node pools
- which Kubernetes object create/delete/view operations a user can perform on all clusters within a compartment or tenancy

See [Policy Configuration for Cluster Creation and Deployment](#).

In addition to IAM, the Kubernetes RBAC Authorizer can enforce additional fine-grained access control for users on specific clusters via Kubernetes RBAC roles and clusterroles. A Kubernetes RBAC role is a collection of permissions. For example, a role might include read permission on pods and list permission for pods. A Kubernetes RBAC clusterrole is just like a role, but can be used anywhere in the cluster. A Kubernetes RBAC rolebinding maps a role to a user or set of users, granting that role's permissions to those users for resources in that namespace. Similarly, a Kubernetes RBAC clusterrolebinding maps a clusterrole to a user or set of users, granting that clusterrole's permissions to those users across the entire cluster.

IAM and the Kubernetes RBAC Authorizer work together to enable users who have been successfully authorized by at least one of them to complete the requested Kubernetes operation.

When a user attempts to perform any operation on a cluster (except for create role and create clusterrole operations), IAM first determines whether the group to which the user belongs has the appropriate and sufficient permissions. If so, the operation succeeds. If the attempted operation also requires additional permissions granted via a Kubernetes RBAC role or clusterrole, the Kubernetes RBAC Authorizer then determines whether the user has been granted the appropriate Kubernetes role or clusterrole.

Typically, you'll want to define your own Kubernetes RBAC roles and clusterroles when deploying a Kubernetes cluster to provide additional fine-grained control. When you attempt to perform a create role or create clusterrole operation, the Kubernetes RBAC Authorizer first determines whether you have sufficient Kubernetes privileges. To create a role or clusterrole, you must have been assigned an existing Kubernetes RBAC role (or clusterrole) that has at least the same or higher privileges as the new role (or clusterrole) you're attempting to create.

By default, users are not assigned any Kubernetes RBAC roles (or clusterroles) by default. So before attempting to create a new role (or clusterrole), you must be assigned an appropriately privileged role (or clusterrole). A number of such roles and clusterroles are always created by default, including the cluster-admin clusterrole (for a full list, see [Default Roles and Role Bindings](#) in the Kubernetes documentation). The cluster-admin clusterrole essentially confers super-user privileges. A user granted the cluster-admin clusterrole can perform any operation across all namespaces in a given cluster.

Note that Oracle Cloud Infrastructure tenancy administrators already have sufficient privileges, and do not require the cluster-admin clusterrole.

Example: Granting the Kubernetes RBAC cluster-admin clusterrole



Note



The following instructions assume:

- You have the required access to create Kubernetes RBAC roles and clusterroles, either because you're in the tenancy's Administrators group, or because you have the Kubernetes RBAC cluster-admin clusterrole.
- The user to which you want to grant the RBAC cluster-admin clusterrole is not an OCI tenancy administrator. If they are an OCI tenancy administrator, they do not require the Kubernetes RBAC cluster-admin clusterrole.

Follow these steps to grant a user who is not a tenancy administrator the Kubernetes RBAC cluster-admin clusterrole on a cluster deployed on Oracle Cloud Infrastructure:

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. In a terminal window, grant the Kubernetes RBAC cluster-admin clusterrole to the user by entering:

```
$ kubectl create clusterrolebinding <my-cluster-admin-binding> --clusterrole=cluster-admin --user=<user_OCID>
```

where:

- `<my-cluster-admin-binding>` is a string of your choice to be used as the name for the binding between the user and the Kubernetes RBAC cluster-admin clusterrole. For example, `jdoe_clst_admin`
- `<user_OCID>` is the user's OCID (obtained from the Console). For example, `ocid1.user.oc1..aaaaa...zutq` (abbreviated for readability).

For example:

```
$ kubectl create clusterrolebinding jdoe_clst_admin --clusterrole=cluster-admin --
user=ocidl.user.oc1..aaaaa...zutq
```

Example: Giving a developer user the ability to read pods in a new cluster



Note

The following instructions assume you're in the tenancy's Administrators group, and therefore have:

- the required permissions to create clusters, and to manage users and groups
- the required access to create Kubernetes RBAC roles and clusterroles

Follow these steps to give a developer the necessary Oracle Cloud Infrastructure and Kubernetes RBAC permissions to use `kubectl` to view pods running on a cluster deployed on Oracle Cloud Infrastructure:

1. Create a new Oracle Cloud Infrastructure user for the developer to use (for example, called `jdoe@acme.com`), and make a note of the new user's OCID (for example, `ocidl.user.oc1..aaaaa...tx5a`, abbreviated for readability). See [To create a user](#).
2. Create a new Oracle Cloud Infrastructure group and add the new user to the group (for example, called `acme-dev-pod-vwr`). See [To create a group](#).
3. Create a new Oracle Cloud Infrastructure policy that grants the new group the `CLUSTER_USE` permission on clusters, with a policy statement like:

```
Allow group acme-dev-pod-vwr to use clusters in <location>
```

In the above policy statement, replace `<location>` with either `tenancy` (if you are creating the policy in the tenancy's root compartment) or `compartment <compartment-`

name> (if you are creating the policy in an individual compartment).

See [To create a policy](#).

4. Create a new cluster in the Console. See [Creating a Kubernetes Cluster](#).
5. Follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
6. In a text editor, create a file (for example, called role-pod-reader.yaml) with the following content. This file defines a Kubernetes RBAC role that enables users to read pod details.

```
kind: Role
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  namespace: default
  name: pod-reader
rules:
- apiGroups: ["" ] # "" indicates the core API group
  resources: ["pods"]
  verbs: ["get", "watch", "list"]
```

7. In a terminal window, create the new role in the cluster using kubectl. For example, if you gave the yaml file that defines the new role the name role-pod-reader.yaml, enter the following:

```
$ kubectl create -f role-pod-reader.yaml
```

8. In a terminal window, bind the Kubernetes RBAC role you just created to the Oracle Cloud Infrastructure user account you created earlier by entering the following to create a new rolebinding (in this case, called pod-reader-binding):

```
$ kubectl create rolebinding pod-reader-binding --role=pod-reader --
user=ocidl.user.oc1..aaaaa...tx5a
```

9. Give the developer the credentials of the new Oracle Cloud Infrastructure user you created earlier, and tell the developer they can now see details of pods running on the cluster deployed on Oracle Cloud Infrastructure by:

- Signing in to the Console using the new user's credentials.
- Following the instructions in [Downloading a kubeconfig File to Enable Cluster Access](#) to obtain their own copy of the cluster's kubeconfig file. If the file does not have the expected default name and location of `$HOME/.kube/config`, the developer will also have to set the KUBECONFIG environment variable to point to the file. Note that the developer must download their own kubeconfig file. They cannot access a cluster using a kubeconfig file that you (or a different user) downloaded.
- Using `kubectl` to see details of the pods by entering:

```
$ kubectl get pods
```

Kubernetes Versions and Container Engine for Kubernetes

When you create a new Kubernetes cluster using Container Engine for Kubernetes, you specify:

- The version of Kubernetes to run on the master nodes in the cluster.
- The version of Kubernetes to run on the worker nodes in each node pool. All worker nodes in the same node pool run the same version of Kubernetes. Different node pools in a cluster can run different versions of Kubernetes.

The version of Kubernetes that you specify for the worker nodes in a node pool must be either the same Kubernetes version as that running on the master nodes, or an earlier Kubernetes version that is still compatible. In other words:

- The master nodes in a new cluster must run the same version of Kubernetes as the version running on worker nodes, or must be no more than two versions ahead.
- The worker nodes in a node pool must not run a more recent version of Kubernetes than the associated master nodes.

About Kubernetes Versions

New versions of Kubernetes are released periodically that contain new features and bug fixes.

Kubernetes version numbers have the format $x.y.z$ where x is a major release, y is a minor release, and z is a patch release. For example, 1.13.5.

Kubernetes itself is supported for three minor versions at a time (the current release version and two previous versions).

As described in the [Kubernetes documentation](#), a certain amount of version variation is permissible between master nodes and worker nodes in a cluster:

- The Kubernetes version on worker nodes can lag behind the version on the master nodes by up to two versions, but no more. If the version on the worker nodes is more than two versions behind the version on the master nodes, the Kubernetes versions on the worker nodes and the master nodes are incompatible.
- The Kubernetes version on worker nodes must never be more recent than the version on the master nodes.

About Upgrading Clusters to Newer Kubernetes Versions

After a new version of Kubernetes has been released and when Container Engine for Kubernetes supports the new version, you can use Container Engine for Kubernetes to upgrade master nodes running older versions of Kubernetes. Because Container Engine for Kubernetes distributes the Kubernetes Control Plane on multiple Oracle-managed master nodes (distributed across different availability domains in a region where supported) to ensure high availability, you're able to upgrade the Kubernetes version running on master nodes with zero downtime.

Having upgraded master nodes to a new version of Kubernetes, you can subsequently create new node pools running the newer version. Alternatively, you can continue to create new node pools that will run older versions of Kubernetes (providing those older versions are compatible with the Kubernetes version running on the master nodes).

Note that you upgrade master nodes by performing an 'in-place' upgrade, but you upgrade worker nodes by performing an 'out-of-place' upgrade. To upgrade the version of Kubernetes running on worker nodes in a node pool, you replace the original node pool with a new node pool that has new worker nodes running the appropriate Kubernetes version. Having 'drained'

existing worker nodes in the original node pool to prevent new pods starting and to delete existing pods, you can then delete the original node pool.

Also note the following:

- Container Engine for Kubernetes only upgrades the Kubernetes version running on master nodes when you explicitly initiate the upgrade operation.
- After upgrading master nodes to a newer version of Kubernetes, you cannot downgrade the master nodes to an earlier Kubernetes version.
- Before you upgrade the version of Kubernetes running on the master nodes, it is your responsibility to test that applications deployed on the cluster are compatible with the new Kubernetes version. For example, before upgrading the existing cluster, you might create a new separate cluster with the new Kubernetes version to test your applications.
- The versions of Kubernetes running on the master nodes and the worker nodes must be compatible (that is, the Kubernetes version on the master nodes must be no more than two minor versions ahead of the Kubernetes version on the worker nodes).
- If the version of Kubernetes currently running on the master nodes is more than one version behind the most recent supported version, you are given a choice of versions to upgrade to. If you want to upgrade to a version of Kubernetes that is more than one version ahead of the version currently running on the master nodes, you must upgrade to each intermediate version in sequence without skipping versions.

Kubernetes Versions Supported by Container Engine for Kubernetes

Container Engine for Kubernetes supports the following versions of Kubernetes:

Kubernetes Version	Supported by Container Engine for Kubernetes?	Notes
1.11.x and earlier	No	N/A
1.12.7	Yes	<p>Note that versions of Kubernetes 1.12.x prior to 1.12.7 (for example, 1.12.6) are no longer supported. You cannot:</p> <ul style="list-style-type: none"> • create clusters running earlier 1.12.x versions • add new node pools to existing clusters running earlier 1.12.x versions <p>If you do have clusters running 1.12.x versions earlier than 1.12.7, Oracle strongly recommends you upgrade those clusters to version 1.12.7.</p>
1.13.5	Yes	N/A

Upgrading the Version of Kubernetes Running on Master Nodes

When Container Engine for Kubernetes supports a newer version of Kubernetes than the version currently running on the master nodes in a cluster, you can upgrade the Kubernetes version running on the master nodes.

Important: After you've upgraded master nodes to a newer version of Kubernetes, you can't downgrade the master nodes to an earlier Kubernetes version. It's therefore important that

before you upgrade the version of Kubernetes running on the master nodes, you test that applications deployed on the cluster are compatible with the new Kubernetes version.

Using the Console

To upgrade the version of Kubernetes running on the master nodes:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster where you want to upgrade the version of Kubernetes running on the master nodes.
If a newer version of Kubernetes is available than the one running on the master nodes in the cluster, the **Upgrade Available** button is enabled at the top of the **Cluster** page.
4. Click **Upgrade Available** to upgrade the master nodes to a newer version.
5. In the **Upgrade Cluster Master** dialog box, select the version of Kubernetes to which to upgrade the master nodes, and click **Confirm**.

The version of Kubernetes running on the master nodes is upgraded. From now on, the new version of Kubernetes will appear as an option when you're defining new node pools for the cluster.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [UpdateCluster](#) operation to upgrade the version of Kubernetes running on the master nodes.

'Upgrading' the version of Kubernetes running on worker nodes by creating a new node pool

To 'upgrade' the version of Kubernetes running on worker nodes in a node pool, you replace the original node pool with a new node pool that has new worker nodes running the appropriate Kubernetes version. Having 'drained' existing worker nodes in the original node pool to prevent new pods starting and to delete existing pods, you can then delete the original node pool.

To 'upgrade' the version of Kubernetes on worker nodes by creating a new node pool:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster where you want to change the Kubernetes version running on worker nodes.
4. On the **Cluster** page, click **Add Node Pool** to create a new node pool and specify the required version of Kubernetes for its worker nodes.
The version of Kubernetes you specify must be compatible with the version that is running on the master nodes.
5. If there are labels attached to worker nodes in the original node pool and those labels are used by selectors (for example, to determine the nodes on which to run pods), then use the `kubectl label nodes` command to attach the same labels to the new worker nodes in the new node pool. See [Assigning Pods to Nodes](#) in the Kubernetes documentation.
6. For each worker node in the original node pool, prevent new pods from starting and delete existing pods by entering `kubectl drain <node_name>` for each worker node.
For more information:
 - about using `kubectl`, see [Accessing a Cluster Using kubectl](#)
 - about the `drain` command, see [drain](#) in the Kubernetes documentation

Recommended: Leverage pod disruption budgets as appropriate for your application to ensure that there's a sufficient number of replica pods running throughout the drain operation.

After all the worker nodes have been drained from the original node pool and pods are running on worker nodes in the new node pool, you can delete the original node pool.

7. On the **Cluster** page, display the **Node Pools** tab and select **Delete Node Pool** from the **Actions** menu.

The original node pool and all its worker nodes are deleted.

'Upgrading' the image running on worker nodes by creating a new node pool

When a node pool is created, the image to use for worker nodes in the node pool is specified. An image is a template of a virtual hard drive that determines the operating system and other software running on the worker nodes.

To 'upgrade' the image running on worker nodes in a node pool (for example, to upgrade to a new version of Oracle Linux), you replace the original node pool with a new node pool that has new worker nodes running the appropriate image. Having 'drained' existing worker nodes in the original node pool to prevent new pods starting and to delete existing pods, you can then delete the original node pool.

To 'upgrade' the image on worker nodes by creating a new node pool:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Container Clusters**.
2. Choose a **Compartment** you have permission to work in.
3. On the **Cluster List** page, click the name of the cluster where you want to change the image running on worker nodes.
4. On the **Cluster** page, click **Add Node Pool** to create a new node pool and specify the required image for its worker nodes.
5. If there are labels attached to worker nodes in the original node pool and those labels are used by selectors (for example, to determine the nodes on which to run pods), then

use the `kubectl label nodes` command to attach the same labels to the new worker nodes in the new node pool. See [Assigning Pods to Nodes](#) in the Kubernetes documentation.

6. For each worker node in the original node pool, prevent new pods from starting and delete existing pods by entering `kubectl drain <node_name>` for each worker node. For more information:

- about using `kubectl`, see [Accessing a Cluster Using kubectl](#)
- about the `drain` command, see [drain](#) in the Kubernetes documentation

Recommended: Leverage pod disruption budgets as appropriate for your application to ensure that there's a sufficient number of replica pods running throughout the drain operation.

After all the worker nodes have been drained from the original node pool and pods are running on worker nodes in the new node pool, you can delete the original node pool.

7. On the **Cluster** page, display the **Node Pools** tab and select **Delete Node Pool** from the **Actions** menu.

The original node pool and all its worker nodes are deleted.

Creating Load Balancers to Distribute Traffic Between Cluster Nodes

When you create a service, you can optionally create a load balancer to distribute service traffic among the nodes assigned to that service. The key fields in the configuration of a load balancer are the **type** of service being created and the **ports** that the load balancer will listen to.

Creating Load Balancers to Distribute HTTP Traffic

Consider the following configuration file, `nginx_lb.yaml`. It defines a deployment (`kind: Deployment`) for the `nginx` app, followed by a service definition with a type of `LoadBalancer` (`type: LoadBalancer`) that balances http traffic on port 80 for the `nginx` app.

```
apiVersion: apps/v1
kind: Deployment
```

CHAPTER 9 Container Engine for Kubernetes

```
metadata:
  name: my-nginx
  labels:
    app: nginx
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:1.7.9
        ports:
        - containerPort: 80
---
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
spec:
  type: LoadBalancer
  ports:
  - port: 80
  selector:
    app: nginx
```

The first part of the configuration file defines an Nginx deployment, requesting that it be hosted on 3 pods running the `nginx:1.7.9` image, and accept traffic to the containers on port 80.

The second part of the configuration file defines the Nginx service, which uses type `LoadBalancer` to balance Nginx traffic on port 80 amongst the available pods.

CHAPTER 9 Container Engine for Kubernetes

To create the deployment and service defined in `nginx_lb.yaml` while connected to your Kubernetes cluster, enter the command:

```
$ kubectl apply -f nginx_lb.yaml
```

This command outputs the following upon successful creation of the deployment and the load balancer:

```
deployment "my-nginx" created
service "my-nginx-svc" created
```

The load balancer may take a few minutes to go from a pending state to being fully operational. You can view the current state of your cluster by entering `kubectl get all`, where your output looks similar to the following:

```
$ kubectl get all
```

```
NAME                                                    READY   STATUS    RESTARTS   AGE
po/my-nginx-431080787-0m4m8                            1/1     Running   0           3m
po/my-nginx-431080787-hqqcr                             1/1     Running   0           3m
po/my-nginx-431080787-n8125                             1/1     Running   0           3m

NAME                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
svc/kubernetes      203.0.113.1     <NONE>           443/TCP          3d
svc/my-nginx-svc    203.0.113.7     192.0.2.22      80:30269/TCP    3m

NAME                DESIRED   CURRENT   UP-TO-DATE   AVAILABLE   AGE
deploy/my-nginx     3         3         3             3           3m

NAME                DESIRED   CURRENT   READY   AGE
rs/my-nginx-431080787 3         3         3       3m
```

The output shows that the `my-nginx` deployment is running on 3 pods (the `po/my-nginx` entries), that the load balancer is running (`svc/my-nginx-svc`) and has an external IP (192.0.2.22) that clients can use to connect to the app that's deployed on the pods.

Creating Load Balancers with SSL Support to Distribute HTTPS Traffic

You can create a load balancer with SSL termination, allowing `https` traffic to an app to be distributed among the nodes in a cluster. This example provides a walkthrough of the configuration and creation of a load balancer with SSL support.

CHAPTER 9 Container Engine for Kubernetes

Consider the following configuration file, `nginx-demo-svc-ssl.yaml`, which defines an Nginx deployment and exposes it via a load balancer that serves http on port 80, and https on port 443. This sample creates an Oracle Cloud Infrastructure load balancer, by defining a service with a type of `LoadBalancer` (`type: LoadBalancer`).

```
apiVersion: apps/v1beta1
kind: Deployment
metadata:
  name: nginx-deployment
spec:
  replicas: 2
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx
        ports:
        - containerPort: 80
---
kind: Service
apiVersion: v1
metadata:
  name: nginx-service
  annotations:
    service.beta.kubernetes.io/oci-load-balancer-ssl-ports: "443"
    service.beta.kubernetes.io/oci-load-balancer-tls-secret: ssl-certificate-secret
spec:
  selector:
    app: nginx
  type: LoadBalancer
  ports:
  - name: http
    port: 80
    targetPort: 80
  - name: https
    port: 443
    targetPort: 80
```

CHAPTER 9 Container Engine for Kubernetes

The Load Balancer's annotations are of particular importance. The ports on which to support https traffic are defined by the value of **oci-load-balancer-ssl-ports**. You can declare multiple SSL ports by using a comma-separated list for the annotation's value. For example, you could set the annotation's value to "443, 3000" to support SSL on ports 443 and 3000.

The required TLS secret, **ssl-certificate-secret**, needs to be created in Kubernetes. This example creates and uses a self-signed certificate. However, in a production environment, the most common scenario is to use a public certificate that's been signed by a certificate authority.

The following command creates a self-signed certificate, `tls.crt`, with its corresponding key, `tls.key`:

```
$ openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.key -out tls.crt -subj  
"/CN=nginxsvc/O=nginxsvc"
```

Now that you created the certificate, you need to store both it and its key as a secret in Kubernetes. The name of the secret must match the name from the **oci-load-balancer-tls-secret** annotation of the load balancer's definition. Use the following command to create a TLS secret in Kubernetes, whose key and certificate values are set by `--key` and `--cert`, respectively.

```
$ kubectl create secret tls ssl-certificate-secret --key tls.key --cert tls.crt
```

You must create the Kubernetes secret before you can create the service, since the service references the secret in its definition. Create the service using the following command:

```
$ kubectl create -f manifests/demo/nginx-demo-svc-ssl.yaml
```

Watch the service and wait for a public IP address (EXTERNAL-IP) to be assigned to the Nginx service (`nginx-service`). This is the load balancer IP to use to connect to the service.

```
$ kubectl get svc --watch
```

NAME	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
nginx-service	192.0.2.1	198.51.100.1	80:30274/TCP	5m

The load balancer is now running, which means the service can now be accessed using either http or https, as demonstrated by the following commands:

```
$ curl http://198.51.100.1
```

CHAPTER 9 Container Engine for Kubernetes

```
$ curl --insecure https://198.51.100.1
```

The "--insecure" flag is used to access the service using https due to the use of self-signed certificates in this example. Do not use this flag in a production environment where the public certificate was signed by a certificate authority.

Note: When a cluster is deleted, a load balancer that's dynamically created when a service is created will not be removed. Before deleting a cluster, delete the service, which in turn will result in the cloud provider removing the load balancer. The syntax for this command is:

```
$ kubectl delete svc SERVICE_NAME
```

For example, to delete the service from the previous example, enter:

```
$ kubectl delete svc nginx-service
```

Creating Internal Load Balancers in Public and Private Subnets

You can create Oracle Cloud Infrastructure load balancers to control access to services running on a cluster:

- When you create a 'custom' cluster, you select an existing VCN that contains the network resources to be used by the new cluster. If you want to use load balancers to control traffic into the VCN, you select existing public or private subnets in that VCN to host the load balancers.
- When you create a 'quick cluster', the VCN that's automatically created contains a public regional subnet to host a load balancer. If you want to host load balancers in private subnets, you can add private subnets to the VCN later.

Alternatively, you can create an internal load balancer service in a cluster to enable other programs running in the same VCN as the cluster to access services in the cluster. You can host internal load balancers in public subnets and private subnets.

To create an internal load balancer hosted on a public subnet, add the following annotation in the metadata section of the manifest file:

```
service.beta.kubernetes.io/oci-load-balancer-internal: "true"
```

CHAPTER 9 Container Engine for Kubernetes

To create an internal load balancer hosted on a private subnet, add both following annotations in the metadata section of the manifest file:

```
service.beta.kubernetes.io/oci-load-balancer-internal: "true"
service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocidl.subnet.oc1..aaaaa...vdfw"
```

where `ocidl.subnet.oc1..aaaaa...vdfw` is the OCID of the private subnet.

For example:

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    service.beta.kubernetes.io/oci-load-balancer-internal: "true"
    service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocidl.subnet.oc1..aaaaa...vdfw"
spec:
  type: LoadBalancer
  ports:
  - port: 8100
  selector:
    app: nginx
```

Specifying Alternative Load Balancer Shapes

The shape of an Oracle Cloud Infrastructure load balancer specifies its maximum total bandwidth (that is, ingress plus egress). By default, load balancers are created with a shape of 100Mbps. Other shapes are available, including 400Mbps and 8000Mbps.

To specify an alternative shape for a load balancer, add the following annotation in the metadata section of the manifest file:

```
service.beta.kubernetes.io/oci-load-balancer-shape: <value>
```

where `value` is the bandwidth of the shape (for example, 100Mbps, 400Mbps, 8000Mbps).

For example:

CHAPTER 9 Container Engine for Kubernetes

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps
spec:
  type: LoadBalancer
  ports:
  - port: 80
  selector:
    app: nginx
```

Note: Sufficient load balancer quota must be available in the region for the shape you specify. Enter the following kubectl command to confirm that load balancer creation did not fail due to lack of quota:

```
$ kubectl describe service <service-name>
```

Specifying Load Balancer Connection Timeout

You can specify the maximum idle time (in seconds) allowed between two successive receive or two successive send operations between the client and backend servers.

To explicitly specify a maximum idle time, add the following annotation in the metadata section of the manifest file:

```
oci-load-balancer-connection-idle-timeout: <value>
```

where `value` is the number of seconds.

For example:

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
```

CHAPTER 9 Container Engine for Kubernetes

```
oci-load-balancer-connection-idle-timeout: 100
spec:
  type: LoadBalancer
  ports:
  - port: 80
  selector:
    app: nginx
```

Note that if you don't explicitly specify a maximum idle time, a default value is used. The default value depends on the type of listener:

- for TCP listeners, the default maximum idle time is 300 seconds
- for HTTP listeners, the default maximum idle time is 60 seconds

Specifying Load Balancer Security List Management Options

You can specify how security lists are managed.

To explicitly specify a security list management mode, add the following annotation in the metadata section of the manifest file:

```
oci-load-balancer-security-list-management-mode: <value>
```

where <value> is one of:

- "All": All required security list rules for load balancer services are managed.
- "Frontend": Only security list rules for ingress to load balancer services are managed. You have to set up a rule that allows inbound traffic to the appropriate ports for node port ranges, the kube-proxy health port, and the health check port ranges.
- "None": No security list management is enabled. You have to set up a rule that allows inbound traffic to the appropriate ports for node port ranges, the kube-proxy health port, and the health check port ranges. Additionally, you have to set up rules to allow inbound traffic to load balancers.

For example:

```
apiVersion: v1
kind: Service
metadata:
```

CHAPTER 9 Container Engine for Kubernetes

```
name: my-nginx-svc
labels:
  app: nginx
annotations:
  oci-load-balancer-security-list-management-mode: "Frontend"
spec:
  type: LoadBalancer
  ports:
  - port: 80
  selector:
    app: nginx
```

Note that if you specify an invalid value for `oci-load-balancer-security-list-management-mode`, the value "All" is used instead.

Specifying Load Balancer Listener Protocol

You can define the type of traffic accepted by the load balancer listener by specifying the protocol on which the listener accepts connection requests.

To explicitly specify the load balancer listener protocol, add the following annotation in the metadata section of the manifest file:

```
oci-load-balancer-backend-protocol: <value>
```

where `<value>` is the protocol that defines the type of traffic accepted by the listener. For example, "HTTP". To get a list of valid protocols, use the [ListProtocols](#) operation.

For example:

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    oci-load-balancer-backend-protocol: "HTTP"
spec:
  type: LoadBalancer
  ports:
  - port: 80
```

```
selector:  
  app: nginx
```

Note that if you don't explicitly specify a protocol, "TCP" is used as the default value.

Creating a Persistent Volume Claim

Container storage via a container's root file system is ephemeral, and can disappear upon container deletion and creation. To provide a durable location to store data and prevent it from being lost, you can create and use persistent volumes to store data outside of containers.

You can define and apply a persistent volume claim to your cluster, which in turn creates a persistent volume that's bound to the claim. A claim is a block storage volume in the underlying IaaS provider that's durable and offers persistent storage, enabling your data to remain intact, regardless of whether the containers that the storage is connected to are terminated.

With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Storage service.

A persistent volume claim (PVC) is a request for storage, similar to how a pod requests compute resources. A PVC provides an abstraction layer to underlying storage. For example, an administrator could create a number of static persistent volumes (PVs) that can later be bound to one or more persistent volume claims. This is analogous to an administrator creating cluster nodes to which pods are later assigned. If none of the static persistent volumes match the user's PVC request, the cluster may attempt to dynamically create a PV that matches the PVC request. This example uses the latter approach, and it assumes that the cluster administrator has not created any suitable PVs that match the PVC request—meaning that the PVCs will dynamically create the PVs for this example.

The minimum amount of persistent storage that a PVC can request is 50 gigabytes. If the request is for less than 50 gigabytes, the request is rounded up to 50 gigabytes.

CHAPTER 9 Container Engine for Kubernetes

The following YAML defines two PVCs that each request 50 gigabytes of persistent storage (storage: 50Gi). You use names of the PVCs (for example, mysqlclaim) when defining which claims to use as the volumes of a deployment.

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysqlclaim
spec:
  storageClassName: "oci"
  selector:
    matchLabels:
      failure-domain.beta.kubernetes.io/zone: "US-ASHBURN-AD-1"
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 50Gi
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wordpressclaim
spec:
  storageClassName: "oci"
  selector:
    matchLabels:
      failure-domain.beta.kubernetes.io/zone: "US-ASHBURN-AD-2"
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 50Gi
```



Note

In the previous example, the PVCs request storage in availability domains in the Ashburn region using `matchLabels:failure-domain.beta.kubernetes.io/zone`. Note that when you specify values for `matchLabels:failure-domain.beta.kubernetes.io/zone`, you must use the shortened versions of availability domain names. For example `US-ASHBURN-AD-1`, `US-ASHBURN-AD-2`. See [Availability by Region Name and Region Code](#) for a list of shortened versions of availability domain names.

Enter the following command to create the PVC from the YAML file:

```
$ kubectl create -f https://raw.githubusercontent.com/wercker/oke_examples/master/kubernetes_
examples/persistent_volume_claims.yaml

persistentvolumeclaim "mysqlclaim" created
persistentvolumeclaim "wordpressclaim" created
```

You can verify that the PVCs have been created and bound to persistent volumes by calling `kubectl get pvc`:

```
$ kubectl get pvc
```

NAME	STATUS	VOLUME	CAPACITY	ACCESSMODES	STORAGECLASS	AGE
mysqlclaim	Bound	abyhq1jrerxpanjto7b5z1xjesy4aedghc5c52f5v43xcrymo77ktdl6ibjq	50Gi			
		oci				4m
wordpressclaim	Bound	abyhq1jt3rzldcclootxn7yrfgv36s7rnggcobennjohevykqpitzkinspka	50Gi			
		oci				4m

CHAPTER 9 Container Engine for Kubernetes

You can use these persistent volumes when creating other objects, such as deployments. For example, the following deployment definition instructs the system to use the `mysqlclaim` PVC as the `mysql-persistent-storage` volume, which is mounted by pods hosting the deployment as `/var/lib/mysql`.

```
#MySQL Deployment
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: mysql
  labels:
    app: mysql
spec:
  replicas: 1
  selector:
    matchLabels:
      app: mysql
  template:
    metadata:
      labels:
        app: mysql
    spec:
      containers:
        - image: mysql:5.6
          name: mysql
          env:
            - name: MYSQL_ROOT_PASSWORD
              valueFrom:
                secretKeyRef:
                  name: mysql
                  key: password
          ports:
            - containerPort: 3306
              name: mysql
          volumeMounts:
            - name: mysql-persistent-storage
              mountPath: /var/lib/mysql
      volumes:
        - name: mysql-persistent-storage
          persistentVolumeClaim:
            claimName: mysqlclaim
```

Adding OCI Service Broker for Kubernetes to Clusters

Service brokers offer a catalog of backing services to workloads running on cloud native platforms. The Open Service Broker API is a commonly-used standard for interactions between service brokers and platforms. The Open Service Broker API specification describes a simple set of API endpoints that platforms use to provision, gain access to, and manage service offerings. For more information about the Open Service Broker API, see resources available online including those at openservicebrokerapi.org.

OCI Service Broker for Kubernetes is an implementation of the Open Service Broker API. OCI Service Broker for Kubernetes is specifically for interacting with Oracle Cloud Infrastructure services from Kubernetes clusters. It includes three service broker adapters to bind to the following Oracle Cloud Infrastructure services:

- Object Storage
- Autonomous Transaction Processing
- Autonomous Data Warehouse

You can add OCI Service Broker for Kubernetes to clusters you've created with Oracle Cloud Infrastructure Container Engine for Kubernetes to interact with the Oracle Cloud Infrastructure services listed above. Having added OCI Service Broker for Kubernetes to a cluster, you don't have to manually provision and de-provision the Oracle Cloud Infrastructure services each time you deploy or un-deploy an application on the cluster. Instead, you interact with the Oracle Cloud Infrastructure services by using `kubectl` to call the Open Service Broker APIs implemented by OCI Service Broker for Kubernetes .

OCI Service Broker for Kubernetes is available as a Helm chart, a Docker container, and as source code from [Github](#).

For more information about OCI Service Broker for Kubernetes, see the OCI Service Broker for Kubernetes documentation in the [Github repository](#).

Adding OCI Service Broker for Kubernetes to a Cluster

To add OCI Service Broker for Kubernetes to a cluster, follow the detailed instructions in the [Github repository](#).

For convenience, here's a high-level summary of the steps involved:

1. Install OCI Service Broker for Kubernetes. During this step, you will typically:
 - Install the Service Catalog.
 - Install the svcat tool.
 - Deploy OCI Service Broker for Kubernetes.
 - Grant RBAC permissions and roles.
 - Register OCI Service Broker for Kubernetes.

For more information about installation, see the OCI Service Broker for Kubernetes documentation in the [Github repository](#).

2. Secure OCI Service Broker for Kubernetes. During this step, you will typically:
 - Restrict access to Service Catalog resources using RBAC permissions and roles.
 - Configure TLS for OCI Service Broker for Kubernetes.
 - Set up an Oracle Cloud Infrastructure user for use by OCI Service Broker for Kubernetes.
 - Set up appropriate policies to control access to resources (according to the Oracle Cloud Infrastructure services to be used).
 - Limit access to the OCI Service Broker for Kubernetes endpoint using NetworkPolicy.
 - Stand up an etcd cluster for Service Catalog and OCI Service Broker for Kubernetes.
 - Protect sensitive values by creating secrets.

The security configuration to choose will depend on your particular requirements. For more information, see the OCI Service Broker for Kubernetes documentation in the [Github repository](#).

3. Provision and bind to the required Oracle Cloud Infrastructure services. During this step, you will typically:
 - Provide service provision request parameters.
 - Provide service binding request parameters.

- Provide service binding response credentials.

The details to provide will depend on the Oracle Cloud Infrastructure service to bind to. For more information, see the OCI Service Broker for Kubernetes documentation in the [Github repository](#).

Example: Setting Up an Ingress Controller on a Cluster

You can set up different open source ingress controllers on clusters you have created with Container Engine for Kubernetes.

This topic explains how to set up an example ingress controller along with corresponding access control on an existing cluster. Having set up the ingress controller, this topic describes how to use the ingress controller with an example hello-world backend, and how to verify the ingress controller is working as expected.

Example Components

The example includes an ingress controller and a hello-world backend.

Ingress Controller Components

The ingress controller comprises:

- An ingress controller deployment called `nginx-ingress-controller`. The deployment deploys an image that contains the binary for the ingress controller and Nginx. The binary manipulates and reloads the `/etc/nginx/nginx.conf` configuration file when an ingress is created in Kubernetes. Nginx upstreams point to services that match specified selectors.
- An ingress controller service called `ingress-nginx`. The service exposes the ingress controller deployment as a `LoadBalancer` type service. Because Container Engine for Kubernetes uses an Oracle Cloud Infrastructure `integration/cloud-provider`, a load

balancer will be dynamically created with the correct nodes configured as a backend set.

Backend Components

The hello-world backend comprises:

- A backend deployment called `docker-hello-world`. The deployment handles default routes for health checks and 404 responses. This is done by using a stock hello-world image that serves the minimum required routes for a default backend.
- A backend service called `docker-hello-world-svc`. The service exposes the backend deployment for consumption by the ingress controller deployment.

Setting Up the Example Ingress Controller

In this section, you create the access rules for ingress. You then create the example ingress controller components, and confirm they are running.

Creating the Access Rules for the Ingress Controller

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the KUBECONFIG environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. If your Oracle Cloud Infrastructure user is a tenancy administrator, skip the next step and go straight to [Creating the Service Account, and the Ingress Controller](#).
3. If your Oracle Cloud Infrastructure user is not a tenancy administrator, in a terminal window, grant the user the Kubernetes RBAC cluster-admin clusterrole on the cluster by entering:

```
$ kubectl create clusterrolebinding <my-cluster-admin-binding> --clusterrole=cluster-admin --user=<user-OCID>
```

where:

CHAPTER 9 Container Engine for Kubernetes

- `<my-cluster-admin-binding>` is a string of your choice to be used as the name for the binding between the user and the Kubernetes RBAC cluster-admin clusterrole. For example, `jdoue_clst_adm`
- `<user-OCID>` is the user's OCID (obtained from the Console). For example, `ocid1.user.oc1..aaaaa...zutq` (abbreviated for readability).

For example:

```
$ kubectl create clusterrolebinding jdoue_clst_adm --clusterrole=cluster-admin --
user=ocid1.user.oc1..aaaaa...zutq
```

Creating the Service Account, and the Ingress Controller

1. Run the following command to create the `nginx-ingress-controller` ingress controller deployment, along with the Kubernetes RBAC roles and bindings:

```
$ kubectl apply -f https://raw.githubusercontent.com/kubernetes/ingress-
nginx/master/deploy/static/mandatory.yaml
```

2. Create and save the file `cloud-generic.yaml` containing the following code to define the `ingress-nginx` ingress controller service as a load balancer service:

```
kind: Service
apiVersion: v1
metadata:
  name: ingress-nginx
  namespace: ingress-nginx
  labels:
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/part-of: ingress-nginx
spec:
  type: LoadBalancer
  selector:
    app.kubernetes.io/name: ingress-nginx
    app.kubernetes.io/part-of: ingress-nginx
  ports:
    - name: http
      port: 80
      targetPort: http
    - name: https
```

CHAPTER 9 Container Engine for Kubernetes

```
port: 443
targetPort: https
```

3. Using the file you just saved, create the `ingress-nginx` ingress controller service by running the following command:

```
$ kubectl apply -f cloud-generic.yaml
```

Verifying the `ingress-nginx` Ingress Controller Service is Running as a Load Balancer Service

1. View the list of running services:

```
$ kubectl get svc -n ingress-nginx
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
ingress-nginx	LoadBalancer	10.96.229.38	<pending>	80:30756/TCP,443:30118/TCP	1h

The `EXTERNAL-IP` for the `ingress-nginx` ingress controller service is shown as `<pending>` until the load balancer has been fully created in Oracle Cloud Infrastructure.

2. Repeat the `kubectl get svc` command until an `EXTERNAL-IP` is shown for the `ingress-nginx` ingress controller service:

```
$ kubectl get svc -n ingress-nginx
```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
ingress-nginx	LoadBalancer	10.96.229.38	129.146.214.219	80:30756/TCP,443:30118/TCP	1h

Creating the TLS Secret

A TLS secret is used for SSL termination on the ingress controller. To generate the secret for this example, a self-signed certificate is used. While this is okay for testing, for production, use a certificate signed by a Certificate Authority.

```
$ openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.key -out tls.crt -subj  
"/CN=nginxsvc/O=nginxsvc"
```

```
$ kubectl create secret tls tls-secret --key tls.key --cert tls.crt
```



Note

Under Windows, you may need to replace `"/CN=nginxsvc/O=nginxsvc"` with `"/CN=nginxsvc\O=nginxsvc"`. For example, this is necessary if you run the `openssl` command from a Git Bash shell.

Setting Up the Example Backend

In this section, you define a hello-world backend service and deployment.

Creating the docker-hello-world Service Definition

1. Create the file `hello-world-ingress.yaml` containing the following code. This code uses a publicly available hello-world image from Docker Hub. You can substitute another image of your choice that can be run in a similar manner.

```
apiVersion: apps/v1beta1
kind: Deployment
metadata:
  name: docker-hello-world
  labels:
    app: docker-hello-world
spec:
  replicas: 3
  template:
    metadata:
      labels:
        app: docker-hello-world
    spec:
      containers:
      - name: docker-hello-world
        image: scottsbaldwin/docker-hello-world:latest
        ports:
```

CHAPTER 9 Container Engine for Kubernetes

```
    - containerPort: 80
---
apiVersion: v1
kind: Service
metadata:
  name: docker-hello-world-svc
spec:
  selector:
    app: docker-hello-world
  ports:
    - port: 8088
      targetPort: 80
  type: ClusterIP
```

Note the `docker-hello-world` service's type is `ClusterIP`, rather than `LoadBalancer`, because this service will be proxied by the `ingress-nginx` ingress controller service. The `docker-hello-world` service does not need public access directly to it. Instead, the public access will be routed from the load balancer to the ingress controller, and from the ingress controller to the upstream service.

2. Create the new hello-world deployment and service on nodes in the cluster by running the following command:

```
$ kubectl create -f hello-world-ingress.yaml
```

Using the Example Ingress Controller to Access the Example Backend

In this section you create an ingress to access the backend using the ingress controller.

Creating the Ingress Resource

1. Create the file `ingress.yaml` and populate it with this code:

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: hello-world-ing
  annotations:
    kubernetes.io/ingress.class: "nginx"
```

CHAPTER 9 Container Engine for Kubernetes

```
spec:
  tls:
  - secretName: tls-secret
  rules:
  - http:
    paths:
    - backend:
        serviceName: docker-hello-world-svc
        servicePort: 8088
```

2. Create the resource:

```
$ kubectl create -f ingress.yaml
```

Verifying that the Example Components are Working as Expected

In this section, you confirm that all of the example components have been successfully created and are operating as expected. The `docker-hello-world-svc` service should be running as a ClusterIP service, and the `ingress-nginx` service should be running as a LoadBalancer service. Requests sent to the ingress controller should be routed to nodes in the cluster.

Obtaining the External IP Address of the Load Balancer

To confirm the `ingress-nginx` service is running as a LoadBalancer service, obtain its external IP address:

```
$ kubectl get svc --all-namespaces
```

NAMESPACE	NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
default	docker-hello-world-svc	ClusterIP	10.96.83.247	<none>	8088/TCP
default	kubernetes	ClusterIP	10.96.0.1	<none>	443/TCP
ingress-nginx	ingress-nginx	LoadBalancer	10.96.229.38	129.146.214.219	80:30756/TCP,443:30118/TCP
kube-system	kube-dns	ClusterIP	10.96.5.5	<none>	53/UDP,53/TCP
kube-system	kubernetes-dashboard	ClusterIP	10.96.208.64	<none>	443/TCP

1h					
kube-system	tiller-deploy	ClusterIP	10.96.28.102	<none>	44134/TCP
1h					

Sending cURL Requests to the Load Balancer

1. Use the external IP address of the `ingress-nginx` service (for example, `129.146.214.219`) to curl an http request:

```
$ curl -I http://129.146.214.219

HTTP/1.1 301 Moved Permanently
Via: 1.1 10.68.69.10 (McAfee Web Gateway 7.6.2.10.0.23236)
Date: Thu, 07 Sep 2017 15:20:16 GMT
Server: nginx/1.13.2
Location: https://129.146.214.219/
Content-Type: text/html
Content-Length: 185
Proxy-Connection: Keep-Alive
Strict-Transport-Security: max-age=15724800; includeSubDomains;
```

The output shows a 301 redirect and a Location header that suggest that http traffic is being redirected to https.

2. Either cURL against the https url or add the `-L` option to automatically follow the location header. The `-k` option instructs cURL to not verify the SSL certificates.

```
$ curl -ikL http://129.146.214.219

HTTP/1.1 301 Moved Permanently
Via: 1.1 10.68.69.10 (McAfee Web Gateway 7.6.2.10.0.23236)
Date: Thu, 07 Sep 2017 15:22:29 GMT
Server: nginx/1.13.2
Location: https://129.146.214.219/
Content-Type: text/html
Content-Length: 185
Proxy-Connection: Keep-Alive
Strict-Transport-Security: max-age=15724800; includeSubDomains;

HTTP/1.0 200 Connection established
```

CHAPTER 9 Container Engine for Kubernetes

```
HTTP/1.1 200 OK
Server: nginx/1.13.2
Date: Thu, 07 Sep 2017 15:22:30 GMT
Content-Type: text/html
Content-Length: 71
Connection: keep-alive
Last-Modified: Thu, 07 Sep 2017 15:17:24 GMT
ETag: "59b16304-47"
Accept-Ranges: bytes
Strict-Transport-Security: max-age=15724800; includeSubDomains;

<h1>Hello webhook world from: docker-hello-world-1732906117-0ztkm</h1>
```

The last line of the output shows the HTML that is returned from the pod whose hostname is `docker-hello-world-1732906117-0ztkm`.

3. Issue the cURL request several times to see the hostname in the HTML output change, demonstrating that load balancing is occurring:

```
$ curl -k https://129.146.214.219

<h1>Hello webhook world from: docker-hello-world-1732906117-61151</h1>

$ curl -k https://129.146.214.219

<h1>Hello webhook world from: docker-hello-world-1732906117-7r89v</h1>

$ curl -k https://129.146.214.219

<h1>Hello webhook world from: docker-hello-world-1732906117-0ztkm</h1>
```

Inspecting `nginx.conf`

The `nginx-ingress-controller` ingress controller deployment manipulates the `nginx.conf` file in the pod within which it is running.

1. Find the name of the pod running the `nginx-ingress-controller` ingress controller deployment and use it with a `kubectl exec` command to show the contents of `nginx.conf`.

CHAPTER 9 Container Engine for Kubernetes

```
$ kubectl get po -n ingress-nginx

NAME                                READY   STATUS    RESTARTS   AGE
nginx-ingress-controller-110676328-h86xg  1/1     Running   0          1h

$ kubectl exec -n ingress-nginx -it nginx-ingress-controller-110676328-h86xg -- cat
/etc/nginx/nginx.conf
```

2. Look for `proxy_pass` in the output. There will be one for the default backend and another that looks similar to:

```
proxy_pass http://upstream_balancer;
```

This shows that Nginx is proxying requests to an upstream called `upstream_balancer`.

3. Locate the upstream definition in the output. It will look similar to:

```
upstream upstream_balancer {
    server 0.0.0.1:1234; # placeholder

    balancer_by_lua_block {
        tcp_udp_balancer.balance()
    }
}
```

The upstream is proxying via Lua.

CHAPTER 10 Data Transfer

This chapter explains how to migrate data to Oracle Cloud Infrastructure using Disk-Based Data Transfer and Data Transfer Appliance.

Overview of Data Transfer Service

Oracle offers offline data transfer solutions that let you migrate data to Oracle Cloud Infrastructure. Moving data over the public internet is not always feasible because of high network costs, unreliable network connectivity, long transfer times, and security concerns. Our transfer solutions address these pain points, are easy to use, and provide faster data upload compared to over-the-wire data transfer.



Note

Data Transfer is not available in Oracle Cloud Infrastructure Government Cloud realms.

Data transfer is supported in US East (Ashburn), US West (Phoenix), Germany Central (Frankfurt), and UK South (London) regions.



Note

To simplify this Data Transfer documentation, we generically refer to Object Storage to mean that you can transfer data into a bucket in either the Object Storage tier or Archive Storage tier.

DISK-BASED DATA TRANSFER

You send your data as files on encrypted commodity disk to an Oracle transfer site. Operators at the Oracle transfer site upload the files into your designated Object Storage bucket in your tenancy.

This transfer solution requires you to source and purchase the disks used to transfer data to Oracle Cloud Infrastructure. The disks are shipped back to you after the data is successfully uploaded.

See [Disk Data Transfer](#) for details.

APPLIANCE-BASED DATA TRANSFER

You send your data as files on secure, high-capacity, Oracle-supplied storage appliances to an Oracle transfer site. Operators at the Oracle transfer site upload the data into your designated Object Storage bucket in your tenancy.

This solution supports data transfer when you are migrating a large volume of data and when using disks is not a practical alternative. You do not need to write any code or purchase any hardware—Oracle supplies the transfer appliance and software required to manage the transfer.

See [Appliance Data Transfer](#) for details.

Limits on Data Transfer Service Resources

When you sign up for Oracle Cloud Infrastructure, a set of service limits is configured for your tenancy. The service limit is the quota or allowance set on a resource. Verify that your service limits are set appropriately before you begin the data transfer process.

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the wanted tags. For general information about applying tags, see Resource Tags.

Data Transfer currently supports applying tags to transfer jobs from the command line (Data Transfer Utility or CLSs). Tagging is not supported using the Console.

What's Next

Now you are ready to prepare for your data transfer. See the following pages for more information on each of the data transfer methods:

- [Disk Data Transfer](#)
- [Appliance Data Transfer](#)

Disk Data Transfer

Disk-Based Data Transfer is one of Oracle's offline data transfer solutions that lets you migrate data to Oracle Cloud Infrastructure. You send your data as files on encrypted disks to an Oracle transfer site. Operators at the Oracle transfer site upload the files into the designated Object Storage bucket in your tenancy. You are then free to move the uploaded data to other Oracle Cloud Infrastructure services as needed.

Disk-Based Data Transfer Concepts

The following concepts are essential to understanding Disk-Based Data Transfer.

DISK

A disk is a user-supplied storage device that is specially prepared to copy and upload data to Oracle Cloud Infrastructure. You copy your data to one or more of these disks and ship

the disks in a parcel to Oracle to upload your data.

The following transfer disks are supported:

- SATA II/III 2.5" or 3.5" hard disk drives
- External USB 2.0/3.0 hard disk drives



Note

Pin-code protected devices and physical-key protected devices are currently not supported.

TRANSFER DISK

A transfer disk is the logical representation of a disk that has been prepared to copy and upload data to Oracle Cloud Infrastructure.

TRANSFER JOB

A transfer job is the logical representation of a data migration to Oracle Cloud Infrastructure. A transfer job consists of one or more transfer packages that each contain one or more transfer disks.

DATA TRANSFER UTILITY

The Data Transfer Utility is the command line software that Oracle provides for you to prepare transfer disks for your data and for shipment to Oracle. In addition, you can use this software to manage transfer jobs and packages.

DATA HOST

The host computer on your site that stores the data you intend to copy to the disk for migration to Oracle Cloud Infrastructure.

TRANSFER PACKAGE

A transfer package is the logical representation of the parcel containing the transfer disks that you ship to Oracle to upload to Oracle Cloud Infrastructure.

BUCKET

The logical container in Oracle Cloud Infrastructure Object Storage where Oracle operators upload your data. A bucket is associated with a single compartment in your tenancy that has policies that determine what actions a user can perform on a bucket and on all the objects in the bucket.

DATA TRANSFER ADMINISTRATOR

A new or existing IAM user that has the authorization and permissions to create and manage transfer jobs. See .

DATA TRANSFER UPLOAD USER

A temporary IAM user that grants Oracle personnel the authorization and permissions to upload the data from your transfer disks to your designated Oracle Cloud Infrastructure Object Storage bucket. Delete this temporary user after your data is uploaded to Oracle Cloud Infrastructure. See .

Roles and Responsibilities

Depending on your organization, the responsibilities of using and managing the data transfer may span multiple roles. Use the following set of roles as a guideline for how you can assign the various tasks associated with the data transfer.

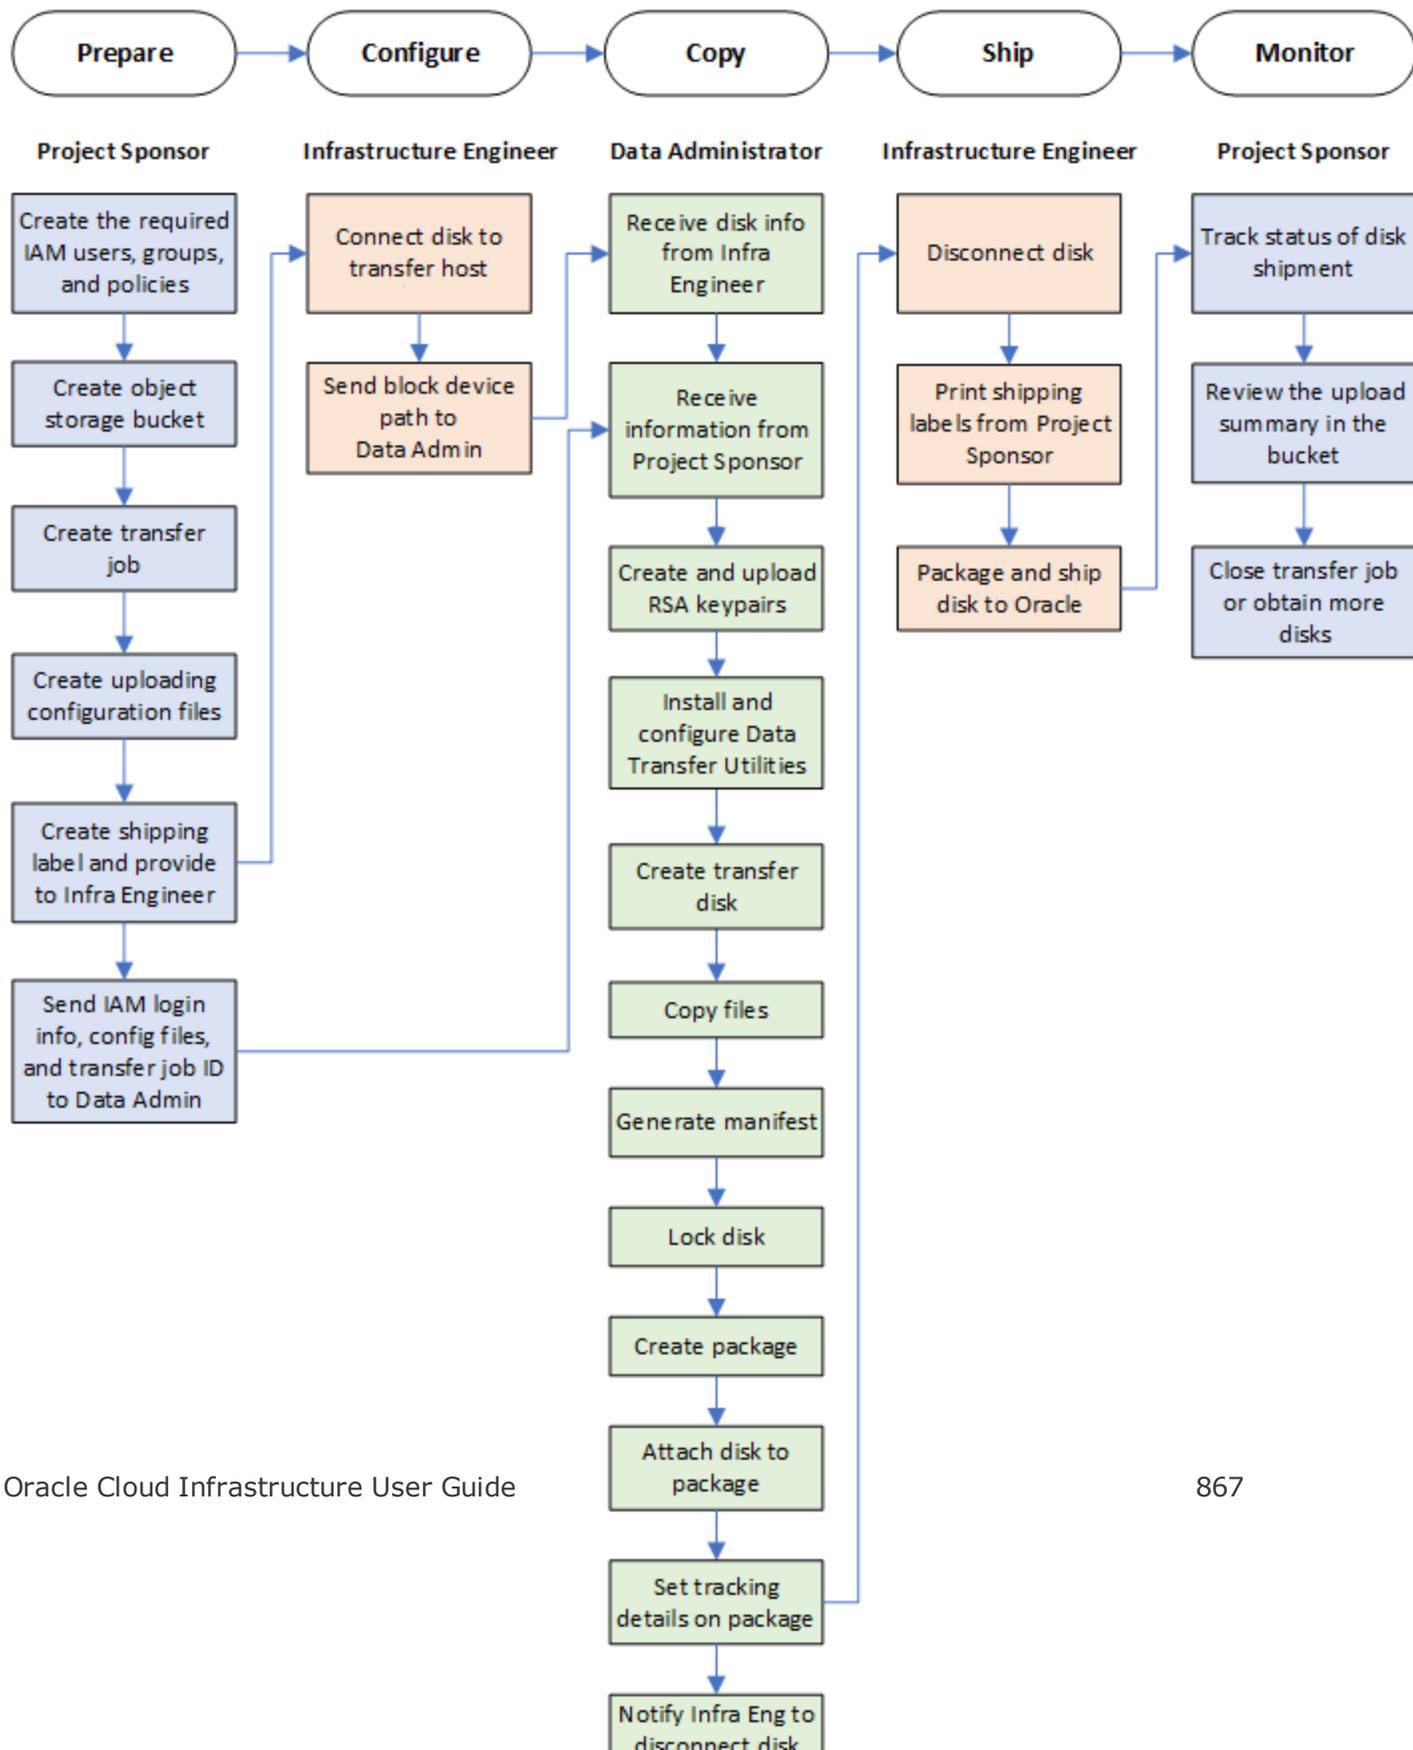
- **Project Sponsor:** Responsible for the overall success of the data transfer. Project Sponsors usually have complete access to their organization's Oracle Cloud Infrastructure tenancy. They coordinate with the other roles in the organization to complete the implementation of data transfer project.
- **Infrastructure Engineer:** Responsible for integrating the transfer appliance into the organization's IT infrastructure from where the data is being transferred. Tasks associated with this role include connecting the transfer appliance to power, placing it within the network, and setting the IP address through a serial console menu using the provided USB-to-Serial adapter.
- **Data Administrator:** Responsible for identifying and preparing the data to be

transferred to Oracle Cloud Infrastructure. This person usually has access to, and expertise with, the data being migrated.

These roles correspond to the various phases of the data transfer described in the following section. A specific role can be responsible for one or more phases.

Task Flow for Disk-Based Data Transfer

Here is a high-level overview of the tasks involved in transferring data to Oracle Cloud Infrastructure using Data Transfer Disk organized by phase. Complete one phase before proceeding to the next one. You can click some of the boxes to get details on how to perform the associated task. Use the roles previously described to distribute the tasks across individuals or groups within your organization.



Secure Disk Data Transfer to Oracle Cloud Infrastructure

This section highlights the security details of the Data Transfer Service process.

- The Data Transfer Utility uses the standard Linux dm-crypt and LUKS utilities to encrypt block devices.
- The dm-crypt software generates a master AES-256 bit encryption key that is used for all data written to or read from the disk. That key is protected by an encryption passphrase that the user must know to access the encrypted data.
- When the data transfer administrator uses the Data Transfer Utility to create disks, Oracle Cloud Infrastructure creates a strong encryption passphrase that is displayed to the user and passed to dm-crypt. The passphrase is displayed to standard output only once and cannot be retrieved again. Copy this passphrase to a durable, secure location for future reference.
- For extra security, you can also encrypt your own data with your own encryption keys. Before copying your data to the transfer disk, you can encrypt your data with a tool and encryption key of your choosing. After the data has been uploaded, you would need to use the same tool and encryption key to access the data.
- All network communication between the Data Transfer Utility and Oracle Cloud Infrastructure is encrypted in-transit using Transport Layer Security (TLS).
- After copying your data to a transfer disk, generate a manifest file using the Data Transfer Utility. The manifest contains an index of all of the copied files and generated data integrity hashes. The Data Transfer Utility copies the `config_upload_user` configuration file and referenced IAM credentials to the encrypted transfer disk. This configuration file describes the temporary IAM data transfer upload user. Oracle uses the credentials and entries defined in the `config_upload_user` file when processing the transfer disk and uploading files to Oracle Cloud Infrastructure Object Storage.



Note

Data Transfer Service Does Not Support Passphrases on Private Keys

While we recommend encrypting a private key with a passphrase when generating API signing keys, Data Transfer does not support passphrases on the key file required for the `config_upload_user`. If you use a passphrase, Oracle personnel cannot upload your data.

Oracle cannot upload data from a transfer disk without the correct credentials defined in this configuration file. See [Data Transfer Utility](#) for more information about the required configuration files.

- When you disconnect or lock a transfer disk using the Data Transfer Utility, the original encryption passphrase is required to once again access the disk. If the encryption passphrase is not known or lost, you cannot access the data on the transfer disk. To reuse a transfer disk, you must reformat the disk. Reformatting a disk removes all the data.
- Oracle retrieves the encryption passphrase for a transfer disk from Oracle Cloud Infrastructure. Oracle uses the passphrase to decrypt, mount the transfer disk, and upload the data to the designated bucket in the tenancy.
- After processing a transfer package, Oracle returns all transfer disks attached to the transfer package using the return shipping label you provide.
- To protect your data, we make the data on the disk unrecoverable before shipping the transfer disks back to you. To comply with customs regulations, we wipe the disks completely before shipping the transfer disks back to international shipping addresses.

Ways to Manage Disk Data Transfers

We provide two ways to manage disk-based data transfers:

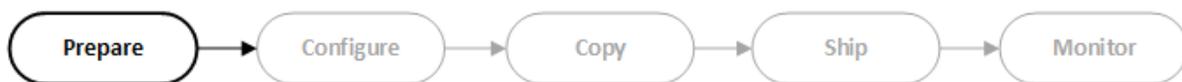
- The Data Transfer Utility is a full-featured command line tool for disk-based data transfers only (appliance-based data transfers use a different command line tool). For more information and installation instructions, see [Data Transfer Utility](#).
- The Console is an easy-to-use, partial-featured browser-based interface. For more information, see [Signing In to the Console](#).



Note

You can perform many data transfer tasks using either the Console or the Data Transfer Utility. However, there are some tasks you can **only** perform using the Data Transfer Utility (for example, creating and locking transfer disks). describes the management tasks in detail and guides you to the appropriate management interface to use for each task.

Preparing for Disk Data Transfers



This topic describes the tasks associated with preparing for the Disk-Based Data Transfer. The Project Sponsor role typically performs these tasks. See [Roles and Responsibilities](#).

Creating the Required IAM Users, Groups, and Policies

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization.

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

Access to resources is provided to groups using policies and then inherited by the users that are assigned to those groups. Data transfer requires the creation of two distinct groups:

- Data transfer *administrators* who can create and manage transfer jobs.
- Data transfer *upload users* who can upload data to Object Storage. For your data security, the permissions for upload users allow Oracle personnel to upload standard and multi-part objects on your behalf and inspect bucket and object metadata. The permissions do not allow Oracle personnel to inspect the actual data.

The Data Administrator is responsible for generating the required RSA keys needed for the temporary upload users. These keys should never be shared between users.

For details on creating groups, see [Managing Groups](#).

An administrator creates these groups with the following policies:

- The data transfer administrator group requires an authorization policy that includes the following:

```
Allow group <group_name> to manage data-transfer-jobs in compartment <compartment_name>
```

```
Allow group <group_name> to manage buckets in compartment <compartment_name>
```

```
Allow group <group_name> to manage objects in compartment <compartment_name>
```

Alternatively, you can consolidate the `manage buckets` and `manage objects` policies into the following:

```
Allow group <group_name> to manage object-family in compartment <compartment_name>
```

- The data transfer upload user group requires an authorization policy that includes the following:

CHAPTER 10 Data Transfer

```
Allow group <group_name> to manage buckets in compartment <compartment_name> where all {
request.permission='BUCKET_READ' }
```

```
Allow group <group_name> to manage objects in compartment <compartment_name> where any {
request.permission='OBJECT_CREATE' , request.permission='OBJECT_OVERWRITE' ,
request.permission='OBJECT_INSPECT' }
```



Important

For security reasons, we recommend that you create a unique IAM data transfer upload user for each transfer job and then delete that user once your data is uploaded to Oracle Cloud Infrastructure.

The Oracle Cloud Infrastructure administrator then adds a user to each of the data transfer groups created. For details on creating users, see [Managing Users](#).

Creating Object Storage Buckets

The Object Storage service is used to upload your data to Oracle Cloud Infrastructure. Object Storage stores objects in a container called a bucket within a compartment in your tenancy. For details on creating the bucket to store uploaded data, see [Managing Buckets](#).

Creating Transfer Jobs

This section describes how to create a transfer job as part of the preparation for the data transfer. See [Transfer Jobs](#) for complete details on all tasks related to transfer jobs.



Tip

You can use the Console or the Data Transfer Utility to create a transfer job.

A transfer job represents the collection of files that you want to transfer and signals the intention to upload those files to Oracle Cloud Infrastructure. A transfer job combines at least

one transfer disk with a transfer package. Identify which compartment and Object Storage bucket that Oracle is to upload your data to. Create the transfer job in the same compartment as the upload bucket and supply a human-readable name for the transfer job. Avoid entering confidential information when providing transfer job names.



Note

It is recommended that you create a compartment for each transfer job to minimize the required access your tenancy.

Creating a transfer job returns a job ID that you specify in other transfer tasks. For example:

```
ocid1.datatransferjob.region1.phx..exampleuniqueID
```

To create a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Select the designated compartment you are to use for data transfers from the drop-down list.
A list of transfer jobs that have already been created is displayed.
3. Click **Create Transfer Job**.
4. In the **Create Transfer Job** dialog, enter a **Job Name**, and select the **Upload Bucket** from the drop-down list.
Avoid entering confidential information in the transfer job name.
5. Select **Disk** for the **Transfer Device Type**.
6. Click **Create Transfer Job**.

To create a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job create` to create a transfer job. The

CHAPTER 10 Data Transfer

`<display_name>` is the name of the transfer job. Avoid entering confidential information in the transfer job name.

```
 dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk
```

Optionally, you can specify one or more free-form or defined tags when you create a transfer job. For more information about tagging, see [Resource Tags](#).

To specify free-form tags when creating a job:

```
 dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --freeform-tags '{ "<tag_key>":"<value>" }'
```

To specify defined tags when creating a job:

```
 dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```



Note

Only users with the required permissions can create tag namespaces and tag keys. Create the tag namespaces and keys before you can specify them when creating a job. See [Working with Defined Tags](#) for details.

To specify multiple tags, comma separate the JSON-formatted key/value pairs:

```
 dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --freeform-tags '{ "<tag_key>":"<value>" }, { "<tag_key>":"<value>" }'
```

Preparing Upload Configuration Files

The Project Sponsor is responsible for creating or obtaining configuration files that allow the uploading of user data to the transfer appliance. Send these configuration files to the Data Administrator where they can be placed in the Data Host. The `config` file is for the data transfer administrator, the IAM user with the authorization and permissions to create and manage transfer jobs. The `config_upload_user` file is for the data transfer upload user, the temporary IAM user that Oracle uses to upload your data on your behalf.

CHAPTER 10 Data Transfer

Create a base Oracle Cloud Infrastructure directory and two configuration files with the required credentials.

CREATING THE DATA TRANSFER DIRECTORY

Create a Oracle Cloud Infrastructure directory (`.oci`) on the same Data Host where the CLI is installed. For example:

```
mkdir /root/.oci/
```

The two configuration files (`config` and `config_upload_user`) are placed in this directory.

CREATING THE DATA TRANSFER ADMINISTRATOR CONFIGURATION FILE

Create the data transfer administrator configuration file `/root/.oci/config` with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer administrator>
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_ path to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

```
[DEFAULT]
user=ocid1.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocid1.user.oc1..exampleuniqueID.pem
tenancy=ocid1.tenancy.oc1..exampleuniqueID
region=us-phoenix-1
```

For the data transfer administrator, you can create a single configuration file that contains different profile sections with the credentials for multiple users. Then use the `--profile` option to specify which profile to use in the command.

Here is an example of a data transfer administrator configuration file with different profile sections:

```
[DEFAULT]
user=ocid1.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocid1.user.oc1..exampleuniqueID.pem
tenancy=ocid1.tenancy.oc1..exampleuniqueID
```

CHAPTER 10 Data Transfer

```
region=us-phoenix-1
[PROFILE1]
user=ocidl.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
tenancy=ocidl.tenancy.oc1..exampleuniqueID
region=us-ashburn-1
```

By default, the `DEFAULT` profile is used for all CLI commands. For example:

```
oci dts job create --compartment-id ocid.compartment.oc1..exampleuniqueID --bucket MyBucket --display-
name MyDisplay --device-type disk
```

Instead, you can issue any CLI command with the `--profile` option to specify a different data transfer administrator profile. For example:

```
oci dts job create --compartment-id ocid.compartment.oc1..exampleuniqueID --bucket MyBucket --display-
name MyDisplay --device-type disk --profile MyProfile
```

Using the example configuration file above, the `<profile_name>` would be `profile1`.

If you created two separate configuration files, use the following command to specify the configuration file to use:

```
oci dts job create --compartment-id <compartment_id> --bucket <bucket_name> --display-name <display_
name>
```

CREATING THE DATA TRANSFER UPLOAD USER CONFIGURATION FILE

The `config_upload_user` configuration file is for the data transfer upload user, the temporary IAM user that Oracle uses to upload your data on your behalf. Create this configuration file with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer upload user>
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_ path to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

```
[DEFAULT]
user=ocidl.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
```

CHAPTER 10 Data Transfer

```
tenancy=ocid1.tenancy.oc1..exampleuniqueID  
region=us-phoenix-1
```



Important

Creating an upload user configuration file with multiple profiles is *not* supported.

CONFIGURATION FILE ENTRIES

The following table lists the basic entries that are required for each configuration file and where to get the information for each entry.



Note

Data Transfer Service does not support passphrases on the key files for both data transfer administrator and data transfer upload user.

Entry	Description and Where to Get the Value	Required?
user	OCID of the data transfer administrator or the data transfer upload user, depending on which profile you are creating. To get the value, see Required Keys and OCIDs .	Yes
fingerprint	Fingerprint for the key pair being used. To get the value, see Required Keys and OCIDs .	Yes

Entry	Description and Where to Get the Value	Required?
key_file	Full path and filename of the private key. Important: The key pair must be in PEM format. For instructions on generating a key pair in PEM format, see Required Keys and OCIDs .	Yes
tenancy	OCID of your tenancy. To get the value, see Required Keys and OCIDs .	Yes
region	An Oracle Cloud Infrastructure region. See Regions and Availability Domains . Data transfer is supported in US East (Ashburn), US West (Phoenix), Germany Central (Frankfurt), and UK South (London).	Yes

You can verify the data transfer upload user credentials using the following command:

```
dts job verify-upload-user-credentials --bucket <bucket_name>
```

Creating Shipping Labels

You can find the shipping address in the transfer package details. Use this information to create a shipping label for the transfer package that is used to send the disk to Oracle.

To get the shipping address for a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to see the details.
3. Click the Actions icon (three dots), and then click **View Details**.

CHAPTER 10 Data Transfer

Alternatively, click the hyper-linked name of the transfer job.

A list of transfer packages that have already been created is displayed.

4. Find the transfer package for which you want to see the details.
5. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.

To get the shipping address for a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package show` to get the shipping address for a transfer package.

```
dts package show --job-id <job_id> --package-label <package_label>
```

Notifying the Data Administrator

When you have completed all the tasks in this topic, provide the Data Administrator of the following:

- IAM login credentials
- Data Transfer Utility configuration files
- Transfer job ID
- Transfer job label

What's Next

You are now ready to configure your system for the data transfer. See [Configuring Disk Data Transfers](#).

Configuring Disk Data Transfers



CHAPTER 10 Data Transfer

This topic describes the tasks associated with configuring the Disk-Based Data Transfer. The Infrastructure Engineer role typically performs these tasks. See [Roles and Responsibilities](#).

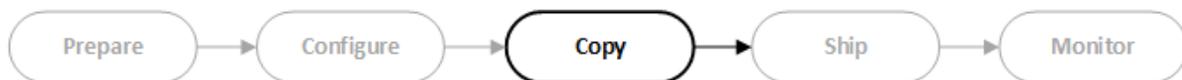
Configuration for the Disk-Based Data Transfer consists of the following tasks:

- Attaching the disk to the Data Host. Remove all partitions and any file systems. To prevent the accidental deletion of data, the Data Transfer Utility does not work with disks that already have partitions or file systems. Disks are visible to the host as block devices and must provide a valid response to the `hdparm -I <device>` Linux command.
- Sending the block device path to the Data Administrator.

What's Next

You are now ready to load your data to the disk. See [Copying the Data](#).

Copying the Data



This topic describes the tasks associated with running the data transfer from the Data Host to the physical disk that has been configured as a transfer disk. The Data Administrator role typically performs these tasks. See [Roles and Responsibilities](#).

Information Prerequisites

Before performing any disk copying tasks, you must obtain the following information:

- Disk block device path. The Infrastructure Engineer typically provides this information.
- IAM login information, Data Transfer Utility configuration files, transfer job ID, and job label. The Project Sponsor typically provides this information.

Generate and Upload RSA Key Pairs

The Data Administrator is responsible for generating and uploading the RSA key pairs. Do not share these generated RSA keys between users. See [Creating a Key Pair](#).

Install and Configure the Data Transfer Utility

The Data Transfer Utility provides a set of command line-based tools for configuring and running disk-based data transfers. Use the Data Transfer Utility as an alternative to running commands from the Console. Sometimes you must use the Data Transfer Utility to complete certain tasks as there is no Console equivalent. See [Data Transfer Utility](#) for details on how to install and configure the Data Transfer Utility for use with disk-based data transfers.

Creating the Transfer Disk

The transfer disk is the logical representation of the physical disk that has been configured for use for receiving data as part of the disk-based data transfer. See [Transfer Disks](#) for complete details on all tasks related to transfer disks.



Tip

You can only use the Data Transfer Utility to create a transfer disk.

When you create a transfer disk for use with the disk on which you are copying your files, the Data Transfer Utility:

- Sets up the disk for encryption using the passphrase
- Creates a file system on the disk
- Mounts the file system at `/mnt/orcdts_<label>`

For example:

```
/mnt/orcdts_DJZNWK3ET
```

When you register a transfer disk, Oracle Cloud Infrastructure generates a strong encryption passphrase that is used to encrypt the contents on the disk. The encryption passphrase is displayed to standard output to the data transfer administrator user and cannot be retrieved again. Create a local, secure copy of the encryption passphrase, so you can reference the passphrase again.

Creating a transfer disk requires the job ID returned from when you [created the transfer job](#) and the path to the attached disk (for example, `/dev/sdb`).

To create a transfer disk using the Data Transfer Utility

At the command prompt on the host, run `dts disk create` to create a transfer disk.

```
dts disk create --job-id <job_id> --block-device <block_device>
```

Copying Files

Attach the disks to the Data Host and copy files to the mount point created by the transfer disk through the Data Transfer Utility.

You can only copy regular files to disks. Special files (links, sockets, pipes, and so forth) cannot be copied directly. To transfer special files, create a tar archive of the files and copy the tar archive to the disk.



Note

Copy all Files Before Disconnecting the Disk

Do not disconnect the disk until you copy all files from the Data Host and generate the manifest file. If you accidentally disconnect the disk before copying all files, you must unlock the disk using the encryption passphrase. The encryption passphrase was generated and displayed when you created the transfer disk. If the generated encryption passphrase is not available, you must delete the transfer disk from the transfer job and re-create the transfer disk. All data previously copied to that disk is lost.

Generating the Manifest



Tip

You can only use the Data Transfer Utility to generate a manifest file.

The amount of time to generate the manifest file depends on the size of the upload files, disk speed, and available processing power.

After copying your data to a transfer disk, generate a manifest file using the Data Transfer Utility. The manifest contains an index of all of the copied files and generated data integrity hashes. The Data Transfer Utility copies the `config_upload_user` configuration file and referenced IAM credentials to the encrypted transfer disk. This configuration file describes the temporary IAM data transfer upload user. Oracle uses the credentials and entries defined in the `config_upload_user` file when processing the transfer disk and uploading files to Oracle Cloud Infrastructure Object Storage.



Note

Data Transfer Service Does Not Support Passphrases on Private Keys

While we recommend encrypting a private key with a passphrase when generating API signing keys, Data Transfer does not support passphrases on the key file required for the `config_upload_user`. If you use a passphrase, Oracle personnel cannot upload your data.

Oracle cannot upload data from a transfer disk without the correct credentials defined in this configuration file. See [Data Transfer Utility](#) for more information about the required configuration files.

To create a manifest file using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk manifest` to create a manifest file.

```
dts disk manifest --job-id <job_id> --disk-label <label>
```



Note

Do You Need to Regenerate the Manifest File?

If you add, remove, or modify any files on the disk after generating the manifest file, you must regenerate the file. If the manifest file does not match the contents of the target bucket, Oracle cannot upload the data.

Locking the Transfer Disk



Tip

You can only use the Data Transfer Utility to lock a transfer disk.

Locking a transfer disk safely unmounts the disk and removes the encryption passphrase from the Data Host.

To lock a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk lock` to lock a transfer disk.

```
dts disk lock --job-id <job_id> --disk-label <label> --block-device <block_device>
```

If you need to unlock the transfer disk , you are prompted for the encryption passphrase that was generated when you created the transfer disk.

To unlock a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk unlock` to unlock a transfer disk.

```
dts disk unlock --job-id <job_id> --disk-label <label> --block-device <block_device> --encryption-passphrase <encryption_passphrase>
```

Creating the Transfer Package

A transfer package is the virtual representation of the physical package of disks that you are shipping to Oracle for upload to Oracle Cloud Infrastructure. See [Transfer Packages](#) for complete details on all tasks related to transfer packages.



Tip

You can use the Console or the Data Transfer Utility to create a transfer package.

Creating a transfer package requires the job ID returned from when you created the transfer job. For example:

```
ocid1.datatransferjob.region1.phx..exampleuniqueID
```

To create a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to create a transfer package.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
A list of transfer packages that have already been created is displayed.
4. Click **Create Transfer Package**.
5. In the **Create Transfer Package** dialog, choose the **Vendor**.
6. Click **Create Transfer Package**.

The Data Transfer Package dialog appears displaying information such as the shipping address, the shipping vendor, and the shipping status.

To create a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package create` to create a transfer package.

```
dts package create --job-id <job_id>
```

The following information is returned:

```
Transfer Package :  
Label :  
TransferSiteShippingAddress :  
DeliveryVendor :  
DeliveryTrackingNumber :  
ReturnDeliveryTrackingNumber :  
Status :  
Devices :
```

Attaching the Transfer Disk to the Transfer Package

Attach a transfer disk to a transfer package after you have performed the following tasks:

1. Copied your data onto the disk
2. Generated the required manifest file
3. Run and reviewed the dry-run report
4. Locked the transfer disk in preparation for shipment



Tip

You can use the Console or the Data Transfer Utility to attach a transfer disk to a transfer package.

A disk can be attached to one package, detached, and then attached to another package.

To attach a transfer disk to a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job associated with the transfer package that you want to attach a disk to.

3. Click the Actions icon (three dots), and then click **View Details**.
A list of transfer packages is displayed.
4. Find the transfer package that you want to attach a disk to.
5. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer package.
A list of transfer disks is displayed.
6. Click **Attach Transfer Disks**.
7. In the **Attach Transfer Disk** dialog, select the **Transfer Disks** that you want to attach to the transfer package.
8. Click **Attach**.

To attach a transfer disk to a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk attach` to attach a disk to a transfer package.

```
dts disk attach --job-id <job_id> --package-label <package_label> --disk-label <label>
```

You have attached a transfer disk to a transfer package, but have changed your mind about shipping that disk with the transfer package. You can also detach a transfer disk from one transfer package and attach that disk to a different transfer package.

To detach a transfer disk to a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer package for which you want to detach a transfer disk.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer package.
A list of transfer disks that have already been attached is displayed.
4. Find the transfer disk that you want to detach.

5. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer disk.
6. Click **Detach Transfer Disk**.

To detach a transfer disk to a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk detach` to detach a disk from a transfer package.

```
dts disk detach --job-id <job_id> --package-label <package_label> --disk-label <label>
```

Setting Tracking Details on the Transfer Package



Tip

You can use the Console or the Data Transfer Utility to update the transfer package with tracking information.

After delivering the transfer package to the shipping vendor, update the transfer package with the tracking information.



Important

Oracle cannot process a transfer package until you update the tracking information.

To update the transfer package with tracking information using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.

2. Find the transfer job for which you want to see the associated transfer packages.
3. Click the Actions icon (three dots), and then click **View Details**.
A list of transfer packages that have already been created is displayed.
4. Find the transfer package that you want to update.
5. Click the Actions icon (three dots), and then click **View Details**.
6. Click **Edit**.
7. Enter the **Tracking ID** and the **Return Tracking ID**.
8. Click **Edit Transfer Package**.

To update the transfer package with tracking information using the Data Transfer Utility

At the command prompt on the host, run `dts package ship` to update the transfer package tracking information.

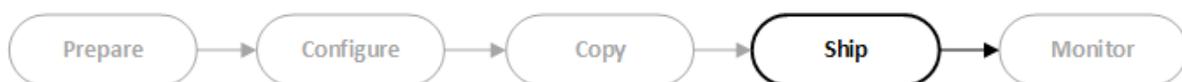
```
dts package ship --job-id <job_id> --package-label <package_label> --package-vendor <vendor_name> --tracking-number <tracking_number> --return-tracking-number <return_tracking_number>
```

Notifying the Infrastructure Engineer

After completing the tasks listed in this topic, notify the Infrastructure Engineer of the following:

- Disconnect the physical disk from the Data Host
- Package the disk for shipment

Shipping the Disk



This topic describes the tasks associated with shipping the physical disk containing the copied data to Oracle. The Infrastructure Engineer role typically performs these tasks. See [Roles and Responsibilities](#).

Disconnecting the Disk from the Data Host

Do not disconnect the disk until you copy all files from the Data Host and generate the [manifest file](#).

Printing Shipping Labels

You should receive the shipping labels electronically from the Project Sponsor. Print them on the appropriate labels for shipping the disk.

Packaging and Shipping the Disk

GENERAL

Include the required return shipping label in the box when packaging transfer disks for shipment.



Note

Return Shipment Label Requirement

If you do not include the return shipping label inside the box, Oracle cannot process the transfer package.

Ensure that the transfer job and transfer package label are clearly readable on the outside of the box containing the transfer disks.



Important

If you are shipping transfer disks to London or Frankfurt, request that the shipping vendor requires a signature delivery.

SHIPPING TRANSFER DISKS INTERNATIONALLY

Create a commercial invoice when shipping transfer disks internationally. To ensure that packages are not held up in customs, follow these guidelines when creating the commercial invoice:

- Show a unique reference number.
- Show the "bill-to party as follows":
 - For shipments to the European Union (Frankfurt) location:
ORACLE Deutschland B.V. & Co. KG
Riesstrasse 25
Munich, 80992
GERMANY
 - For shipments to the United States location:
Oracle America, Inc.
500 Oracle Parkway
REDWOOD CITY CA 94065
UNITED STATES
- Show the "ship-to party" as the address provided in the transfer package details. See [Creating Shipping Labels](#) for details.
- State that "The value shown includes the value of software and data recorded onto the hard drive unit."
- State that the "Goods are free of charge - no payment required."
- State that the type of export is "Temporary."

- Ensure that the commodity code shows the correct HS code for a hard drive unit as specified in the source country's HS code list.
- State the description as the manufacture's description of the hard drive unit and include the words "Hard Disk Drive."
- Ensure that the invoice is signed and includes the printed name of the signer.

Monitoring the Disk



This topic describes the tasks to be done after the data transfer is complete the disk has been shipped to Oracle. The Project Sponsor role typically performs these tasks. See [Roles and Responsibilities](#).

Tracking the Disk Shipment

When Oracle has processed the transfer disks associated with a transfer package, the status of the transfer package changes to **Processed**. When Oracle has shipped the transfer disks associated with a transfer package, the status of the transfer package changes to **Returned**.



Tip

You can use the Console or the Data Transfer Utility to check the status of a transfer package.

To check the status of a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Choose the data transfer package for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.
4. Look at the **Status**.

To check the status of a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package show` to show the status of a transfer package.

```
dts package show --job-id <job_id> --package-label <package_label>
```

Reviewing the Upload Summary

Oracle creates upload summary log files for each uploaded disk. These logs are placed in the bucket where the data was uploaded to Oracle Cloud Infrastructure. The upload summary file compares the disk's manifest file to the contents of the target Oracle Cloud Infrastructure Object Storage bucket after file upload.

The top of the log report summarizes the overall file processing status:

```
P - Present: The file is present in both the disk and the target bucket
M - Missing: The file is present in the disk but not the target bucket. It was likely uploaded and then
deleted by another user before the summary was generated.
C - Name Collision: The file is present in the manifest but a file with the same name but different
contents is present in the target bucket.
U - Unreadable: The file is not readable from the disk
N - Name Too Long: The file name on disk is too long and could not be uploaded
```

Complete file upload details follow the summary.

CHAPTER 10 Data Transfer

```
#####
##### SUMMARY FOR DISK [WDH887L] #####
Generated 2017-09-08 19:46:36
TOTAL : 1110
### P present: 1110
### M missing: 0
### C name collision: 0
### U unreadable: 0
### N nameTooLong: 0
#####
| STATUS | NAME | LAST_MODIFIED | SIZE (MB) | MD5 | ETag |
| present | small/01/T6QKX | Fri Sep 08 19:04:54 UTC 2017 | 10.00 | 8c1kXbWU793H2KH1aF8m6v== | 58B2F70D1C874AAA8053824318ACAC52 |
| present | small/01/FAFKU | Fri Sep 08 19:12:20 UTC 2017 | 10.00 | 8c1kXbWU793H2KH1aF8m6v== | 58B2EC4616E4608E053824318AC6DF6 |
| present | small/01/EVPPD | Fri Sep 08 19:02:42 UTC 2017 | 10.00 | 8c1kXbWU793H2KH1aF8m6v== | 58B2ECDFa7494606E053824318AC383A |
| present | small/01/2U02H | Fri Sep 08 19:13:06 UTC 2017 | 10.00 | 8c1kXbWU793H2KH1aF8m6v== |
```

Closing the Transfer Job

Typically, you would close a transfer job when no further transfer job activity is required or possible. Closing a transfer job requires that the status of all associated transfer packages be returned, canceled, or deleted. In addition, the status of all associated transfer disks must be complete, in error, missing, canceled, or deleted.



Tip

You can use the Console or the Data Transfer Utility to close a transfer job.

To close a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer package for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
4. Click **Close Transfer Job**.

To close a transfer job using the Data Transfer Utility

At the command prompt on the host, run `dts job close` to close a transfer job.

```
dts job close --job-id <job_id>
```

What's Next

You have completed the process of setting up, running, and monitoring the disk-based data transfer. After the disk contents is successfully migrated to Oracle Cloud Infrastructure, your physical disk is erased and returned to you.

If you determine that another disk-based data transfer is required, repeat the procedure from the beginning.

Data Transfer Utility

This topic describes how to install and configure the Data Transfer Utility for use in disk-based data transfers. In addition, this topic describes the syntax for the Data Transfer Utility commands.



Important

With this release, the Data Transfer Utility only supports disk-based data transfers. Use of the Data Transfer Utility for appliance-based transfers has been replaced with the Oracle Cloud Infrastructure command line interface (CLI). See for more information.

The Data Transfer Utility is licensed under the Universal Permissive License 1.0 and the Apache License 2.0. Third-party content is separately licensed as described in the code.

Prerequisites

To install and use the Data Transfer Utility, obtain the following:

- An Oracle Cloud Infrastructure account.
- Required Oracle Cloud Infrastructure users and groups with the required IAM policies. See for details.
- A Data Host machine with the following installed:
 - Oracle Linux 6 or greater, Ubuntu 14.04 or greater, or SUSE 11 or greater
 - Java 1.8 or Java 1.11
 - hdparm 9.0 or later
 - Cryptsetup 1.2.0 or greater
- Firewall access: If you have a restrictive firewall in the environment where you are using the Data Transfer Utility, you may need to open your firewall configuration to the following IP address ranges: 140.91.0.0/16.
You also need to open access to the object storage IP address ranges: 134.70.0.0/17.

Installing the Data Transfer Utility

Download and install the Data Transfer Utility installer that corresponds to your Data Host's operating system.

To install the Data Transfer Utility on Debian or Ubuntu

1. Download the [installation .deb file](#).
2. Issue the `apt install` command as the `root` user that has write permissions to the `/opt` directory.

```
sudo apt install ./dts-X.Y.Z.x86_64.deb
```

`X.Y.Z` represents the version numbers that match the installer you downloaded.

3. Confirm that the Data Transfer Utility installed successfully.

```
sudo dts --version
```

Your Data Transfer Utility version number is returned.

To install the Data Transfer Utility on Oracle Linux or Red Hat Linux

1. Download the [installation .rpm file](#).
2. Issue the `yum install` command as the `root` user that has write permissions to the `/opt` directory.

```
sudo yum localinstall ./dts-X.Y.Z.x86_64.rpm
```

X.Y.Z represents the version numbers that match the installer you downloaded.

3. Confirm that the Data Transfer Utility installed successfully.

```
sudo dts --version
```

Your Data Transfer Utility version number is returned.

Configuring the Data Transfer Utility

Before using the Data Transfer Utility, you must create a base Oracle Cloud Infrastructure directory and two configuration files with the required credentials. One configuration file is for the data transfer administrator, the IAM user with the authorization and permissions to create and manage transfer jobs. The other configuration file is for the data transfer upload user, the temporary IAM user that Oracle uses to upload your data on your behalf.

BASE DATA TRANSFER DIRECTORY

Create a base Oracle Cloud Infrastructure directory:

```
mkdir /root/.oci/
```

CONFIGURATION FILE FOR THE DATA TRANSFER ADMINISTRATOR

Create a data transfer administrator configuration file `/root/.oci/config` with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer administrator>
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_ path to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

CHAPTER 10 Data Transfer

```
[DEFAULT]
user=ocidl.user.oc1..<unique_ID>
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..<unique_ID>.pem
tenancy=ocidl.tenancy.oc1..<unique_ID>
region=us-phoenix-1
```

For the data transfer administrator, you can create a single configuration file that contains different profile sections with the credentials for multiple users. Then use the `--profile` option to specify which profile to use in the command. Here is an example of a data transfer administrator configuration file with different profile sections:

```
[DEFAULT]
user=ocidl.user.oc1..<unique_ID>
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..<unique_ID>.pem
tenancy=ocidl.tenancy.oc1..<unique_ID>
region=us-phoenix-1
[PROFILE1]
user=ocidl.user.oc1..<unique_ID>
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..<unique_ID>.pem
tenancy=ocidl.tenancy.oc1..<unique_ID>
region=us-ashburn-1
```

By default, the `DEFAULT` profile is used for all Data Transfer Utility commands. For example:

```
dts job create --compartment-id <compartment_id> --bucket <bucket_name> --display-name <display_name>
--device-type <disk>
```

Instead, you can issue any Data Transfer Utility command with the `--profile` option to specify a different data transfer administrator profile. For example:

```
dts job create --compartment-id <compartment_id> --bucket <bucket_name> --display-name <display_name>
--device-type <disk> --profile <profile_name>
```

Using the example configuration file above, the `<profile_name>` would be `profile1`.

CONFIGURATION FILE FOR THE DATA TRANSFER UPLOAD USER

Create a data transfer upload user `/root/.oci/config_upload_user` configuration file with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer upload user>
```

CHAPTER 10 Data Transfer

```
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_path to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

```
[DEFAULT]
user=ocidl.user.oc1..<unique_ID>
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..<unique_ID>.pem
tenancy=ocidl.tenancy.oc1..<unique_ID>
region=us-phoenix-1
```



Important

Creating an upload user configuration file with multiple profiles is *not* supported.

CONFIGURATION FILE ENTRIES

The following table lists the basic entries that are required for each configuration file and where to get the information for each entry.



Note

Data Transfer Service does not support passphrases on the key files for both data transfer administrator and data transfer upload user.

Entry	Description and Where to Get the Value	Required?
user	OCID of the data transfer administrator or the data transfer upload user, depending on which profile you are creating. To get the value, see Required Keys and OCIDs .	Yes
fingerprint	Fingerprint for the key pair being used. To get the value, see Required Keys and OCIDs .	Yes
key_file	Full path and filename of the private key. Important: The key pair must be in PEM format. For instructions on generating a key pair in PEM format, see Required Keys and OCIDs .	Yes
tenancy	OCID of your tenancy. To get the value, see Required Keys and OCIDs .	Yes
region	An Oracle Cloud Infrastructure region. See Regions and Availability Domains . Data transfer is supported in US East (Ashburn), US West (Phoenix), Germany Central (Frankfurt), and UK South (London).	Yes

You can verify the data transfer upload user credentials using the following command:

```
dts job verify-upload-user-credentials --bucket <bucket_name>
```

CONFIGURATION FILE LOCATION

The location of the configuration files is `/root/.oci/config`.

Using the Data Transfer Utility

This section provides an overview of the syntax for the Data Transfer Utility.



Important

The Data Transfer Utility must be run as the `root` user.

You can specify Data Transfer Utility command options using the following commands:

- `--option <value>` *or*
- `--option=<value>`

SYNTAX

The basic Data Transfer Utility syntax is:

```
dts <resource> <action> <options>
```

This syntax is applied to the following:

- `dts` is the shortened utility command name
- `job` is an example of a `<resource>`
- `create` is an example of an `<action>`
- Other utility strings are `<options>`

The following commands to create a transfer job shows a typical Data Transfer Utility construct.

```
dts job create --compartment-id ocidl.compartment.oc1..<unique_ID> --display-name "mycompany transfer1" --bucket mybucket --device-type disk
```

Or:

```
dts job create --compartment-id=ocidl.compartment.oc1..<unique_ID> --display-name="mycompany transfer1" --bucket=mybucket --device-type=disk
```



Note

In the previous examples, provide a friendly name for the transfer job using the `--display-name` option. Avoid entering confidential information when providing resource names or descriptions.

GETTING HELP WITH COMMANDS

You can get help with the different commands associated with Data Transfer Utility using `dts` by itself. For example:

```
dts
```

FINDING OUT THE INSTALLED VERSION OF THE DATA TRANSFER UTILITY

You can get the installed version of the Data Transfer Utility using `--version` or `-v`. For example:

```
dts --version
```

What's Next

You are now ready to perform disk-based data transfers. See .

Disk Data Transfer Reference

This topic provides complete task details for certain components associated with Disk-Based Data Transfers. Use this topic as a reference to learn and use commands associated with components included in the Disk-Based Data Transfer procedure.

Transfer Jobs

A transfer job represents the collection of files that you want to transfer and signals the intention to upload those files to Oracle Cloud Infrastructure. A transfer job combines at least

one transfer disk with a transfer package. Identify which compartment and Object Storage bucket that Oracle is to upload your data to.



Note

It is recommended that you create a compartment for each transfer job to minimize the required access your tenancy.



Tip

You can use the Console or the Data Transfer Utility to create a transfer job.

CREATING TRANSFER JOBS

Create the transfer job in the same compartment as the upload bucket and supply a human-readable name for the transfer job. Avoid entering confidential information when providing transfer job names.

Creating a transfer job returns a job ID that you specify in other transfer tasks. For example:

```
ocidl.datatransferjob.region1.phx..<unique_ID>
```

To create a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Select the designated compartment you are to use for data transfers from the drop-down list.
A list of transfer jobs that have already been created is displayed.

3. Click **Create Transfer Job**.
4. In the **Create Transfer Job** dialog, enter a **Job Name**, and select the **Upload Bucket** from the drop-down list.
Avoid entering confidential information in the transfer job name.
5. Select **Disk** for the **Transfer Device Type**.
6. Click **Create Transfer Job**.

To create a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job create` to create a transfer job. The `<display_name>` is the name of the transfer job. Avoid entering confidential information in the transfer job name.

```
dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk
```

Optionally, you can specify one or more free-form or defined tags when you create a transfer job. For more information about tagging, see [Resource Tags](#).

To specify free-form tags when creating a job:

```
dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --freeform-tags '{ "<tag_key>": "<value>" }'
```

To specify defined tags when creating a job:

```
dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --defined-tags '{ "<tag_namespace>": { "<tag_key>": "<value>" } }'
```



Note

Users create tag namespaces and tag keys with the required permissions. These items must exist before you can specify them when creating a job. See [Working with Defined Tags](#) for details.

To specify multiple tags, comma separate the JSON-formatted key/value pairs:

```
 dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name>
--device-type disk --freeform-tags '{ "<tag_key>":"<value>" }', '{ "<tag_key>":"<value>" }'
```

DISPLAYING TRANSFER JOBS

To display the list of transfer jobs using the Console

Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.

To display the list of transfer jobs using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job list` to display the list of transfer jobs.

```
 dts job list --compartment-id <compartment_id>
```

When you use the Data Transfer Utility to list jobs, tagging details are also included in the output if you specified tags.

DISPLAYING TRANSFER JOB DETAILS

To display the details of a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.

To display the details of a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job show` to display the details of a

transfer job.

```
dts job show --job-id <job_id>
```

When you use the Data Transfer Utility to display the details of a job, tagging details are also included in the output if you specified tags.

EDITING TRANSFER JOBS

To edit the name of a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer job that you want to edit.
3. Click the Actions icon (three dots), and then click **Edit**.
4. Edit the name of the transfer job.
Avoid entering confidential information in the transfer job name.
5. Click **Save**.

To edit the name of a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job update` to edit the name (`--display-name`) of a transfer job. The `<display_name>` is the new name of the transfer job. Avoid entering confidential information in the transfer job name.

```
dts job update --job-id <job_id> --display-name <display_name>
```

To edit the tags associated with a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job update` to edit the tags associated with a transfer job. The Data Transfer Utility **replaces** any existing tags with the new key/value pairs you specify.

To edit free-form tags, provide the replacement key/value pairs:

```
dts job update --job-id <job_id> --freeform-tags '{ "<tag_key>":"<value>" }'
```

To edit defined tags, provide the replacement key value pairs:

```
dts job update --job-id <job_id> --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```

To delete the tags associated with a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job update` to delete the tags associated with a transfer job. The Data Transfer Utility **replaces** any existing tags with the new key/value pairs you specify. If you want to delete some of the tags, you would specify new tag string that does not contain the key/value pair you want to delete.

Partial tag deletion is handled in the same way as you edit tags:

- To edit free-form tags, provide the replacement key/value pairs:

```
dts job update --job-id <job_id> --freeform-tags '{ "<tag_key>":"<value>" }'
```

- To edit defined tags, provide the replacement key value pairs:

```
dts job update --job-id <job_id> --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```

To delete all free-form tags:

```
dts job update --job-id <job_id> --freeform-tags '{}'
```

To delete all defined tags:

```
dts job update --job-id <job_id> --defined-tags '{}'
```

DELETING TRANSFER JOBS

Typically, you would delete a transfer job early in the transfer process and before you create any transfer packages or disks. For example, you initiated the data transfer by creating a transfer job, but changed your mind. If you want to delete a transfer job later in the transfer process, you must first delete all transfer packages and disks associated with the transfer job.

To delete a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer job that you want to delete.
3. Click the Actions icon (three dots), and then click **Delete**.
Alternatively, you can delete a transfer job from the **View Details** page.
4. Confirm the deletion when prompted.

To delete a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job delete` to delete a transfer job.

```
dts job delete --job-id <job_id>
```

CLOSING TRANSFER JOBS



Tip

You can use the Console or the Data Transfer Utility to close a transfer job.

Typically, you would close a transfer job when no further transfer job activity is required or possible. Closing a transfer job requires that the status of all associated transfer packages be returned, canceled, or deleted. In addition, the status of all associated transfer disks must be complete, in error, missing, canceled, or deleted.

To close a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer package for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
4. Click **Close Transfer Job**.

To close a transfer job using the Data Transfer Utility

At the command prompt on the Data Host, run `dts job close` to close a transfer job.

```
dts job close --job-id <job_id>
```

Transfer Disks

The section describes the creation and management transfer disks.



Important

Before creating a transfer disk from an attached disk, remove all partitions and any file systems. To prevent the accidental deletion of data, the Data Transfer Utility does not work with disks that already have partitions or file systems. Disks are visible to the Data Host as block devices and must provide a valid response to the `hdparm -I <device>` Linux command.

CREATING TRANSFER DISKS



Tip

You can only use the Data Transfer Utility to create a transfer disk.

When you create a transfer disk, the Data Transfer Utility:

- Sets up the disk for encryption using the passphrase
- Creates a file system on the disk
- Mounts the file system at `/mnt/orcdts_<label>`

For example:

```
/mnt/orcdts_DJZWNK3ET
```

When you register a transfer disk, Oracle Cloud Infrastructure generates a strong encryption passphrase that is used to encrypt the transfer disk. The encryption passphrase is displayed to standard output to the data transfer administrator user and cannot be retrieved again. Create a local, secure copy of the encryption passphrase, so you can reference the passphrase again.

Creating a transfer disk requires the job ID returned from when you [created the transfer job](#) and the path to the attached disk (for example `/dev/sdb`).

To create a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk create` to create a transfer disk.

```
dts disk create --job-id <job_id> --block-device <block_device>
```

DELETING TRANSFER DISKS



Tip

You can only use the Data Transfer Utility to delete a transfer disk.

Typically, you would delete a transfer disk during the disk preparation process. You created, attached, and copied data to the transfer disk, but have changed your mind about shipping the disk. If you want to reuse the disk, remove all file systems and create the disk again.

To delete a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dtc disk delete` to delete a transfer disk.

```
dtc disk delete --job-id <job_id> --disk-label <label>
```

CANCELING TRANSFER DISKS



Tip

You can only use the Data Transfer Utility to cancel a transfer disk.

If you shipped a disk to Oracle, but have changed your mind about uploading the files, you can cancel the transfer disk. You can cancel a disk in a transfer package, while allowing the file upload from other disks.

Oracle cannot process canceled transfer disks. Oracle returns canceled transfer disks to the sender.

To cancel a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dtc disk cancel` to cancel a transfer disk.

CHAPTER 10 Data Transfer

```
dts disk cancel --job-id <job_id> --disk-label <label>
```

LOCKING TRANSFER DISKS



Tip

You can only use the Data Transfer Utility to lock a transfer disk.

Locking a transfer disk safely unmounts the disk and removes the encryption passphrase from the Data Host.

To lock a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk lock` to lock a transfer disk.

```
dts disk lock --job-id <job_id> --disk-label <label> --block-device <block_device>
```

UNLOCKING TRANSFER DISKS

If you need to unlock the transfer disk, you are prompted for the encryption passphrase that was generated when you created the transfer disk.

To unlock a transfer disk using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk unlock` to unlock a transfer disk.

```
dts disk unlock --job-id <job_id> --disk-label <label> --block-device <block_device> --encryption-passphrase <encryption_passphrase>
```

Transfer Packages

A transfer package is the virtual representation of the physical package of disks that you are shipping to Oracle for upload to Oracle Cloud Infrastructure.



Tip

You can use the Console or the Data Transfer Utility to create a transfer package.

CREATING TRANSFER PACKAGES

Creating a transfer package requires the job ID returned from when you created the transfer job. For example:

```
ocid1.datatransferjob.region1.phx..exampleuniqueID
```

To create a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to create a transfer package.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
A list of transfer packages that have already been created is displayed.
4. Click **Create Transfer Package**.
5. In the **Create Transfer Package** dialog, choose the **Vendor**.
6. Click **Create Transfer Package**.

The Data Transfer Package dialog appears displaying information such as the shipping address, the shipping vendor, and the shipping status.

To create a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package create` to create a transfer package.

CHAPTER 10 Data Transfer

```
dts package create --job-id <job_id>
```

The following information is returned:

```
Transfer Package :  
Label :  
TransferSiteShippingAddress :  
DeliveryVendor :  
DeliveryTrackingNumber :  
ReturnDeliveryTrackingNumber :  
Status :  
Devices :
```

DISPLAYING TRANSFER PACKAGE DETAILS

To display the details of a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to see the details.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
A list of transfer packages that have already been created is displayed.

To display the details of a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package show` to display the details of a transfer package.

```
dts package show --job-id <job_id> --package-label <package_label>
```

The following information is returned:

```
Transfer Package :  
Label :  
TransferSiteShippingAddress :  
DeliveryVendor :  
DeliveryTrackingNumber :
```

CHAPTER 10 Data Transfer

```
ReturnDeliveryTrackingNumber :  
Status :  
Devices :
```

EDITING TRANSFER PACKAGES

Edit the transfer package and supply the tracking information when you ship the package.

To edit a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to see the associated transfer packages.
3. Click the Actions icon (three dots), and then click **View Details**.
4. Find the transfer package that you want to edit.
5. Click the Actions icon (three dots), and then click **View Details**.
6. Click **Edit**.
Change the vendor and supply the tracking information as needed.
7. Click **Edit Transfer Package**.

To edit a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package update` to edit the details of a transfer package.

```
dts package update --job-id <job_id> --package-label <package_label> [--package-vendor <vendor_name>]  
[--tracking-number <tracking_number>] [--return-tracking-number <return_tracking_number>]
```

DELETING TRANSFER PACKAGES

Typically, you would delete a transfer package early in the transfer process and before you created any transfer disks. You initiated the transfer job and package, but have changed your mind. If you delete a transfer package later in the transfer process, you must first detach all

associated transfer disks. You cannot delete a transfer package once the package has been shipped to Oracle.

To delete a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to see the associated transfer packages.
3. Click the Actions icon (three dots), and then click **View Details**.
4. Find the transfer package that you want to edit.
5. Click the Actions icon (three dots), and then click **View Details**.
6. Click **Edit**.
Change the vendor and supply the tracking information as needed.
7. Click **Edit Transfer Package**.

To delete a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package delete` to delete a transfer package.

```
dts package delete --job-id <job_id> --package-label <package_label>
```

CANCELING TRANSFER PACKAGES

If you shipped a transfer package, but have changed your mind about uploading the data, you can cancel a transfer package. Before canceling a transfer package, you must first cancel all transfer disks associated with that transfer package. Oracle cannot process canceled transfer packages. Oracle returns canceled transfer packages to the sender.

To cancel a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and

click **Data Transfer**.

2. Find the transfer job for which you want to see associated transfer packages.
3. Click the Actions icon (three dots), and then click **View Details**.
4. Find the transfer package that you want to cancel.
5. Click the Actions icon (three dots), and then click **View Details**.
6. Click **Cancel Transfer Package**.

To cancel a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts package cancel` to cancel a transfer package.

```
dts package cancel --job-id <job_id> --package-label <package_label>
```

ATTACHING TRANSFER DISKS TO A TRANSFER PACKAGE



Tip

You can use the Console or the Data Transfer Utility to attach a transfer disk to a transfer package.

Attach a transfer disk to a transfer package after you have done the following tasks in order:

- Copied your data onto the disk
- Generated the required manifest file
- Run and reviewed the dry-run report
- Locked the transfer disk in preparation for shipment

A disk can be attached to one package, detached, and then attached to another package.

To attach a transfer disk to a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job associated with the transfer package that you want to attach a disk to.
3. Click the Actions icon (three dots), and then click **View Details**.
A list of transfer packages is displayed.
4. Find the transfer package that you want to attach a disk to.
5. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer package.
A list of transfer disks is displayed.
6. Click **Attach Transfer Disks**.
7. In the **Attach Transfer Disk** dialog, select the **Transfer Disks** that you want to attach to the transfer package.
8. Click **Attach**.

To attach a transfer disk to a transfer package using the Data Transfer Utility

At the command prompt on the Data Host, run `dts disk attach` to attach a disk to a transfer package.

```
dts disk attach --job-id <job_id> --package-label <package_label> --disk-label <label>
```

You have attached a transfer disk to a transfer package, but have changed your mind about shipping that disk with the transfer package. You can also detach a transfer disk from one transfer package and attach that disk to a different transfer package.

To detach a transfer disk to a transfer package using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and

click **Data Transfer**.

2. Find the transfer package for which you want to detach a transfer disk.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer package.
A list of transfer disks that have already been attached is displayed.
4. Find the transfer disk that you want to detach.
5. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer disk.
6. Click **Detach Transfer Disk**.

To detach a transfer disk to a transfer package using the Data Transfer Utility

At the command prompt on the host, run `dts disk detach` to detach a disk from a transfer package.

```
dts disk detach --job-id <job_id> --package-label <package_label> --disk-label <label>
```

Appliance Data Transfer

Appliance-Based Data Transfer is one of Oracle's offline data transfer solutions that lets you migrate petabyte-scale datasets to Oracle Cloud Infrastructure. You send your data as files on one or more secure, high-capacity, Oracle-supplied storage appliances to an Oracle transfer site. Operators at the Oracle transfer site upload the files into the designated Object Storage bucket in your tenancy. You are then free to move the uploaded data to other Oracle Cloud Infrastructure services as needed.



Note

Appliance-Based Data Transfer is not available for free trial or Pay As You Go accounts.

Appliance-Based Data Transfer Concepts

The following concepts are essential to understanding Appliance-Based Data Transfer.

TRANSFER JOB

A transfer job is the logical representation of a data migration to Oracle Cloud Infrastructure. A transfer job is associated with one or more appliances.

APPLIANCE

An appliance is high storage capacity device that is specially prepared to copy and upload data to Oracle Cloud Infrastructure. You request an appliance from Oracle, copy your data to the appliance, and then ship the appliance back to Oracle to upload your data.

COMMAND LINE INTERFACE

The command line interface (CLI) is a small footprint tool that you can use on its own or with the Console to complete Oracle Cloud Infrastructure tasks, including Appliance-Based Data Transfer jobs.



Note

You can only run Oracle Cloud Infrastructure CLI commands from a Linux host. This differs from running CLI commands for other Oracle Cloud Infrastructure Services on a variety of host operating systems. Appliance-based commands require validation that is only available on Linux hosts.

HOST

A physical computer on which one or more of the logical hosts (Control, Data, Terminal Emulation) is running. Depending on your computing environment, you can have a separate physical host for each logical host, consolidate all three logical hosts onto a single physical host, or have two logical hosts on one physical host and the third logical

host on a separate physical host. All physical hosts must be on network used for the data transfer.

CONTROL HOST

The logical representation of the host computer at your site from which you perform Data Transfer Service tasks. Depending on your needs, you may use one or more separate hosts (Control and Data) to run your Appliance-Based Data Transfer job.

DATA HOST

The logical representation of the host computer on your site that stores the data you intend to copy to Oracle Cloud Infrastructure.

TERMINAL EMULATION HOST

The logical representation of the host computer that uses terminal emulation software to communicate with, and allow you to command, the appliance.

BUCKET

The logical container in Oracle Cloud Infrastructure Object Storage where Oracle operators upload your data. A bucket is associated with a single compartment in your tenancy that has policies that determine what actions a user can perform on a bucket and on all the objects in the bucket.

DATA TRANSFER ADMINISTRATOR

A new or existing IAM user that has the authorization and permissions to create and manage transfer jobs.

DATA TRANSFER UPLOAD USER

A temporary IAM user that grants Oracle personnel the authorization and permissions to upload the data from the appliance to your designated Oracle Cloud Infrastructure Object Storage bucket. Delete this temporary user after your data is uploaded to Oracle Cloud Infrastructure.

APPLIANCE MANAGEMENT SERVICE

Software running on the appliance that provides management functions. Users interact with this service through the Oracle Cloud Infrastructure CLI.

Appliance Specifications

Use NFS versions 3, 4, or 4.1 to copy your data onto the appliance. Here are some details about the appliance:

Item Description	Specification
Storage Capacity	150 TB of protected usable space
Network Interfaces	<ul style="list-style-type: none">- 10 GbE - RJ45- 10 GbE - SFP+ You are responsible for providing all network cables. If you want to use SFP+, your transceivers must be compatible with Intel X520 NICs.
Provided Cables	<ul style="list-style-type: none">- NEMA 5–15 type B to C13- C13 - 14 power- USB - DB9 serial
Environmental	<ul style="list-style-type: none">- Operational temperature: 50–95°F (10–35°C)- Operational relative humidity: 8–90% non-condensing- Acoustics: < 75 dB @ 73°F (23° C)- Operational altitude: -1,000 ft - 10,000 ft (approx. -300–3048 m))

Item Description	Specification
Power	<ul style="list-style-type: none"> - Consumption: 554 W - Voltage: 100–240 VAC - Frequency: 47–63 Hz - Conversion efficiency: 89%
Weight	<ul style="list-style-type: none"> - Unit: 38 lbs (approx. 17 kg) - Unit + Transit Case: 64 lbs (approx. 29 kg)
Height	3.5" (approx. 9 cm) (2U)
Width	17" (approx. 43 cm)
Depth	24" (approx. 61 cm)
Shipping Case	11" x 25" x 28" (approx. 28 x 63.5 x 71 cm)

Roles and Responsibilities

Depending on your organization, the responsibilities of using and managing the data transfer may span multiple roles. Use the following set of roles as a guideline for how you can assign the various tasks associated with the data transfer.

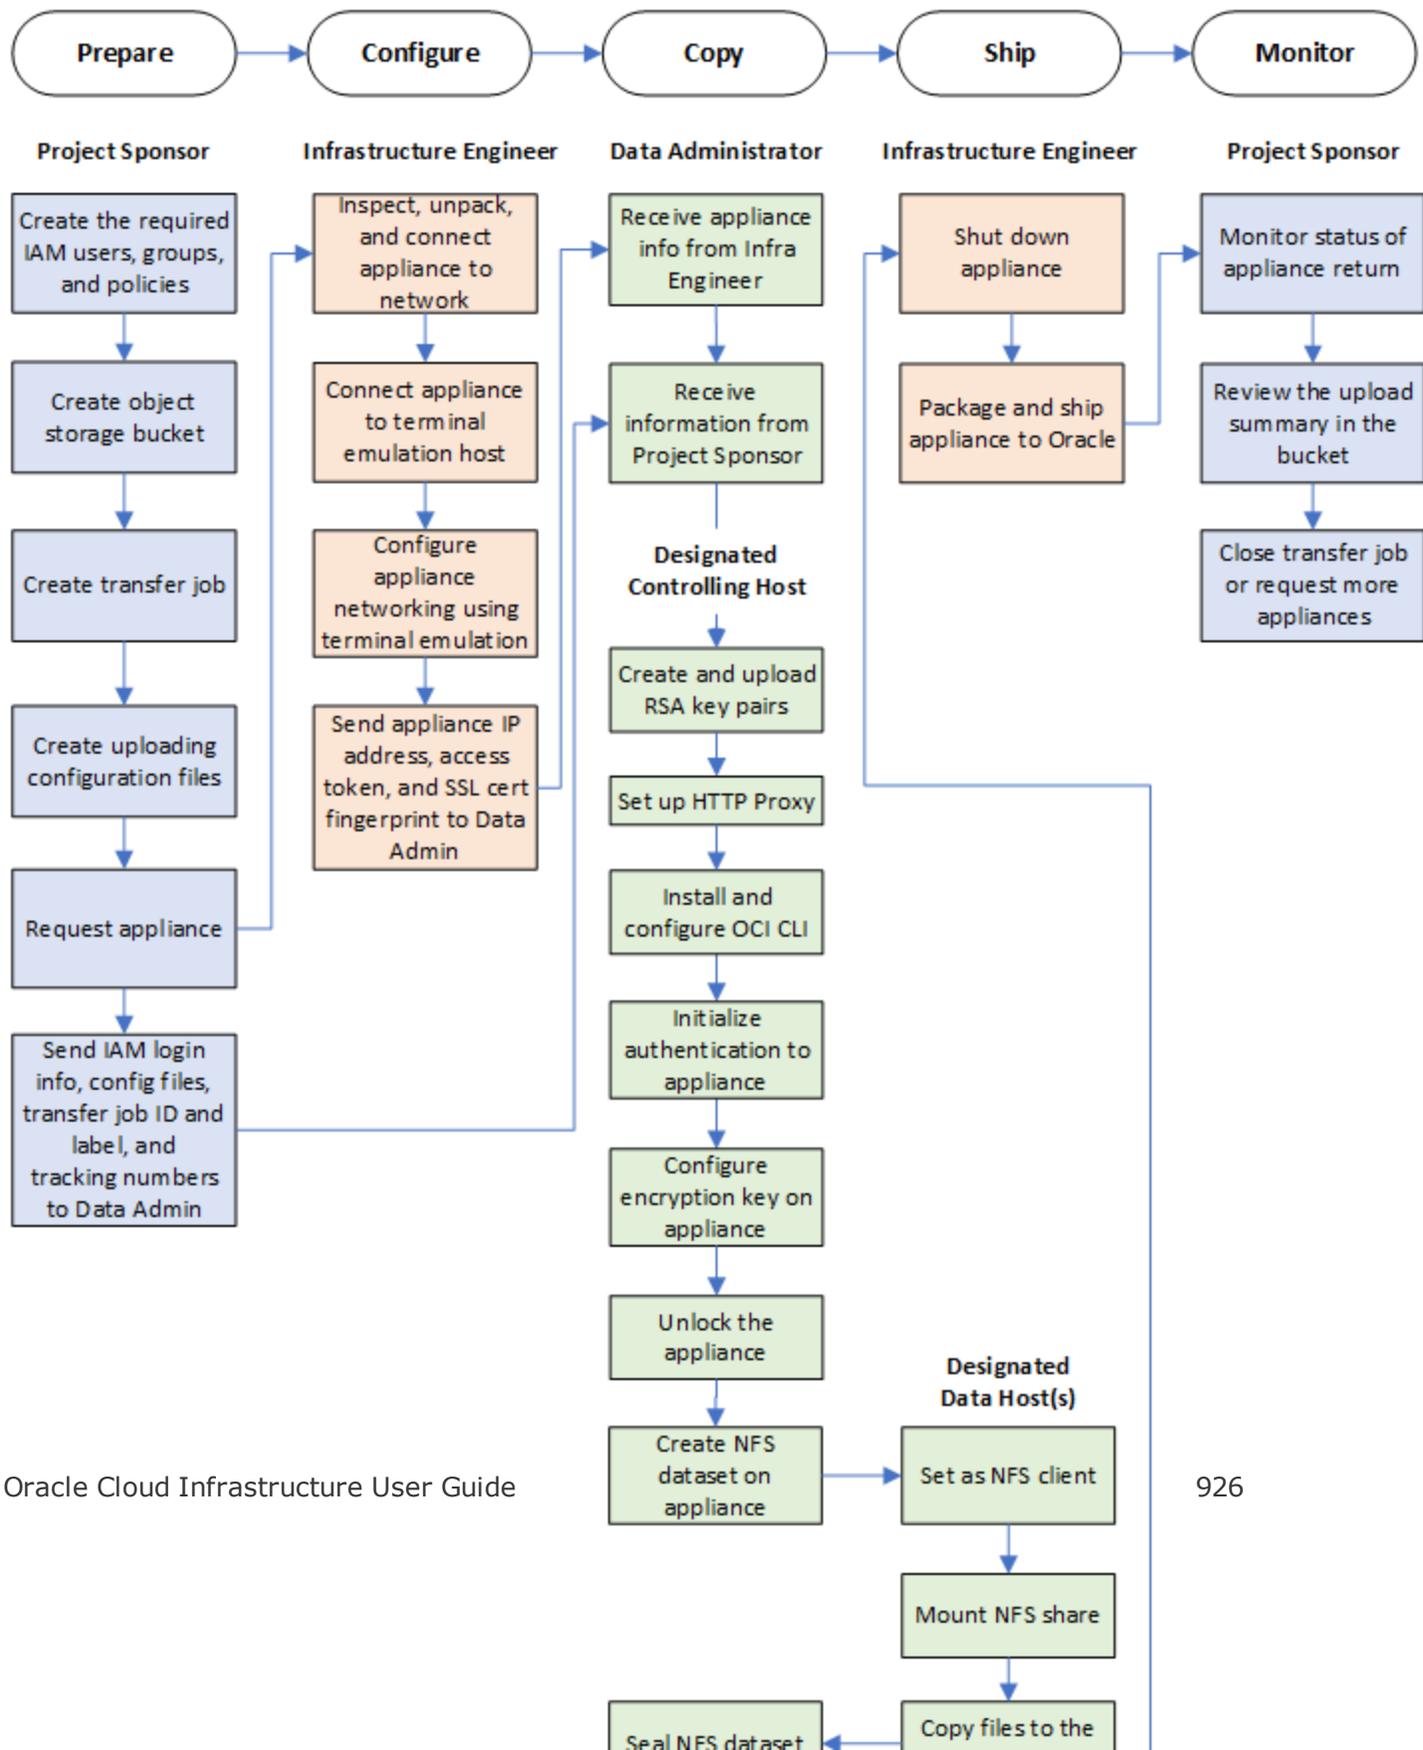
- **Project Sponsor:** Responsible for the overall success of the data transfer. Project Sponsors usually have complete access to their organization's Oracle Cloud Infrastructure tenancy. They coordinate with the other roles in the organization to complete the implementation of data transfer project.
- **Infrastructure Engineer:** Responsible for integrating the transfer appliance into the organization's IT infrastructure from where the data is being transferred. Tasks associated with this role include connecting the transfer appliance to power, placing it within the network, and setting the IP address through a serial console menu using the provided USB-to-Serial adapter.

- **Data Administrator:** Responsible for identifying and preparing the data to be transferred to Oracle Cloud Infrastructure. This person usually has access to, and expertise with, the data being migrated.

These roles correspond to the various phases of the data transfer described in the following section. A specific role can be responsible for one or more phases.

Task Flow for Appliance-Based Data Transfer

Here is a high-level overview of the tasks involved in the Appliance-Based Data Transfer to Oracle Cloud Infrastructure using organized by phase. Complete one phase before proceeding to the next one. Use the roles previously described to distribute the tasks across individuals or groups within your organization.



Secure Appliance Data Transfer to Oracle Cloud Infrastructure

This section highlights the security details of the Data Transfer Appliance process.

- Appliances are shipped from Oracle to you with a tamper-evident security tie on the transit case. A second tamper-evident security tie is included in the appliance transit case for you to secure the case when you ship the case back to Oracle. The number on the physical security ties must match the numbers logged by Oracle in the appliance details.
- When you configure the appliance for the first time:
 - The appliance generates a master AES-256 bit encryption key that is used for all data written to or read from the device. The encryption key never leaves the device.
 - The encryption key is protected by an encryption passphrase that you must know to access the encrypted data. The system securely fetches a provided encryption passphrase from Oracle Cloud Infrastructure and registers that passphrase on the appliance.



Note

The encryption passphrase is never stored on the appliance

- All data is encrypted as the data is copied to an appliance.
- For more security, you can also encrypt your own data with your own encryption keys. Before copying your data to the transfer appliance, you can encrypt your data with a tool and encryption key of your choosing. After the data has been uploaded, you would need to use the same tool and encryption key to access the data.
- All network communication between your appliance-based data transfer environment and Oracle Cloud Infrastructure is encrypted in-transit using Transport Layer Security (TLS).

- After copying your data to a transfer appliance, the data transfer system generates a manifest file. The manifest contains an index of all of the copied files and generated data integrity hashes. The system also encrypts and copies the `config_upload_user` configuration file to the transfer appliance. This configuration file describes the temporary IAM data transfer upload user. Oracle uses the credentials and entries defined in the `config_upload_user` file when processing the transfer appliance and uploading files to Oracle Cloud Infrastructure Object Storage.



Note

Data Transfer Service Does Not Support Passphrases on Private Keys

While we recommend encrypting a private key with a passphrase when generating API signing keys, the Data Transfer Service does not support passphrases on the key file required for the `config_upload_user` configuration file. If you use a passphrase, Oracle personnel cannot upload your data.

Oracle cannot upload data from a transfer appliance without the correct credentials defined in this configuration file. See [Preparing Upload Configuration Files](#) for more information about the required configuration files.

- Oracle erases all of your data from the transfer appliance after it has been processed. The erasure process follows the NIST 800-88 standards.
- Keep possession of the security tie after you have finished unpacking and connecting the appliance. Include it when returning the appliance to Oracle. Failure to include the security tie can result in a delay in the data migration process.

Preparing for Appliance Data Transfers



This topic describes the tasks associated with preparing for the Appliance-Based Data Transfer. The Project Sponsor role typically performs these tasks. See [Roles and Responsibilities](#).



Note

You can only run Oracle Cloud Infrastructure CLI commands from a Linux host. This differs from running CLI commands for other Oracle Cloud Infrastructure Services on a variety of host operating systems. Appliance-based commands require validation that is only available on Linux hosts.

Creating the Required IAM Users, Groups, and Policies

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization.

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

Access to resources is provided to groups using policies and then inherited by the users that are assigned to those groups. Data transfer requires the creation of two distinct groups:

- Data transfer *administrators* who can create and manage transfer jobs.
- Data transfer *upload users* who can upload data to Object Storage. For your data security, the permissions for upload users allow Oracle personnel to upload standard and multi-part objects on your behalf and inspect bucket and object metadata. The permissions do not allow Oracle personnel to inspect the actual data.

The Data Administrator is responsible for generating the required RSA keys needed for the temporary upload users. These keys should never be shared between users.

For details on creating groups, see [Managing Groups](#).

An administrator creates these groups with the following policies:

- The data transfer administrator group requires an authorization policy that includes the following:

```
Allow group <group_name> to manage data-transfer-jobs in compartment <compartment_name>
```

```
Allow group <group_name> to manage buckets in compartment <compartment_name>
```

```
Allow group <group_name> to manage objects in compartment <compartment_name>
```

Alternatively, you can consolidate the `manage buckets` and `manage objects` policies into the following:

```
Allow group <group_name> to manage object-family in compartment <compartment_name>
```

- The data transfer upload user group requires an authorization policy that includes the following:

```
Allow group <group_name> to manage buckets in compartment <compartment_name> where all {  
  request.permission='BUCKET_READ' }
```

```
Allow group <group_name> to manage objects in compartment <compartment_name> where any {  
  request.permission='OBJECT_CREATE' , request.permission='OBJECT_OVERWRITE' ,  
  request.permission='OBJECT_INSPECT' }
```



Important

For security reasons, we recommend that you create a unique IAM data transfer upload user for each transfer job and then delete that user once your data is uploaded to Oracle Cloud Infrastructure.

The Oracle Cloud Infrastructure administrator then adds a user to each of the data transfer groups created. For details on creating users, see [Managing Users](#).

REQUESTING THE DATA TRANSFER APPLIANCE ENTITLEMENT

If your tenancy has not been entitled to perform Appliance-Based Data Transfers, you are required to request it before creating an appliance-based transfer job. The Data Transfer Appliance Entitlement is a tenancy-wide entitlement that you need to request once for each tenancy.

Use the following policy to enable users in a specific group to request a Data Transfer Appliance Entitlement in your tenancy.

```
Allow group <group_name> to {DTA_ENTITLEMENT_CREATE} in tenancy
```

Creating Object Storage Buckets

The Object Storage service is used to upload your data to Oracle Cloud Infrastructure. Object Storage stores objects in a container called a bucket within a compartment in your tenancy. For details on creating the bucket to store uploaded data, see [Managing Buckets](#).

Creating Transfer Jobs

This section describes how to create a transfer job as part of the preparation for the data transfer. See [Transfer Jobs](#) for complete details on all tasks related to transfer jobs.



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to create a transfer job.

A transfer job represents the collection of files that you want to transfer and signals the intention to upload those files to Oracle Cloud Infrastructure. Identify which compartment and Object Storage bucket that Oracle is to upload your data to. Create the transfer job in the same compartment as the upload bucket and supply a human-readable name for the transfer job. Avoid entering confidential information when providing transfer job names.



Note

It is recommended that you create a compartment for each transfer job to minimize the required access your tenancy.

Creating a transfer job returns a job ID that you specify in other transfer tasks. For example:

```
ocid1.datatransferjob.region1.phx..<unique_ID>
```

To create a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Select the designated compartment you are to use for data transfers from the drop-down list.
A list of transfer jobs that have already been created is displayed.
3. Click **Create Transfer Job**.
4. In the **Create Transfer Job** dialog, enter a **Job Name**, and select the **Upload Bucket** from the drop-down list.
Avoid entering confidential information in the transfer job name.
5. Select **Appliance** for the **Transfer Device Type**.
6. Click **Create Transfer Job**.

To create a transfer job using the CLI

At the command prompt on the Control Host, run `dts job create` to create a transfer job. The `<display_name>` is the name of the transfer job. Avoid entering confidential information in the transfer job name.

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type appliance
```

CHAPTER 10 Data Transfer

Optionally, you can specify one or more free-form or defined tags when you create a transfer job. For more information about tagging, see [Resource Tags](#).

To specify free-form tags when creating a job:

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type appliance --freeform-tags '{ "<tag_key>":"<value>" }'
```

To specify defined tags when creating a job:

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type appliance --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```



Note

Users with the required permissions create tag namespaces and tag keys. Namespaces and keys must be present before you can specify them when creating a job. See [Working with Defined Tags](#) for details.

To specify multiple tags, comma separate the JSON-formatted key/value pairs:

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type appliance --freeform-tags '{ "<tag_key>":"<value>","<tag_key>":"<value>" }'
```

REQUESTING THE DATA TRANSFER APPLIANCE ENTITLEMENT



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to request the Data Transfer Appliance Entitlement.

If your tenancy is not entitled to use the Data Transfer Appliance, you must request the Data Transfer Appliance Entitlement before creating an appliance-based transfer job.

To request the Data Transfer Appliance Entitlement using the Console

Open the Transfer Job page and click **Request** at the top. Otherwise, you are prompted to request the entitlement when attempting to create your first appliance-based transfer job.

Once requested, the status of your request is visible at the top of the Transfer Job page. It can take a while to get the Data Transfer Appliance Entitlement approved. After Oracle receives your request, a Terms and Conditions Agreement is sent to the account owner via DocuSign to use the appliance. The entitlement request is approved once the signature is received. The Data Transfer Appliance Entitlement is a tenancy-wide entitlement that you need to request once for each tenancy.

To request the Data Transfer Appliance Entitlement using the CLI

```
oci dts appliance request-entitlement --compartment-id <compartment_id> --profile <profile> --name <your_name> --email <your_email>
```

```
oci dts appliance show-entitlement --compartment-id <compartment_id> --profile <profile>
```

Preparing Upload Configuration Files

The Project Sponsor is responsible for creating or obtaining configuration files that allow the uploading of user data to the transfer appliance. Send these configuration files to the Data Administrator where they can be placed in the Control Host (if there are separate Control and Data Hosts). The `config` file is for the data transfer administrator, the IAM user with the authorization and permissions to create and manage transfer jobs. The `config_upload_user` file is for the data transfer upload user, the temporary IAM user that Oracle uses to upload your data on your behalf.

Create a base Oracle Cloud Infrastructure directory and two configuration files with the required credentials.

CREATING THE DATA TRANSFER DIRECTORY

Create a Oracle Cloud Infrastructure directory (`.oci`) on the same Control Host machine where the Oracle Cloud Infrastructure CLI is installed. For example:

```
mkdir /root/.oci/
```

The two configuration files (`config` and `config_upload_user`) are placed in what ever location you choose.



Note

You can store the configuration files anywhere on your Control Host. The `root` directory is only given as an example.

CREATING THE DATA TRANSFER ADMINISTRATOR CONFIGURATION FILE

Create the data transfer administrator configuration file `/root/.oci/config` with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer administrator>
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_path_ to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

```
[DEFAULT]
user=ocidl.user.oc1..<unique_ID>
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
tenancy=ocidl.tenancy.oc1..<unique_ID>
region=us-phoenix-1
```

For the data transfer administrator, you can create a single configuration file that contains different profile sections with the credentials for multiple users. Then use the `--profile` option to specify which profile to use in the command. Here is an example of a data transfer administrator configuration file with different profile sections:

```
[DEFAULT]
user=ocidl.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
tenancy=ocidl.tenancy.oc1..exampleuniqueID
region=us-phoenix-1
[PROFILE1]
```

CHAPTER 10 Data Transfer

```
user=ocidl.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
tenancy=ocidl.tenancy.oc1..exampleuniqueID
region=us-ashburn-1
```



Important

Creating an upload user configuration file with multiple profiles is *not* supported.

By default, the `DEFAULT` profile is used for all CLI commands. For example:

```
oci dts job create --compartment-id ocid.compartment.oc1..exampleuniqueID --bucket MyBucket --display-name MyDisplay --device-type appliance
```

Instead, you can issue any CLI command with the `--profile` option to specify a different data transfer administrator profile. For example:

```
oci dts job create --compartment-id ocid.compartment.oc1..exampleuniqueID --bucket MyBucket --display-name MyDisplay --device-type appliance --profile MyProfile
```

Using the example configuration file above, the `<profile_name>` would be `profile1`.

If you created two separate configuration files, use the following command to specify the configuration file to use:

```
oci dts job create --compartment-id <compartment_id> --bucket <bucket_name> --display-name <display_name>
```

CREATING THE DATA TRANSFER UPLOAD USER CONFIGURATION FILE

The `config_upload_user` configuration file is for the data transfer upload user, the temporary IAM user that Oracle uses to upload your data on your behalf. Create this configuration file with the following structure:

```
[DEFAULT]
user=<The OCID for the data transfer upload user>
fingerprint=<The fingerprint of the above user's public key>
key_file=<The _absolute_path to the above user's private key file on the host machine>
tenancy=<The OCID for the tenancy that owns the data transfer job and bucket>
region=<The region where the transfer job and bucket should exist. Valid values are:
us-ashburn-1, us-phoenix-1, eu-frankfurt-1, and uk-london-1.>
```

For example:

CHAPTER 10 Data Transfer

```
[DEFAULT]
user=ocidl.user.oc1..exampleuniqueID
fingerprint=4c:1a:6f:a1:5b:9e:58:45:f7:53:43:1f:51:0f:d8:45
key_file=/home/user/ocidl.user.oc1..exampleuniqueID.pem
tenancy=ocidl.tenancy.oc1..exampleuniqueID
region=us-phoenix-1
```

CONFIGURATION FILE ENTRIES

The following table lists the basic entries that are required for each configuration file and where to get the information for each entry.



Note

Data Transfer Service does not support passphrases on the key files for both data transfer administrator and data transfer upload user.

Entry	Description and Where to Get the Value	Required?
user	OCID of the data transfer administrator or the data transfer upload user, depending on which profile you are creating. To get the value, see Required Keys and OCIDs .	Yes
fingerprint	Fingerprint for the key pair being used. To get the value, see Required Keys and OCIDs .	Yes
key_file	Full path and filename of the private key. Important: The key pair must be in PEM format. For instructions on generating a key pair in PEM format, see Required Keys and OCIDs .	Yes

Entry	Description and Where to Get the Value	Required?
tenancy	OCID of your tenancy. To get the value, see Required Keys and OCIDs .	Yes
region	An Oracle Cloud Infrastructure region. See Regions and Availability Domains . Data transfer is supported in US East (Ashburn), US West (Phoenix), Germany Central (Frankfurt), and UK South (London).	Yes

You can verify the data transfer upload user credentials using the following command:

```
oci dts job verify-upload-user-credentials --bucket <bucket_name>
```

Requesting the Transfer Appliance

This section describes how to request a transfer appliance from Oracle for copying your data to Oracle Cloud Infrastructure. See [Appliances](#) for complete details on all tasks related to transfer jobs.



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to request a transfer appliance.

Oracle Cloud Infrastructure customers can use data transfer appliances to migrate data for free. You are only charged for Object Storage usage once the data is successfully transferred to your designated bucket. All appliance requests still require approval from Oracle.



Tip

We recommend that you identify the data you intend to upload and make data copy preparations before requesting the transfer appliance.

Creating a transfer appliance request returns an Oracle-assigned appliance label. For example:

XA8XM27EVH

To request a transfer appliance using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
Choose the transfer job that you want to request a transfer appliance for.
2. Under **Transfer Appliances**, click **Request Transfer Appliance**.
3. In the **Request Transfer Appliance** dialog box, provide the shipping address details where you want the appliance sent.
 - **Company Name:** Required. Specify the name of the company that owns the data being migrated to Oracle Cloud Infrastructure.
 - **Recipient Name:** Required. Specify the name of the recipient to send the appliance to.
 - **Recipient Phone Number:** Required. Specify the recipient's phone number.
 - **Recipient Email Address:** Required. Specify the recipient's email address.
 - **Care Of:** Optional intermediary party responsible for transferring the appliance shipment from the delivery vendor to the intended recipient.
 - **Address Line 1:** Required. Specify the street address to send the appliance to.
 - **Address Line 2:** Optional identifying address details like building, suite, unit, or floor information.

- **City/Locality:** Required. Specify the city or locality.
 - **State/Province/Region:** Required. Specify the state, province, or region.
 - **Zip/Postal Code:** Specify the zip code or postal code.
 - **Country:** Required. Select the country.
4. Click **Request Transfer Appliance**.

To request a transfer appliance using the CLI

At the command prompt on the Control Host, run `oci dts appliance request` to request a data transfer appliance.

Here are the minimum requirements for the transfer appliance request:

```
oci dts appliance request --job-id <job_id> --addressee <addressee> --address1 <address_line1> --city-or-locality <city_or_locality> --state-or-region <state_or_region>--country<country>--zip-code<zip>
```

In addition, you can specify these optional fields in the request:

```
--care-of <care_of>  
--address2 <address_line2>, --address3 <address_line3>, and --address4 <address_line4>  
--email <email_address>  
--phone-number <phone_number>  
--profile <profile>
```

When you submit an appliance request, Oracle generates a unique label (name) to identify the transfer appliance and your request is sent to Oracle for approval and processing.

Notifying the Data Administrator

When you have completed all the tasks in this topic, provide the Data Administrator of the following:

- IAM login credentials
- Oracle Cloud Infrastructure CLI configuration files
- Transfer job ID
- Transfer job label

What's Next

You are now ready to configure your system for the data transfer. See [Configuring Appliance Data Transfers](#).

Configuring Appliance Data Transfers



This topic describes the tasks associated with configuring the Appliance-Based Data Transfer. The Infrastructure Engineer role typically performs these tasks. See [Roles and Responsibilities](#).

Unpacking and Connecting the Appliance to the Network

When the shipping vendor delivers your transfer appliance, Oracle updates the status as **Delivered** and provides the date and time the appliance was received in the **Transfer Appliance Details**.



Important

Your transfer appliance arrives in a transit case with a telescoping handle and wheels. The case amenities allow for easy movement to the location where you intend to place the appliance to upload your data.

Retain all packaging materials! When shipping the transfer appliance back to Oracle, you must package the appliance in the same manner and packaging in which the appliance was received.

Here are the tasks involved in unpacking and getting your transfer appliance ready to configure.

1. Inspect the tamper-evident security tie on the transit case.

If the appliance was tampered with during transit, the tamper-evident security tie serves to alert you.



Warning

If the security tie is damaged or is missing, do not plug the appliance into your network! Immediately file a Service Request (SR).

2. Remove and compare the number on the security tie with the number logged by Oracle.

To see the security tie number logged by Oracle using the Console

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
- b. Find the transfer job and transfer appliance associated with the removed security tie.
- c. Click the **Actions** icon (⋮), and then click **View Details**.
- d. Look at the contents of the **Send Security Tie ID** field in the **Transfer Appliance Details** and compare that number with the number on the physical tag.

To see the security tie number logged by Oracle using the CLI

At the command prompt on the Control Host, run `oci dts appliance show` to delete a transfer appliance.

```
oci dts appliance show --job-id <job_id> --appliance-label <label>
```



Warning

If the number on the physical security tie does not match the number logged by Oracle, do not plug the appliance into your network! Immediately file a Service Request (SR).



Note

Keep possession of the security tie after you have finished unpacking and connecting the appliance. Include it when returning the appliance to Oracle. Failure to include the security tie can result in a delay in the data migration process.

3. Open the transit case and ensure that the case contains the following items:
 - Appliance unit and power cable (two types of power cables provided: C14 and C13 to 14)
 - USB to DB-9 serial cable
 - Return shipping instructions (retain these instructions)
 - Return shipping label, label sleeve, tie-on tag, and zip tie
 - Return shipment tamper-evident security tie (use this tie to ensure secure transit case back to Oracle)
4. Compare the number on the return shipment security tie with the number logged by Oracle.

To see the security tie number logged by Oracle using the Console

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
- b. Find the transfer job and transfer appliance associated with the return shipment security tie.
- c. Click the Actions icon (three dots), and then click **View Details**.
- d. Look at the contents of the **Return Security Tie ID** field in the **Transfer Appliance Details** and compare that number with the number on the physical tag.

To see the security tie number logged by Oracle using the CLI

At the command prompt on the Control Host, run `oci dts appliance show` to the security tie number associated with the transfer appliance.

```
oci dts appliance show --job-id <job_id> --appliance-label <label>
```



Warning

If the number on the return security tie does not match the number logged by Oracle, file a Service Request (SR). These security tie numbers must match or Oracle cannot upload data from your returned transfer appliance.

5. Remove the transfer appliance from the case and place the appliance on a solid surface or in a rack.



Warning

We recommend assistance lifting the transfer appliance out of the transit case and placing the appliance in a rack or on a desk top. The total shipping weight is about 64 lbs (29.0299 kg) and appliance weight is 38 lbs (17.2365 kg).

6. Connect the appliance to your local network using one of the following:
 - 10GBase-T: Standard RJ-45
 - SFP+: The transceiver must be compatible with Intel X520 NICs.
7. Attach one of the provided power cords to the appliance and plug the other end into a grounded power source.
8. Turn on the appliance by flipping the power switch on the back of the appliance.

Connecting the Appliance to the Terminal Emulation Host

Connect the appliance to your designated Terminal Emulation Host computer using the provided USB to DB-9 serial cable.



Note

You might need to download the driver for this cable on your Terminal Emulation Host:
<https://www.cablestogo.com/product/26887/5ft-usb-to-db9-male-serial-rs232-adapter-cable#support>

SETTING UP TERMINAL EMULATION

Appliance-based transfers require you to set up your host for terminal emulation so you can communicate with the appliance device through the appliance's serial console. This

CHAPTER 10 Data Transfer

communication requires installing serial console terminal emulator software. We recommend using the following:

- PuTTY for Windows
- ZOC for OS X
- PuTTY or Minicom for Linux

Configure the following terminal emulator software settings:

- Baud Rate: 115200
- Emulation: VT102
- Handshaking: Disabled/off
- RTS/DTS: Disabled/off



Note

PuTTY does not allow you to configure all of these settings individually. However, you can configure the PuTTY default settings by selecting the **Serial** connection type and specifying "115200" for the **Serial Line** baud speed. This is sufficient to use PuTTY as a terminal emulator for the appliance.

Configuring the Transfer Appliance Networking

When the appliance boots up, an appliance serial console configuration menu is displayed on the Terminal Emulation Host to which the appliance is connected.

```
Oracle Cloud Data Transfer Appliance
- For use with minimum dts version: dts-0.4.140
- See "Help" for determining your dts version

1) Configure Networking
2) Show Networking
3) Reset Authentication
4) Show Authentication
```

CHAPTER 10 Data Transfer

```
5) Show Status
6) Collect Appliance Diagnostic Information
7) Generate support bundle
8) Shutdown Appliance
9) Reboot Appliance
10) Help
```

Select a command:



Note

It can take up to 5 minutes for the serial console menu to display. Press **Enter** if you do not see the serial console configuration menu after this amount of time.

The appliance supports a single active network interface on any of the 10-Gbps network ports. If only one interface is cabled and active, that interface is chosen automatically. If multiple interfaces are active, you are given the choice to select the interface to use.

To configure your transfer appliance networking

1. From the Terminal Emulation Host, select **Configure Networking** from the appliance serial console menu.
2. Provide the required networking information when prompted:
 - **IP Address:** IP address of the transfer appliance.
 - **Subnet Mask Length:** The count of leading 1 bit in the subnet mask. For example, if the subnet mask is 255.255.255.0 then the length is 24.
 - **Default Gateway:** Default gateway for network communications.

For example:

```
Configure Networking:
^C to cancel

Configuring IP address, subnet mask length, gateway
Example:
```

CHAPTER 10 Data Transfer

```
IP Address : 10.0.0.2
Subnet Mask Length : 24
Gateway : 10.0.0.1

Address: 10.0.0.1
Subnet Mask Length: 24
Gateway: 10.0.0.1

Configuring IP address 10.0.0.1 netmask 255.255.255.0 default gateway 10.0.0.1
Enabling enp0s3
Now trying to restart the network

Network configuration is complete

New authentication material created.

Client access token           : 4iH1gw1okPJO
Appliance certificate MD5 fingerprint : BF:C6:49:9B:25:FE:9F:64:06:7E:DF:F5:F9:E5:C6:56
Press ENTER to return...
```

When you configure a network interface, the appliance software generates a new client access token and appliance X.509/SSL certificate. The access token is used to authorize your Control Host to communicate with the Data Transfer Appliance's Management Service. The x.509/SSL certificate is used to encrypt communications with the Data Transfer Appliance's Management Service over the network. Provide the access token and SSL certificate fingerprint values displayed here when you use the CLI commands to [initialize authentication on your host machine](#).

You can change the selected interface, network information, and reset the authentication material at any time by selecting **Configure Networking** again from the appliance serial console menu.

Notify the Data Administrator

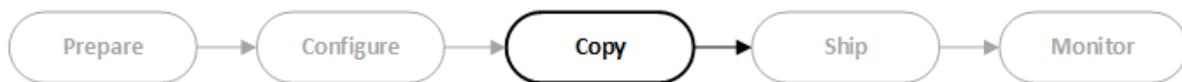
After completing the tasks in this topic, send the following appliance information IP address of the transfer appliance to the Data Administrator:

- Appliance IP address
- Access token
- SSL certificate fingerprint

What's Next

You are now ready to load your data to the disk. See [Copying Data to the Appliance](#).

Copying Data to the Appliance



This topic describes the tasks associated with copying data from the Data Host to the appliance using the Control Host. The Data Administrator role typically performs these tasks. See [Roles and Responsibilities](#).



Note

You can only run Oracle Cloud Infrastructure CLI commands from a Linux host. This differs from running CLI commands for other Oracle Cloud Infrastructure Services on a variety of host operating systems. Appliance-based commands require validation that is only available on Linux hosts.

Information Prerequisites

Before performing any disk copying tasks, you must obtain the following information:

- Appliance IP address - typically is provided by the Infrastructure Engineer.
- IAM login information, Data Transfer Utility configuration files, transfer job ID, and job label - typically is provided by the Project Sponsor.

Generate and Upload RSA Key Pairs

The Data Administrator is responsible for generating and uploading the RSA key pairs. Do not share these generated RSA keys between users. See [Creating a Key Pair](#).

Setting Up an HTTP Proxy Environment

You might need to set up an HTTP proxy environment on the Control Host to allow access to the public internet. This proxy environment allows the Oracle Cloud Infrastructure CLI to communicate with the Data Transfer Appliance Management Service and the appliance over a local network connection. If your environment requires internet-aware applications to use network proxies, configure the Control Host to use your environment's network proxies by setting the standard Linux environment variables on your Control Host.

Assume that your organization has a corporate internet proxy at `http://www-proxy.myorg.com` and that the proxy is an HTTP address at port 80. You would set the following environment variable:

```
export HTTPS_PROXY=http://www-proxy.myorg.com:80
```

If you configured a proxy on the Control Host and the transfer appliance is directly connected to that host, the Control Host tries unsuccessfully to communicate with the transfer appliance using a proxy. Set a `no_proxy` environment variable for the appliance. For example, if the appliance is on a local network at `10.0.0.1`, you would set the following environment variable:

```
export NO_PROXY=10.0.0.1
```

Install and Using the Oracle Cloud Infrastructure Command Line Interface

The Oracle Cloud Infrastructure Command Line Interface (CLI) provides a set of command line-based tools for configuring and running Appliance-Based Data Transfer. Use the Oracle

Cloud Infrastructure CLI as an alternative to running commands from the Console. Sometimes you must use the CLI to complete certain tasks as there is no Console equivalent.

Installation and configuration of the CLIs is described in detail in [Command Line Interface \(CLI\)](#).

USING THE CLI

The CLI must be run as the `root` user.

You can specify CLI options using the following commands:

- `--option <value>` *or*
- `--option=<value>`

The basic CLI syntax is:

```
oci dts <resource> <action> <options>
```

This syntax is applied to the following:

- `oci dts` is the shortened CLI command name
- `job` is an example of a `<resource>`
- `create` is an example of an `<action>`
- Other strings are `<options>`

The following command to create a transfer job shows a typical CLI command construct.

```
oci dts job create --compartment-id ocid1.compartment.oc1..exampleuniqueID --bucket MyBucket --device-type appliance --display-name transfer1
```



Note

In the previous examples, provide a friendly name for the transfer job using the `--display-name` option. Avoid entering confidential information as part of the display name.

Firewall Access

If you have a restrictive firewall in the environment where you are using the Oracle Cloud Infrastructure CLI , you may need to open your firewall configuration to the following IP address ranges: 140.91.0.0/16.

Initializing Authentication to the Appliance



Tip

You can only use the Oracle Cloud Infrastructure CLI to initialize authentication.

Initialize authentication to allow the host machine to communicate with the appliance. Use the values returned from the **Configure Networking** command. See [Configuring the Transfer Appliance Networking for details](#).

To initialize authentication using the CLI

Perform this task using [CLI commands](#). There is no Console equivalent.

At the command prompt on the host, run `oci dts physical-appliance initialize-authentication` to initialize authentication.

```
oci dts physical-appliance initialize-authentication --job-id <job-id> --appliance-cert-fingerprint <fingerprint> --appliance-ip <ip_address> --appliance-label <appliance-label>
```

For example:

```
oci dts physical-appliance initialize-authentication --job-id
ocid1.datatransferjob.region1.phx..exampleuniqueID --appliance-cert-fingerprint
F7:1B:D0:45:DA:04:0C:07:1E:B2:23:82:E1:CA:1A:E9 --appliance-ip 10.0.0.1 --appliance-label XA8XM27EVH
```

When prompted, supply the access token and system. For example:

```
oci dts physical-appliance initialize-authentication --appliance-certfingerprint
86:CA:90:9E:AE:3F:0E:76:E8:B4:E8:41:2F:A4:2C:38 --applianceip 10.0.0.5 --jobid
ocid1.datatransferjob.oc1..exampleuniqueID --appliance-label XAKKJAO9KT
```

CHAPTER 10 Data Transfer

```
Retrieving the Appliance serial id from Oracle Cloud Infrastructure.
Access token ('q' to quit):
Found an existing appliance. Is it OK to overwrite it? [y/n]y
Registering and initializing the authentication between the dts CLI and the appliance
Appliance Info :
  encryptionConfigured : false
  lockStatus            : NA
  finalizeStatus        : NA
  totalSpace            : Unknown
  availableSpace        : Unknown
```

The Control Host can now communicate with the appliance.

To show the status of and storage details about the connected appliance using the CLI

At the command prompt on the host, run `oci dts physical-appliance show` to show the status of the connected appliance.

```
oci dts physical-appliance show
```

For example:

```
Appliance Info :
  encryptionConfigured : false
  lockStatus            : NA
  finalizeStatus        : NA
  totalSpace            : Unknown
  availableSpace        : Unknown
```

Configuring Appliance Encryption

Configure the appliance to use encryption. Oracle Cloud Infrastructure creates a strong passphrase for each appliance. The command securely collects the strong passphrase from Oracle Cloud Infrastructure and sends that passphrase to the Data Transfer service.

If your environment requires Internet-aware applications to use network proxies, ensure that you set up the required Linux environment variables. See for more information.



Important

If you are working with multiple appliances at the same time, be sure the job ID and appliance label you specify in this step matches the physical appliance you are currently working with. You can get the serial number associated with the job ID and appliance label using the Console or the Oracle Cloud Infrastructure CLI. You can find the serial number of the physical appliance on the back of the device on the agency label.



Tip

You can only use the Oracle Cloud Infrastructure CLI to configure encryption.

To configure appliance encryption using the CLI

At the command prompt on the host, run `oci dts physical-appliance configure-encryption` to configure appliance encryption.

```
oci dts physical-appliance configure-encryption --job-id <job_id> --appliance-label <label>
```

For example:

```
oci dts physical-appliance configure-encryption --job-id  
ocid1.datatransferjob.region1.phx..exampleuniqueID --appliance-label XA8XM27EVH
```

Unlocking the Appliance

Before you can write data to the transfer appliance, you must unlock the appliance. Unlocking the transfer appliance requires the strong passphrase that is created by Oracle Cloud Infrastructure for each appliance. Unlocking can be accomplished in two different ways:

- If you provide the `--job-id` and `--appliance-label` when running the `unlock` command, the data transfer system retrieves the passphrase from Oracle Cloud Infrastructure and sends it to the transfer appliance during the unlock operation.
- You can query Oracle Cloud Infrastructure for the passphrase and provide that passphrase when prompted during the unlock operation.



Important

It can take up to 10 minutes to unlock an appliance the first time. Subsequent unlocks are not as time consuming.

To retrieve the passphrase to unlock the appliance using the CLI

At the command prompt on the host, run `dts physical-appliance unlock` with `--job-id` and `--appliance-label` to unlock the appliance.

```
oci dts physical-appliance unlock --job-id <job_id> --appliance-label <label>
```

For example:

```
oci dts physical-appliance unlock --job-id ocid1.datatransferjob.region1.phx..exampleuniqueID --  
appliance-label XA8XM27EVH
```

To query Oracle Cloud Infrastructure for the passphrase to provide to unlock the appliance using the CLI

At the command prompt on the host, run `dts appliance get-passphrase` to obtain the passphrase from the Oracle Cloud Infrastructure.

```
oci dts appliance get-passphrase --job-id <job_id> --appliance-label <label>
```

Then, run `dts physical-appliance unlock` without `--job-id` and `--appliance-label` and supply the passphrase when prompted.

```
oci dts physical-appliance unlock
```

Creating NFS Datasets

A dataset is a collection of files that are treated similarly. You can write up to 100 million files onto the appliance for migration to Oracle Cloud Infrastructure. We currently support one dataset per appliance. Appliance-Based Data Transfer supports NFS versions 3, 4, and 4.1 to write data to the appliance. In preparation for writing data, create and configure a dataset to write to. See [Datasets](#) for complete details on all tasks related to datasets.

To create a dataset using the CLI

At the command prompt on the host, run `oci dts nfs-dataset create` to create a dataset.

```
oci dts nfs-dataset create --name <dataset_name>
```

For example:

```
oci dts nfs-dataset create --name nfs-ds-1
```

ACTIVATING THE DATASET

Activation creates the NFS export, making the dataset accessible to NFS clients.

To activate the dataset

At the command prompt on the host, run `oci dts nfs-dataset activate` to activate the NFS dataset.

```
oci dts nfs-dataset activate --name <dataset_name>
```

For example:

```
oci dts nfs-dataset activate --name nfs-ds-1
```

CHAPTER 10 Data Transfer

CONFIGURING EXPORT SETTINGS ON THE DATASET

To configure export settings on a dataset

At the command prompt on the host, run `oci dts nfs-dataset set-export` to configure export settings on an NFS dataset.

```
oci dts nfs-dataset set-export --name <dataset_name> --rw=true --world=true
```

For example:

```
oci dts nfs-dataset set-export --name nfs-ds-1 --rw=true --world=true
```

Here is another example of creating the export to give read/write access to a subnet:

```
oci dts nfs-dataset set-export --name nfs-ds-1 --ip 10.0.0.0 --subnet-mask-length 24 --rw true --world false
```

Setting Your Data Host as an NFS Client

Set up your Data Host as an NFS client:

- For Debian or Ubuntu, install the `nfs-common` package. For example:

```
sudo apt-get install nfs-common
```

- For Oracle Linux or Red Hat Linux, install the `nfs-utils` package. For example:

```
sudo yum install nfs-utils
```

Mounting the NFS Share

To mount the NFS share

At the command prompt on the Data Host, use the `mount` command to mount the NFS share.

```
mount -t nfs <appliance_ip>:/data/<dataset_name><mountpoint>
```

For example:

```
mount -t nfs 10.0.0.1:/data/nfs-ds-1 /mnt/nfs-ds-1
```

After the NFS share is mounted, you can write data to the share.

Copying Files to the NFS Share

Copy your file to the appliance using normal file system tools.



Important

You can only copy regular files to transfer appliances. Special files (for example, symbolic links, device special, sockets, and pipes) cannot be copied directly. To transfer special files, create a tar archive of these files and copy the tar archive to the transfer appliance.

Deactivating the Dataset



Note

Deactivating the dataset is only required if you are running appliance commands using the Data Transfer Utility. If you are using the Oracle Cloud Infrastructure CLI to run your Appliance-Based Data Transfer, you can skip this step and proceed to [Sealing the Dataset](#).

After you are done writing data, deactivate the dataset. Deactivation removes the NFS export on the dataset, disallowing any further writes.

To deactivate the dataset

At the command prompt on the host, run `dts nfs-dataset deactivate` to deactivate the NFS dataset.

CHAPTER 10 Data Transfer

```
dts nfs-dataset deactivate --name <dataset_name>
```

For example:

```
dts nfs-dataset deactivate --name nfs-ds-1
```

Sealing the Dataset

Sealing a dataset stops all writes to the dataset. Sealing a dataset is a long running process that can take some time to complete. The completion time depends upon the number of files and total amount of data that was copied to the appliance.

If you issue the `seal` command without the `--wait` option, the seal operation is triggered and runs in the background. You are returned to the command prompt and can use the `seal-status` command to monitor the sealing status. If you issue the `seal` command with the `--wait` option, the seal operation is triggered and continues to provide status updates until sealing completion.



Important

You can only copy regular files to transfer appliances. Special files (for example, symbolic links, device special, sockets, and pipes) cannot be copied directly. To transfer special files, create a tar archive of these files and copy the tar archive to the transfer appliance.

The sealing operation generates a manifest across all files in the dataset. The manifest contains an index of the copied files and generated data integrity hashes.

TO SEAL THE DATASET USING THE CLI

At the command prompt on the host, run `oci dts nfs-dataset seal` to seal the NFS dataset.

```
oci dts nfs-dataset seal --name <dataset_name> [--wait]
```

For example:

CHAPTER 10 Data Transfer

```
oci dts nfs-dataset seal --name nfs-ds-1
Seal initiated. Please use seal-status command to get progress.
```

TO MONITOR THE DATASET SEALING PROCESS USING THE CLI

At the command prompt on the host, run `oci dts nfs-dataset seal-status` to monitor the dataset sealing process.

```
oci dts nfs-dataset seal-status --name <dataset_name>
```

For example, here is the status that is issued upon sealing completion:

```
oci dts nfs-dataset seal-status --name nfs-ds-1

Seal Status :
  success      : true
  failureReason : *** none ***
  startTime    : 2018/07/10 18:24:05 EDT
  endTime     : 2018/07/10 18:24:06 EDT
  numFilesToProcess : 2000
  numFilesProcessed : 2000
  bytesToProcess  : 1.95 GB
  bytesProcessed  : 1.95 GB
```



Note

If changes are necessary after sealing a dataset or finalizing an appliance, you must reopen the dataset to modify the contents. See [Reopening a Dataset](#).

DOWNLOADING THE DATASET SEAL MANIFEST

After sealing the dataset, you can optionally download the dataset's seal manifest to a user-specified location. The manifest file contains the checksum details of all the files. The transfer site uploader consults the manifest file to determine the list of files to upload to object storage. For every uploaded file, it validates that the checksum reported by object storage matches the checksum in manifest. This validation ensures that no files got corrupted in transit.

To download the dataset seal manifest file using the CLI

At the command prompt on the host, run `oci dts nfs-dataset get-seal-manifest` to download the seal manifest.

```
oci dts nfs-dataset get-seal-manifest --name <dataset_name> --output-file <file_path>
```

For example:

```
oci dts nfs-dataset get-seal-manifest --name nfs-ds-1 --output-file ~/Downloads/seal-manifest
```

Finalizing the Appliance



Tip

You can only use the CLI commands to finalize the appliance.

Finalizing an appliance tests and copies the following to the appliance:

- [Upload user configuration credentials](#)
- Private PEM key details
- Name of the upload bucket

The credentials, API key, and bucket are required for Oracle to be able to upload your data to Oracle Cloud Infrastructure Object Storage. When you finalize an appliance, you can no longer access the appliance for dataset operations unless you unlock the appliance. See [Reopening a Dataset](#) if you need to unlock an appliance that was finalized.



Important

If you are working with multiple appliances at the same time, be sure the job ID and appliance label you specify in this step matches the physical appliance you are currently working with. You can get the serial number associated with the job ID and appliance label using the Console or the Oracle Cloud Infrastructure CLI. You can find the serial number of the physical appliance on the back of the device on the agency label.

To finalize the appliance

1. [Seal](#) the dataset before finalizing the appliance.
2. At the command prompt on the host, run `oci dts physical-appliance finalize` to finalize an appliance.

```
oci dts physical-appliance finalize --job-id <job_id> --appliance-label <label>
```

For example:

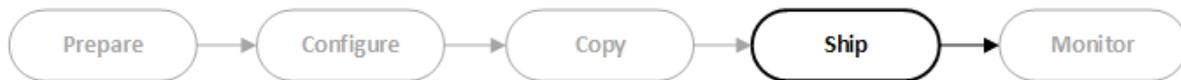
```
oci dts physical-appliance finalize --job-id ocidl.datatransferjob.region1.phx..exampleuniqueID --  
appliance-label XA8XM27EVH
```



Note

If changes are necessary after sealing a dataset or finalizing an appliance, you must reopen the dataset to modify the contents. See [Reopening a Dataset](#).

Shipping the Appliance



This topic describes the tasks associated with shipping the transfer appliance containing the copied data to Oracle. The Infrastructure Engineer role typically performs these tasks. See [Roles and Responsibilities](#).



Note

You can only run Oracle Cloud Infrastructure CLI commands from a Linux host. This differs from running CLI commands for other Oracle Cloud Infrastructure Services on a variety of host operating systems. Appliance-based commands require validation that is only available on Linux hosts.

Shutting Down the Transfer Appliance

Shut down the appliance before packing up and shipping the appliance back to Oracle.

To shut down the appliance

Using the terminal emulator on the host machine, select **Shutdown** from the appliance serial console.

Packing and Shipping Transfer Appliance to Oracle

Return the appliance to Oracle within 30 days. If you need the transfer appliance beyond the standard 30-day window, you can file a Service Request (SR) to ask for an extension of up to 60 days.



Important

Review and follow the instructions that were provided in the transit case with the appliance.

To pack and ship the appliance

1. Unplug the power cord from the power source and detach the other end of the cord from the appliance.
2. Disconnect the appliance from your network.
3. Remove the return shipment tamper-evident security tie from the transit case.
4. Place the transfer appliance, power cord, and serial cable in the transit case.



Warning

We recommend assistance lifting and placing the transfer appliance back into the transit case. The total shipping weight is about 64 lbs (29.0299 kg) and appliance weight is 38 lbs (17.2365 kg).

5. Close and secure the transit case with the return tamper-evident security tie.
6. Loop the top of the plastic tie-on tag with return shipping label through the handle of the transit case. Remove the protective tape from the back of the tie-on tag, exposing the adhesive area on which to secure the tag onto itself. Use the provided zip tie to secure the tie-on tag to the handle.
7. Return the transit case to FedEx by doing one of the following:
 - Drop off the packed, sealed, and labeled transit case to an FedEx Authorized ShipCenter location or a nearby FedEx Office location. *Obtain a receipt from the vendor to certify transfer of custody.*

- Schedule a pickup with FedEx at your location. Ensure that the transit case is packed, sealed, and labeled before FedEx arrives for pickup.

The shipping vendor notifies Oracle when the transfer appliance is shipped back to Oracle for upload to Oracle Cloud Infrastructure Object Storage.

Monitoring the Appliance Return



This topic describes the tasks to be done after the data transfer is complete the transfer appliance has been returned to Oracle. The Project Sponsor role typically performs these tasks. See [Roles and Responsibilities](#).



Note

You can only run Oracle Cloud Infrastructure CLI commands from a Linux host. This differs from running CLI commands for other Oracle Cloud Infrastructure Services on a variety of host operating systems. Appliance-based commands require validation that is only available on Linux hosts.

Monitoring the Status of Your Transfer Appliance Return Shipment

The shipping vendor notifies Oracle when your transfer appliance is picked up and shipped back for upload to Oracle Cloud Infrastructure Object Storage.

To monitor the status of your transfer appliance return shipment using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job and transfer appliance that you shipped back to Oracle for data upload.
3. Under **Transfer Appliances**, look at the **Status** field.

To monitor the status of your transfer appliance using the CLI

At the command prompt on the host, run `oci dts appliance show` to monitor a transfer appliance status.

```
oci dts appliance show --job-id <job_id> -appliance <label>
```

TRANSFER APPLIANCE STATUS VALUES

Here are the transfer appliance status values, listed in alphabetic order:

CANCELED

You can change your mind about uploading your data to Oracle Cloud Infrastructure Object Storage and cancel your transfer appliance. Ship the appliance back to Oracle and then cancel the appliance. Oracle always uses secure wipe tools on the boot and data areas whenever a transfer appliance is returned.

COMPLETE

Oracle completed your transfer appliance data upload. Your data is available in your designated bucket in Oracle Cloud Infrastructure Object Storage.

CUSTOMER LOST

You have not returned a data transfer appliance within the required 90 days.

DELIVERED

Oracle received a delivery confirmation from the shipping vendor that your transfer appliance was delivered. When the appliance is delivered, Oracle provides the date and time the appliance was received in the transfer appliance details. Appliance usage tracking begins.

ERROR

Oracle encountered an unrecoverable error trying to process your transfer appliance. Oracle cannot upload your data from the appliance. To protect your data, Oracle uses secure wipe tools on the boot and data areas any transfer appliance that cannot be processed.

Complete another request for a transfer appliance.

ORACLE PREPARING

Oracle approved your transfer appliance request. The status displays preparing until the transfer appliance is shipped to you.

ORACLE RECEIVED

Oracle received your transfer appliance shipment. The status displays Oracle received until Oracle begins processing and uploading your transfer appliance.

ORACLE RECEIVED CANCELED

You canceled your transfer appliance after you shipped the appliance back to Oracle. Oracle received your canceled transfer appliance. Oracle does *not* upload the appliance data.

PREPARING

You activated your transfer appliance. You can now copy your data onto the transfer appliance. The status displays preparing until you ship the transfer appliance back to Oracle.

PROCESSING

Oracle is processing and uploading the data from your transfer appliance. The status displays the processing status until Oracle completes uploading your data from your transfer appliance.

REJECTED

Oracle denied your transfer appliance request.



Important

If your appliance request is denied and you have questions, contact your Sales Representative or file a Service Request (SR).

REQUESTED

You successfully completed your request for a transfer appliance. The status displays requested until Oracle approves your transfer appliance request.

RETURN SHIPPED

Oracle received confirmation from the shipping vendor that you shipped your transfer appliance back to Oracle. The status displays return shipped until Oracle receives your transfer appliance.

RETURN SHIPPED CANCELED

You canceled your transfer appliance after the appliance was delivered to you or after you shipped the appliance back to Oracle. Oracle received confirmation from the shipping vendor that your canceled transfer appliance is on the way back to Oracle. The status displays return shipped canceled until Oracle receives your transfer appliance.

SHIPPING

Oracle completed the necessary preparations and shipped your transfer appliance. When the appliance is shipped, Oracle provides the serial number of the appliance, the shipping

vendor, and the tracking number in the appliance details. The status displays shipping until the appliance is delivered to you.

Reviewing the Upload Summary

Oracle creates upload summary log files for each uploaded appliance. These log files are placed in the bucket where data was uploaded to Oracle Cloud Infrastructure. The upload summary file compares the appliance's manifest file to the contents of the target Oracle Cloud Infrastructure Object Storage bucket after file upload.



Note

If you chose to upload your data to an Archive Storage bucket, you must first restore the log file object before you can download that file for review.

The top of the log report summarizes the overall file processing status:

```
P - Present: The file is present in both the device and the target bucket
M - Missing: The file is present in the device but not the target bucket. It was likely uploaded and then deleted by another user before the summary was generated.
C - Name Collision: The file is present in the manifest but a file with the same name but different contents is present in the target bucket.
U - Unreadable: The file is not readable from the disk
N - Name Too Long: The file name on disk is too long and could not be uploaded
```

Complete file upload details follow the summary.

```
#####
##### SUMMARY FOR DISK [WDH0B87L] #####
Generated 2017-09-08 19:46:36
TOTAL : 1110
### P present: 1110
### M missing: 0
### C name collision: 0
### U unreadable: 0
### N nameTooLong: 0
#####
| STATUS | NAME | LAST_MODIFIED | SIZE (MB) | MD5 | ETag |
| present | small/01/T6QKX | Fri Sep 08 19:04:54 UTC 2017 | 10.00 | 8c1kXbWU793H2KHif8m6w== | 58B2F70D1C874AAA0E953824310ACAC52 |
| present | small/01/FAFKU | Fri Sep 08 19:12:20 UTC 2017 | 10.00 | 8c1kXbWU793H2KHif8m6w== | 58B2EC461E4608E0E953824310AC6DF6 |
| present | small/01/EVVPD | Fri Sep 08 19:02:42 UTC 2017 | 10.00 | 8c1kXbWU793H2KHif8m6w== | 58B2ECDFa749460E0E953824310AC385A |
| present | small/01/2U02H | Fri Sep 08 19:13:06 UTC 2017 | 10.00 | 8c1kXbWU793H2KHif8m6w== |
```

If you upload more than 100,000 files, the upload details are broken into multiple pages. You can only download the first page from the Console. Download the rest of the pages directly

from the Object Storage bucket. The subsequent pages have the same object name as the first page, but have an enumerated suffix.

Closing a Transfer Job



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to close a transfer job.

Typically, you would close a transfer job when no further transfer job activity is required or possible. Closing a transfer job requires that the status of all associated transfer appliances be returned, canceled, or deleted.

To close a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer package for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
4. Click **Close Transfer Job**.

To close a transfer job using the CLI

At the command prompt on the host, run `oci dts job close` to close a transfer job.

```
oci dts job close --job-id <job_id>
```

What's Next

You have completed the process of setting up, running, and monitoring the Appliance-Based Data Transfer. If you determine that another appliance-based data transfers is required,

repeat the procedure from the beginning.

Appliance Data Transfer Reference

This topic provides complete task details for certain components associated with Appliance-Based Data Transfers. Use this topic as a reference to learn and use commands associated with components included in the Appliance-Based Data Transfer procedure.

Transfer Jobs

A transfer job represents the collection of files that you want to transfer and signals the intention to upload those files to Oracle Cloud Infrastructure. A transfer job combines at least one transfer disk with a transfer package. Identify which compartment and Object Storage bucket that Oracle is to upload your data to.



Note

It is recommended that you create a compartment for each transfer job to minimize the required access your tenancy.



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to create a transfer job.

CREATING TRANSFER JOBS

Create the transfer job in the same compartment as the upload bucket and supply a human-readable name for the transfer job. Avoid entering confidential information when providing transfer job names.

Creating a transfer job returns a job ID that you specify in other transfer tasks. For example:

```
ocid1.datatransferjob.region1.phx..<unique_ID>
```

To create a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Select the designated compartment you are to use for data transfers from the drop-down list.
A list of transfer jobs that have already been created is displayed.
3. Click **Create Transfer Job**.
4. In the **Create Transfer Job** dialog, enter a **Job Name**, and select the **Upload Bucket** from the drop-down list.
Avoid entering confidential information in the transfer job name.
5. Select **Disk** for the **Transfer Device Type**.
6. Click **Create Transfer Job**.

To create a transfer job using the CLI

At the command prompt on the host, run `oci dts job create` to create a transfer job. The `<display_name>` is the name of the transfer job. Avoid entering confidential information in the transfer job name.

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type disk
```

Optionally, you can specify one or more free-form or defined tags when you create a transfer job. For more information about tagging, see [Resource Tags](#).

To specify free-form tags when creating a job:

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type disk --freeform-tags '{ "<tag_key>":"<value>" }'
```

To specify defined tags when creating a job:

CHAPTER 10 Data Transfer

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type disk --defined-tags '{ "<tag_namespace>": { "<tag_key>": "<value>" } }'
```



Note

Users create tag namespaces and tag keys with the required permissions. These items must exist before you can specify them when creating a job. See [Working with Defined Tags](#) for details.

To specify multiple tags, comma separate the JSON-formatted key/value pairs:

```
oci dts job create --bucket <bucket_name> --compartment-id <compartment_id> --display-name <display_name> --device-type disk --freeform-tags '{ "<tag_key>": "<value>" }, { "<tag_key>": "<value>" }'
```

DISPLAYING TRANSFER JOBS

To display the list of transfer jobs using the Console

Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.

To display the list of transfer jobs using the CLI

At the command prompt on the host, run `dts job list` to display the list of transfer jobs.

```
oci dts job list --compartment-id <compartment_id>
```

When you use the CLI command to list jobs, tagging details are also included in the output if you specified tags.

DISPLAYING TRANSFER JOB DETAILS

To display the details of a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.

To display the details of a transfer job using the CLI

At the command prompt on the host, run `oci dts job show` to display the details of a transfer job.

```
oci dts job show --job-id <job_id>
```

When you use the CLI command to display the details of a job, tagging details are also included in the output if you specified tags.

EDITING TRANSFER JOBS

To edit the name of a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer job that you want to edit.
3. Click the Actions icon (three dots), and then click **Edit**.
4. Edit the name of the transfer job.
Avoid entering confidential information in the transfer job name.
5. Click **Save**.

To edit the name of a transfer job using the CLI

At the command prompt on the host, run `oci dts job update` to edit the name (`--display-name`) of a transfer job. The `<display_name>` is the new name of the transfer job. Avoid entering confidential information in the transfer job name.

```
oci dts job update --job-id <job_id> --display-name <display_name>
```

To edit the tags associated with a transfer job using the CLI

At the command prompt on the host, run `oci dts job update` to edit the tags associated with a transfer job. The CLI command **replaces** any existing tags with the new key/value pairs you specify.

To edit free-form tags, provide the replacement key/value pairs:

```
oci dts job update --job-id <job_id> --freeform-tags '{ "<tag_key>":"<value>" }'
```

To edit defined tags, provide the replacement key value pairs:

```
oci dts job update --job-id <job_id> --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```

To delete the tags associated with a transfer job using the CLI

At the command prompt on the host, run `oci dts job update` to delete the tags associated with a transfer job. The CLI command **replaces** any existing tags with the new key/value pairs you specify. If you want to delete some of the tags, you would specify new tag string that does not contain the key/value pair you want to delete.

Partial tag deletion is handled in the same way as you edit tags:

- To edit free-form tags, provide the replacement key/value pairs:

```
oci dts job update --job-id <job_id> --freeform-tags '{ "<tag_key>":"<value>" }'
```

- To edit defined tags, provide the replacement key value pairs:

```
oci dts job update --job-id <job_id> --defined-tags '{ "<tag_namespace>": { "<tag_key>":"<value>" } }'
```

To delete all free-form tags:

```
oci dts job update --job-id <job_id> --freeform-tags '{}'
```

To delete all defined tags:

```
oci dts job update --job-id <job_id> --defined-tags '{}'
```

DELETING TRANSFER JOBS

You can delete transfer jobs when they are in the Initiated, Preparing, and Close states.

To delete a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer job that you want to delete.
3. Click the Actions icon (three dots), and then click **Delete**.
Alternatively, you can delete a transfer job from the **View Details** page.
4. Confirm the deletion when prompted.

To delete a transfer job using the CLI

At the command prompt on the host, run `oci dts job delete` to delete a transfer job.

```
oci dts job delete --job-id <job_id>
```

CLOSING TRANSFER JOBS



Tip

You can use the Console or the Oracle Cloud Infrastructure CLI to close a transfer job.

Typically, you would close a transfer job when no further transfer job activity is required or possible. Closing a transfer job requires that the status of all associated transfer packages be returned, canceled, or deleted. In addition, the status of all associated transfer disks must be complete, in error, missing, canceled, or deleted.

To close a transfer job using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer package for which you want to display the details.
3. Click the Actions icon (three dots), and then click **View Details**.
Alternatively, click the hyper-linked name of the transfer job.
4. Click **Close Transfer Job**.

To close a transfer job using the CLI

At the command prompt on the host, run `oci dts job close` to close a transfer job.

```
oci dts job close --job-id <job_id>
```

Appliances

This section describes tasks associated with the Oracle-provided appliance.



Tip

You can use the Console or the Oracle Cloud Infrastructure CLIs to request an appliance.



Tip

We recommend that you identify the data you intend to upload and make data copy preparations before requesting the appliance.

Creating an appliance request returns an Oracle-assigned appliance label. For example:

```
XA8XM27EVH
```

To request an appliance using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
Choose the transfer job for which you want to request an appliance.
2. Under **Transfer Appliances**, click **Request Transfer Appliance**.
3. In the **Request Transfer Appliance** dialog box, provide the shipping address details where you want the appliance sent.
 - **Company Name:** Required. Specify the name of the company that owns the data being migrated to Oracle Cloud Infrastructure.
 - **Recipient Name:** Required. Specify the name of the recipient to send the appliance to.

- **Recipient Phone Number:** Required. Specify the recipient's phone number.
- **Recipient Email Address:** Required. Specify the recipient's email address.
- **Care Of:** Optional intermediary party responsible for transferring the appliance shipment from the delivery vendor to the intended recipient.
- **Address Line 1:** Required. Specify the street address to send the appliance to.
- **Address Line 2:** Optional identifying address details like building, suite, unit, or floor information.
- **City/Locality:** Required. Specify the city or locality.
- **State/Province/Region:** Required. Specify the state, province, or region.
- **Zip/Postal Code:** Specify the zip code or postal code.
- **Country:** Required. Select the country.

4. Click **Request Transfer Appliance**.

To request an appliance using the CLIs

At the command prompt on the host, run `oci dts appliance request` to request an appliance.

Here are the minimum requirements for the appliance request:

```
oci dts appliance request --job-id <job_id> --addressee <addressee> --address1 <address_line1> --city-or-locality <city_or_locality> --state-or-region <state_or_region>--country<country>--zip-code<zip>
```

In addition, you can specify these optional fields in the request:

```
--care-of <care_of>  
--address2 <address_line2>, --address3 <address_line3>, and --address4 <address_line4>  
--email <email_address>  
--phone-number <phone_number>  
--profile <profile>
```

When you submit an appliance request, Oracle generates a unique label (name) to identify the appliance and your request is sent to Oracle for approval and processing.

MONITORING THE APPLIANCE REQUEST STATUS

The time it takes to approve, prepare, and ship your appliance request varies and depends on various factors, including current available inventory. Oracle provides status updates daily throughout the appliance request and ship process.

To monitor the status of your appliance request using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find and select the transfer job for which you want to monitor associated appliance requests.
3. Under **Transfer Appliances**, find the appliance label Oracle assigned to your appliance request and look at the **Status** field.

Here are the key status values to look for when monitoring your appliance request:

- **Requested:** You successfully completed your request for an appliance. The status displays **Requested** until Oracle approves your appliance request.
- **Rejected:** Oracle denied your appliance request.



Important

If your appliance request is denied and you have questions, contact your Sales Representative or file a Service Request (SR).

- **Oracle Preparing:** Oracle approved your appliance request. The status displays **Oracle Preparing** until the appliance is shipped to you.
- **Shipping:** Oracle completed the necessary preparations and shipped your appliance. When the appliance is shipped, Oracle provides the serial number of the appliance, the shipping vendor, and the tracking number in the **Transfer Appliance Details**. The status displays **Shipping** until the appliance is delivered to you.

- **Delivered:** The shipping vendor delivered your appliance. When the appliance is delivered, Oracle provides the date and time the appliance was received in the **Transfer Appliance Details**. The status displays **Delivered**.

To monitor the status of your appliance request using the CLI

At the command prompt on the host, run `oci dts appliance show` to monitor your appliance request.

```
oci dts appliance show --job-id <job_id> --appliance-label <label>
```

DISPLAYING THE LIST OF APPLIANCES



Tip

You can use the Console to get a list of appliances that are associated with a job.

To display the list of appliances using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Choose the transfer job for which you want to display the list of associated appliances. The list of appliances is displayed below the transfer job details.

To display the list of appliances using the CLI commands

At the command prompt on the host, run `oci dts appliance list` to display the list of transfer jobs.

```
oci dts appliance list --job-id <job_id>
```

DISPLAYING THE DETAILS OF AN APPLIANCE

To display the details of an appliance using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the transfer job for which you want to display the details of an associated appliance.
The list of appliances is displayed below the transfer job details.
3. Find the appliance for which you want to display the details.
4. Click the Actions icon (three dots), and then click **View Details**.

To display the details of an appliance using the CLI commands

At the command prompt on the host, run `oci dts appliance show` to display the details of an appliance.

```
oci dts appliance show --job-id <job_id> --appliance-label <label>
```

EDITING THE APPLIANCE REQUEST SHIPPING INFORMATION

You can only edit the shipping information when the status is **Requested**.

To edit the appliance request shipping information using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the **Requested** appliance that you want to edit the shipping information.
3. Click the Actions icon (three dots), and then click **Edit**.
4. Edit the shipping information for the appliance.

5. Click **Save**.

To edit the appliance request shipping information using the CLI

At the command prompt on the host, run `oci dts appliance update` to edit the shipping information for the appliance.

```
oci dts appliance update-shipping-address --job-id <job_id>--appliance-label<label><changed_fields>
```

The `<changed_fields>` variable represents one or more of the following shipping address fields that you want to update:

```
--addressee <addressee> --careOf <care_of> --address1 <street_address> --city <city> --state <addressee> --zip <zip> --country <country> --phone <phone>
```

DELETING AN APPLIANCE REQUEST

You can delete an appliance request before Oracle approves the request—the status must be **Requested**. For example, you initiated the transfer by creating a transfer job and requested an appliance, but changed your mind.

To delete an appliance request using the Console

1. Open the navigation menu. Under **Core Infrastructure**, go to **Object Storage** and click **Data Transfer**.
2. Find the data transfer job and appliance request that you want to delete.
3. Click the Actions icon (three dots), and then click **Delete**.
Alternatively, you can delete an appliance request from the **Transfer Appliance Details** page.
4. Confirm the deletion when prompted.

To delete an appliance request using the CLI commands

At the command prompt on the host, run `oci dts appliance delete` to delete an appliance.

```
oci dts appliance delete --job-id <job_id>--appliance-label<label>
```

Datasets

A dataset is a collection of files that are treated similarly. You can write up to 100 million files onto the appliance for transfer. We currently support one dataset per appliance.

CREATING DATASETS

Appliance data transfer supports NFSversion 3, 4, and 4.1 to write data to the appliance. In preparation for writing data, create and configure a dataset to write to.

To create a dataset using the CLI

At the command prompt on the host, run `oci dts nfs-dataset create` to create a dataset.

```
oci dts nfs-dataset create --name <dataset_name>
```

For example:

```
oci dts nfs-dataset create --name nfs-ds-1
```

ACTIVATING THE DATASET

Activation creates the NFS export, making the dataset accessible to NFS clients.

To activate the dataset

At the command prompt on the host, run `oci dts nfs-dataset activate` to activate the NFS dataset.

```
oci dts nfs-dataset activate --name <dataset_name>
```

For example:

CHAPTER 10 Data Transfer

```
oci dts nfs-dataset activate --name nfs-ds-1
```

CONFIGURING EXPORT SETTINGS ON THE DATASET

To configure export settings on a dataset

At the command prompt on the host, run `oci dts nfs-dataset set-export` to configure export settings on an NFS dataset.

```
oci dts nfs-dataset set-export --name <dataset_name> --rw=true --world=true
```

For example:

```
oci dts nfs-dataset set-export --name nfs-ds-1 --rw=true --world=true
```

Here is another example of creating the export to give read/write access to a subnet:

```
oci dts nfs-dataset set-export --name nfs-ds-1 --ip=10.0.0.0 --subnet-mask-length=24 --rw=true --world=false
```

DEACTIVATING THE DATASET



Note

Deactivating the dataset is only required if you are running appliance commands using the Data Transfer Utility. If you are using the Oracle Cloud Infrastructure CLI to run your Appliance-Based Data Transfer, you can skip this step and proceed to [Sealing the Dataset](#).

After you are done writing data, deactivate the dataset. Deactivation removes the NFS export on the dataset, disallowing any further writes.

To deactivate the dataset

At the command prompt on the host, run `dts nfs-dataset deactivate` to deactivate the

NFS dataset.

```
dts nfs-dataset deactivate --name <dataset_name>
```

For example:

```
dts nfs-dataset deactivate --name nfs-ds-1
```

SEALING THE DATASET

Sealing a dataset stops all writes to the dataset. Sealing a dataset is a long running process that can take some time to complete. The completion time depends upon the number of files and total amount of data that was copied to the appliance.

If you issue the `seal` command without the `--wait` option, the seal operation is triggered and runs in the background. You are returned to the command prompt and can use the `seal-status` command to monitor the sealing status. If you issue the `seal` command with the `--wait` option, the seal operation is triggered and continues to provide status updates until sealing completion.



Important

You can only copy regular files to transfer appliances. Special files (for example, symbolic links, device special, sockets, and pipes) cannot be copied directly. To transfer special files, create a tar archive of these files and copy the tar archive to the transfer appliance.

The sealing operation generates a manifest across all files in the dataset. The manifest contains an index of the copied files and generated data integrity hashes.

To seal the dataset using the CLI

At the command prompt on the host, run `oci dts nfs-dataset seal` to seal the NFS dataset.

CHAPTER 10 Data Transfer

```
oci dts nfs-dataset seal --name <dataset_name> [--wait]
```

For example:

```
oci dts nfs-dataset seal --name nfs-ds-1
Seal initiated. Please use seal-status command to get progress.
```

To monitor the dataset sealing process using the CLI

At the command prompt on the host, run `oci dts nfs-dataset seal-status` to monitor the dataset sealing process.

```
oci dts nfs-dataset seal-status --name <dataset_name>
```

For example, here is the status that is issued upon sealing completion:

```
oci dts nfs-dataset seal-status --name nfs-ds-1

Seal Status :
success      : true
failureReason : *** none ***
startTime    : 2018/07/10 18:24:05 EDT
endTime      : 2018/07/10 18:24:06 EDT
numFilesToProcess : 2000
numFilesProcessed : 2000
bytesToProcess  : 1.95 GB
bytesProcessed  : 1.95 GB
```

DOWNLOADING THE DATASET SEAL MANIFEST

After sealing the dataset, you can optionally download the dataset's seal manifest to a user-specified location. The manifest file contains the checksum details of all the files. The transfer site uploader consults the manifest file to determine the list of files to upload to object storage. For every uploaded file, it validates that the checksum reported by object storage matches the checksum in manifest. This validation ensures that no files got corrupted in transit.

To download the dataset seal manifest file using the CLI

At the command prompt on the host, run `oci dts nfs-dataset get-seal-manifest` to download the seal manifest.

CHAPTER 10 Data Transfer

```
oci dts nfs-dataset get-seal-manifest --name <dataset_name> --output-file <file_path>
```

For example:

```
oci dts nfs-dataset get-seal-manifest --name nfs-ds-1 --output-file ~/Downloads/seal-manifest
```

REOPENING A DATASET



Tip

You can only use the CLI command to reopen a dataset.

If changes are necessary after sealing a dataset or finalizing an appliance, you must reopen the dataset to modify the contents. Make the required changes and again [seal the dataset](#). Resealing the dataset generates a new manifest.



Note

If an appliance is rebooted or power cycled, follow the instructions in this topic to reopen the dataset.

STEP 1: UNLOCKING THE APPLIANCE

Before you can write data to the transfer appliance, you must unlock the appliance. Unlocking the transfer appliance requires the strong passphrase that is created by Oracle Cloud Infrastructure for each appliance. Unlocking can be accomplished in two different ways:

- If you provide the `--job-id` and `--appliance-label` when running the `unlock` command, the data transfer system retrieves the passphrase from Oracle Cloud Infrastructure and sends it to the transfer appliance during the unlock operation.
- You can query Oracle Cloud Infrastructure for the passphrase and provide that passphrase when prompted during the unlock operation.

To retrieve the passphrase to unlock the appliance

At the command prompt on the host, run `oci dts physical-appliance unlock` with `--job-id` and `--appliance-label` to unlock the appliance.

```
oci dts physical-appliance unlock --job-id <job_id> --appliance-label <label>
```

For example:

```
oci dts physical-appliance unlock --job-id ocid1.datatransferjob.region1.phx..exampleuniqueID --  
appliance-label XA8XM27EVH
```

To query Oracle Cloud Infrastructure for the passphrase to provide to unlock the appliance

At the command prompt on the host, run `oci dts appliance get-passphrase` to obtain the passphrase from the Oracle Cloud Infrastructure.

```
oci dts appliance get-passphrase --job-id <job_id> --appliance-label <label>
```

Then, run `oci dts physical-appliance unlock` without `--job-id` and `--appliance-label` and supply the passphrase when prompted.

```
oci dts physical-appliance unlock
```

STEP 2: REOPENING THE APPLIANCE

Reopen the dataset to write data to the appliance again.

To reopen an NFS dataset

At the command prompt on the host, run `oci dts nfs-dataset reopen` to reopen an NFS dataset.

```
oci dts nfs-dataset reopen --name <dataset_name>
```

STEP 3: REPEAT STEPS TO WRITE DATA TO THE APPLIANCE

Repeat the same tasks you performed when you originally wrote data to the appliance beginning with [activating the dataset](#) in the [Copying Files to the NFS Share](#) section.

Troubleshooting

This topic describes various troubleshooting issues related to the Data Transfer Service.

Troubleshooting the Appliance

You can generate performance information for troubleshooting issues with the appliance through the terminal emulator on the host machine. Select **Collect Appliance Diagnostic Information** from the [serial console configuration menu](#). The diagnostic tool generates system, network, storage, and performance data while the transfer job is running. It then forwards the data to the appliance serial console. Here you can scroll through the terminal to view it.

You can also use the log capture feature of the serial port emulator to capture the output. Serial port emulators often support the ability to copy the session to a file. Refer to the documentation of your serial port emulation package for instructions. Copying to a log file is useful if you need assistance from Oracle or if your emulation session does not allow you to scroll back and see all the output.

For each operation, the display shows exactly what command was executed and all the options.

Here is an example of the diagnostic output:

```
-----  
- systemctl -l --type=service --state=active -  
-----  
UNIT LOAD ACTIVE SUB DESCRIPTION  
auditd.service loaded active running Security Auditing Service  
blk-availability.service loaded active exited Availability of block devices  
chronyd.service loaded active running NTP client/server  
console-diags@39-3147-1001.service loaded active running Diagnostic Collection Server for the XA (PID  
3147/UID 1001)  
crond.service loaded active running Command Scheduler  
data-transfer-appliance.service loaded active running Data Transfer Appliance  
data-transfer-console.service loaded active running Data Transfer Serial Console
```

Any problem with the diagnostic data collection results in the console output being written to the log file of the service. Failure of the commands is indicative of a serious problem, perhaps requiring the return of the appliance. Here is an example of the log:

```
Mar 6 17:55:33 localhost console-diags: {"Module": "main", "Type": "Info", "Message": "Received message\n{\"cmd\": \"collect\"}"}\nMar 6 17:55:33 localhost console-diags: {"Module": "main", "Type": "Info", "Message": "Setting up output\nfile. First to remove all /tmp/xa-diags-results"}\nMar 6 17:55:33 localhost console-diags: {"Module": "main", "Type": "Info", "Message": "Removing /tmp/xa-\ndiags-results.2019-03-06T17:54:56.000471"}
```

Initializing Appliance Fails Because of IP Address Issues

Initializing the Appliance can fail because of using the incorrect IP address. The IP address for `initialize-auth` can differ from the IP address obtained when running `ping` or `ssl connect`. If you experience an initialization failure, ensure that you are using the correct IP address for your Appliance and try initializing again.

Dataset Sealing Process Fails

The dataset sealing process can fail sometimes because there are special files in the dataset:

```
dts nfs-dataset seal-status --name nfs-ds-1\n\nSeal Status :\n  success          : false\n  failureReason    :\nNumber of special files : 5\n  startTime        : 2019/03/26 11:52:37 PDT\n  endTime          : 2019/03/26 11:52:39 PDT\n  numFilesToProcess : 0\n  numFilesProcessed : 0\n  bytesToProcess   : 0.00 KB\n  bytesProcessed   : 0.00 KB\n  bytesToProcess   : 0.00 KB
```

At the command prompt on the host, reactivate the NFS dataset.

```
oci dts nfs-dataset activate --name <dataset_name>
```

Then run `find` to get the full list of all special files and the specific type of each one.

CHAPTER 10 Data Transfer

```
find <mountpoint> \! -type f \! -type d | xargs file
```

For example:

```
$ find /mnt/nfs-ds-1 \! -type f \! -type d | xargs file
/mnt/nfs-ds-1/myfile1: symbolic link to `/home/user1/myfile1'
/mnt/nfs-ds-1/myfile2: symbolic link to `/home/user1/myfile2'
```

Next, review the list and remove all special files from the NFS mount point.

```
find <mountpoint> \! -type f \! -type d | xargs rm
```

Deactivate the NFS dataset.

```
oci dts nfs-dataset deactivate --name <dataset_name>
```

Finally, reseal the dataset.

```
oci dts nfs-dataset seal --name <dataset_name>[--wait]
```

Monitor the seal progress. Wait for it to complete successfully and continue with the subsequent steps.

Initialize Authentication Fails with "connection refused" or "connection timed out"

If you try to configure networking using the appliance serial console but fail with a "connection refused" or "connection timed out" message, follow these troubleshooting steps.

Run the following command at the command prompt on the host:

```
ping <appliance_ip>
```

If a failure occurs, run the following command to verify appliance IP and the path to appliance.

```
ping -I <local_interface> <appliance_ip>
```

To determine expected interface, run `ip route` or an equivalent command. Verify that routing table is sane. Try running `traceroute` if you're not sure to see the network path to the appliance IP.

Run the following command:

```
curl -k https://<appliance_ip>
```

You should receive the response "Not found." This failure can indicate the IP address may be wrong. For example, nothing is listening on port 443. If you receive a failure message, run the following command:

```
openssl s_client -showcerts -connect <appliance_ip>:443
```

You should see a certificate issued for "Oracle Cloud Infrastructure" / "Data Transfer Appliance."

This command is similar to `curl` but does not use HTTPS and so proxies do not affect it. If this command works, and `curl` fails, then verify there are no proxy environmental variables.

Data Transfer Utility Fails with "invalid configuration file"

If you attempt to run Data Transfer commands and receive the error message "invalid configuration file," verify that the following files are present on your host and are correctly set up:

- `~/oci/config` *and*
- `~/oci/config_upload_user`

Both files must have "[DEFAULT]" as the first line. Use of the "~" character in a path is not valid in the file's contents.

Data Transfer Utility Fails with "Processing exception..." while communicating to Oracle Cloud Infrastructure

Check if your environment has proxies to the internet. If so, update them to the latest version and set "https_proxy." If you are using the appliance, set "no_proxy" environmental variables. See [Prerequisites](#) for more information on proxies.

Data Transfer Utility Fails Because of Lack of Exclusive Access to Disk

The Data Transfer Utility requires exclusive access to the disk. If you have any drivers that already claim exclusive access to the disk, then the Data Transfer Utility fails. For example, if you employ a devicemapper multipath driver over all your disk devices, you must first

remove the disk used for the data transfer from the list of devices managed by the multipath driver.

Be sure that access to the disk is not done through any devicemapper or volume manager. During the data transfer, the expectation is that the file system is created on a "raw" device. Any layering or mapping through intermediate drivers or abstraction layers makes it impossible for the disk to be uploaded at the transfer site. The source of these failures can include drivers like multipath, md, striping, logical volume managers, and potentially others as well.

You can confirm that the Data Transfer Utility has exclusive access by attempting to manually format the disk being used for your data transfer. The Data Transfer Utility uses the `cryptsetup` utility to create an encrypted device. You can run `cryptsetup` from the command line (root privileges required):

```
cryptsetup luksFormat -c aes-xts-plain64 -s 512 -h sha512 --iter-time 2000 --use random /dev/<sdXX>
```

`<sdXX>` is the name of the disk being used for the data transfer.

When prompted, respond that you do want to encrypt the device. You are required to provide a passphrase. Any passphrase is acceptable as the `cryptsetup` utility can run on a disk repeatedly without any problems.

If the command succeeds, then you know that the Data Transfer Utility can gain exclusive ownership of the disk to do the necessary for the data transfer.

CHAPTER 11 Database

This chapter explains how to launch a DB System and manage databases on DB Systems.

Overview of the Database Service

The Database service offers autonomous and user-managed Oracle Database cloud solutions. Autonomous databases are preconfigured, fully-managed environments that are suitable for either transaction processing or for data warehouse workloads. User-managed solutions are bare metal, virtual machine, and Exadata DB systems that you can customize with the resources and settings that meet your needs.

You can quickly provision an autonomous database or user-managed DB system. You have full access to the features and operations available with the database, but Oracle owns and manages the infrastructure.

You can also extend user-managed database services into your data center by using Exadata Cloud at Customer, which applies the combined power of Exadata and Oracle Cloud Infrastructure while enabling you to meet your organization's data-residency requirements.

For details about each offering, start with the following overview topics:

Autonomous Databases

The Database service offers Oracle's [Autonomous Database](#) with transaction processing and data warehouse workload types.

User-managed Systems

- [Bare Metal and Virtual Machine DB Systems](#)
- [Exadata DB Systems](#)
- [Exadata Cloud at Customer](#)

License Types and Bring Your Own License (BYOL) Availability

Oracle Cloud Infrastructure supports a licensing model with two license types. With **License included**, the cost of the cloud service includes a license for the Database service. With **Bring Your Own License (BYOL)**, Oracle Database customers can use existing licenses with Oracle Cloud Infrastructure. Note that Oracle Database customers remain responsible for complying with license restrictions applicable to their BYOL licenses, as defined in their program order for those licenses.

You do not need separate on-premises licenses and cloud licenses. BYOL databases support all advanced Database service manageability functionality, including backing up and restoring a DB system, patching, and Oracle Data Guard.

You can choose BYOL when you launch an Oracle Cloud Infrastructure database or DB system. Choosing BYOL impacts how the usage data for the instance is metered and subsequent billing.

Note that on some provisioning dialogs in the Console, the BYOL option is labeled **My Organization Already Owns Oracle Database Software Licenses**.

For additional information about license pricing and features, see [Oracle Cloud Database Services](#).

Always Free Database Resources

The Database service is one of the Oracle Cloud Infrastructure services that provides you with Always Free resources as a part of Oracle's Free Tier. For an introduction to the Free Tier, see [Oracle Cloud Infrastructure's Free Tier](#). For details about the Always Free Autonomous Database, see [Always Free Availability](#) in the Autonomous Database overview topic. To provision an Always Free Autonomous Database, see [To create an Always Free Autonomous Database](#).

Moving Database Resources to a Different Compartment

You can move DB systems, Autonomous Database resources, and Exadata Cloud at Customer resources from one compartment to another. When you move a Database resource to a new

compartment, its dependent resources move with it. After you move the resource to the new compartment, inherent policies apply immediately and affect access to that resource and its dependent resources through the Console.



Important

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see [Details for the Database Service](#).

Dependent Resource Details

Details about dependent resources are as follows:

- **Bare metal, virtual machine, and Exadata DB systems:** Dependent resources that move with these DB systems include Database Homes and databases, as well as the metadata for automatic backups. To verify the compartment of a dependent resource, check the compartment of the DB system.
- **Autonomous Database:** Autonomous Database dependent resources are limited to its automatic backups. Autonomous Exadata Infrastructure instances and Autonomous Container Databases have no dependent resources that move with them. Associated (non-dependent) resources remain in their current compartments.
- **Exadata Cloud at Customer:** Resources that can be moved are Exadata Infrastructure, VM clusters, and backup destinations. VM cluster networks are dependent resources of Exadata Infrastructure instances, so they move with them. VM clusters have the following dependent resources: Database Homes, and databases and their automatic backups. Backup destinations have no dependent resources.

For more information about moving resources to other compartments, see [Moving Resources to a Different Compartment](#).

Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about available Database service metrics and how to view them, see [Database Metrics](#).

Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

See [Database](#) for details about Database resources that emit events.

Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For more information on tenancies and compartments, see "Key Concepts and Terminology" in the *Oracle Cloud Infrastructure Getting Started Guide*. For general information about using the API, see [REST APIs](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to write policies that provide stricter access to database resources, see [Details for the Database Service](#).

Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

For common policies used to authorize Oracle Cloud Infrastructure Database users, see [Common Policies](#).

For in-depth information on granting users permissions for the Database service, see [Details for the Database Service](#) in the IAM policy reference.

Limits on the Database Service

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).



Note

Service limits and compartment quotas do not apply to Exadata Cloud at Customer.

Many Database API operations are subject to [throttling](#).

Work Requests Integration

The Database service is integrated with the Oracle Cloud Infrastructure Work Requests API. Work requests allow you to monitor long-running operations such as the provisioning of DB systems. A work request is an activity log that enables you to track each step in the operation's progress. Each work request has an OCID that allows you to interact with it programmatically and use it for automation. For general information on using work requests in Oracle Cloud Infrastructure, see [Work Requests](#) and the [Work Requests API](#).

Database service operations that create work requests

The following Database operations result in the creation of a work request:

Autonomous Databases



Note

In the Database service, work requests are currently supported only for the Autonomous Database resources in the list that follows.

- Creating or terminating the following resource types:
 - Autonomous Databases
 - Autonomous Container Databases
 - Autonomous Exadata Infrastructure instances
- Starting or stopping an Autonomous Database instance
- Restoring an Autonomous Database instance
- Cloning an Autonomous Database instance
- Creating or deleting manual backups
- Scaling database storage or CPU
- Updating the database license type
- Updating a database's network access control list (ACL)
- Registering or deregistering an Autonomous Database with Data Safe.

Bare Metal, Virtual Machine, and Exadata Databases

Work requests are not currently available for these products.

Overview of Autonomous Database

Oracle Cloud Infrastructure's Autonomous Database is a fully managed, preconfigured database environment with two workload types available, Autonomous Transaction Processing and Autonomous Data Warehouse. You do not need to configure or manage any hardware, or install any software. After provisioning, you can scale the number of CPU cores or the storage capacity of the database at any time without impacting availability or performance.

Autonomous Database handles creating the database, as well as the following maintenance tasks:

- Backing up the database
- Patching the database

- Upgrading the database
- Tuning the database

Always Free Availability

Autonomous Database can be used for free as part of Oracle Cloud Infrastructure's suite of Always Free resources. Users of both paid and free Oracle Cloud Infrastructure accounts have access to two Always Free instances of Autonomous Database. Always Free Autonomous Databases have a fixed 8 GB of memory, 20 GB of storage, 1 OCPU, and can be configured for either Autonomous Transaction Processing or Autonomous Data Warehouse workloads.

For an introduction to the Free Tier, see [Oracle Cloud Infrastructure's Free Tier](#). For details of the Always Free Autonomous Database, see [Overview of the Always Free Autonomous Database](#). To provision an Always Free Autonomous Database, see [To create an Always Free Autonomous Database](#).

Currently, Always Free Autonomous Databases are available in all commercial [regions](#).

Available Workload Types

Autonomous Database offers two workload types:

- The **Autonomous Transaction Processing** workload type configures the database for a transactional workload, with a bias towards high volumes of random data access. For a complete product overview of Autonomous Transaction Processing, see [Autonomous Transaction Processing](#). For Autonomous Transaction Processing tutorials, see [Quick Start tutorials](#).
- The **Autonomous Data Warehouse** workload type configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations. For a complete product overview of Autonomous Data Warehouse, see [Autonomous Data Warehouse](#). For Autonomous Data Warehouse tutorials, see [Quick Start tutorials](#).

Deployment Types

Autonomous Databases have the following deployment options:

- **Dedicated deployment.** Using the dedicated deployment option, you have exclusive use of the Exadata infrastructure and hardware. Dedicated deployment offers multitenant database architecture, allowing you to create and manage multiple Autonomous Databases within a single database system. For an overview of Dedicated Deployments, see [Overview of Autonomous Database Dedicated Deployments](#). Both workload types (transaction processing and warehouse) can be provisioned with dedicated deployment.
- **Serverless deployment.** With the serverless deployment option, you provision and manage only the Autonomous Database, while Oracle handles the infrastructure deployment and management tasks. Both workload types (transaction processing and warehouse) can be provisioned with serverless deployment.

CPU Scaling

For serverless deployments, the **auto scaling** feature enables the system to automatically adjust the number of CPU cores as load demand fluctuates, allowing the system to use cores more efficiently. As demand increases, auto scaling gradually increases the number of cores, up to a maximum of three times the assigned number. Likewise, it gradually decreases cores as demand drops. You can also scale the database's assigned number of CPU cores up or down at any time. CPU scaling does not impact Autonomous Database availability or performance. Note the following points regarding the auto scaling feature:

- The maximum number of cores that can be made available to a database remains 128, regardless of whether auto scaling is enabled or not. This means that database with a CPU core count of 64 could auto scale up to two times the assigned number of cores ($2 \times 64 = 128$). A database with 42 cores (or fewer) could auto scale up to three times the assigned number ($3 \times 42 = 126$).
- Auto scaling can be enabled or disabled at any time.

- The auto scaling status for a database (enabled or disabled) is displayed on the database details page.
- You can view hourly snapshots of the database's actual CPU usage over the most recent 8 days. This information is available in the [Service Console](#), in the Overview page graph "Number of OCPUs Allocated". For more information, see [To view CPU allocation hourly snapshot data for an Autonomous Database](#).
- For billing purposes, the Autonomous Database service determines the average CPU utilization per hour.

Storage Scaling

Autonomous Database allows you to scale the storage capacity of the database at any time without impacting availability or performance.

Performance Monitoring Using Oracle Performance Hub

You can monitor and diagnose the performance of an Autonomous Database in the Oracle Cloud Infrastructure Console using the [Performance Hub](#) ASH Analytics and SQL Monitoring features. These features provide the same information as the ASH Analytics and SQL Monitoring tools found in Oracle's EM Express, Oracle Management Cloud (OMC), and SQL Developer Web applications. For more information about using these features in the Oracle Cloud Infrastructure Console, see [Using Performance Hub to Analyze Database Performance in Oracle Cloud Infrastructure](#).

Oracle Database Preview Version Availability

Oracle Cloud Infrastructure periodically offers Autonomous Database preview versions of Oracle Database for testing purposes. You can provision an Autonomous Database using preview version software to test applications before the general availability of the software in Autonomous Database. Oracle will notify Autonomous Database customers when preview versions are available. Preview version software is available for a limited time. Databases provisioned with preview version software will display the end date of the preview period at

the top of the database details page in the Console. If you are using the Console, you can also see the end date of the preview period in the Create Database provisioning dialog before the database is created.

Preview version software should not be used for production databases or for databases that need to persist beyond the limited preview period. Note that preview databases and their associated resources (including backups) are terminated automatically at the conclusion of the preview period. Oracle will notify customers prior to the conclusion of the preview period regarding the end date of the preview.

Any existing Autonomous Database (including those provisioned with preview version software) can be cloned using a preview version of Autonomous Database. However, preview version databases cannot be cloned using the regular (general-availability) Autonomous Database software.

See [Creating an Autonomous Database](#) for details on provisioning a preview version of Autonomous Database.

Availability

Autonomous Database is currently available in all regions of the [commercial realm](#). Autonomous Database is currently not available in regions within the Government Cloud realm.

Security Considerations

Data Safe Integration

Oracle Data Safe is a cloud service that enables you to monitor the security posture of your Autonomous Databases. Data Safe helps you discover, protect, and mask sensitive and regulated data. It also helps you assess database users and their authentication information for security risks, and provides auditing and reporting features need for compliance.

Autonomous Databases using serverless deployment can be registered with the Data Safe instance in the region containing the Autonomous Database. For information on registering or

deregistering a database, see [To register or deregister an Autonomous Database with Data Safe](#).

For information on creating and using Data Safe instances, see the [Data Safe Overview](#).

Service Gateway

Oracle Autonomous Database is one of the Oracle Cloud services that can be privately accessed through a service gateway within a VCN. This means you do not need a public IP or NAT to access your Autonomous Database instance from any of the cloud services within the [Oracle Services Network](#). For example, if you have a Compute instance that uses a VCN with a service gateway, you can route traffic between your Compute instance and an Autonomous Database in the same region without the traffic going over the internet. For information on setting up a VCN service gateway and configuring it to access all supported Oracle Service Network services (which include Autonomous Database), see [Access to Oracle Services: Service Gateway](#).

Access Control Lists (ACLs) for Serverless Deployments

For Autonomous Databases using serverless deployment, an access control list (ACL) provides additional protection for your Autonomous Database by allowing only specified IP addresses and VCNs in the list to connect to the database. Specified IP addresses can include private IP addresses from your on-premises network that connect to your database using [transit routing](#) and allow traffic to move directly from your on-premises network to your Autonomous Database without going over the internet. See [Transit Routing: Private Access to Oracle Services](#) for more information on this method of access.

You can add the following to your ACL:

- Public IP addresses (individually, or in CIDR blocks)
- An entire VCN (specified by OCID)
- Private IP addresses within a specified VCN (individually, or in CIDR blocks)
- Private IP addresses within an on-premises network that have access using a [transit routing](#)

You can create an ACL during database provisioning, or at any time thereafter. You can also edit an ACL at any time. Removing all entries from the list makes the database accessible to all clients with the applicable credentials. See [To manage the access control list of an Autonomous Database with Serverless Deployment](#) to learn how to create, update, or delete an ACL.



Important

If you are using a service gateway and you configure an access control list, you must add the CIDR range 240.0.0.0/4 to the ACL to enable clients accessing the database through the service gateway to connect to it.

Security Tools for Dedicated Deployments

Network security groups (NSGs), an optional Networking feature available with dedicated deployment, act as a virtual firewall for your Autonomous Exadata Infrastructure resources. An NSG consists of a set of ingress and egress [security rules](#) that apply only to *a set of VNICs of your choice within a single VCN*. For more information, see the following topics:

- [Network Security Groups](#)
- [To edit the network security groups \(NSGs\) for your Autonomous Exadata Infrastructure resource](#)

Development and Administration Tools

Oracle's SQL Developer Web, Application Express (APEX), and Machine Learning applications are available for Autonomous Databases. For information on how to use these applications and access them from the Console, see [Autonomous Database Development and Administration Tools](#).

Using the Oracle Cloud Infrastructure Console to Manage Autonomous Databases

For information on provisioning, managing, and backing up an Autonomous Database in the Oracle Cloud Infrastructure Console, see the following topics:

- [Creating an Autonomous Database](#)
- [Managing an Autonomous Database](#)
- [Connecting to an Autonomous Database](#)
- [Backing Up an Autonomous Database Manually](#)
- [Restoring an Autonomous Database](#)

Additional Autonomous Database Product Information

Autonomous Transaction Processing

Information for dedicated deployments

For in-depth documentation on using and managing your Autonomous Transaction Processing dedicated deployment, see the following topics:

- [Getting Started with Autonomous Transaction Processing](#)
- [Connecting to Autonomous Transaction Processing](#)
- [Loading Data into Autonomous Transaction Processing](#)
- [Starting, Stopping and Scaling Autonomous Transaction Processing](#)
- [Managing Database Users](#)
- [Managing and Monitoring Performance](#)
- [Backing Up and Restoring Autonomous Transaction Processing](#)
- [Cloud Object Storage URI Formats](#)
- [Using Oracle Database Features in Autonomous Transaction Processing](#)

For information on how application developers connect their applications to Autonomous Transaction Processing databases, see [Developer's Guide to Oracle Autonomous Transaction Processing Dedicated Deployments](#).

For known issues, see [Known Issues for Oracle Autonomous Transaction Processing Dedicated Deployments](#).

Information for serverless deployments

For in-depth documentation on using and managing your Autonomous Transaction Processing database, see the following topics:

- [Getting Started with Autonomous Transaction Processing](#)
- [Connecting to Autonomous Transaction Processing](#)
- [Loading Data with Autonomous Transaction Processing](#)
- [Querying External Data with Autonomous Transaction Processing](#)
- [Creating Dashboards, Reports, and Notebooks with Autonomous Transaction Processing](#)
- [Managing Users on Autonomous Transaction Processing](#)
- [Managing and Monitoring Performance of Autonomous Transaction Processing](#)

For information on using a database client to manage your database, see [Connect Autonomous Transaction Processing Using a Client Application](#).

Autonomous Data Warehouse

Information for serverless deployments

For in-depth documentation on using and managing your Autonomous Data Warehouse database, see the following topics:

- [Getting Started with Autonomous Data Warehouse Cloud](#)
- [Connecting to Autonomous Data Warehouse Cloud](#)

- [Loading Data with Autonomous Data Warehouse Cloud](#)
- [Migrating Data from Amazon Redshift](#)
- [Querying External Data with Autonomous Data Warehouse Cloud](#)
- [Creating Dashboards, Reports, and Notebooks with Autonomous Data Warehouse Cloud](#)
- [Managing Users on Autonomous Data Warehouse Cloud](#)
- [Managing and Monitoring Performance of Autonomous Data Warehouse Cloud](#)

For information on using a database client to manage your database, see the following topic:

- [Connect Autonomous Data Warehouse Using a Client Application](#)

Creating an Autonomous Database

This topic describes how to provision a new Autonomous Database using the Oracle Cloud Infrastructure Console or the [API](#). Autonomous Databases can be provisioned using either the [dedicated deployment](#) option or the [serverless deployment](#) option. Your database can be optimized for either [transaction processing](#) or [data warehouse](#) workloads.

If you want to provision an Always Free Autonomous Database, see the [To create an Always Free Autonomous Database](#) task instructions in this topic. For more information on the Free Tier, see [Oracle Cloud Infrastructure's Free Tier](#).

For *Oracle By Example* tutorials on provisioning Autonomous Databases, see [Provisioning Autonomous Transaction Processing](#) and [Provisioning Autonomous Data Warehouse Cloud](#).

Prerequisites

- To create an Autonomous Database, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See

[Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service.



Tip

See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

- For information on additional prerequisites for provisioning an Autonomous Transaction Processing database, see [What Do You Need?](#) Likewise, for information on additional prerequisites for provisioning an Autonomous Data Warehouse, see [What Do You Need?](#)
- To create an Autonomous Transaction Processing database with the dedicated deployment option, you must first provision the infrastructure and at least one Autonomous Container Database. For more information, see [Creating an Autonomous Exadata Infrastructure Resource](#) and [Creating an Autonomous Container Database](#). Note that Oracle Cloud Infrastructure does not currently offer Autonomous Data Warehouse databases with the dedicated deployment option.

Using the Oracle Cloud Infrastructure Console to Create an Autonomous Database

When provisioning an Autonomous Database using the serverless deployment option, you can choose either the transaction processing or the data warehouse workload types.

For Autonomous Transaction Processing databases, you can choose the dedicated deployment option. Dedicated deployment provides you with exclusive use of the Exadata infrastructure and hardware.

To create an Autonomous Database using the serverless deployment option

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.

In the **Create Autonomous Database** dialog, enter the following:

BASIC DATABASE INFORMATION

- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. Avoid entering confidential information.
- **Compartment:** Select the compartment of the Autonomous Database.
- **Database Name:** The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.



Note

You cannot use the same database name concurrently for both an Autonomous Data Warehouse and an Autonomous Transaction Processing database.

WORKLOAD TYPE

Select the desired workload type. See [About Autonomous Transaction Processing](#) and [About Autonomous Data Warehouse](#) for information about each workload type.

INFRASTRUCTURE TYPE

Select **Serverless**.

DATABASE CPU CORE COUNT AND STORAGE CONFIGURATION

- **Always Free:** Use this selector to show only Always Free configuration options if you are provisioning an Always Free Autonomous Database. See [Overview of the Always Free Autonomous Database](#) for more information.
- **CPU Core Count:** You can enable up to 128 cores for your Autonomous Transaction Processing database. The actual number of available cores is subject to your tenancy's [service limits](#). Select the number of cores you wish to assign to your database. The total number of number of cores available to all database within the Autonomous Exadata Infrastructure depends on the infrastructure shape being used. For information on the CPU core resources available in the current Oracle Cloud Infrastructure Exadata offerings, see [Exadata X7 Shapes](#).
Auto Scaling: Auto scaling allows Autonomous Database to automatically increase the number of CPU cores by up to three times the assigned CPU core count value, depending on demand for processing. The auto scaling feature reduces the number of CPU cores when additional cores are not needed. For databases with up to 42 assigned cores, you can increase the maximum number of cores available through auto scaling by increasing the CPU core count value.



Note

The maximum number of cores that are available to any Autonomous Database database not using dedicated deployment is 128, regardless of whether auto scaling is enabled or not. This means that database with a CPU core count of 64 could auto scale up to two times the assigned number of cores ($2 \times 64 = 128$). A database with 42 cores (or fewer) could auto scale up to three times the assigned number ($3 \times 42 = 126$). For billing purposes, the database service determines the average number of CPUs used per hour.

- **Storage:** Specify the storage you wish to make available to your Autonomous Database, in terabytes. You can make up to 128 TB available.
- **Enable Preview Version:** (*This option only displays during periods when a preview version of Autonomous Database is available*) Select this option to provision the database with an Autonomous Database preview version. Preview versions of Autonomous Database are made available for limited periods for testing purposes. Do not select this option if you are provisioning a database for production purposes or if you will need the database to persist beyond the limited availability period of the preview version.

ADMINISTRATOR CREDENTIALS

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. You use this password when accessing the Autonomous Database service console and when using an SQL client tool.

- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

NETWORK ACCESS

Optionally, you can create an Access Control List (ACL) for your database. ACLs provide additional protection for your Autonomous Database by allowing only the public and VCN IP addresses in the list to connect to the database. Click **Configure Access Control Rules** to create an ACL for your database.

You can specify the following types of addresses in your list by using the **Source** drop-down selector:

- **IP Address** allows you to specify one or more individual public IP address. Use commas to separate your addresses in the input field.
- **CIDR Block** allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.
- **Virtual Cloud Network** allows you to specify an existing VCN. The drop-down listing in the input field allows you to choose from the VCNs in your current compartment for which you have access permissions. Click the **Change Compartment** link to display the VCNs of a different compartment.
- **Virtual Cloud Network (OCID)** allows you input the OCID of a VCN in a text box. You can use this input method if the VCN you are specifying is in a compartment which you do not have permission to access.

If you add a **Virtual Cloud Network** to your ACL, you can limit further by specifying allowed VCN IP addresses or CIDR ranges. Enter those addresses or CIDR blocks in the

IP Addresses or CIDRs field that is displayed below your **Virtual Cloud Network** choice. Use commas to separate your VCN addresses and CIDR blocks in the input field. You can specify the following types of IP addresses at the VCN level:

- Private IP addresses within your Oracle Cloud Infrastructure VCN
- Private IP addresses within an on-premises network that have access to your Autonomous Database using a [transit routing and a private connection via FastConnect or VPN Connect](#).

Click **+Additional Entry** to add additional access rules to your list.

LICENSE TYPE

Specify the license type setting you want to use. Your choice affects metering for billing. You have the following options:

- **Bring your own license:** Bring my existing database software licenses to the database cloud service.
- **License included:** Subscribe to new database software licenses and the Database cloud service.

ADVANCED OPTIONS

Tags: Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

2. Click **Create Autonomous Database**.



Note

The following naming restrictions apply to Autonomous Transaction Processing and Autonomous Data Warehouse databases:

- Names associated with databases terminated within the last 60 days cannot be used when creating a new database.
- A database name cannot be used concurrently for both an Autonomous Data Warehouse and an Autonomous Transaction Processing database.

To create an Autonomous Database database using the dedicated deployment option



Note

To provision an Autonomous Database using the dedicated deployment option, you must have already provisioned an Autonomous Exadata Infrastructure resource and at least one container database on that resource. See the following for more information:

- [Overview of Autonomous Database Dedicated Deployments](#)
- [Creating an Autonomous Exadata Infrastructure Resource](#)
- [Creating an Autonomous Container Database](#)

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.

In the **Create Autonomous Database** dialog, enter the following:

BASIC DATABASE INFORMATION

- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. Avoid entering confidential information.
- **Compartment:** Select the compartment of the Autonomous Database.
- **Database Name:** The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.

WORKLOAD TYPE

Select the desired workload type. See [About Autonomous Transaction Processing](#) and [About Autonomous Data Warehouse](#) for information about each workload type.

INFRASTRUCTURE TYPE

Select **Dedicated**.

DATABASE CPU CORE COUNT AND STORAGE CONFIGURATION

- **CPU Core Count:** Select the number of cores you wish to assign to your database. The total number of number of cores available to all database within the Autonomous Exadata Infrastructure depends on the infrastructure shape being used. For information on the CPU core resources available in the current Oracle Cloud Infrastructure Exadata offerings, see [Exadata X7 Shapes](#).



Tip

Autonomous Database with dedicated deployment allows for the over-subscription of CPU cores when provisioning databases. This means that for a given infrastructure instance, the sum of all CPU cores provisioned to Autonomous Databases can exceed the number of CPU cores on the Exadata rack. This approach allows for better CPU core utilization. Note that if all databases are heavily loaded at the same time, there will be some contention for CPU resources, and degraded performance.

- **Storage:** Specify the storage you wish to make available to your Autonomous Database, in terabytes. The available storage depends on the infrastructure shape being used. For information on storage available in the current Oracle Cloud Infrastructure Exadata offerings, see [Exadata X7 Shapes](#).

AUTONOMOUS CONTAINER DATABASE

- **Compartment:** Specify the compartment containing the Autonomous Container Database you wish to use.
- **High Availability Database Container:** Specify the Autonomous Container Database you wish to use for your Autonomous Database. This selection determines which Autonomous Exadata Infrastructure the database runs within.

See [Creating an Autonomous Container Database](#) for information on provisioning a container database.

ADMINISTRATOR CREDENTIALS

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. You use this password when accessing the Autonomous Database service console and when using an SQL client tool.

- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

ADVANCED OPTIONS

Tags: Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

2. Click **Create Autonomous Database**.



Note

The following naming restrictions apply to Autonomous Transaction Processing and Autonomous Data Warehouse databases:

- Names associated with databases terminated within the last 60 days cannot be used when creating a new database.
- A database name cannot be used concurrently for both an Autonomous Data Warehouse and an Autonomous Transaction Processing database.

To create an Always Free Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.

In the **Create Autonomous Database** dialog, enter the following:

BASIC DATABASE INFORMATION

- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. Avoid entering confidential information.
- **Compartment:** Select the compartment of the Autonomous Database.
- **Database Name:** The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.

CHAPTER 11 Database

WORKLOAD TYPE

Select the desired workload type. See [About Autonomous Transaction Processing](#) and [About Autonomous Data Warehouse](#) for information about each workload type.

INFRASTRUCTURE TYPE

Applies to Autonomous Transaction Processing only.

Select **Serverless**.

DATABASE CPU CORE COUNT AND STORAGE CONFIGURATION

- **Always Free:** Move this selector to the right so that the provisioning workflow shows only the Always Free configuration options. Note that the **Core CPU count** and **Storage** configuration fields are disabled when provisioning an Always Free Autonomous Database. Your database will have 1 OCPU, 8 GB of memory, and 20 GB of storage.

ADMINISTRATOR CREDENTIALS

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria:

- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

Use this password when accessing the Autonomous Database service console and when using an SQL client tool.

LICENSE TYPE

When you provision an Always Free Autonomous Database, the license type is set to **License included** and cannot be adjusted.

ADVANCED OPTIONS

Tags: Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

2. Click **Create Autonomous Database**.



Note

The following naming restrictions apply to Autonomous Transaction Processing and Autonomous Data Warehouse databases:

- Names associated with databases terminated within the last 60 days cannot be used when creating a new database.
- A database name cannot be used concurrently for two Autonomous Databases, regardless of workload type.

Using the API

Use the [CreateAutonomousDatabase](#) API operation to create Autonomous Databases of either the transaction processing or warehouse workload types.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For More Information

AUTONOMOUS TRANSACTION PROCESSING

- [Using Oracle Autonomous Transaction Processing](#) (the Autonomous Transaction Processing database administrator guide)
- [Using Oracle Autonomous Transaction Processing Dedicated Deployments](#) (the Autonomous Transaction Processing Dedicated Deployment database administrator guide)
- [Autonomous Transaction Processing: Tutorials](#) (Oracle By Example tutorials)
- [Autonomous Transaction Processing: Videos](#) (video tutorials)

AUTONOMOUS DATA WAREHOUSE

- [Using Oracle Autonomous Data Warehouse](#) (the Autonomous Data Warehouse user guide)
- [Autonomous Data Warehouse: Tutorials](#) (Oracle By Example tutorials)
- [Autonomous Data Warehouse: Videos](#) (video tutorials)

Managing an Autonomous Database

This topic describes the database management tasks for Autonomous Databases that you complete using the Oracle Cloud Infrastructure Console or the [API](#). Note that some database management tasks not covered here are performed using the [Autonomous Transaction Processing service console](#) or the [Autonomous Data Warehouse service console](#).

Prerequisites

To perform the management tasks in this topic, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure

Database service. See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

Using the Oracle Cloud Infrastructure Console

You can perform basic administrative tasks for Autonomous Databases in the Oracle Cloud Infrastructure Console including stopping, starting, and scaling your databases. You can also use the Console to [back up](#) or to [restore](#) the database.

To set the Admin password

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Go to **Actions**, and then click **Admin Password**. The Admin Password dialog opens.
5. Enter a password for the Autonomous Database. The password must meet the following criteria:
 - Between 12 and 30 characters long
 - Contains at least one lowercase letter
 - Contains at least one uppercase letter
 - Contains at least one number
 - Does not contain the double quotation mark (")
 - Does not contain the string "admin", regardless of casing
 - Is not one of the last four passwords used for the database
 - Is not a password you previously set within the last 24 hours
6. Enter the password again in the Confirm Password field.

7. Click **Update**.

To scale the CPU core count or storage of an Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Click **Scale Up/Down**.
5. Enter a new value for **CPU Core Count** or **Storage** between 1 and 128. The number you enter represents the desired total (final) value for your database's CPU core count or storage.

The number of available cores is subject to your tenancy's [service limits](#). An Autonomous Database database can have a maximum of 128 cores and 128 TB of storage. Scaling the CPU core count affects your CPU billing.

6. Click **Update**.

To enable or disable auto scaling for an Autonomous Database with serverless deployment

Note the following points regarding the auto scaling feature:

- If auto scaling is disabled while more CPU cores are in use than the database's currently assigned number of cores, then Autonomous Database scales the number of CPU cores in use down to the assigned number.
- Enabling auto scaling does not change the concurrency and parallelism settings for the predefined services. See [Managing Concurrency and Priorities on Autonomous Data Warehouse](#) and [Managing Priorities on Autonomous Transaction Processing](#) for more information.

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Click **Scale Up/Down**.
5. Check **Auto Scaling** to enable the auto scaling feature, or uncheck **Auto Scaling** to disable the feature. Auto scaling allows Autonomous Database to automatically increase the number of CPU cores by up to three times the assigned CPU core count value, depending on demand for processing. The auto scaling feature reduces the number of CPU cores when additional cores are not needed. You can enable or disable auto scaling at any time. For databases with up to 42 assigned cores, you can increase the maximum number of cores available through auto scaling by increasing the CPU core count value.



Note

The maximum number of cores that are available to any Autonomous Database database not using dedicated deployment is 128, regardless of whether auto scaling is enabled or not. This means that database with a CPU core count of 64 could auto scale up to two times the assigned number of cores ($2 \times 64 = 128$). A database with 42 cores (or fewer) could auto scale up to three times the assigned number ($3 \times 42 = 126$). For billing purposes, the database service determines the average number of CPUs used per hour.

6. Click **Update**.

To view CPU allocation hourly snapshot data for an Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database for which you wish to view CPU usage data.
4. Click the **Service Console** button. The [Service Console](#) opens in a new tab or window.
5. In the Overview screen, the **Number of OCPUs allocated** graph shows hourly snapshot data of CPU allocation over the last eight days. Place your cursor over the graph and move it to the left or right to see data for a specific day and hour.

To move an Autonomous Database to another compartment



Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see [Details for the Database Service](#).

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to move.

4. Go to **Actions**, and then click **Move Resource**.
5. Select the new compartment.
6. Click **Move Resource**.

For information about dependent resources for Database resources, see [Moving Database Resources to a Different Compartment](#).

To stop or start an Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Go to **Actions**, and then click **Stop** (or **Start**). When you stop your Autonomous Database, billing stops for CPU usage. Billing for storage continues when the database is stopped.
5. Confirm that you wish to stop or start your Autonomous Database in the confirmation dialog.



Note

Stopping your database has the following consequences:

- On-going transactions are rolled back.
- CPU billing is halted based on full-hour cycles of usage.
- You will not be able to connect to your database using database clients or tools.

To terminate an Autonomous Database



Warning

Terminating an Autonomous Database permanently deletes it. The database data, including automatic backups, will be lost when the system is terminated. Manual backups remain in Object Storage and are not automatically deleted when you terminate an Autonomous Database. Oracle recommends that you create a manual backup prior to terminating.

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Go to **Actions**, and then click **Terminate**.
5. Confirm that you wish to terminate your Autonomous Database in the confirmation dialog.

To check the lifecycle state of your Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.

4. In the Information tab, note the value displayed for **Lifecycle State**. For some lifecycle states, an information icon (i) is displayed to provide additional details regarding the lifecycle state or ongoing operations such as backups, restores, or terminations. The database has one of the following lifecycle states:
 - Available
 - Available needs attention
 - Backup in progress
 - Provisioning
 - Restore in progress
 - Scaling in progress
 - Starting
 - Stopping
 - Stopped
 - Terminating
 - Terminated
 - Unavailable

To view a work request for your Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. In the **Resources** section, click **Work Requests**. The status of all work requests appears on the page.
5. To see the log messages, error messages, and resources that are associated with a

specific work request, click the operation name. Then, select an option in the **More information** section.

For associated resources, you can click the the Actions icon (three dots) next to a resource to copy the resource's OCID.

For more information, see [Work Requests](#).

To manage tags for your Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Go to **Actions**, and then click **Apply Tag(s)** to add new tags. Or click the **Tags** tab to view or edit the existing tags.

For more information, see [Resource Tags](#).

To access the Autonomous Database service console (serverless deployments only)

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Click **Service Console**.

For information on using the Autonomous Transaction Processing service console features, see [Managing and Monitoring Performance of Autonomous Transaction Processing](#). For

information on using the Autonomous Data Warehouse service console features, see [Managing and Monitoring Performance of Autonomous Data Warehouse Cloud](#).

To change the license type of an Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. Go to **Actions**, and then click **Update License Type**.
The dialog displays the options with your current license type selected.
5. Select the new license type.
6. Click **Update**.

See [Known Issue](#).

To manage the access control list of an Autonomous Database with Serverless Deployment

An access control list (ACL) provides additional protection for your Autonomous Database by allowing only the IP addresses in the list to connect to the database. An ACL must contain at least one entry representing an IP address or a range of addresses. To create or edit an ACL for an existing [serverless](#) Autonomous Database, do the following:

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.

4. Under **Network** in the database details, find the **Access Control List** field and click **add** (if no ACL currently exists) or **edit** (to update an existing ACL).
5. In the **Access Control List dialog**, add or modify entries, as applicable.
You can specify the following types of addresses in your list by using the **Source** drop-down selector:
 - **IP Address** allows you to specify one or more individual public IP address. Use commas to separate your addresses in the input field.
 - **CIDR Block** allows you to specify one or more ranges of public IP addresses using CIDR notation. Use commas to separate your CIDR block entries in the input field.
 - **Virtual Cloud Network** allows you to specify an existing VCN. The drop-down listing in the input field allows you to choose from the VCNs in your current compartment for which you have access permissions. Click the **Change Compartment** link to display the VCNs of a different compartment.
 - **Virtual Cloud Network (OCID)** allows you input the OCID of a VCN in a text box. You can use this input method if the VCN you are specifying is in a compartment which you do not have permission to access.

If you add a **Virtual Cloud Network** to your ACL, you can limit further by specifying allowed VCN IP addresses or CIDR ranges. Enter those addresses or CIDR blocks in the **IP Addresses or CIDRs** field that is displayed below your **Virtual Cloud Network** choice. Use commas to separate your VCN addresses and CIDR blocks in the input field. You can specify the following types of IP addresses at the VCN level:

- Private IP addresses within your Oracle Cloud Infrastructure VCN
- Private IP addresses within an on-premises network that have access to your Autonomous Database using a [transit routing and a private connection via FastConnect or VPN Connect](#).

Click **+Additional Entry** to add additional access rules to your list.



Important

If you are using a service gateway, ensure that the CIDR range 240.0.0.0/4 is included in the list to allow clients accessing the database through the service gateway to connect to it.

To remove the ACL, simply delete all entries in the list. This action allows all clients to connect to the database.

6. Click **Update**.

For information about access control lists, see [Access Control Lists \(ACLs\) for Serverless Deployments](#).

To register or deregister an Autonomous Database with Data Safe

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to administer.
4. In the Autonomous Database Information tab on the database details page, click the **register** or **deregister** link under **Data Safe**, depending on the operation you are performing.
5. In the confirmation dialog, click **Confirm** to complete the registration or deregistration. You can monitor the progress of the registration or deregistration using the [work request](#) created by the system.

If you are registering your Autonomous Database, you can click the **View Console** link to display the Data Safe user interface for the registered database. For information on using Data Safe, see the [Data Safe Overview](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Autonomous Databases:

- [ListAutonomousDatabases](#)
- [GetAutonomousDatabase](#)
- [UpdateAutonomousDatabase](#)
- [ChangeAutonomousDatabaseCompartment](#)
- [StartAutonomousDatabase](#)
- [StopAutonomousDatabase](#)
- [DeleteAutonomousDatabase](#)

For More Information

Autonomous Transaction Processing

- [Managing Users on Autonomous Transaction Processing](#)
- [Managing and Monitoring Performance of Autonomous Transaction Processing](#)
- Autonomous Transaction Processing [Quickstart Tutorials](#)
- [Autonomous Transaction Processing](#) (*complete user guide*)

Autonomous Data Warehouse

- [Managing Users on Autonomous Data Warehouse Cloud](#)
- [Managing and Monitoring Performance of Autonomous Data Warehouse Cloud](#)
- Autonomous Data Warehouse [Quickstart Tutorials](#).

- [Autonomous Data Warehouse](#) (*complete user guide*)

Connecting to an Autonomous Database

This topic describes the following actions related to connecting client applications to an Autonomous Database:

- Connecting a client to an Autonomous Database
- Obtaining the credentials and information you need to create a connection
- Rotating the keys and credentials needed for a connection (wallet rotation)
- Obtaining access URLs for Oracle Application Express (APEX) and Oracle SQL Developer Web

About Connecting to Autonomous Databases

Applications and tools connect to Autonomous Databases by using Oracle Net Services (also known as SQL*Net). SQL*Net supports a variety of connection types to Autonomous Databases, including Oracle Call Interface (OCI), ODBC drivers, JDBC OC, and JDBC Thin Driver.

To support connections of any type, you'll need to download the client security credentials and network configuration settings required to access your database. You'll also need to supply the applicable TNS names or connection strings for a connection, depending on the client application or tool, type of connection, and service level. You can view or copy the TNS names and connection strings in the DB Connection dialog for your Autonomous Database. For detailed information about the TNS names, see [Predefined Database Service Names for Autonomous Transaction Processing](#) and [Predefined Database Service Names for Autonomous Data Warehouse](#).

CONNECTING FROM A VCN

To connect to Autonomous Databases from a VCN, the VCN must be configured with one of the following gateways:

- [internet gateway](#): For access from a public subnet in the VCN
- [service gateway](#): For access from a private subnet in the VCN

Make sure to configure the subnet's [route table](#) with a rule that sends the desired traffic to the specific gateway. Also configure the subnet's [security lists](#) to allow the desired traffic.

You can also connect to your database from a private IP addresses in your on-premises network by using transit routing with an Oracle Cloud Infrastructure VCN. This allows traffic to move directly from your on-premises network to your Autonomous Database without going over the internet. See [Transit Routing: Private Access to Oracle Services](#) for more information on this method of access.

About Downloading Client Credentials

The client credentials .zip that you download contains the following files:

- cwallet.sso - Oracle auto-login wallet
- ewallet.p12 - PKCS #12 wallet file associated with the auto-login wallet
- sqlnet.ora - SQL*Net profile configuration file that includes the wallet location and TNSNAMES naming method
- tnsnames.ora - SQL*Net configuration file that contains network service names mapped to connect descriptors for the local naming method
- Java Key Store (JKS) files - Key store files for use with JDBC Thin Connections



Important

Wallet files, along with the database user ID and password, provide access to data in your Autonomous Database. Store wallet files in a secure location. Share wallet files only with authorized users. If wallet files are transmitted in a way that might be accessed by unauthorized users (for example, over public email), transmit the wallet password separately and securely.

For Autonomous Databases using [serverless deployment](#), you have the choice of downloading an **instance wallet** file or a **regional wallet** file. The instance wallet contains only credentials and keys for a single Autonomous Database. The regional wallet contains credentials and keys for all Autonomous Databases in a specified region. For security purposes, Oracle recommends that regional wallets be used only by database administrators, and that instance wallets be supplied to other users whenever possible.

For Autonomous Databases using [dedicated deployment](#), the wallet file contains only credentials and keys for a single Autonomous Database.

About Rotating Your Autonomous Database Wallet

For Autonomous Databases using [serverless deployment](#), you can rotate an instance or regional wallet for security purposes. When your wallet rotation is complete, you will have a new set of certificate keys and credentials, and the old wallet's keys and credentials will be invalid. Rotating an instance wallet does not invalidate the regional wallet that covers the same database instance. Rotating a regional wallet affects all databases in the specified region. User session termination begins after wallet rotation completes, however this process does not happen immediately.



Important

If you are rotating a wallet to address a security breach and need to reestablish all database connections immediately using the keys and credentials of your newly rotated wallet, stop and restart the database instance.

Before You Begin

The Autonomous Database is preconfigured to support Oracle Net Services (a TNS listener is installed and configured to use secure TCPS and client credentials.) The client computer must be prepared to use Oracle Net Services to connect to the Autonomous Database. Preparing your client includes downloading the client credentials. See the following links for steps you might have to perform before you access the client credentials and connection information for your Autonomous Database

Autonomous Transaction Processing

- [Preparing for Oracle Call Interface \(OCI\), ODBC, and JDBC OCI Connections](#)
- [Preparing for JDBC Thin Connections](#)

Autonomous Data Warehouse

- [Preparing for Oracle Call Interface \(OCI\), ODBC, and JDBC OCI Connections](#)
- [Preparing for JDBC Thin Connections](#)

Using the Oracle Cloud Infrastructure Console

To download a wallet for an Autonomous Database with serverless deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you are interested in.
4. Click **DB Connection**.
5. In the **Download Client Credentials (Wallet)** section, select the **Wallet Type**. You can choose to download an instance wallet or a regional wallet.
6. To obtain the client credentials, click **Download Wallet**.
You will be prompted to provide a password to encrypt the keys inside the wallet. The password must be at least 8 characters long and must include at least 1 letter and either 1 numeric character or 1 special character.
Save the client credentials zip file to a secure location. See [About Downloading Client Credentials](#) for information about the files included in the download.
7. Take note of or copy the TNS names or connection strings you need for your connection. See [About Connecting to Autonomous Databases](#) for information about making connections.

To download a wallet for an Autonomous Database with dedicated deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you are interested in.

4. Click **DB Connection**.
5. Select the **DB Connection** option.
6. Click the **DB Connection** tab.
7. To obtain the client credentials, click **Download**.
You will be prompted to provide a password to encrypt the keys inside the wallet. The password must be at least 8 characters long and must include at least 1 letter and either 1 numeric character or 1 special character.
Save the client credentials zip file to a secure location. See [About Downloading Client Credentials](#) for information about the files included in the download.
8. Take note of or copy the TNS names or connection strings you need for your connection. See [About Connecting to Autonomous Databases](#) for information about making connections.

To rotate an Autonomous Database wallet (serverless deployment only)

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you are interested in.
4. Click **DB Connection**.
5. In the **Download Client Credentials (Wallet)** section, select the **Wallet Type**. You can choose to rotate an instance wallet or a regional wallet.
6. Click **Rotate Wallet**. A confirmation dialog will prompt you to enter the database name to confirm the rotation.
7. Enter the name of the database, then click **Rotate Wallet**.
The rotation takes a few minutes to complete.

To obtain access URLs for Oracle Application Express (APEX) and Oracle SQL Developer Web (dedicated deployment only)

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you are interested in.
4. Select the **Application Connection** option.
5. Application URLs are displayed in plain text in the **Application URL** field. Copy the URL string using the **Copy** link.
6. Paste the URL into a browser running on a Compute instance that is inside of the VCN of the Autonomous Database. Alternately, you can use the URL with a compute instance that has a direct connection to the VCN of the Autonomous Database.

Using the API

Use the [GenerateAutonomousDatabaseWallet](#) API operation to download the client credentials for your Autonomous Database.

Use the [UpdateAutonomousDatabaseWalletDetails](#) API operation to rotate the wallet for your Autonomous Database.

Use the [AutonomousDatabase](#) API operation to get the access URLs for Application Express (APEX) and SQL Developer Web.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

What's Next

For information and instructions on making secure connections to your database, see [Connecting to Autonomous Transaction Processing](#) and [Connecting to Autonomous Data](#)

[Warehouse](#).

Backing Up an Autonomous Database Manually

This topic describes how to create manual backups of Autonomous Databases. You can use the Oracle Cloud Infrastructure Console or the [API](#) to perform these tasks.

Oracle Cloud Infrastructure automatically backs up your Autonomous Databases and retains these backups for 60 days. Automatic backups are weekly full backups and daily incremental backups. You can also create manual backups to supplement your automatic backups. Manual backups are stored in an Object Storage bucket that you create, and are retained for 60 days.



Note

During the backup operation, the database remains available. However, lifecycle management operations such as stopping the database, scaling it, or terminating it are disabled.

Prerequisites

- To create or manage Autonomous Database backups, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service. See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

- To create a manual backup for an Autonomous Database, you must first configure an Object Storage bucket to serve as a destination for your manual backups. See [Setting Up a Bucket to Store Manual Backups](#) for instructions.

Using the Oracle Cloud Infrastructure Console

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, find the database that you want to back up.
4. Click the name to display the Autonomous Database details.
5. Click **Create Manual Backup**.



Tip

If this step does not successfully complete, confirm that you have an Object Storage bucket set up and configured to store manual backups. See [Setting Up a Bucket to Store Manual Backups](#) for instructions.

6. In the **Create Manual Backup** dialog, enter a name for your backup. Avoid entering confidential information.
7. Click **Create**. Your backup may take several hours to complete, depending on the size of your database.
8. Optionally, you can check the state of your backup in the list of backups on the database details page. For some states, an information icon (ⓘ) is displayed to provide additional details regarding the state or ongoing operations like deletions. The backup has one of the following states:
 - Creating
 - Active

- Deleting
- Deleted
- Failed

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Autonomous Database backups:

- [ListAutonomousDatabaseBackups](#)
- [GetAutonomousDatabaseBackup](#)
- [CreateAutonomousDatabaseBackup](#)

Setting Up a Bucket to Store Manual Backups

You must create an Oracle Cloud Infrastructure Object Storage bucket to hold your Autonomous Database manual backups and configure your database to connect to it. This is a one-time operation.

To set up an object store and user credentials for your manual backups

Some of the steps in this procedure require you to connect to the database by using an Oracle Database client such as SQL Developer. See [Connecting with Oracle SQL Developer \(18.2 or later\)](#) for information and instructions on connecting to an Autonomous Transaction Processing database. See [Connecting with Oracle SQL Developer \(18.2 or later\)](#) for information and instructions on connecting to an Autonomous Data Warehouse database.

1. If you have not already done so, generate an auth token for the Oracle Cloud Infrastructure Object Storage user to access the bucket you create in the next step. See [To create an auth token](#) to learn how to do this. (You will need this auth token for the

database credential you create in step 4.)

2. In the Oracle Cloud Infrastructure Console, [create a bucket](#) in your designated Object Storage [Swift compartment](#) to hold the backups. The format of the bucket name is `backup_databasename`, where *databasename* is lowercase. For example, if you provision a database named DATABASE1, the bucket name should be `backup_database1`.



Note

When you create your bucket:

- Pick **Standard** as the storage tier. Manual backups are only supported with buckets created in the [standard storage tier](#).
- Ensure that you use the [database name](#), and not the display name, as the bucket name.

3. Using an Oracle Database client, log in to the database as the administrator set the database `default_bucket` property to your Oracle Cloud Infrastructure Object Storage tenancy URL. The format of the tenancy URL is `https://swiftobjectstorage.region.oraclecloud.com/v1/object_storage_namespace`. For example:

```
ALTER DATABASE PROPERTY SET default_bucket='https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/ansh8lvru1zp';
```

In the example, the Object Storage namespace is **ansh8lvru1zp**.

**Tip**

- To determine the Object Storage namespace string to use, click **View Bucket** in the actions (three dots) menu of the bucket you created in the previous step.
- Do not include the bucket name in the URL.
- Ensure that you follow the format indicated. Do not use a pre-authenticated request URL.

4. With the tenancy user and the auth token referenced in step 1, create the credential for your Oracle Cloud Infrastructure Object Storage account. Use `DBMS_CLOUD.CREATE_CREDENTIAL` to create the credential.

For example:

```
BEGIN
DBMS_CLOUD.CREATE_CREDENTIAL(
  credential_name => 'DEF_CRED_NAME',
  username => 'db1_user@oracle.com',
  password => '<auth_token>'
);
END;
/
```

For more information on creating this credential, see [CREATE_CREDENTIAL Procedure](#).

5. Set the database property `default_credential` to the credential you created in the previous step.

For example:

```
ALTER DATABASE PROPERTY SET default_credential = 'ADMIN.DEF_CRED_NAME';
```

To list the current value for the default bucket, run the following command:

```
SELECT PROPERTY_VALUE from database_properties WHERE PROPERTY_NAME='DEFAULT_BUCKET';
```

After completing these steps you can take manual backups any time you want.

Restoring an Autonomous Database

This topic describes how to restore an Autonomous Database from a backup. You can use the Oracle Cloud Infrastructure Console or the [API](#) to perform this task.

You can use any existing manual or automatic backup to restore your database, or you can restore and recover your database to any point in time in the 60-day retention period of your automatic backups. For point-in-time restores, you specify a timestamp, and your Autonomous Database decides which backup to use for the fastest restore.



Note

Restoring Autonomous Database puts the database in the unavailable state during the restore operation. You cannot connect to a database in that state. The only lifecycle management operation supported in the unavailable state is terminate.

Prerequisites

To restore Autonomous Databases, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service. See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

Using the Oracle Cloud Infrastructure Console

To restore an Autonomous Database from a backup

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, find the database that you wish to restore.
4. Click the name of the Autonomous Database to display the database details.
5. Click the **Restore** button to open the restore dialog.
6. Click **Select Backup**.
7. Specify the date range for a list of backups to display.
8. Select the backup.
9. Click **Restore**.

To restore an Autonomous Database using point-in-time restore

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, find the database that you wish to restore.
4. Click the name of the Autonomous Database to display the database details.
5. Click the **Restore** button to open the restore dialog.
6. Click **Specify Timestamp**.
7. Enter a timestamp. Your Autonomous Database decides which backup to use for faster recovery. The timestamp input allows you to specify precision to the seconds level (YYYY-MM-DD HH:MM:SS GMT).

8. Click **Restore**.

Using the API

Use the [RestoreAutonomousDatabase](#) API operation to restore your Autonomous Database from a backup.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Cloning an Autonomous Database

This topic describes how to clone an existing Autonomous Database using the Oracle Cloud Infrastructure Console or the [API](#). You may wish to use the cloning feature to create a point-in-time copy of your Autonomous Database for purposes such as testing, development or analytics. If you need to clone only the database schema of your source database, the Console's cloning feature is a quick and easy way to accomplish this task.



Note

Any existing Autonomous Database (including those provisioned with preview version software) can be cloned using a preview version of Autonomous Database. However, preview version databases cannot be cloned using the regular (general-availability) Autonomous Database software.

Clone Types

The clone feature offers the following two types of Autonomous Database clones:

- The full clone option creates a new database that includes all of the source database's metadata and data.
- The metadata clone option creates a new database that includes the source database's metadata, but not the source database's data.

Prerequisites

- To clone an Autonomous Database, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service.

Password Requirement for New Database with Dedicated Deployment

When cloning a database that uses [dedicated deployment](#), the password you set for the target database cannot be one of the three most recently used passwords of the source database.

Using the Oracle Cloud Infrastructure Console

To clone an Autonomous Database with serverless deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to clone.
4. Go to **Actions**, and then click **Create Clone**.

In the **Create Autonomous Database Clone** dialog, enter the following:

CLONE TYPE

Select the [type of clone](#) you wish to create. Choose either **Full Clone** or **Metadata Clone**.

DATABASE INFORMATION

- **Compartment:** Your current compartment is the default selection.
- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. Avoid entering confidential information.
- **Database Name:** The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.
- **CPU Core Count:** You can enable up to 128 cores for your Autonomous Database. The actual number of available cores is subject to your tenancy's [service limits](#).

Auto Scaling: allows Autonomous Database to automatically increase the number of CPU cores by up to three times the assigned CPU core count value, depending on demand for processing. The auto scaling feature reduces the number of CPU cores when additional cores are not needed. For databases with up to 42 assigned cores, you can increase the maximum number of cores available through auto scaling by increasing the CPU core count value.



Note

The maximum number of cores that are available to any Autonomous Database database not using dedicated deployment is 128, regardless of whether auto scaling is enabled or not. This means that database with a CPU core count of 64 could auto scale up to two times the assigned number of cores ($2 \times 64 = 128$). A database with 42 cores (or fewer) could auto scale up to three times the assigned number ($3 \times 42 = 126$). For billing purposes, the database service determines the average number of CPUs used per hour.

- **Storage:** Specify the storage you wish to make available to your Autonomous Database database, in terabytes. You can make up to 128 TB available. For full clones, the size of the source database determines the minimum amount of storage you can make available.
- **Enable Preview Version:** (*This option only displays during periods when a preview version of Autonomous Database is available*) Select this option to provision the database with an Autonomous Database preview version. Preview versions of Autonomous Database are made available for limited periods for testing purposes. Do not select this option if you are provisioning a database for production purposes or if you will need the database to persist beyond the limited availability period of the preview version.

ADMINISTRATOR CREDENTIAL

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. You use this password when accessing the Autonomous

Database service console and when using an SQL client tool.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

LICENSE TYPE

The type of license you want to use for the Autonomous Transaction Processing database. Your choice affects metering for billing. You have the following options:

- **My Organization Already Owns Oracle Database Software Licenses:** This choice is used for the Bring Your Own License (BYOL) license type. If you choose this option, make sure you have proper entitlements to use for new service instances that you create.
- **Subscribe to New Database Software Licenses and the Database Cloud Service:** This is used for the License Included license type. With this choice, the cost of the cloud service includes a license for the Database service.

5. Click **Create Autonomous Database Clone**.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.
- Backups associated with the source database are not cloned for either the full clone or the metadata clone option.

- Oracle recommends that you evaluate the security requirements for the new database and implement them, as applicable. See [Security Considerations](#) for details.

To clone an Autonomous Database with dedicated deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to clone.
4. Go to **Actions**, and then click **Create Clone**.

In the **Create Autonomous Database Clone** dialog, enter the following:

CLONE TYPE

Select the [type of clone](#) you wish to create. Choose either **Full Clone** or **Metadata Clone**.

DATABASE INFORMATION

- **Compartment:**Your current compartment is the default selection.
- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique, and you can change it whenever you like. Avoid entering confidential information.
- **Database Name:** The database name must consist of letters and numbers only, starting with a letter. The maximum length is 14 characters. Avoid entering confidential information.
- **Autonomous Container Database:** (*Read only*) Currently, cloned target databases must be created within the same Autonomous Container Database as the source database.
- **CPU Core Count:** You can enable up to 92 cores for your target Autonomous

Database.

- **Storage:** Specify the storage you wish to make available to your Autonomous Database database, in terabytes. You can make up to 128 TB available. For full clones, the size of the source database determines the minimum amount of storage you can make available.

ADMINISTRATOR CREDENTIAL

Set the password for the Autonomous Database Admin user by entering a password that meets the following criteria. You use this password when accessing the Autonomous Database service console and when using an SQL client tool.

- Password cannot be one of the three most recently used passwords of the source database
- Between 12 and 30 characters long
- Contains at least one lowercase letter
- Contains at least one uppercase letter
- Contains at least one number
- Does not contain the double quotation mark (")
- Does not contain the string "admin", regardless of casing

5. Click **Create Autonomous Database Clone**.

The Console displays the details page for the new clone of your database and the service begins provisioning the Autonomous Database. Note the following:

- The new clone displays the **Provisioning** lifecycle state until the provisioning process completes.
- The source database remains in the **Available** lifecycle state.
- Backups associated with the source database are not cloned for either the full clone or the metadata clone option.

- Oracle recommends that you evaluate the security requirements for the new database and implement them, as applicable. See [Security Tools for Dedicated Deployments](#) for details.

Using the API

Use the [CreateAutonomousDatabase](#) API operation to clone an Autonomous Database.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For More Information

For information about optimizer statistics, resource management rules and performance data for a cloned database, see the [Using Oracle Autonomous Data Warehouse](#) and [Using Oracle Autonomous Transaction Processing](#) user guides.

Maintenance Updates for Autonomous Databases with Serverless Deployment

Autonomous Databases perform maintenance updates and database patching for you. Your database remains available throughout the maintenance process. This topic describes Autonomous Database maintenance for [serverless deployments](#).

For information on maintenance for dedicated deployments, see [Overview of Dedicated Deployment Maintenance](#).

Maintenance Duration

For Autonomous Databases with [serverless deployments](#), Oracle performs regular maintenance updates that generally take no more than two hours.

Checking the Scheduling of Maintenance Updates

To see when your next scheduled maintenance update is, navigate to the Autonomous Database details page in the Console for the database you are interested in. The **Next Maintenance** metadata field displays the beginning and ending times of the next database maintenance . You can also use the [GetAutonomousDatabase](#) API operation to determine the time of your next maintenance update.

Overview of Autonomous Database Dedicated Deployments

This topic describes the database architecture, features, and user roles of Autonomous Database dedicated deployments. For a general overview of Autonomous Databases that covers the basics common to both deployment options, see [Overview of Autonomous Database](#).

Dedicated Deployment Database Architecture

Autonomous Databases with dedicated deployment have a three-level database architecture model that makes use of [Oracle multitenant database architecture](#).

RESOURCE TYPES

Each level of the architecture model corresponds to one of the following resources types:

- An **Autonomous Exadata Infrastructure** resource, a hardware rack which includes compute nodes and storage servers, tied together by a high-speed, low-latency InfiniBand network and intelligent Exadata software. With dedicated deployment, you have exclusive use of the Exadata infrastructure and hardware on which your Autonomous Transaction Processing databases run.

For a list of the hardware and Oracle Cloud resource characteristics of Autonomous Exadata Infrastructure resources, see [Characteristics of Autonomous Exadata Infrastructure Resources](#).

- An **Autonomous Container Database**, which provides a container for multiple user databases. This resource is sometimes referred to as a CDB, and is functionally

equivalent to the multitenant container databases found in Oracle 12c and higher databases.

Multitenant architecture offers many advantages over non-CDB architecture. For example, it does the following:

- Allows you to easily manage multiple individual user databases
 - Makes more efficient use of database hardware, as individual databases may use only a fraction of the server hardware capacity
 - Allows for easier and more rapid movement of data and code
 - Allows for easier testing, as development databases can be housed within the same container as production databases
 - Allows for the separation of duties between database administrators, who manage only the individual Autonomous Database instances to which they are granted privileges, and fleet managers, who manage infrastructure resources and container databases.
- An **Autonomous Database**. You can create multiple Autonomous Databases within the same container database. This level of the database architecture is analogous to the pluggable databases (PDBs) found in non-Autonomous Exadata systems. Your Autonomous Database can be configured for either transaction processing or data warehouse workloads.

DEPLOYMENT ORDER

You must create the dedicated deployment resources in the following order:

1. Autonomous Exadata Infrastructure. See [Creating an Autonomous Exadata Infrastructure Resource](#) for more information.
2. Autonomous Container Database. See [Creating an Autonomous Container Database](#) for more information.
3. Autonomous Database. See [Creating an Autonomous Database](#) for more information.

RELATED RESOURCES

Related resources and prerequisites include:

- A **Virtual Cloud Network (VCN)** and a **Subnet**, which you create using Oracle Cloud Infrastructure's Networking service. You must have at least one VCN and one subnet available to provision an Autonomous Database with dedicated deployment.

For more information, see the following topics:

- [Network Isolation](#) (from *Fleet Administrator's Guide to Oracle Autonomous Transaction Processing Dedicated Deployments*)
 - [Overview of Networking](#)
 - [To create a VCN](#)
 - [To create a subnet](#)
- **Autonomous Backups**, created for you automatically by the Autonomous Database service. You do not have to provision storage for your backups. Backups are stored in Object Storage that is managed by Oracle Cloud Infrastructure's Database service. Note that automatic backups incur Object Storage usage costs. By default, backups are stored for 60 days. Using the Console, you can choose to change the retention period to 7, 15, or 30 days.
 - **Manual Backups**. Optionally, you can configure Autonomous Database to create on-demand manual backups. See [Setting Up a Bucket to Store Manual Backups](#) and [Backing Up an Autonomous Database Manually](#) for more information on manual backups. Manual backups are subject to the retention policy you have in place for the Autonomous Container Database. By default, manual backups are stored for 60 days.

User Roles

Your organization may choose to split the administration of the Autonomous Database with dedicated deployment into the following roles:

- **Fleet Administrator**. Fleet administrators create, monitor and manage Autonomous Exadata Infrastructure and Autonomous Container Database resources. A fleet administrator must have permissions for using the networking resources required by the dedicated deployment, and permissions to manage the infrastructure and container database resources.

See [Fleet Administrator's Guide to Oracle Autonomous Transaction Processing Dedicated Deployments](#) for a complete overview of the fleet administrator role.

- **Database Administrator.** Database administrators create, monitor and manage Autonomous Databases. They also create and manage users within the database. Database administrators must have permissions for using container databases, for managing Autonomous Transaction Processing databases and backups, and for using the related networking resources. For manual backups, they must have permissions to use the designated Object Storage bucket. At the time of provisioning an Autonomous Database, the administrator provides user credentials for the automatically created ADMIN account, which provides administrative rights to the new database. See [Using Oracle Autonomous Transaction Processing Dedicated Deployments](#) for a complete overview of the database administrator role.
- **Database User.** Database users are the developers who write applications that connect to and use an Autonomous Database to store and access the data. Database users do not need Oracle Cloud Infrastructure accounts. They gain network connectivity to and connection authorization information for the database from the database administrator.

CPU Provisioning, CPU Scaling, and Storage Scaling

Autonomous Database with dedicated deployment allows for the over-subscription of CPU cores when provisioning databases. This means that for a given infrastructure resource, the sum of all CPU cores provisioned to Autonomous Databases can exceed the number of CPU cores on the Exadata rack. This approach allows for better CPU core utilization. Note that if all databases are heavily loaded at the same time, there will be some contention for CPU resources, and degraded performance.

Additionally, you can scale the CPU count and the storage capacity of the database at any time without impacting availability or performance.

Overview of Dedicated Deployment Maintenance

Autonomous Database dedicated deployment systems have separate regularly scheduled maintenance runs for both Autonomous Exadata Infrastructure resources and Autonomous

Container Databases. You can choose to set the scheduling for your maintenance runs, or let the system handle maintenance scheduling. You can view the maintenance history for infrastructure instances and container databases in the Oracle Cloud Infrastructure Console.



Tip

Oracle recommends that you define the acceptable maintenance times for your Autonomous Exadata Infrastructure resources and Autonomous Container Databases. Doing so will prevent maintenance runs from occurring at times that would be disruptive to regular database operations.

AUTONOMOUS EXADATA INFRASTRUCTURE MAINTENANCE

Infrastructure maintenance takes place at least once each quarter and is mandatory. You can schedule the time your infrastructure maintenance will begin. Infrastructure maintenance runs are for infrastructure patching (including patching of the Exadata grid infrastructure code and operating systems updates), and do not include database patching. Oracle will notify you in the weeks leading up to the quarterly infrastructure patching date about the upcoming maintenance. You can also view scheduled maintenance runs in the Oracle Cloud Infrastructure console. The following tasks explain how to view scheduled and past maintenance updates, and to edit the maintenance schedule for an infrastructure instance:

- [To configure the automatic maintenance schedule for an Autonomous Exadata Infrastructure resource](#)
- [To view the next scheduled maintenance for an Autonomous Exadata Infrastructure resource](#)
- [To view the maintenance history of an Autonomous Exadata Infrastructure resource](#)

You can use the [GetMaintenanceRun](#), [ListMaintenanceRun](#), and [UpdateAutonomousExadataInfrastructure](#) API operations to view details about scheduled and

past maintenance updates, and to update the maintenance schedule of your infrastructure instance.

AUTONOMOUS CONTAINER DATABASE MAINTENANCE

Container database maintenance takes place at the time of the quarterly infrastructure maintenance run. Container database maintenance includes Oracle Database software patches. Autonomous Database offers two CDB maintenance type choices:

- Release Update (RU): Autonomous Database installs only the most current release update.
- Release Update Revision (RUR): Autonomous Database installs the release update plus additional fixes.

The following tasks explain how to view and edit maintenance update information for Autonomous Container Databases:

- [To view the maintenance history of an Autonomous Container Database](#)
- [To skip a scheduled maintenance run for an Autonomous Container Database](#)
- [To configure the type of maintenance patching for an Autonomous Container Database](#)

Use the [UpdateAutonomousContainerDatabase](#) API operation to change the patching type for an Autonomous Container Database. Use the [ListMaintenanceRun](#) API operation to see past maintenance update information. Use the [UpdateMaintenanceRun](#) API operations to skip a container database maintenance update. You can skip maintenance runs for up to 2 consecutive quarters if needed.

Using the Oracle Cloud Infrastructure Console to Manage Dedicated Deployments

For information on provisioning, managing, and backing up dedicated deployments in the Oracle Cloud Infrastructure Console, see the following topics:

FOR FLEET ADMINISTRATORS

- [Creating an Autonomous Exadata Infrastructure Resource](#)
- [Creating an Autonomous Container Database](#)

- [Managing an Autonomous Exadata Infrastructure Resource](#)
- [Managing an Autonomous Container Database](#)
- [Fleet Administrator's Guide to Oracle Autonomous Transaction Processing Dedicated Deployments](#) (*complete fleet administrator guide*)

FOR DATABASE ADMINISTRATORS

- [Creating an Autonomous Database](#)
- [Managing an Autonomous Database](#)
- [Connecting to an Autonomous Database](#)
- [Backing Up an Autonomous Database Manually](#)
- [Restoring an Autonomous Database](#)
- [Using Oracle Autonomous Transaction Processing Dedicated Deployments](#) (*complete database administrator guide*)

Additional Information

For known issues, see [Known Issues for Oracle Autonomous Transaction Processing Dedicated Deployments](#).

Creating an Autonomous Exadata Infrastructure Resource

This topic describes how to provision an [Autonomous Exadata Infrastructure](#) resource for Autonomous Databases with dedicated deployment, using the Oracle Cloud Infrastructure Console or the [API](#). For an overview of the dedicated deployment option, see [Overview of Autonomous Database Dedicated Deployments](#).

The infrastructure resource includes the physical Exadata hardware and intelligent Exadata software. Once you have provisioned an infrastructure instance, you can provision one or more [Autonomous Container Databases](#) to run on your infrastructure. To provision an Autonomous Database, you must have both an infrastructure resource and at least one container database available.



Note

This topic is not applicable to Autonomous Databases using the serverless deployment option.

Prerequisites

- To create an Autonomous Exadata Infrastructure resource, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service.



Tip

See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

- You will also need a **Virtual Cloud Network** and a **Subnet**, which you create using Oracle Cloud Infrastructure's [Networking service](#). For information on creating and managing these resources, see [VCNs and Subnets](#).

Using the Oracle Cloud Infrastructure Console

TO CREATE AN AUTONOMOUS EXADATA INFRASTRUCTURE RESOURCE

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Click **Autonomous Exadata Infrastructure**.
4. Click **Create Autonomous Exadata Infrastructure**.
5. In the **Create Autonomous Exadata Infrastructure** dialog, enter the following general information:
 - **Display Name:** A user-friendly description or other information that helps you easily identify the infrastructure resource. The display name does not have to be unique. Avoid entering confidential information.
 - **Compartment:** Specify the compartment in which the Autonomous Exadata Infrastructure will be created.
 - **Availability Domain:** Select an availability domain for the Autonomous Exadata Infrastructure.
 - **Shape:** *Read only.* Oracle Cloud Infrastructure currently offers the Exadata.Quarter2.92 shape for provisioning an Autonomous Exadata Infrastructure resource. This shape has 92 cores, 106 TB storage, and 1440 GB RAM. See [Exadata X7 shapes](#) for more information on Exadata shape specifications.
6. Enter the following network information:
 - **Virtual cloud network compartment:** The compartment containing the VCN you wish to use for the Autonomous Exadata Infrastructure. The default value is the user's current compartment. Click **change compartment** to select a VCN in a different compartment.
 - **Virtual cloud network:** The VCN in which to launch the Autonomous Exadata Infrastructure.

- **Subnet compartment:** The compartment containing the subnet you wish to use for the Autonomous Exadata Infrastructure. The default value is the user's current compartment. Click **change compartment** to select a subnet in a different compartment.
- **Subnet:** The subnet to which the Autonomous Exadata Infrastructure should attach. Do not use a subnet that overlaps with 141.144.75.0/24.
- **Use network security groups to control traffic:** *Optional.* You can specify up to five network security groups (NSGs) for your Autonomous Exadata Infrastructure resource by selecting this option. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress [security rules](#) to your infrastructure resource. A maximum of five NSGs can be specified. To add an NSG, select the compartment containing the NSG using the **Network security group compartment** selector, then select the NSG itself using the **Network security group** selector.

For more information on creating and working with NSGs, see [Network Security Groups](#).

Note that if you choose a subnet with a [security list](#), the security rules for the infrastructure resource will be a union of the rules in the security list and the NSGs.

7. Optionally, you can specify the date and start time for the Autonomous Exadata Infrastructure quarterly maintenance.



Tip

Oracle recommends that you define the acceptable maintenance times for your Autonomous Exadata Infrastructure resources and Autonomous Container Databases. Doing so will prevent maintenance runs from occurring at times that would be disruptive to regular database operations.

To change the Autonomous Exadata Infrastructure maintenance schedule

- a. In the Create Autonomous Exadata Infrastructure dialog, click **Modify Schedule**.
 - b. In the Automatic Maintenance Schedule dialog, select **Specify a Schedule**
 - c. In the **Maintenance months** selector, specify at least one month for each quarter during which infrastructure maintenance will take place.
 - d. For **Week of the Month**, select a week during the month that the maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. For example, to allow maintenance during the 2nd week of the month (from the 8th day to the 14th day of the month), use the value 2. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days.
 - e. For **Day of the Week**, select the day of the week that the maintenance will take place.
 - f. For **Start Hour**, select one of the six start time windows available. The maintenance will begin during the 4 hour time window that you specify and may continue beyond the end of the period chosen. The start time window is specified in universal coordinated time (UTC).
 - g. Click **Update Maintenance Schedule**.
8. Choose the license type you wish to use. Your choice affects metering for billing. You have the following options:
- **Bring your own license**: If you choose this option, make sure you have proper entitlements to use for new service instances that you create.
 - **License included**: With this choice, the cost of the cloud service includes a license for the Database service.
9. The following **Advanced Options** are available:
- Tags** - Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined

tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

10. Click **Create Autonomous Exadata Infrastructure**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [LaunchAutonomousExadataInfrastructure](#) API operation to create an Autonomous Exadata Infrastructure resource.

What's Next

After creating an Autonomous Exadata Infrastructure resource, you can [create one or more Autonomous Container Databases](#) within your infrastructure. You must have provisioned both an infrastructure resource and at least one container database before you can [create your first Autonomous Database](#) in Oracle Cloud Infrastructure.

Managing an Autonomous Exadata Infrastructure Resource

This topic describes the Autonomous Exadata Infrastructure management tasks for Autonomous Databases that you complete using the Oracle Cloud Infrastructure Console or the [API](#). Autonomous Exadata Infrastructure resources are used by Autonomous Databases with dedicated deployment. For an overview of the dedicated deployment option, see [Overview of Autonomous Database Dedicated Deployments](#).



Note

This topic is not applicable to Autonomous Databases using the serverless deployment option.

Using the Oracle Cloud Infrastructure Console

The following management operations can be performed on Autonomous Exadata Infrastructure resources in Oracle Cloud Infrastructure:

- Configuring the automatic maintenance schedule
- Viewing the next scheduled maintenance date and maintenance history
- Copying the Autonomous Exadata Infrastructure endpoint
- Terminating the Autonomous Exadata Infrastructure resource



Tip

Oracle recommends that you define the acceptable maintenance times for your Autonomous Exadata Infrastructure resources and Autonomous Container Databases. Doing so will prevent maintenance runs from occurring at times that would be disruptive to regular database operations.

To configure the automatic maintenance schedule for an Autonomous Exadata Infrastructure resource

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.

2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.
5. On the Autonomous Exadata Infrastructure details page, under **Maintenance**, click the **edit** link in the **Maintenance Schedule** field.
6. In the **Automatic Maintenance Schedule** dialog, select **Specify a schedule**.
7. Under **Maintenance months**, specify at least one month for each quarter during which Autonomous Exadata Infrastructure maintenance will take place.
8. Under **Week of the month**, specify which week of the month maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days.
9. Under **Day of the week**, specify the day of the week on which the maintenance will occur.
10. Under **Start hour**, specify the hour during which the maintenance run will begin.
11. Click **Update Maintenance Schedule**.

To view the next scheduled maintenance for an Autonomous Exadata Infrastructure resource

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.

5. On the Autonomous Exadata Infrastructure details page, under **Maintenance**, click the **view** link in the **Next Maintenance** field.
6. On the **Maintenance** page, scheduled maintenance events are listed under the **Regular Autonomous Exadata Infrastructure maintenance** heading.

To view the maintenance history of an Autonomous Exadata Infrastructure resource

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.
5. On the Autonomous Exadata Infrastructure details page, under **Maintenance**, click the **view** link in the **Next Maintenance** field.
6. On the **Maintenance** page, under **Autonomous Database Maintenance**, click **History**. In the list of past maintenance events, you can click on an individual event title to read the details of the maintenance that took place. Maintenance event details include the following:
 - The category of maintenance (quarterly software maintenance, hardware maintenance, or a critical patch)
 - Whether the maintenance was scheduled or unplanned
 - The OCID of the maintenance event. (Go to **Actions**, then choose **Copy OCID**.)
 - The start time and date of the maintenance

To view or copy the Autonomous Exadata Infrastructure endpoint

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.
5. On the Autonomous Exadata Infrastructure Information tab, click **Show** or **Copy** in the DB Infrastructure Endpoint Name field.

To edit the network security groups (NSGs) for your Autonomous Exadata Infrastructure resource

Your Autonomous Exadata Infrastructure instance can use up to five network security groups (NSGs). Note that if you choose a subnet with a [security list](#), the security rules for the infrastructure instance will be a union of the rules in the security list and the NSGs. For more information, see [Network Security Groups](#).

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.
5. In the **Network** details, click the **Edit** link to the right of the **Network Security Groups** field.
6. In the **Edit Network Security Groups** dialog, click **+ Another Network Security Group** to add an NSG to the Autonomous Exadata Infrastructure resource.

To change an assigned NSG, click the drop-down menu displaying the NSG name, then select a different NSG.

To remove an NSG from your DB system, click the **X** icon to the right of the displayed NSG name.

7. Click **Save**.

To move an Autonomous Exadata Infrastructure resource to another compartment



Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see [Details for the Database Service](#).

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you wish to move.
5. Click **Move Resource**.
6. Select the new compartment.
7. Click **Move Resource**.

For information about dependent resources for Database resources, see [Moving Database Resources to a Different Compartment](#).

To terminate an Autonomous Exadata Infrastructure resource



Warning

Terminating an Autonomous Exadata Infrastructure resource permanently deletes it and removes associated resources such as Autonomous Container Databases, Autonomous Databases, and Autonomous Database backups. You cannot recover a terminated Autonomous Exadata Infrastructure resource.

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Exadata Infrastructure**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the resource you are interested in.
5. Go to **Actions**, and then click **Terminate**.
6. Confirm that you wish to terminate your Autonomous Exadata Infrastructure in the confirmation dialog.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [UpdateAutonomousExadataInfrastructure](#) API operation to configure the automatic maintenance schedule for your infrastructure resource.

Use the [GetMaintenanceRun](#) API to view the details of a maintenance run that is scheduled, in progress, or that has ended.

Use the [ListMaintenanceRun](#) API to get a list of maintenance runs in a specified compartment.

Use the [ChangeAutonomousExadataInfrastructureCompartment](#) API operation to move an Autonomous Exadata Infrastructure resource to another compartment.

Use the [TerminateAutonomousExadataInfrastructure](#) API operation to delete an Autonomous Exadata Infrastructure resource.

For More Information

[Create and Manage Autonomous Exadata Infrastructure Resources](#) (*Fleet Administrator's Guide to Oracle Autonomous Transaction Processing Dedicated Deployments*)

Creating an Autonomous Container Database

This topic describes how to provision a new Autonomous Container Database using the Oracle Cloud Infrastructure Console or the [API](#). Container databases are only necessary for Autonomous Databases that use the [dedicated deployment](#) option. For a brief overview of the dedicated deployment option, see [Overview of Autonomous Database Dedicated Deployments](#).



Note

This topic is not applicable to Autonomous Databases using the serverless deployment option.

Prerequisites

- To create an Autonomous Container Database, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service.



Tip

See [Let database and fleet administrators manage Autonomous Databases](#) for sample Autonomous Database policies. See [Details for the Database Service](#) for detailed information on policy syntax.

- To create an Autonomous Container Database, you must have an available Autonomous Exadata Infrastructure instance. For information on creating an infrastructure instance, see [Creating an Autonomous Exadata Infrastructure Resource](#).

Using the Oracle Cloud Infrastructure Console

TO CREATE AN AUTONOMOUS CONTAINER DATABASE

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Click **Autonomous Container Database**.
3. Click **Create Autonomous Container Database**.
4. In the **Create Autonomous Container Database** dialog, enter the following database information:

- **Display Name:** A user-friendly description or other information that helps you easily identify the resource. The display name does not have to be unique. Avoid entering confidential information.
 - **Compartment:** Specify the compartment in which the container database will be created.
5. Choose the Autonomous Exadata Infrastructure you wish to use to create your container database:
- **Compartment:** Specify the compartment containing the Autonomous Exadata Infrastructure you wish to use for your container database.
 - **Autonomous Exadata Infrastructure:** Choose an Autonomous Exadata Infrastructure for your container database. See [Creating an Autonomous Exadata Infrastructure Resource](#) for more information.
6. Optionally, you can change the maintenance type for your Autonomous Container Database. The maintenance type choices are Release Update (RU) and Release Update Revision (RUR). The Release Update setting installs only the most current release update, while the Release Update Revision installs the release update plus additional fixes. For more information, see [Release Update Introduction and FAQ \(Doc ID 2285040.1\)](#) in the My Oracle Support online help portal.

To change the Autonomous Container Database maintenance type

- a. In the Create Autonomous Container Database dialog, click **Modify Schedule**.
 - b. In the **Automatic Maintenance Type** dialog, select your desired maintenance type.
 - c. Click **Update Automatic Maintenance**.
7. The following **Advanced Options** are available:
- **Tags** - Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you

should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.

- **Management** - Optionally, you can specify the backup retention policy, which controls the length of time you retain backups in the Autonomous Container Database. The choices are 7 days, 15 days, 30 days, and 60 days. The default setting is 60 days.
8. Click **Create Autonomous Container Database**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [CreateAutonomousContainerDatabase](#) API operation to create an Autonomous container database.

What's Next

After creating an Autonomous Container Database, you can [create one or more Autonomous Databases](#) within the container database.

Managing an Autonomous Container Database

This topic describes the database management tasks for Autonomous Container Databases that you complete using the Oracle Cloud Infrastructure Console or the [API](#). Container databases are used by Autonomous Databases using the dedicated deployment option. For an overview of the dedicated deployment option, see [Overview of Autonomous Database Dedicated Deployments](#).



Note

This topic is not applicable to Autonomous Databases using the serverless deployment option.

The following management operations can be performed on Autonomous Container Databases in Oracle Cloud Infrastructure:

- Edit the backup retention policy. By default, database backups are retained for 60 days. You have the option of retaining backups for 7, 15, 30, or 60 days. The current backup retention policy for an Autonomous Container Database is displayed on the Autonomous Container Database details page.
- Configure the type of database maintenance. You can choose to use Release Update (RU) or Release Update Revision (RUR) updates for your Autonomous Container Database maintenance. For more information, see [Release Update Introduction and FAQ \(Doc ID 2285040.1\)](#) in the My Oracle Support online help portal.
- View the Autonomous Container Database next scheduled maintenance and maintenance history.
- Skip a scheduled maintenance run. For container databases, you can skip maintenance runs for up to 2 consecutive quarters if needed.
- Perform a rolling restart of databases within an Autonomous Container Database. You can perform a "rolling restart" on all the Autonomous Databases in an Autonomous Container Database to ensure that the current memory allocation is optimized. During a rolling restart, each node of an Autonomous Database is restarted separately while the remaining nodes continue to be available. No interruption of service occurs during a rolling restart. You cannot perform a container database restart if a backup is in progress.
- Terminate an Autonomous Container Database. Note that you must terminate all Autonomous Databases within a container database before you can terminate the container database itself.

Using the Oracle Cloud Infrastructure Console

To set the backup retention policy for an Autonomous Container Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Container Databases, click on the display name of the container database you are interested in.
5. On the Autonomous Container Database details page, under **Backup**, click the **Edit** link in the **Backup Retention Field**.
6. Specify a backup retention period from the list of choices.
7. Click **Save Changes**.

To configure the type of maintenance patching for an Autonomous Container Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Container Databases, click on the display name of the container database you are interested in.
5. On the Autonomous Container Database details page, under **Maintenance**, click the **Edit** link in the **Maintenance Schedule** field.
6. In the **Automatic Maintenance Schedule** dialog, under **Maintenance Type**, select either **Release Update (RU)** or **Release Update Revision (RUR)**.

7. Click **Update Maintenance Schedule**.

To view the maintenance history of an Autonomous Container Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Container Databases, click on the display name of the container database you are interested in.
5. On the Autonomous Container Database details page, under **Maintenance**, click the **Edit** link in the **Maintenance Schedule** field.
6. On the **Maintenance** page, under **Autonomous Database Maintenance**, click **History**. In the list of past maintenance events, you can click on an individual event title to read the details of the maintenance that took place. Maintenance event details include the following:
 - The category of maintenance (quarterly software maintenance or a critical patch)
 - Whether the maintenance was scheduled or unplanned
 - The OCID of the maintenance event. (Go to **Actions**, then choose **Copy OCID**.)
 - The start time and date of the maintenance

To skip a scheduled maintenance run for an Autonomous Container Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.

4. In the list of Autonomous Container Databases, click on the display name of the container database you are interested in.
5. On the Autonomous Container Database details page, under **Maintenance**, click the **View** link in the **Next Maintenance** field.
6. On the **Maintenance** page, any container database maintenance events planned for the next 15 days will appear in the list of maintenance events. To skip a container database maintenance run, click the **Skip Maintenance** button for the scheduled maintenance event.

To perform a rolling restart of databases within an Autonomous Container Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Container Databases, click on the display name of the container database you are interested in.
5. On the Autonomous Container Database details page, click **Restart**.
6. In the confirmation dialog, type the name of the Autonomous Container Database.
7. Click **Restart**.

To move an Autonomous Container Database to another compartment



Note

To move resources between compartments, resource users must have sufficient access permissions on the



compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see [Details for the Database Service](#).

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Container Databases, click on the display name of the container database you wish to move.
5. Click **Move Resource**.
6. Select the new compartment.
7. Click **Move Resource**.

For information about dependent resources for Database resources, see [Moving Database Resources to a Different Compartment](#).

To terminate an Autonomous Container Database



Warning

Terminating an Autonomous Container Database permanently deletes it and removes associated resources such as Autonomous Databases and Autonomous Database backups. You cannot recover a terminated Autonomous Container Database.

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. Under **Autonomous Database**, click **Autonomous Container Database**.
4. In the list of Autonomous Exadata Infrastructure resources, click on the display name of the infrastructure resource you are interested in.
5. Click **Terminate**.
6. Confirm that you wish to terminate your Autonomous Exadata Infrastructure in the confirmation dialog.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [UpdateAutonomousContainerDatabase](#) API operation to perform the following management actions:

- Set the backup retention period for an Autonomous Container Database.
- Set the maintenance patching type of an Autonomous Container Database.

Use the [UpdateMaintenanceRun](#) API operation to skip a container database maintenance run.

Use the [ListMaintenanceRun](#) API to get a list of maintenance runs in a specified compartment. Can be used to see maintenance history and scheduled maintenance runs.

Use the [RestartAutonomousContainerDatabase](#) API operation to perform a rolling restart on a container database.

Use the [ChangeAutonomousContainerDatabaseCompartment](#) API operation to move a container database to another compartment.

Use the [TerminateAutonomousContainerDatabase](#) API operation to terminate a container database.

Autonomous Database Development and Administration Tools

This topic describes the Oracle Database tools available for Autonomous Database using the Console and how to access them using the Console. The following tools can be accessed directly from the Oracle Cloud Infrastructure Console:

- [Oracle SQL Developer Web](#)
- [Oracle Application Express](#)
- [Oracle Machine Learning User Administration](#) (available for serverless deployments only)



Tip

Autonomous Database supports a range of other Oracle and third-party tools and applications. See [Autonomous Data Warehouse Tools and Application Test Matrix](#) to learn about other tools you can use with your Autonomous Database.

For Autonomous Databases with [serverless deployment](#), additional tools can be accessed through the [Service Console](#).

Oracle SQL Developer Web

Oracle SQL Developer Web in Autonomous Data Warehouse provides a development environment and a data modeler interface for Autonomous Databases. SQL Developer Web is available for both [dedicated](#) and [serverless](#) Autonomous Database deployments.

The main features of SQL Developer Web are:

- Run SQL statements and scripts in the worksheet
- Export data

- Design Data Modeler diagrams using existing objects

SQL Developer Web is a browser-based interface of Oracle SQL Developer and provides a subset of the features of the desktop version.

For more information, see the Autonomous Transaction Processing and Autonomous Data Warehouse user guides. Each guide contain a section on using SQL Developer Web with Autonomous Databases. Use the following links:

- [Using Oracle Autonomous Transaction Processing](#)
- [Using Oracle Autonomous Data Warehouse](#)

Complete product information can be found in [About Oracle SQL Developer Web](#).

Oracle Application Express

Oracle Application Express (APEX) is a low-code development platform that enables you to build scalable, secure enterprise applications with world-class features that can be deployed anywhere. APEX provides you with an easy-to-use browser-based environment to load data, manage database objects, develop REST interfaces, and rapidly build applications for both desktop and mobile devices.

Oracle Application Express is available for both [dedicated](#) and [serverless](#) Autonomous Database deployments.

See [Oracle Application Express](#)

For complete information, see the APEX topics in the Autonomous Transaction Processing and the Autonomous Data Warehouse user guides. Use the following links:

- [Using Oracle Autonomous Transaction Processing](#)
- [Using Oracle Autonomous Data Warehouse](#)

Oracle Machine Learning User Administration

[Oracle Machine Learning](#) is a collaborative web-based interface that provides a development environment to create data mining notebooks where you can perform data analytics, data

discovery and data visualizations. Using the Oracle Cloud Infrastructure Console, you can quickly get to the Oracle Machine Learning User Administration interface to create and manage users.

Machine Learning is currently available for [serverless](#) Autonomous Database deployments only.

Using the Oracle Cloud Infrastructure Console

FOR AUTONOMOUS DATABASES WITH SERVERLESS DEPLOYMENTS

To access Oracle SQL Developer Web

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to work with.
4. Click the **Tools** tab on the Autonomous Database Details page.
5. Click **Open SQL Developer Web**

To access Oracle Application Express (APEX)

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to work with.
4. Click the **Tools** tab on the Autonomous Database Details page.
5. Click **Open APEX**

To access Oracle Machine Learning's User Administration Interface

To use Oracle Machine Learning with your Autonomous Database, you must first create a user account within the application. The following steps explain how to navigate to the User Administration interface for Machine Learning from the Autonomous Database details page within the Console.

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to work with.
4. Click the **Tools** tab on the Autonomous Database Details page.
5. Click **Open Oracle ML User Administration**

FOR AUTONOMOUS DATABASES WITH DEDICATED DEPLOYMENTS

For dedicated deployments, the Console provides access URLs for Application Express (APEX) and SQL Developer Web that you can use to connect to these applications. The URLs only work from browsers within the same VCN as the Autonomous Database being accessed by the applications. Therefore, to use these URLs, you will need to open a browser running on a computer that meets one of the following conditions:

- The computer is a Compute instance is provisioned in the VCN of the Autonomous Database.
- The computer has a direct connection to the VCN of the Autonomous Database

To access APEX or SQL Developer Web, paste the appropriate access URL into the browser's address field, and then provide the Autonomous Database username and password when prompted. For more information on APEX, see the [APEX documentation](#). For more information on SQL Developer Web, see [Oracle SQL Developer Web](#).

The following tasks explain how to obtain an access URLs for APEX and SQL Developer Web.

To obtain the access URL for Oracle SQL Developer Web for an Autonomous Database with dedicated deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to work with.
4. Click the **Tools** tab on the Autonomous Database Details page.
5. Click **Open SQL Developer Web**

To obtain the access URLs for Oracle Application Express (APEX) for an Autonomous Database with dedicated deployment

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to work with.
4. Click the **Tools** tab on the Autonomous Database Details page.
5. Click **Open APEX**

Overview of the Always Free Autonomous Database

Oracle Cloud Infrastructure's Always Free Autonomous Database is part of Oracle Cloud Infrastructure's [Free Tier](#) of services. You can provision up to two Always Free Autonomous Databases in the [home region](#) of your tenancy. These databases are provided free of charge, and they are available to users of both free and paid accounts. You can use these Autonomous Databases for small-scale applications, for development or testing purposes, or for learning about and exploring Oracle Cloud Infrastructure.

Always Free Autonomous Database Specifications

- **Processor:** 1 Oracle CPU processor (cannot be scaled)
- **Memory:** 8 GB RAM
- **Database Storage:** 20 GB storage (cannot be scaled)
- **Workload Type:** Your choice of either the [transaction processing](#) or [data warehouse](#) workload type
 - The **Autonomous Transaction Processing** workload type configures the database for a transactional workload, with a bias towards high volumes of random data access.
For a complete product overview of Autonomous Transaction Processing, see [Autonomous Transaction Processing](#). For Autonomous Transaction Processing tutorials, see [Quick Start tutorials](#).
 - The **Autonomous Data Warehouse** workload type configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations.
For a complete product overview of Autonomous Data Warehouse, see [Autonomous Data Warehouse](#). For Autonomous Data Warehouse tutorials, see [Quick Start tutorials](#).
- **Deployment Type:** [Serverless deployment](#).
- **Maximum Simultaneous Database Sessions:** 20



Tip

Always Free Autonomous Databases can be [upgraded to regular paid instances](#) after provisioning in order to access all Autonomous Database features.

For free accounts, you will need to [upgrade your account](#) before upgrading an Always Free Autonomous Database to a paid instance.

Lifecycle for Always Free Autonomous Databases

After provisioning, you can continue using your Always Free Autonomous Database for as long as you want at no charge. You can terminate the database at any time.

If your Always Free Autonomous Database has no activity for a period of 7 consecutive days, the Database service will stop the database automatically. If this happens, you are allowed to restart the database and continue using it. If your Always Free Autonomous Database remains in a stopped state for 3 consecutive months, the resource will be reclaimed by the Database service.

Exadata Cloud at Customer

Exadata Cloud at Customer enables you to apply the combined power of Exadata and Oracle Cloud Infrastructure inside your own data center. You have full access to the features and capabilities of Oracle Database along with the intelligent performance and scalability of Exadata, but with Oracle owning and managing the Exadata infrastructure. You can use the Oracle Cloud Infrastructure console and APIs to manage Exadata Cloud at Customer like any other cloud resource, while maintaining absolute sovereignty over your data.

Each Exadata Cloud at Customer system configuration contains compute nodes (database servers) and Exadata Storage Servers that are interconnected using a high-speed, low-latency InfiniBand network and intelligent Exadata software. Each configuration is equipped with a fixed amount of memory, storage, and network resources.

Exadata Cloud at Customer uses virtual machine (VM) technology to separate the customer-managed and Oracle-managed components on each compute node. You have root privilege for the Exadata compute node VMs, so you can manage the Oracle Database, Oracle Grid Infrastructure, and Exadata system software. However, you do not have administrative access to the physical compute node hardware, which Oracle administers.

Exadata Cloud at Customer uses Exadata Storage Servers for database storage. The storage is allocated to disk groups managed by Oracle Automatic Storage Management (ASM). You have full administrative access to the ASM disk groups, but Oracle administers the Exadata Storage Server hardware and software.

In addition to the compute node hardware and Exadata Storage Servers, Oracle also manages other Exadata Cloud at Customer infrastructure components, including the network switches, power distribution units (PDUs), and integrated lights-out management (ILOM) interfaces.

Subscription to Exadata Cloud at Customer can include all of the required Oracle Database software licenses, or you can choose to bring Oracle Database software licenses that you already own to Exadata Cloud at Customer. If you choose to include Oracle Database software licenses in your Exadata Cloud at Customer subscription, then the included licenses contain all of the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (RAC). Exadata Cloud at Customer also comes with cloud-specific software tools that assist with administration tasks, such as backup, recovery, and patching.

On each Exadata Cloud at Customer system, you can create one or more databases. Apart from the inherent storage and processing capacity of your Exadata system, there is no set maximum for the number of databases that you can create.

System Configuration

Exadata Cloud at Customer is offered in the following system configurations:

- **Base System:** Containing two compute nodes and three Exadata Storage Servers. A Base System is an entry-level configuration. Compared to other configurations, a Base System contains Exadata Storage Servers with significantly less storage capacity and

compute nodes with significantly less memory and processing power.

- Quarter Rack: Containing two compute nodes and three Exadata Storage Servers.
- Half Rack: Containing four compute nodes and six Exadata Storage Servers.
- Full Rack: Containing eight compute nodes and 12 Exadata Storage Servers.

Each system configuration is equipped with a fixed amount of memory, storage, and network resources, and all system configurations are based on Oracle Exadata X8 systems.

The following table outlines the technical specifications for each Exadata Cloud at Customer system configuration.

Exadata Cloud at Customer X8 System Specifications

Property	Base System	Quarter Rack	Half Rack	Full Rack
Number of Compute Nodes	2	2	4	8
Total Maximum Number of Enabled CPU Cores	48	100	200	400
Total RAM Capacity	720 GB	1440 GB	2880 GB	5760 GB
Number of Exadata Storage Servers	3	3	6	12
Total Raw Flash Storage Capacity	38.4 TB	76.8 TB	153.6 TB	307.2 TB
Total Raw Disk Storage Capacity	252 TB	504 TB	1008 TB	2016 TB
Total Usable Storage Capacity	74.8 TB	149.7 TB	299.4 TB	598.7 TB

Storage Configuration

As part of configuring each Exadata Cloud at Customer VM cluster, the storage space inside the Exadata Storage Servers is configured for use by Oracle Automatic Storage Management (ASM). By default, the following ASM disk groups are created:

- The DATA disk group is primarily intended for the storage of Oracle Database data files. Also, a small amount of space is allocated from the DATA disk group to support the shared file systems that are used to store software binaries (and patches) and files associated with the cloud-specific tooling. You should not store your own data, including Oracle Database data files, backups, trace files, and so on, inside the system-related ACFS file systems.
- The RECO disk group is primarily used for storing the Fast Recovery Area (FRA), which can be used to provide a local store for files related to backup and recovery. By default, the FRA is used to store archived redo log files and the backup control file. If you configure your VM cluster with the option to allocate storage for local backups, then you can use the FRA as a database backup destination. Finally, if you enable flashback features on a database, then the FRA is used to store the flashback logs.

In addition, you can optionally create the SPARSE disk group. The SPARSE disk group is required to support Exadata snapshot functionality. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily. Snapshot clones are often used for development, testing, or other purposes that require a transient database. For more information about Exadata snapshot functionality, see [Setting up Oracle Exadata Storage Snapshots](#) in *Oracle Exadata System Software User's Guide*.

Impact of Configuration Settings on Storage

As an input to the VM cluster creation process, you must choose options that determine how storage space in the Exadata Storage Servers is allocated to the ASM disk groups:

- **Allocate Storage for Exadata Snapshots:** If you select this option, the SPARSE disk group is created and less space is allocated to the DATA and RECO disk groups. If you do not select this option, then the SPARSE disk group is not created and you cannot use Exadata snapshot functionality.
- **Allocate Storage for Local Backups:** If you select this option, more space is allocated to the RECO disk group to accommodate local backups to Exadata storage. If

CHAPTER 11 Database

you do not select this option, more space is allocated to the DATA disk group but you cannot use local Exadata storage as a backup destination for any databases in the VM cluster.

Your choices profoundly affect how storage space in the Exadata Storage Servers is allocated to the ASM disk groups. The following table outlines the proportional allocation of storage among the DATA, RECO, and SPARSE disk groups for each possible configuration:

Configuration Settings	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Allocate Storage for Exadata Snapshots: No Enable Backups on Local Exadata Storage: No	80%	20%	0% The SPARSE disk group is not created.
Allocate Storage for Exadata Snapshots: No Enable Backups on Local Exadata Storage: Yes	40%	60%	0% The SPARSE disk group is not created.
Allocate Storage for Exadata Snapshots: Yes Enable Backups on Local Exadata Storage: No	60%	20%	20%
Allocate Storage for Exadata Snapshots: Yes Enable Backups on Local Exadata Storage: Yes	35%	50%	15%

Preparing for Exadata Cloud at Customer

This topic describes the site and network requirements to deploy Exadata Cloud at Customer in a customer data center. Checklists are also provided to help you prepare for Exadata Cloud at Customer.

Site Requirements

The following outlines the site requirements for Exadata Cloud at Customer.

Space

The space requirements for each Exadata Cloud at Customer X8 rack are as follows:

Description	Millimeters (mm)	Inches (")
Height	2000 mm	78.74"
Width	601 mm	23.66"
Depth	1197 mm	47.13"

Weight

The following table lists the weight of each Exadata Cloud at Customer X8 rack:

Model (X8)	Kilograms (kg)	Pounds (lbs)
Base System	435.9 kg	961.1 lbs
Quarter Rack	449.0 kg	989.8 lbs

Model (X8)	Kilograms (kg)	Pounds (lbs)
Half Rack	591.5 kg	1304.1 lbs
Full Rack	883.9 kg	1948.7 lbs

Receiving, Unpacking, and Access

Before your Exadata Cloud at Customer rack arrives, ensure that the receiving area is large enough for the package.

Use the following package dimensions for each Exadata Cloud at Customer rack:

Description	Millimeters (mm)	Inches (")
Shipping Height	2159 mm	85 inches
Shipping Width	1219 mm	48 inches
Shipping Depth	1575 mm	62 inches

If your loading dock meets the height and ramp requirements for a standard freight carrier truck, then you can use a pallet jack to unload the rack. If the loading dock does not meet the requirements, then you must provide a standard forklift or other means to unload the rack. You can also request that the rack is shipped in a truck with a lift gate.

Use a conditioned space to remove the packaging material to reduce particles before entering the data center. Allow enough space for unpacking it from its shipping cartons.

Use the information in the following table to ensure that there is a clear pathway for moving the Exadata Cloud at Customer rack. Also, the entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.

Access Route Item	With Shipping Pallet	Without Shipping Pallet
Minimum door height	2184 mm (86 inches)	2040 mm (80.32 inches)
Minimum door width	1270 (50 inches)	640 mm (25.19 inches)
Minimum elevator depth	1625.6 mm (64 inches)	1240 mm (48.82 inches)
Maximum incline	6 degrees	6 degrees
Minimum elevator, pallet jack, and floor loading capacity	1134 kg (2500 lbs)	1134 kg (2500 lbs)

Flooring

Oracle recommends that the Exadata Cloud at Customer system is installed on raised flooring. The site floor and the raised flooring must be able to support the total weight of the Exadata Cloud at Customer rack. See [Weight](#).

Electrical Power

Exadata Cloud at Customer can operate effectively over a wide range of voltages and frequencies. However, each rack must have a reliable power source.

Damage may occur if the ranges are exceeded. Electrical disturbances such as the following may damage Exadata Cloud at Customer:

- Fluctuations caused by brownouts
- Wide and rapid variations in input voltage levels or in input power frequency
- Electrical storms
- Faults in the distribution system, such as defective wiring

To protect Exadata Cloud at Customer from such disturbances, you should have a dedicated power distribution system, power-conditioning equipment, and lightning arresters or power cables to protect from electrical storms.

POWER DISTRIBUTION UNIT SPECIFICATIONS

Each rack has two pre-installed power distribution units (PDUs). The PDUs accept different power sources. You must choose the type of PDU that is correct for your data center and the Exadata Cloud at Customer rack.

The following table outlines the minimum PDU rating requirement for each rack type:

Model (X8)	Minimum PDU Rating (kVA)
Base System	15 kVA
Quarter Rack	15 kVA
Half Rack	15 kVA
Full Rack	22 kVA

The following list outlines the available PDUs for Exadata Cloud at Customer depending on your region. Follow each of the links to access detailed specifications for each PDU type:

- Americas, Japan, and Taiwan
 - [Low-Voltage 15 kVA Single-Phase](#)
 - [Low-Voltage 15 kVA Three-Phase](#)

- [Low-Voltage 22 kVA Single-Phase](#)
- [Low-Voltage 24 kVA Three-Phase](#)
- Europe, the Middle East and Africa (EMEA), and Asia Pacific (APAC), except for Japan and Taiwan
 - [High-Voltage 15 kVA Three-Phase](#)
 - [High-Voltage 22 kVA Single-Phase](#)
 - [High-Voltage 24 kVA Three-Phase](#)

FACILITY POWER REQUIREMENTS

To prevent catastrophic failures, design the input power sources to ensure that adequate power is provided to the PDUs.

Use dedicated AC breaker panels for all power circuits that supply power to the PDU. When planning for power distribution requirements, balance the power load between available AC supply branch circuits. In the United States of America and Canada, ensure that the overall system AC input current load does not exceed 80 percent of the branch circuit AC current rating.



Note

Electrical work and installations must comply with applicable local, state, or national electrical codes.

PDU power cords are 4 meters (13.12 feet) long, and 1–1.5 meters (3.3–4.9 feet) of the cord is routed within the rack cabinet. The installation site AC power receptacle must be within 2 meters (6.6 feet) of the rack.

CIRCUIT BREAKER REQUIREMENTS

If computer equipment is subjected to repeated power interruptions and fluctuations, then it is susceptible to a higher rate of component failure.

You are responsible for supplying the circuit breakers. One circuit breaker is required for each power cord. In addition to circuit breakers, provide a stable power source, such as an uninterruptible power supply (UPS) to reduce the possibility of component failures.

Use dedicated AC breaker panels for all power circuits that supply power to the server. Servers require grounded electrical circuits.



Note

Electrical work and installations must comply with applicable local, state, or national electrical codes.

ELECTRICAL GROUNDING GUIDELINES

The cabinets for Oracle Exadata Rack are shipped with grounding-type power cords.

- Always connect the cords to grounded power outlets.
- Check the grounding type, because different grounding methods may be used depending on your location.
- Refer to documentation such as IEC documents for the correct grounding method.
- Ensure that the facility administrator or qualified electrical engineer verifies the grounding method for the building, and performs the grounding work.

Temperature and Humidity

Excessive internal temperatures may result in full or partial shutdown of Exadata Cloud at Customer system components.

**Note**

Studies have shown that temperature increases of 10 degrees Celsius (15 degrees Fahrenheit) above 20 degrees Celsius (70 degrees Fahrenheit) reduce long-term electronics reliability by 50 percent.

The following table lists the temperature, humidity, and altitude requirements for operating and non-operating machines.

Condition	Operating Requirement	Non-operating Requirement	Optimal Requirement
Temperature	5–32 degrees Celsius (59–89.6 degrees Fahrenheit)	-40–70 degrees Celsius (-40–158 degrees Fahrenheit)	21–23 degrees Celsius (70–74 degrees Fahrenheit)
Relative Humidity	10–90 percent relative humidity, non-condensing	Up to 93 percent relative humidity	45–50 percent, non-condensing
Altitude	3048 meters (10000 feet) maximum	12,000 meters (40000 feet) maximum	Maximum ambient temperature is reduced by 1 degree Celsius for every 300 meters of altitude over 900 meters above sea level.

To minimize the chance of downtime because of component failure, set conditions to the optimal temperature and humidity ranges. Maintaining an Exadata Cloud at Customer system for extended periods at or near the operating limits can significantly increase the potential for hardware component failure.

The ambient temperature range of 21–23 degrees Celsius (70–74 degrees Fahrenheit) is optimal for server reliability and operator comfort. Most computer equipment can operate in a wide temperature range, but near 22 degrees Celsius (72 degrees Fahrenheit) is desirable because it is easier to maintain safe humidity levels. Operating in this temperature range provides a safety buffer in case the air conditioning system fails for some time.

The ambient relative humidity range of 45–50 percent is suitable for safe data processing operations. Most computer equipment can operate in a wide range (20–80 percent), but the range of 45–50 percent is recommended for the following reasons:

- The optimal range helps protect computer systems from corrosion problems associated with high humidity levels.
- The optimal range provides the greatest operating time buffer in case the air conditioning system fails for some time.
- The optimal range avoids failures or temporary malfunctions caused by interference from static discharges that may occur when relative humidity is too low. Electrostatic discharge (ESD) is easily generated, and hard to dissipate in areas of low relative humidity, such as below 35 percent. ESD becomes critical when humidity drops below 30 percent.

Ventilation

To allow for proper ventilation, always provide adequate space in front and behind the rack.

Do not obstruct the front or rear of the rack with equipment or objects that might prevent air from flowing through the rack. Each Exadata Cloud at Customer rack draws cool air in through the front of the rack and discharges warm air out the rear of the rack. There is no air flow requirement for the left and right sides because of front-to-back cooling.

Each Exadata Cloud at Customer rack is designed to function while installed in a natural convection air flow. To ensure adequate air flow, allow a minimum clearance of 1219.2 mm (48 inches) at the front of the server, and 914 mm (36 inches) at the rear of the server for ventilation.

Use perforated tiles, approximately 400 CFM/tile, in front of the rack for cold air intake. The tiles can be arranged in any order in front of the rack, as long as cold air from the tiles can flow into the rack. Inadequate cold air flow could result in a higher inlet temperature in the servers because of exhaust air recirculation. The following is the recommended number of floor tiles:

- Four floor tiles for an Exadata Cloud at Customer Full Rack.
- Three floor tiles for an Exadata Cloud at Customer Half Rack.
- One floor tile for an Exadata Cloud at Customer Quarter Rack or Base System.

Network Requirements

Networks and Network Services

NETWORK REQUIREMENTS

Exadata Cloud at Customer utilizes various different networks to provide secure and reliable network connectivity for different application and management functions. The following list outlines the minimum network requirements to install an Exadata Cloud at Customer system:

- Client network
This network connects the Exadata Cloud at Customer database servers to your existing client network and is used for client access to the database servers. Applications access databases on Exadata Cloud at Customer through this network using Single Client Access Name (SCAN) and Oracle RAC Virtual IP (VIP) interfaces.
The client access network uses a pair of network interfaces on each database server, which are connected to the customer network.

The database servers support channel bonding to provide higher bandwidth or availability for client connections to Exadata Cloud at Customer. Oracle recommends channel bonding for the client access network. For the connection to your corporate network, you must provide network switches capable of supporting your chosen bonding mode. For example, if mode 4 (IEEE 802.3ad Link Aggregation) is configured, then you must supply and configure network switches capable of supporting this bonding mode.

- Backup network

This network is similar to the client access network, as it connects the Exadata Cloud at Customer database servers to your existing network. It can be used for access to the database servers for various purposes, including backups and bulk data transfers.

Like the client network, the backup network uses a pair of network interfaces on each database server, which are connected to the customer network.

Channel bonding is supported for the backup network to provide higher bandwidth or availability, and Oracle recommends channel bonding for the backup network. For the connection to your corporate network, you must provide network switches capable of supporting your chosen bonding mode. For example, if mode 4 (IEEE 802.3ad Link Aggregation) is configured, then you must supply and configure network switches capable of supporting this bonding mode.

- Control plane network

This virtual private network (VPN) connects the two control plane servers that are located in the Exadata Cloud at Customer rack to Oracle Cloud Infrastructure. It facilitates customer-initiated operations using the Oracle Cloud Infrastructure Console and APIs. It also facilitates monitoring and administration of the Oracle-managed infrastructure components in Exadata Cloud at Customer.

- Administration network

This network connects Exadata Cloud at Customer servers and switches to the two control plane servers that are located in the Exadata Cloud at Customer rack. It facilitates customer-initiated operations using the Oracle Cloud Infrastructure Console and APIs. It also facilitates monitoring and administration of the Oracle-managed infrastructure components in Exadata Cloud at Customer.

This network is fully contained within the Exadata Cloud at Customer rack, and does not connect to your corporate network. However, the Exadata infrastructure is indirectly connected to your corporate network through the control plane servers. This connection is required to provide Domain Name System (DNS) and Network Time Protocol (NTP) services to the Exadata infrastructure. Therefore, the IP addresses that are allocated to the administration network must not exist elsewhere in your corporate network.

Each database server and Exadata Storage Server has two network interfaces connected to the administration network. One provides management access to the server through one of the embedded Ethernet ports (NET0). The other provides access to the Integrated Lights-Out Management (ILOM) subsystem through a dedicated ILOM Ethernet port. Exadata Cloud at Customer is delivered with the ILOM and NET0 ports connected to the Ethernet switch in the rack. Cabling or configuration changes to these interfaces are not permitted.

- **InfiniBand network**

This network connects the database servers, Exadata Storage Servers, and control plane servers using the InfiniBand switches on the rack. Each server contains two InfiniBand network interfaces (IB0 and IB1) that are connected to separate InfiniBand switches in the rack. Primarily, Oracle Database uses this network for Oracle RAC cluster interconnect traffic and for accessing data on Exadata Storage Servers.

This non-routable network is fully contained within the Exadata Cloud at Customer rack, and does not connect to your corporate network. However, because the control plane servers are connected to the InfiniBand network and to your corporate network, the IP addresses that are allocated to the InfiniBand network must not exist elsewhere in your corporate network.

DATA CENTER NETWORK SERVICES

Exadata Cloud at Customer requires the following data center network services:

- **Domain Name System (DNS)**

As part of the deployment process, you must decide on the host names and IP addresses to be used for various Exadata Cloud at Customer network interfaces. Oracle recommends that you register the host names and IP addresses for the Exadata Cloud

at Customer network interfaces in your corporate DNS. At least one reliable DNS server is required, which must be accessible to the control plane servers and to all of the servers on the client network. Up to three DNS servers can be registered in Exadata Cloud at Customer to ensure coverage in case a server is unavailable.

- Network Time Protocol (NTP) Services

Exadata Cloud at Customer uses NTP to ensure that all system components are synchronized to the same time. At least one reliable NTP server is required, which must be accessible to the control plane servers and to all of the servers on the client network. Up to three NTP servers can be registered in Exadata Cloud at Customer to ensure coverage in case a server is unavailable.

IP Addresses and Subnets

You must allocate a range of IP addresses to the administration network and another range of IP addresses to the InfiniBand network. No overlap is permitted between the address ranges for the administration network and the InfiniBand network, and all IP addresses should be unique within your corporate network. You must also allocate IP addresses from your corporate network to the control plane servers. These network configuration details are specified when you create the Exadata infrastructure.

When you create the Exadata infrastructure, the Console pre-populates default values for the administration network CIDR block and the InfiniBand network CIDR block. You can use the suggested CIDR blocks if there is no overlap with existing IP addresses in your corporate network.

The following table outlines the IP address requirements for each of these networks. The table specifies the maximum and minimum CIDR block prefix length that are allowed for each network. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, you can choose a smaller CIDR block prefix length, within the allowable range, which reserves more IP addresses for the network.

Network Type	IP Address Requirements
Administration network	Maximum CIDR block prefix length: /23 Minimum CIDR block prefix length: /16
InfiniBand network	Maximum CIDR block prefix length: /22 Minimum CIDR block prefix length: /19
Control plane network	2 IP addresses, 1 for each control plane server

To connect to your corporate network, Exadata Cloud at Customer requires several host names and IP addresses for network interfaces on the client network and the backup network. The precise number of IP addresses depends on the Exadata system shape. These network configuration details, including host names and IP addresses, are specified when you create a VM cluster network. All IP addresses must be statically assigned IP addresses, not dynamically assigned (DHCP) addresses. The client network and the backup network require separate subnets.

The following table outlines the IP address requirements for the client and backup networks. The table specifies the maximum and recommended CIDR block prefix length for each network. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller CIDR block prefix length is recommended, which reserves more IP addresses for the network.

Network Type	IP Address Requirements for Base System, Quarter Rack, or Half Rack	IP Address Requirements for Full Rack
Client network	Maximum CIDR block prefix length: /28 Recommended CIDR block prefix length: /27	Maximum CIDR block prefix length: /27 Recommended CIDR block prefix length: /26
Backup network	Maximum CIDR block prefix length: /29 Recommended CIDR block prefix length: /28	Maximum CIDR block prefix length: /28 Recommended CIDR block prefix length: /27

Uplinks

Exadata Cloud at Customer has the following uplink requirements:

- Two uplinks are required to connect the control plane servers to your corporate network and the control plane virtual private network (VPN).
- Typically, four uplinks are required for each compute node to connect to your corporate network. Using this configuration, two uplinks support the client network and the other two uplinks support the backup network.

On Quarter Rack, Half Rack, or Full Rack systems, you can choose to use 10 Gbps RJ45 copper or 10/25 Gbps SFP28 fiber network connections to your corporate network. However, you cannot have a mixture. For example, you cannot use fiber for the client network and copper for the backup network.

On Base System configurations, the options are more limited because of the physical network interfaces that are available on each compute node. On Base Systems, you can choose to use copper or fiber network connections only for the client network, while the backup network uses a fiber connection.

There is also an option to use shared network interfaces for the client network and the backup network, which reduces the uplink requirement to two uplinks for each compute node. Using shared network interfaces also enables you to use copper network connections to support both the client and backup networks on Base System configurations. However, using shared network interfaces is not generally recommended because it compromises the bandwidth and availability of both networks.

Network Cabling

Every Exadata Cloud at Customer rack is shipped with all of the network equipment and cables that are required to interconnect all hardware in the Exadata Cloud at Customer rack. Oracle also supplies small form-factor pluggable (SFP) network interfaces to enable connectivity to your corporate network. However, you are responsible to provide the required cabling to connect the Exadata compute nodes and control plane servers to your corporate network.

Checklists

System Components Checklist

Use this checklist to ensure that the system component considerations are addressed.

The cells in the second column of the following table are intentionally left blank so that the site survey team can fill in the requested information.

CHAPTER 11 Database

System Components Checklist Items	Customer Response or Comment
How many racks will be installed?	
Will additional equipment be attached to or installed in the rack? If so, ensure that the additional equipment falls within Oracle guidelines and there is sufficient power and cooling.	

Data Center Room Checklist

Use this checklist to ensure that the data center room requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Data Center Room Checklist Items	Yes	No	Not Applicable	Comment
Has the rack location been allocated and is it vacant?				
Does the floor layout meet the equipment maintenance access requirements?				
Will the rack be positioned so that the exhaust air of one rack does not enter the air inlet of another rack?				

CHAPTER 11 Database

Data Center Room Checklist Items	Yes	No	Not Applicable	Comment
Have cabinet stabilization measures been considered?				
If the data center has a raised floor: <ul style="list-style-type: none">• Does the raised floor satisfy the weight requirements for the rack?• Is permission required to remove floor tiles for cabling and servicing below the floor?				
Will the rack location require any non-standard cable lengths?				
Is the floor-to-ceiling height a minimum of 2914 mm (114.72 inches)?				
Is the depth of the raised floor a minimum of 46 cm (18 inches)?				

RELATED TOPICS:

- [Space](#)
- [Weight](#)
- [Flooring](#)
- [Ventilation](#)

Data Center Environment Checklist

Use this checklist to ensure that the data center environment requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Data Center Environment Checklist Items	Yes	No	Not Applicable	Comment
Does the computer room air conditioning meet temperature and humidity requirements?				
Does the installation floor layout satisfy the ventilation requirements?				
If the room cooling is from a raised floor: <ul style="list-style-type: none"> • Are the perforated floor tiles each rated at 400 CFM or greater? • Can additional perforated floor tiles be obtained if required for additional cooling? 				
Does the data center air conditioning provide sufficient front-to-back airflow?				

Data Center Environment Checklist Items	Yes	No	Not Applicable	Comment
Is airflow adequate to prevent hot spots?				
Can the data center continuously satisfy the environmental requirements?				

RELATED TOPICS:

- [Temperature and Humidity](#)
- [Ventilation](#)

Access Route Checklist

Use this checklist to ensure that the access route requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

CHAPTER 11 Database

Access Route Checklist Items	Yes	No	Not Applicable	Comment
Has the access route been checked for clearance of the rack, including the minimum width and height requirements for all doors on the route?				
Are there any stairs, ramps, or thresholds that are of concern? If yes, then provide details.				
Are all access route incline angles within the permitted range (under 6 degrees)?				
Are all the surfaces acceptable for rolling the new unpacked and packed equipment?				
If a pallet jack is to be used: <ul style="list-style-type: none"> • Can the pallet jack support the weight of the rack? Are the pallet jack tines compatible with the shipping pallet?				
If there are stairs, is a loading elevator available for the equipment?				

Access Route Checklist Items	Yes	No	Not Applicable	Comment
If an elevator is to be used: <ul style="list-style-type: none"> • Does the elevator car meet the height, width, and depth requirements for carrying the rack? • Do the elevator doors meet the height and width requirements for moving the rack? • Does the elevator meet the weight requirements for transporting the rack? 				
Can the complete access route support the weight of the rack?				
Is the access route onto the raised floor rated for dynamic loading of the rack?				

RELATED TOPICS:

- [Space](#)
- [Weight](#)
- [Flooring](#)
- [Receiving, Unpacking, and Access](#)

Facility Power Checklist

Use this checklist to ensure that the facility power requirements are addressed.

CHAPTER 11 Database

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Facility Power Checklist Items	Yes	No	Not Applicable	Comment
Have the operating voltage and electric current requirements been reviewed?				
What type of power supply will be used? <ul style="list-style-type: none">• Single-phase or 3-phase.• Low-voltage or High-voltage.				
Are enough power outlets provided within 2 meters for each rack?				
Do the power outlets have appropriate socket receptacles for the planned Power Distribution Units (PDUs)?				
Will optional ground cables be attached to the rack?				
Are the electrical circuits suitable in terms of voltage and current-carrying capacities?				
Does the power frequency meet the equipment specifications?				

Facility Power Checklist Items	Yes	No	Not Applicable	Comment
Are power outlets available for the new equipment at the designated location?				
Will system power be delivered from two separate grids?				
Is there a UPS to power the equipment?				
Are the minimum required power sources available to support the power load (kW or kVA) for the new hardware?				

RELATED TOPICS:

- [Electrical Power](#)

Safety Checklist

Use this checklist to ensure that the safety requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

CHAPTER 11 Database

Safety Checklist Items	Yes	No	Not Applicable	Comment
Is there an emergency power shut off?				
Is there a fire protection system in the data center room?				
Is the computer room adequately equipped to extinguish a fire?				
Is antistatic flooring installed?				
Is the area below the raised floor free of obstacles and blockages?				

Logistics Checklist

Use this checklist to ensure that the logistics requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

CHAPTER 11 Database

Logistics Checklist Items	Yes	No	Not Applicable	Comment
Is contact information for the data center personnel available?				
Is there security or access control for the data center?				
Are there any security background checks or security clearances required for Oracle personnel to access the data center? If yes, then provide the process for Oracle to follow.				
How many days in advance must background checks be completed?				
Are there any additional security access issues?				
Is computer room access available for installation personnel?				
Are laptops allowed in the data center?				
Are cell phones allowed in the data center?				

CHAPTER 11 Database

Logistics Checklist Items	Yes	No	Not Applicable	Comment
Are cameras allowed in the data center?				
Does the building have a delivery dock?				
Is there a delivery/ unpacking/ staging area?				
Is inside delivery planned (direct to the final rack location in the data center room)?				
If the delivery is not inside, then is the site prepared for uncrating?				
Is the delivery/ unpacking/ staging area protected from the elements?				
Does the building have adequate receiving space?				
Is the unpacking area air-conditioned to avoid thermal shock for various hardware components?				
Will sufficient moving personnel be available to transport the rack?				

CHAPTER 11 Database

Logistics Checklist Items	Yes	No	Not Applicable	Comment
Is union labor required for any part of the delivery or installation?				
Is the site prepared for uncrating and packaging removal? Package removal should take place outside the data center room.				
Is uncrating of cabinet and packaging removal required?				
Are there any restrictions on delivery truck length, width, or height?				
Is there storage space (cabinet) for the ride along spares? If not, does the customer allow cardboard boxes and other packing material in the computer room, since the spares are packed in cardboard boxes?				
Is there a time constraint on dock access? If yes, provide time constraints.				

Logistics Checklist Items	Yes	No	Not Applicable	Comment
Is a tail or side lift required on the delivery carrier to unload the equipment at the delivery dock?				
<p>Will any special equipment be required to place the rack in the data center room?</p> <p>For example:</p> <ul style="list-style-type: none"> • Stair walkers • Lifters • Ramps • Steel plates • Floor covers 				
Does the delivery carrier require any special equipment, such as non-floor damaging rollers, transport dollies, pallet jacks, or fork lifts?				

RELATED TOPICS:

- [Space](#)
- [Receiving, Unpacking, and Access](#)

Network Configuration Checklist

Use this checklist to ensure that the network configuration requirements are addressed.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Network Configuration Checklist Items	Yes	No	Not Applicable	Comment
Will the required network cables be laid from the network equipment to the location where the Oracle Exadata Rack will be installed?				
Will the network cables that will connect to the Oracle Exadata Rack be labeled?				
Will the 10 GbE or 25 GbE interfaces be used for the client access network? If so, has the customer ordered the appropriate cables to their switch?				
Will the Cisco Ethernet switch have IP routing disabled (recommended)?				

RELATED TOPICS:

- [Network Requirements](#)

Reracking Checklist

Use this checklist to ensure that the reracking requirements are addressed.



Note

- Reracking requires prior approval. Check that reracking has been approved.
- Customer must purchase the Oracle Reracking service.
- Oracle does not provide technical support for customer-supplied equipment.

The cells in the second and subsequent columns of the following table are intentionally left blank so that the site survey team can fill in the requested information.

Reracking Checklist Items	Yes	No	Not Applicable	Comment
Has the customer purchased the Oracle Reracking Service?				
Is there a cart capable of carrying the weight of the servers to move the components and associated cabling from the supplied rack to the customer supplied rack?				
Is the target rack empty?				
Attach pictures of the target rack (inside and outside).				

Reracking Checklist Items	Yes	No	Not Applicable	Comment
<p>Does the target rack meet the following requirements?</p> <ul style="list-style-type: none"> • Height: 42 RU Width: 600 mm (23.62 inches) Depth: 1112 mm (43.78 inches) without front and rear doors <p>If the rack is less than 42 RU tall, then the rack must be at least 30 RU tall and the customer must provide compatible PDUs to install in the target rack.</p>				
<p>Is the distance between the front and rear mounting planes between the minimum of 610 mm and the maximum 915 mm (24–36 inches)?</p>				
<p>Is the clearance depth in the front of the front mounting plane (distance to the front cabinet door) at least 25.4 mm (1 inch)?</p>				
<p>Does the target rack meet the following minimum load capacity?</p> <ul style="list-style-type: none"> • 19 kg (41.89 lb) per RU • 785 kg (1730.63 lb) total 				

Reracking Checklist Items	Yes	No	Not Applicable	Comment
<p>Is the rack a four-post rack (mounting at both front and rear)?</p> <div data-bbox="290 594 660 978" style="border: 1px solid #0070C0; background-color: #D9E1F2; padding: 10px;">  <p style="text-align: center;">N o t e</p> <p>Two-post racks are not compatible.</p> </div>				
<p>Does the target rack's horizontal opening and unit vertical pitch conform to ANSI/EIA 310-D-1992 or IEC 60297 standards?</p>				

CHAPTER 11 Database

Reracking Checklist Items	Yes	No	Not Applicable	Comment
Does the target rack have RETMA rail support?				

Reracking Checklist Items	Yes	No	Not Applicable	Comment
<div style="background-color: #e1f5fe; padding: 10px; border: 1px solid #00796b;">  <p style="text-align: center;">N o t e</p> <p>Oracle Exadata rack requires 19 inches (483 mm) for RETMA rail spacing width. The minimum rack width of 600 mm (23.63 inches) is recommended to accommodate the PDU and cable harnesses on the side. If the rack is less than 600 mm wide, then it must have additional depth to accommodate mounting behind the server cable management arms.</p> </div>				

CHAPTER 11 Database

Reracking Checklist Items	Yes	No	Not Applicable	Comment
Does the target rack support Oracle cable management arms?				
Does the target rack support installation of Oracle vented and solid filler panels?				
Can the target rack provide tie-downs along the left rear side of the rack (as viewed from the front of the rack) to support the InfiniBand cables?				
Can the target rack provide tie-downs for the Ethernet wiring harness?				
Is there sufficient space for the cable harnesses and the PDUs in the target rack?				
Can a label with the Oracle Exadata Rack serial number be printed and attached to the target rack?				
Does the target rack support installation of standard Oracle PDUs? If not, then complete the following checklist items:				

CHAPTER 11 Database

Reracking Checklist Items	Yes	No	Not Applicable	Comment
<ul style="list-style-type: none">• Can the customer provide an equivalent pair of PDUs?				
<ul style="list-style-type: none">• Can the customer provide two PDUs, each with a capacity of 10 kVA?				
<ul style="list-style-type: none">• Can the customer provide at least 17 x IOA C13 plugs per PDU?				
<ul style="list-style-type: none">• Can the customer provide a single PDU and its circuits to support the Oracle Exadata Rack power requirements in case one PDU fails?				
<ul style="list-style-type: none">• Can the customer ensure that power loads are evenly distributed across all circuits of a single PDU?				
<ul style="list-style-type: none">• Can the customer provide appropriate power drops for the PDUs?				

Provisioning Exadata Cloud at Customer Systems

This topic explains how to provision an Exadata Cloud at Customer system.

Provisioning an Exadata Cloud at Customer system is a collaborative process that involves you and Oracle. The process is outlined as follows:

- You create the Exadata Cloud at Customer infrastructure.
- You generate a file containing the infrastructure configuration details and provide it to Oracle.
- The Exadata Cloud at Customer system is physically installed in your data center.
- Oracle uses the infrastructure configuration file to perform initial system configuration. At the end of this task, Oracle supplies you with an activation file.
- You activate the Exadata Cloud at Customer infrastructure by using the supplied activation file.

Following the provisioning process, the Exadata Cloud at Customer system is ready for you to use. You can then create a virtual machine (VM) cluster and later create some databases.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Prerequisites

- Before you can provision Exadata Cloud at Customer infrastructure, your Oracle Cloud Infrastructure tenancy must be enabled to use Exadata Cloud at Customer. See [Welcome to Oracle Cloud Infrastructure](#). Contact Oracle for further details.
- In preparation for the provisioning process, ensure that your corporate data center and network infrastructure meet the requirements described in [Preparing for Exadata Cloud at Customer](#).

Using the Console

To create Exadata Cloud at Customer infrastructure

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** that you want to associate with the Exadata infrastructure.
The region that is associated with your Exadata infrastructure cannot be changed after the Exadata infrastructure is created. Therefore, ensure that you select the most appropriate region for your Exadata infrastructure. Consider the following factors:
 - Consider any business policies or regulations that preclude the use of a particular region. For example, you may be required to maintain all operations within national boundaries.
 - Consider the physical proximity of the region to your data center. Needless extra physical separation adds unnecessary latency to network communications between Oracle Cloud Infrastructure and your corporate data center.
3. Click **Exadata Infrastructure**.
4. Click **Create Exadata Infrastructure**.

5. Provide the requested information in the **Create Exadata Infrastructure** page:
 - **Oracle Cloud Infrastructure region:** The region that is associated with your Exadata infrastructure cannot be changed after the Exadata infrastructure is created. Therefore, check the displayed region to ensure that you are using the most appropriate region for your Exadata infrastructure.
See step 2 (earlier in this procedure) for further considerations. To switch regions now, use the Region menu at the top of the Console.
 - **Choose a compartment:** From the list of available compartments, choose the compartment that you want to contain the Exadata infrastructure.
See also [Understanding Compartments](#).
 - **Provide the display name:** The display name is a user-friendly name that you can use to identify the Exadata infrastructure. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the Exadata infrastructure.
 - **Select the Exadata system model:** From the list, choose the model of the Exadata hardware that is being used.
The Exadata system model and the Exadata system shape combine to define the amount of CPU, memory, and storage resources that are available in the Exadata infrastructure. For more details, see [System Configuration](#).
 - **Select an Exadata system shape:** Together with the Exadata system model, the Exadata system shape defines the amount of CPU, memory, and storage resources that are available in the Exadata infrastructure.

Exadata system shapes

- **Base System:** includes two compute nodes and three Exadata Storage Servers. A Base System is an entry-level configuration. Compared to other configurations, a Base System contains Exadata Storage Servers with significantly less storage capacity and compute nodes with significantly less memory and processing power.

- **Quarter Rack:** includes two compute nodes and three Exadata Storage Servers.
- **Half Rack:** includes four compute nodes and six Exadata Storage Servers.
- **Full Rack:** includes eight compute nodes and 12 Exadata Storage Servers.

For more details, see [System Configuration](#).

CONFIGURE THE CLOUD CONTROL PLANE NETWORK

Each Exadata Cloud at Customer system contains two control plane servers, which enable connectivity to Oracle Cloud Infrastructure. The control plane servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the required network parameters:

- **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first control plane server to your corporate network using the control plane network.
- **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP is address for the network interface that connects the second control plane server to your corporate network using the control plane network.
- **Netmask:** Specify the IP netmask for the control plane network.
- **Gateway:** Specify the IP address of the control plane network gateway.
- **HTTPS Proxy:** Specify your corporate HTTPS proxy. The expected format is `http://server.domain:port`. For example, `http://proxy.example.com:80`.

CONFIGURE THE EXADATA SYSTEM NETWORKS

Each Exadata Cloud at Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

- **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units. You can accept the suggested default, or specify a custom value.

The maximum CIDR block prefix length is $/23$, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller CIDR block prefix length is recommended, which reserves more IP addresses for the network. The minimum CIDR block prefix length is $/16$.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the InfiniBand network CIDR block.

- **InfiniBand Network CIDR Block:** Specifies the IP address range for the Exadata InfiniBand network using CIDR notation. The Exadata InfiniBand network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components. You can accept the suggested default, or specify a custom value.

The maximum CIDR block prefix length is $/22$, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller CIDR block prefix length is recommended, which reserves more IP addresses for the network. The minimum CIDR block prefix length is $/19$.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the administration network CIDR block.

CONFIGURE DNS AND NTP SERVICES

Each Exadata Cloud at Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.
- **Select the time zone:** Choose the geographic region and time zone for the Exadata infrastructure.

SHOW ADVANCED OPTIONS

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Create Exadata Infrastructure**.
If all of your inputs are valid, the Infrastructure Details page appears. The page outlines the next steps in the provisioning process. Initially after creation, the state of the Exadata infrastructure is **Requires-Activation**.

To edit Exadata Cloud at Customer infrastructure networking

You can only edit Exadata Cloud at Customer infrastructure networking if the current state of the Exadata infrastructure is **Requires Activation**. Also, ensure that you do not edit the Exadata infrastructure after you download the configuration file and provide it to Oracle.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that you want to edit.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

The Infrastructure Details page displays information about the selected Exadata infrastructure.

5. Click **Edit Infrastructure Networking**.
6. Use the Edit Infrastructure Networking dialog to edit the Exadata infrastructure networking:

CONFIGURE THE CLOUD CONTROL PLANE NETWORK

Each Exadata Cloud at Customer system contains two control plane servers, which enable connectivity to Oracle Cloud Infrastructure. The control plane servers are connected to the control plane network, which is a subnet on your corporate network. The following settings define the required network parameters:

- **Control Plane Server 1 IP Address:** Provide the IP address for the first control plane server. This IP address is for the network interface that connects the first control plane server to your corporate network using the control plane network.
- **Control Plane Server 2 IP Address:** Provide the IP address for the second control plane server. This IP address is for the network interface that connects the second control plane server to your corporate network using the control plane network.
- **Netmask:** Specify the IP netmask for the control plane network.
- **Gateway:** Specify the IP address of the control plane network gateway.
- **HTTPS Proxy:** Specify your corporate HTTPS proxy. The expected format is `http://server.domain:port`. For example, `http://proxy.example.com:80`.

CONFIGURE THE EXADATA SYSTEM NETWORKS

Each Exadata Cloud at Customer system contains two system networks, which are not connected to your corporate network. The following settings define IP address allocations for these networks:

- **Administration Network CIDR Block:** Specifies the IP address range for the administration network using CIDR notation. The administration network provides

connectivity that enables Oracle to administer the Exadata system components, such as the Exadata compute servers, storage servers, network switches, and power distribution units.

The maximum CIDR block prefix length is $/23$, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller CIDR block prefix length is recommended, which reserves more IP addresses for the network. The minimum CIDR block prefix length is $/16$.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the InfiniBand network CIDR block.

- **InfiniBand Network CIDR Block:** Specifies the IP address range for the Exadata InfiniBand network using CIDR notation. The Exadata InfiniBand network provides the high-speed low-latency interconnect used by Exadata software for internal communications between various system components.

The maximum CIDR block prefix length is $/22$, which defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller CIDR block prefix length is recommended, which reserves more IP addresses for the network. The minimum CIDR block prefix length is $/19$.

Ensure that the IP address range does not conflict with other hosts your corporate network, and does not overlap with the administration network CIDR block.

CONFIGURE DNS AND NTP SERVICES

Each Exadata Cloud at Customer system requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services to the Exadata infrastructure:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the control plane network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the control plane network. You may specify up to three NTP servers.
- **Select the time zone:** Choose the geographic region and time zone for the

Exadata infrastructure.

7. Click **Save Changes**.

To download a file containing the configuration details for Exadata Cloud at Customer infrastructure

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure for which you want to download a file containing the infrastructure configuration details.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Download Configuration**.

Your browser downloads a file containing the infrastructure configuration details.

When you provide the generated infrastructure configuration file to Oracle, ensure that it has not been altered in any way. Also, ensure that you do not edit the Exadata infrastructure after you download the configuration file and provide it to Oracle.

To activate Exadata Cloud at Customer infrastructure

To activate Exadata Cloud at Customer infrastructure, ensure that you that have the activation file. This file is supplied to you by Oracle after installation and initial configuration of your Exadata Cloud at Customer system. You can only activate the Exadata infrastructure if its current state is **Requires Activation**.

1. Download the activation file.
2. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.

3. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that you want to activate.
4. Click **Exadata Infrastructure**.
5. Click the name of the Exadata infrastructure that you want to activate.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
6. Click **Activate**.
The Activate button is only available if the Exadata infrastructure requires activation. You cannot activate Exadata infrastructure multiple times.
7. Use the Activate dialog to upload the activation file and then click **Activate Now**.
After activation, the state of the Exadata infrastructure changes to **Active**.

To check the status of Exadata Cloud at Customer infrastructure

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you are interested in.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Check the icon on the Infrastructure Details page. The color of the icon and the text below it indicates the status of the Exadata infrastructure.
 - **Creating**: Yellow icon. The Exadata infrastructure definition is being created in the control plane.
 - **Requires Activation**: Yellow icon. The Exadata infrastructure is defined in the control plane but it must be provisioned and activated before it can be used.
 - **Active**: Green icon. The Exadata infrastructure is successfully provisioned and activated.

- **Deleting:** Gray icon. The Exadata infrastructure is being deleted by using the Console or API.
- **Deleted:** Gray icon. The Exadata infrastructure is deleted and is no longer available. This state is transitory and is displayed for a short time, after which the Exadata infrastructure is no longer displayed.
- **Activation Failed:** Red icon. An error condition currently prevents the activation of the Exadata infrastructure. Typically, this state is auto-correcting and does not require user intervention.

To move Exadata Cloud at Customer infrastructure to another compartment

You can change the compartment that contains your Exadata Cloud at Customer infrastructure by moving it.

When you move Exadata infrastructure, the compartment change is also applied to the associated VM cluster networks. However, the compartment change does not affect any other associated resources, such as the VM clusters, which remain in their current compartment.

To move Exadata Cloud at Customer infrastructure:

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that you want to move.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to move.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the Exadata infrastructure and click **Move Resource**.

To delete Exadata Cloud at Customer infrastructure

Deleting Exadata Cloud at Customer infrastructure removes it from the Cloud Control Plane.

If you are deleting Exadata infrastructure before activation, then if required you can create replacement Exadata infrastructure without any input from Oracle.

In you are deleting active Exadata infrastructure, then to create replacement Exadata infrastructure you must repeat the full provisioning process, including the tasks that Oracle performs.

Before you can delete active Exadata infrastructure, you must:

- Terminate all of the resources that it contains, including the databases, VM cluster, and VM cluster network.
- Lodge a service request (SR) with Oracle indicating your intention to delete the Exadata infrastructure. In response to the SR, Oracle flags the Exadata infrastructure as ready for deletion, which enables you to delete the Exadata infrastructure by using the following process.

To delete Exadata Cloud at Customer infrastructure:

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that you want to delete.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to delete.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Delete**.
6. In the resulting dialog, enter the Exadata infrastructure name and click **Delete Exadata Infrastructure** to confirm the action.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Exadata Cloud at Customer infrastructure:

- [ActivateExadataInfrastructure](#)
- [CreateExadataInfrastructure](#)
- [DeleteExadataInfrastructure](#)
- [DownloadExadataInfrastructureConfigFile](#)
- [GenerateRecommendedVmClusterNetwork](#)
- [GetExadataInfrastructure](#)
- [ListExadataInfrastructure](#)
- [UpdateExadataInfrastructure](#)

For the complete list of APIs, see [Database Service API](#).

Managing VM Clusters on Exadata Cloud at Customer

This topic explains how to manage virtual machine (VM) clusters on Exadata Cloud at Customer.

Before you can create any databases on your Exadata Cloud at Customer infrastructure, you must create a VM cluster network and associate it with a VM cluster. Each Exadata Cloud at Customer infrastructure can support one VM cluster network and associated VM cluster.

The VM cluster network specifies network resources, such as IP addresses and host names, that reside in your corporate data center and are allocated to Exadata Cloud at Customer. The VM cluster network includes definitions for the Exadata client network and the Exadata backup network. The client network and backup network contain the network interfaces that you use to connect to the VM cluster compute nodes, and ultimately the databases that reside on those compute nodes.

The VM cluster provides a link between your Exadata Cloud at Customer infrastructure and databases. The VM cluster contains an installation of Oracle Clusterware, which supports the databases in the cluster. In the VM cluster definition, you also specify the number of enabled CPU cores, which determines the amount of CPU resources that are available to your databases.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Prerequisites

The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the VM cluster compute nodes via SSH. The following shows a sample public key, which is abbreviated for readability.

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA...lo/gKMLVM2xzclxJr/Hc26biw3TXWGEakrK10Q== rsa-key-20160304
```

For more information, see [Managing Key Pairs on Linux Instances](#).

Using the Console

To create a VM cluster network

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure for which you want to create a VM cluster network.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure for which you want to create a VM cluster network.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click **Create VM Cluster Network**.
6. Provide the requested information in the **Data Center Network Details** page:
 - **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster network. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster network.

PROVIDE CLIENT NETWORK DETAILS

The client network is the primary channel for application connectivity to Exadata Cloud at Customer resources. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network.
- **CIDR Block:** Using CIDR notation, provide the IP address range for the client network.

The following table specifies the maximum and recommended CIDR block prefix lengths for each Exadata system shape. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller

CIDR block prefix length is recommended, which reserves more IP addresses for the network.

Exadata System Shape	Base System, Quarter Rack, or Half Rack	Full Rack
Maximum CIDR block prefix length	/28	/27
Recommended CIDR block prefix length	/27	/26

- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the client network.
- **Domain Name:** Specify the domain name for the client network.

PROVIDE BACKUP NETWORK DETAILS

The backup network is the secondary channel for connectivity to Exadata Cloud at Customer resources. It is typically used to segregate application connections on the client network from other network traffic. The following settings define the required network parameters:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network.
- **CIDR Block:** Using CIDR notation, provide the IP address range for the backup network.

The following table specifies the maximum and recommended CIDR block prefix lengths for each Exadata system shape. The maximum CIDR block prefix length defines the smallest block of IP addresses that are required for the network. To allow for possible future expansion within Exadata Cloud at Customer, a smaller

CIDR block prefix length is recommended, which reserves more IP addresses for the network.

Exadata System Shape	Base System, Quarter Rack, or Half Rack	Full Rack
Maximum CIDR block prefix length	/29	/28
Recommended CIDR block prefix length	/28	/27

- **Netmask:** Specify the IP netmask for the backup network.
- **Gateway:** Specify the IP address of the backup network gateway.
- **Hostname Prefix:** Specify the prefix that is used to generate the hostnames in the backup network.
- **Domain Name:** Specify the domain name for the backup network.

PROVIDE DNS AND NTP SERVER DETAILS

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. The following settings specify the servers that provide these services:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
- **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.

SHOW ADVANCED OPTIONS

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To

apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Review Configuration**.

The Review Configuration page displays detailed information about the VM cluster network, including the hostname and IP address allocations. These allocations are initially system-generated and are based on your inputs to the Specify Parameters page.

8. Optionally, you can adjust the system-generated network definitions on the Review Configuration page.

a. Click **Edit IP Allocation**.

b. Use the **Edit** dialog to adjust the system-generated network definitions to meet your requirements.

c. Click **Save Changes**.

9. Click **Create VM Cluster Network**.

The VM Cluster Network Details page is now displayed. Initially after creation, the state of the VM cluster network is **Requires Validation**.

To edit a VM cluster network

You can only edit a VM cluster network that is not associated with a VM cluster.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to edit.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.

The Infrastructure Details page displays information about the selected Exadata infrastructure.

5. Click the name of the VM cluster network that you want to edit.
The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Edit VM Cluster Network**.
7. Use the **Edit** dialog to edit the VM cluster network attributes:

CLIENT NETWORK

The client network is the primary channel for application connectivity to Exadata Cloud at Customer resources. You can edit the following client network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the client network.
- **Netmask:** Specify the IP netmask for the client network.
- **Gateway:** Specify the IP address of the client network gateway.
- **Hostname:** Specify the hostname for each address in the client network.
- **IP Address:** Specify the IP address for each address in the client network.

BACKUP NETWORK

The backup network is the secondary channel for connectivity to Exadata Cloud at Customer resources. It is typically used to segregate application connections on the client network from other network traffic. You can edit the following backup network settings:

- **VLAN ID:** Provide a virtual LAN identifier (VLAN ID) for the backup network.
- **Hostname:** Specify the hostname for each address in the backup network.
- **IP Address:** Specify the IP address for each address in the backup network.

CONFIGURE DNS & NTP SERVERS

The VM cluster network requires access to Domain Names System (DNS) and Network Time Protocol (NTP) services. You can edit the following settings:

- **DNS Servers:** Provide the IP address of a DNS server that is accessible using the client network. You may specify up to three DNS servers.
 - **NTP Servers:** Provide the IP address of an NTP server that is accessible using the client network. You may specify up to three NTP servers.
8. Click **Save Changes**.
After editing, the state of the VM cluster network is **Requires Validation**.

To download a file containing the VM cluster network configuration details

You can download a file containing the VM cluster network configuration details. Your network administrator can use the details to configure your corporate DNS and other network devices to work along with Exadata Cloud at Customer.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network for which you want to download a file containing the VM cluster network configuration details.
The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Download Network Configuration**.
Your browser downloads a file containing the VM cluster network configuration details.

To validate a VM cluster network

You can only validate a VM cluster network if its current state is **Requires Validation**, and if

the underlying Exadata infrastructure is activated.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to validate.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to validate.
The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Validate VM Cluster Network**.
Validation performs a series of automated checks on the VM cluster network. The Validate VM Cluster Network button is only available if the VM cluster network requires validation.
7. In the resulting dialog, click **Validate** to confirm the action.
After successful validation, the state of the VM cluster network changes to **Validated** and the VM cluster network is ready to use. If validation fails for any reason, examine the error message and resolve the issue before repeating validation.

To terminate a VM cluster network

Terminating a VM cluster network removes it from the Cloud Control Plane. Before you can terminate a VM cluster network, you must first terminate the associated VM cluster, if one exists, and all the databases it contains.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the Exadata infrastructure that is associated with the VM cluster network that you want to terminate.

3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that is associated with the VM cluster network that you are interested in.
The Infrastructure Details page displays information about the selected Exadata infrastructure.
5. Click the name of the VM cluster network that you want to terminate.
The VM Cluster Network Details page displays information about the selected VM cluster network.
6. Click **Terminate**.
7. In the resulting dialog, enter the name of the VM cluster network and click **Terminate VM Cluster Network** to confirm the action.

To create a VM cluster

To create a VM cluster, ensure that you that have:

- Active Exadata infrastructure available to host the VM cluster.
 - A validated VM cluster network available for the VM cluster to use.
1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
 2. Choose the **Region** that contains your Exadata infrastructure.
 3. Click **VM Clusters**.
 4. Click **Create VM Cluster**.
 5. Provide the requested information in the **Create VM Cluster** page:
 - **Choose a compartment:** From the list of available compartments, choose the compartment that you want to contain the VM cluster.
 - **Provide the display name:** The display name is a user-friendly name that you can use to identify the VM cluster. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the VM cluster.

- **Select Exadata Cloud at Customer Infrastructure:** From the list, choose the Exadata infrastructure to host the VM cluster. You are not able to create a VM cluster without available and active Exadata infrastructure.
- **Choose the Oracle Grid Infrastructure version:** From the list, choose the version of Oracle Grid Infrastructure to install on the VM cluster.
The Grid Infrastructure version determines the Oracle Database versions that can be supported on the VM cluster. You cannot run an Oracle Database version that is higher than the version of the Oracle Grid Infrastructure software.
- **Specify the OCPU count:** Specify the total number of CPU cores that are allocated to the VM cluster.
If you specify a value of zero, then the VM cluster compute nodes are all shut down at the end of the cluster creation process. In this case, you can later start the compute nodes by scaling the CPU resources. See [To scale the CPU resources on a VM cluster](#).
Otherwise, this value must be a multiple of the number of compute nodes so that every compute node has the same number of CPU cores enabled.
- **Add SSH Key:** Specify the public key portion of an SSH key pair that you want to use to access the VM cluster compute nodes. You can upload a file containing the key, or paste the SSH key string.
To provide multiple keys, upload multiple key files or paste each key into a separate field. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
- **Select a VM Cluster Network:** From the list, choose a VM cluster network definition to use for the VM cluster. You must have an available and validated VM cluster network before you can create a VM cluster.

CONFIGURE THE EXADATA STORAGE

The following settings define how the Exadata storage is configured for use with the VM cluster. These settings cannot be changed after creating the VM cluster. See also [Storage Configuration](#).

- **Allocate Storage for Exadata Snapshots:** Check this option to create a sparse disk group, which is required to support Exadata snapshot functionality. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily.
- **Allocate Storage for Local Backups:** Check this option to configure the Exadata storage to enable local database backups. If you select this option, more space is allocated to the RECO disk group to accommodate the backups. If you do not select this option, you cannot use local Exadata storage as a backup destination for any databases in the VM cluster.
- **Choose a license type:** Either:
 - **Bring Your Own License (BYOL):** Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.
 - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Cloud at Customer.

SHOW ADVANCED OPTIONS

- **Select the time zone:** Choose the geographic region and time zone for the VM cluster. The default time zone is inherited from the time zone setting of the underlying Exadata infrastructure.
 - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Create VM Cluster**.
- The VM Cluster Details page is now displayed. While the creation process is running, the state of the VM cluster is **Pending**. When the VM cluster creation process completes, the state of the VM cluster changes to **Available**.

To scale the CPU resources on a VM cluster

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to scale the CPU resources.
The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Scale Up/Down**.
6. In the dialog box, adjust the **OCPU Count** and then click **Save Changes**.
The OCPU Count value must be a multiple of the number of compute nodes so that every compute node has the same number of CPU cores enabled.
If you set the OCPU Count to zero, then the VM cluster compute nodes are all shut down.
If you change from a zero setting, then the VM cluster compute nodes are all started.
Otherwise, modifying the number of enabled CPU cores is an online operation, and compute nodes are not rebooted because of this operation. See also [System Configuration](#).



Note

If you have explicitly set the `CPU_COUNT` database initialization parameter, that setting is not affected by modifying the number of CPU cores that are allocated to the VM cluster. Therefore, if you have enabled the Oracle Database instance caging feature, the database instance does not use extra CPU cores until you alter the `CPU_COUNT` setting. If `CPU_COUNT` is set to 0 (the default setting), then Oracle Database continuously monitors the number of CPUs reported by the operating system and uses the current count.

To stop, start, or reboot a VM cluster compute node

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the compute node that you want to stop, start, or reboot.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the compute node that you want to stop, start, or reboot.
The VM Cluster Details page displays information about the selected VM cluster.
5. In the Resources list, click **Nodes**.
The list of compute nodes displays.
6. In the list of nodes, click the Actions icon (three dots) for a node and then click one of the following actions:
 - **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.

- **Stop:** Shuts down the node. After the node is stopped, the **Start** action is enabled.
- **Reboot:** Shuts down the node, and then restarts it.

To check the status of a VM cluster compute node

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the compute node that you are interested in.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the compute node that you are interested in.
The VM Cluster Details page displays information about the selected VM cluster.
5. In the Resources list, click **Nodes**.
The list of compute nodes displays. For each compute node in the VM cluster, the name, state, and client IP address are displayed.
6. In the node list, find the compute node that you are interested in and check its state.
The color of the icon and the associated text it indicates its status.
 - **Available:** Green icon. The node is operational.
 - **Starting:** Yellow icon. The node is starting because of a start or reboot action in the Console or API.
 - **Stopping:** Yellow icon. The node is stopping because of a stop or reboot action in the Console or API.
 - **Stopped:** Yellow icon. The node is stopped.
 - **Failed:** Red icon. An error condition prevents the continued operation of the compute node.

To update the license type on a VM cluster

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.

2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to update the license type.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to update the license type. The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Update License Type**.
6. In the dialog box, choose one of the following license types and then click **Save Changes**.
 - **Bring Your Own License (BYOL):** Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.
 - **License Included:** Select this option to subscribe to Oracle Database software licenses as part of Exadata Cloud at Customer.

Updating the license type does not change the functionality or interrupt the operation of the VM cluster.

To move a VM cluster to another compartment

You can change the compartment that contains your VM cluster by moving it.

When you move a VM cluster, the compartment change is also applied to the compute nodes and databases that are associated with the VM cluster. However, the compartment change does not affect any other associated resources, such as the Exadata infrastructure, which remains in its current compartment.

To move a VM cluster:

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to move.
3. Click **VM Clusters**.

4. Click the name of the VM cluster that you want to move.
The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the VM cluster and click **Move Resource**.

To terminate a VM cluster

Terminating a VM cluster removes it from the Cloud Control Plane. In the process, the compute node VMs and their contents are destroyed.

Before you can terminate a VM cluster, you must first terminate the databases it contains.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to terminate.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that you want to terminate.
The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Terminate**.
6. In the resulting dialog, enter the name of the VM cluster and click **Terminate VM Cluster** to confirm the action.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Exadata Cloud at Customer VM cluster networks and VM clusters:

VM cluster networks:

- [GenerateRecommendedVmClusterNetwork](#)
- [CreateVmClusterNetwork](#)
- [DeleteVmClusterNetwork](#)
- [GetVmClusterNetwork](#)
- [ListVmClusterNetwork](#)
- [UpdateVmClusterNetwork](#)
- [ValidateVmClusterNetwork](#)

VM clusters:

- [CreateVmCluster](#)
- [DeleteVmCluster](#)
- [GetVmCluster](#)
- [ListVmCluster](#)
- [UpdateVmCluster](#)

For the complete list of APIs, see [Database Service API](#).

Managing Backup Destinations for Exadata Cloud at Customer

This topic explains how to manage backup destinations for Exadata Cloud at Customer on Oracle Zero Data Loss Recovery Appliance or Network File System (NFS).

Exadata Cloud at Customer provides a backup facility, which can be individually configured on each database. See [Managing Databases on Exadata Cloud at Customer](#) and [Managing Database Backup and Recovery on Exadata Cloud at Customer](#).

However, if you want to store backups on a Recovery Appliance or NFS location that you manage, you must first create a backup destination. Each backup destination defines the properties that are required to connect to the Recovery Appliance or NFS location, and each backup destination must be accessible in your data center from the VM cluster nodes.

The Exadata Cloud at Customer backup facility can also store backups on Oracle Cloud Infrastructure object storage or local Exadata storage on your Exadata Cloud at Customer

system. However, you do not need to create a backup destination for any of these other locations. Instead, applicable options for backup to cloud object storage or local Exadata storage are available directly when you create a database.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Prerequisites

- For a Recovery Appliance backup destination:
 - The Recovery Appliance must be configured with a virtual private catalog (VPC) user, which is used for taking the backups.
 - The Recovery Appliance must be configured with the unique database name of the database being backed up, and a mapping to the VPC user.

- The Recovery Appliance must be accessible from the Exadata Cloud at Customer system using the Oracle Net Services connection string, which is provided by the Recovery Appliance administrator.
- For an NFS backup destination:
 - You must mount the NFS server location to a local mount point directory on each node in the VM cluster.
 - The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
 - You must ensure that the NFS mount is maintained continuously on all of the VM cluster nodes.
 - The NFS-mounted file system must be readable and writable by the `oracle` OS user on all of the VM cluster nodes.

Using the Console

To create a backup destination

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **Backup Destinations**.
4. Click **Create Backup Destination**.
5. Provide the requested information in the **Create Backup Destination** page:
 - **Choose a compartment:** From the list of available compartments, choose the compartment that you want to contain the backup destination.
 - **Name your backup destination:** Specify a user-friendly name that you can use to identify the backup destination. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the backup destination.
 - **Choose the backup destination type:** Select **Recovery Appliance** or **Network Storage (NFS)**.

- If you select **Recovery Appliance**, then you must also specify the following:
 - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the Recovery Appliance. This information is typically provided by the Recovery Appliance administrator.
 - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Recovery Appliance. You can specify multiple VPC user names in case you want to use the Recovery Appliance as a backup destination for multiple databases. This information is typically provided by the Recovery Appliance administrator.
- If you select **Network Storage (NFS)**, then you must also specify the following:
 - **Provide the local NFS mount point path:** Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.

SHOW ADVANCED OPTIONS

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Create Backup Destination**.
- The Backup Destination Details page displays the newly created backup destination.

To edit a backup destination

You can only edit a backup destination if it is not currently associated with database.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to edit.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to edit.
The Backup Destination Details page displays information about the selected backup destination.
5. Click **Edit**.
6. Use the **Edit Backup Destination** dialog to edit the backup destination attributes:
 - If you are editing a Recovery Appliance backup destination:
 - **Provide the Recovery Appliance connection string:** Specify the Oracle Net Services connection string that connects to the Recovery Appliance. This information is typically provided by the Recovery Appliance administrator.
 - **Provide the Virtual Private Catalog (VPC) Users:** Provide a VPC user name for connecting to the Recovery Appliance. You can specify multiple VPC user names in case you want to use the Recovery Appliance as a backup destination for multiple databases. This information is typically provided by the Recovery Appliance administrator.
 - If you are editing an NFS backup destination:
 - **Provide the local NFS mount point path:** Specify the local directory path on each VM cluster node where the NFS server location is mounted. The local directory path and the NFS server location must each be the same across all of the VM cluster nodes.
7. Click **Save Changes**.

To move a backup destination to another compartment

You can change the compartment that contains your backup destination by moving it.

When you move a backup destination, the compartment change does not affect other associated resources. These other resources, such as the associated databases, remain in their current compartment.

To move a backup destination:

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to move.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to move.
The Backup Destination Details page displays information about the selected backup destination.
5. Click **Move Resource**.
6. In the resulting dialog, choose the new compartment for the backup destination and click **Move Resource**.

To terminate a backup destination

Terminating a backup destination removes it from the Cloud Control Plane. Before you can terminate a backup destination, you must ensure that it is not associated with any databases.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the backup destination that you want to terminate.
3. Click **Backup Destinations**.
4. Click the name of the backup destination that you want to terminate.
The Backup Destination Details page displays information about the selected backup destination.

5. Click **Terminate**.
6. In the resulting dialog, enter the backup destination name and click **Terminate Backup Destination** to confirm the action.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Exadata Cloud at Customer backup destinations:

- [CreateBackupDestination](#)
- [DeleteBackupDestination](#)
- [GetBackupDestination](#)
- [ListBackupDestination](#)
- [UpdateBackupDestination](#)

For the complete list of APIs, see [Database Service API](#).

Managing Databases on Exadata Cloud at Customer

This topic explains how to manage Oracle databases on Exadata Cloud at Customer.

Before you can create and use Oracle databases on Exadata Cloud at Customer, you must:

- Provision Exadata Cloud at Customer infrastructure.
- Configure a VM cluster.
- Create any required backup destinations.

You can create one or more databases on each Exadata Cloud at Customer system. Apart from the inherent storage and processing capacity of your Exadata system, there is no set maximum for the number of databases that you can create.

By default, databases on Exadata Cloud at Customer use Oracle Database Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all the database enterprise management packs and all the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (RAC). If you are using your own Oracle Database licenses, then your ability to use various features is limited by your license holdings.

Exadata Cloud at Customer supports the following Oracle Database software releases:

- Oracle Database 19c
- Oracle Database 18c
- Oracle Database 12c Release 2
- Oracle Database 12c Release 1
- Oracle Database 11g Release 2

When you provision a database, you can associate it with a backup destination and enable automatic backups. For more information, see [Managing Backup Destinations for Exadata Cloud at Customer](#) and [Managing Database Backup and Recovery on Exadata Cloud at Customer](#).

Each Oracle database is configured as follows:

- Each database is configured with Oracle RAC database instances running on every node in the VM cluster.
- Each database uses a separate set of Oracle binaries in a separate Oracle home location.
- Each database is configured with default instance parameter settings. While the defaults are reasonable for many cases, you should review the instance parameter settings to ensure that they meet your specific application needs.
In particular, consider the Oracle Database system global area (SGA) and program global area (PGA) instance parameter settings, especially if your VM cluster supports multiple databases. And, ensure that the sum of all Oracle Database memory allocations never exceeds the available physical memory on each compute node.
- Each database using Oracle Database 12c Release 1, or later, is configured as a

container database (CDB) and one pluggable database (PDB) is created inside the CDB. By default:

- The first PDB is configured with a local PDB administration user account named `PDBADMIN`.
- The `PDBADMIN` user account is initially configured with the same administration password as the CDB `SYS` and `SYSTEM` users.
- The `PDBADMIN` user account is initially configured with basic privileges assigned through two roles; `CONNECT` and `PDB_DBA`. However, for most practical administration purposes you must assign extra privileges to the `PDBADMIN` user account or the `PDB_DBA` role.

You can use native Oracle Database facilities to create extra PDBs and to manage all of your PDBs. The `dbaascli` utility also provides a range of convenient PDB management functions. See [Using the dbaascli Utility on Exadata Cloud at Customer](#).

- Depending on the Oracle Database version in use, each database contains web-based monitoring and management tools provided by Enterprise Manager Database Express (EM Express) or Enterprise Manager Database Control (Database Control). See [Using EM Express and Database Control on Exadata Cloud at Customer](#). The `dbaascli` utility also provides a range of convenient database management functions. See [Using the dbaascli Utility on Exadata Cloud at Customer](#).



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Using the Console

To create a database

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **VM Clusters**.
4. Click the name of a VM cluster.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click **Create Database**.
7. Provide the requested information in the **Create Database** page:
 - **Provide the database name:** Specify a user-friendly name that you can use to identify the database. The name doesn't need to be unique because an Oracle Cloud Identifier (OCID) uniquely identifies the database.
 - **Provide a unique name for the database:** Optionally specify a unique name for the database. This attribute defines the value of the `DB_UNIQUE_NAME` database parameter. The value is case insensitive, it can be up to 30 characters in length, and include alphanumeric characters, underscore (`_`), number sign (`#`), and dollar sign (`$`).

If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

- **Select a VM cluster:** From the list of available VM clusters, choose the VM cluster that you want to host the database. You must have an available VM cluster before you can create a database.
- **Select a database version:** From the list, choose the Oracle Database software version to use for the database.
- **Provide the name of the first PDB:** Optionally specify the name for the first PDB, which is created along with the database.

To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used. However, the system-generated name is not guaranteed to be unique across all of the databases in the VM cluster.

- **Provide the administration password:** Provide and confirm the administration password. This password is used for administration accounts and functions in the database, including:
 - The password for the Oracle Database `SYS` and `SYSTEM` users.
 - The Transparent Data Encryption (TDE) keystore password.
 - For databases using Oracle Database 12c Release 1, or later, the password for the PDB administration user in the first PDB (`PDBADMIN`).

The password must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.

- **Choose the database workload type:** Select the workload type that best suits your application:
 - **Online Transactional Processing (OLTP)** configures the database for a transactional workload, with a bias toward high volumes of random data access.

- **Decision Support System (DSS)** configures the database for a decision support or data warehouse workload, with a bias toward large data scanning operations.

CONFIGURE BACKUPS

Use the following settings to define the backup configuration for the database:

- **Backup Destination Type:** From the list, choose an option.
 - **None** does not define a backup configuration for the database.
 - **Local** stores backups locally in the Exadata Storage Servers on your Exadata Cloud at Customer system.
This option is available only if you enabled backups on local Exadata storage in the VM cluster that you want to host the database.
 - **Object Storage** stores backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.
To use this option, your Exadata Cloud at Customer system must have egress connectivity to Oracle Cloud Infrastructure Object Storage.
 - **NFS** stores backups in one of your previously defined backup destinations that uses Network File System (NFS) storage. See [Managing Backup Destinations for Exadata Cloud at Customer](#).
If you select this option, you must also choose from the list of NFS **Backup Destinations**.
 - **Recovery Appliance** stores backups in one of your previously defined backup destinations that uses Oracle Zero Data Loss Recovery Appliance. See [Managing Backup Destinations for Exadata Cloud at Customer](#).
If you select this option, you must also:
 - Choose from the list of Recovery Appliance **Backup Destinations**.
 - Choose from the **VPC User** list, which contains the list of virtual private catalog (VPC) user names that are defined in the Recovery

Appliance backup destination.

- Provide the **Password** for the VPC user.



Note

If you select a backup destination, you cannot change it after the database is created. However, if you select **None** now, you can select a backup destination after the database is created.

- **Enable automatic backups:** Select this option to enable daily backups using the policy for automatic backups.
This option is only enabled when you select a **Backup Destination Type** other than **None**. You can change this setting after database creation.

SHOW ADVANCED OPTIONS

- **Backup retention period:** From the list, you can choose the length of time that automatic backups are retained.
For backups to local Exadata storage, you can choose a retention period of 7 days or 14 days. The default retention period is 7 days.
For backups to Oracle Cloud Infrastructure Object Storage or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
This option does not apply to Recovery Appliance backup destinations. For backups to a Recovery Appliance, the retention policy that is implemented in the Recovery Appliance controls the retention period.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For

more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click **Create Database**.

The Database Details page is now displayed. While the creation process is running, the state of the database is **Pending**. When the database creation process completes, the state of the database changes to **Active**.

To terminate a database

Terminating a database removes it from the Cloud Control Plane. In the process, all of the associated data files and backups are destroyed.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database that you want to terminate.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the database that you want to terminate.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you want to terminate.
The Database Details page displays information about the selected database.
7. Click **Terminate**.
8. In the resulting dialog, enter the name of the database and click **Terminate Database** to confirm the action.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage databases.

Database homes:

- [CreateDbHome](#)
- [DeleteDbHome](#)
- [GetDbHome](#)
- [ListDbHomes](#)

Databases:

- [GetDatabase](#)
- [ListDatabases](#)

Nodes:

- [GetDbNode](#)
- [ListDbNodes](#)

For the complete list of APIs, see [Database Service API](#).

Managing Database Backup and Recovery on Exadata Cloud at Customer

This topic explains how to work with the backup and recovery facilities provided by Exadata Cloud at Customer.

Exadata Cloud at Customer provides automatic database backup facilities that use Oracle Recovery Manager (RMAN). When you create a database on Exadata Cloud at Customer, you can specify a backup destination and enable automatic backups. See [Managing Databases on Exadata Cloud at Customer](#).

After database creation, you can also:

- View a list of available backups.
- Enable or disable automatic backups.

- Edit backup settings.
- Restore a database.

You can perform these operations by using the Console or the API.

Automatic database backups are configured as follows:

- Automatic backups are scheduled daily. The automatic backup process can run at any time within the daily backup window, which is between midnight and 6:00 AM in the time zone of the VM cluster that hosts the database.
- Automatic backups use a combination of full (RMAN level 0) and incremental (RMAN level 1) database backups:
 - For backups to a Recovery Appliance, after an initial full backup is performed the Recovery Appliance creates and validates virtual full backups from each daily incremental backup.
 - For backups to a Network File System (NFS) backup destination, incremental backups are always performed after an initial full backup is taken. Also, the incremental backups are merged into the full backup when they become older than the retention period.
 - For backups to all other storage types, the default interval between full backups is seven days.
- The retention period defines the period for which automatic backups are maintained:
 - For backups to a Recovery Appliance, the retention policy that is implemented in the Recovery Appliance controls the retention period.
 - For backups to local Exadata storage, you can choose a retention period of 7 days or 14 days. The default retention period is 7 days.
 - For backups to Oracle Cloud Infrastructure Object Storage or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
- By default, the database runs in `ARCHIVELOG` mode, and archived redo log files are backed up every 60 minutes.
- Regardless of the backup destination, backups of user data are encrypted by default.

While a backup is in progress, Oracle recommends that you avoid performing actions that could interfere with availability, such as restarting compute nodes or applying patches. If an automatic backup operation fails, the backup is deferred until the next day's backup window.

When required, you can restore a database to:

- The latest available restore point.
- A specific point in time by providing a timestamp.
- An Oracle Database System Change Number (SCN).



Note

The backup and recovery facilities described in this topic cater only for database backup and recovery, which includes Oracle Database data files, log files, control files, and the server parameter (SP) file. You are responsible for backing up other files on your compute nodes. In particular, Oracle strongly recommends that you back up the Transparent Data Encryption (TDE) keystore (wallet). Without the TDE keystore, the database backups are effectively useless because you cannot read the data contained therein.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Using the Console

To view a list of available backups

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database that you are interested in.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database that you are interested in.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you are interested in.
The Database Details page displays information about the selected database, which includes a list of the available backups.

To edit backup settings

Use the following procedure to change the available backup settings.

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database for which you want to edit backup settings.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database for which you want to edit backup settings.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database for which you want to edit backup settings.
The Database Details page displays information about the selected database.
7. Click **Backup Settings**.
8. Your current backup configuration determines the changes that you can make in the

Backup Settings dialog, as follows:

- If automatic backups are not configured (**Backup Destination Type** is set to **None**), then you can use the following settings to define the backup configuration for the database:
 - **Backup Destination Type:** From the list, choose an option.
 - **None** does not define a backup configuration for the database.
 - **Local** stores backups locally in the Exadata Storage Servers on your Exadata Cloud at Customer system.

This option is available only if you enabled backups on local Exadata storage in the VM cluster that you want to host the database.
 - **Object Storage** stores backups in an Oracle-managed object storage container on Oracle Cloud Infrastructure.

To use this option, your Exadata Cloud at Customer system must have egress connectivity to Oracle Cloud Infrastructure Object Storage.
 - **NFS** stores backups in one of your previously defined backup destinations that uses Network File System (NFS) storage. See [Managing Backup Destinations for Exadata Cloud at Customer](#).

If you select this option, you must also choose from the list of NFS **Backup Destinations**.
 - **Recovery Appliance** stores backups in one of your previously defined backup destinations that uses Oracle Zero Data Loss Recovery Appliance. See [Managing Backup Destinations for Exadata Cloud at Customer](#).

If you select this option, you must also:

 - Choose from the list of Recovery Appliance **Backup Destinations**.

- Choose from the **VPU User** list, which contains the list of virtual private catalog (VPC) user names that are defined in the Recovery Appliance backup destination.
- Provide the **Password** for the VPC user.



Note

If you select a backup destination (other than **None**), you cannot change it later.

- **Enable automatic backups:** Select this option to enable daily backups using the policy for automatic backups.
This option is only enabled when you select a **Backup Destination Type** other than **None**. You can change this setting later.
- **Backup retention period:** From the list, you can choose the length of time that automatic backups are retained.
For backups to local Exadata storage, you can choose a retention period of 7 days or 14 days. The default retention period is 7 days.
For backups to Oracle Cloud Infrastructure Object Storage or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
This option does not apply to Recovery Appliance backup destinations. For backups to a Recovery Appliance, the retention policy that is implemented in the Recovery Appliance controls the retention period.
- If automatic backups were previously configured, then you can make the following changes:
 - For Recovery Appliance backup destinations, you can update the **Password** for the virtual private catalog (VPC) user that is used to access the Recovery Appliance.

- For backup destinations that do not use a Recovery Appliance, you can update the **Backup retention period** for automatic backups:
 - For backups to local Exadata storage, you can choose a retention period of 7 days or 14 days. The default retention period is 7 days.
 - For backups to Oracle Cloud Infrastructure Object Storage or to an NFS backup destination, you can choose one of the following preset retention periods: 7 days, 14 days, 30 days, 45 days, or 60 days. The default retention period is 30 days.
 - For backups to a Recovery Appliance, the retention policy that is implemented in the Recovery Appliance controls the retention period.
 - You can set the option to **Enable automatic backups**. Select this option to enable automatic database backups. Deselect this option to suspend automatic database backups.
9. Click **Save Changes**.

To restore a database

1. Open the navigation menu. Under Database, click **Exadata Cloud at Customer**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that hosts the database that you want to restore.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that hosts the database that you want to restore.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you want to restore.
The Database Details page displays information about the selected database.
7. Click **Restore Database**.
8. In the resulting dialog box, select one of the following options, and click **Restore Database**:

- **Restore to latest:** The database is restored and recovered with zero, or least possible, data loss.
- **Restore to a timestamp:** The database is restored and recovered to the specified timestamp.
- **Restore to SCN:** The database is restored and recovered to the specified Oracle Database System Change Number (SCN). The specified SCN must be valid otherwise the operation fails.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage database backup and recovery:

- [GetBackup](#)
- [ListBackups](#)
- [RestoreDatabase](#)
- [UpdateDatabase](#) - To enable and disable automatic backups.

For the complete list of APIs, see [Database Service API](#).

Connecting to an Exadata Cloud at Customer System

This topic explains how to connect to an Exadata Cloud at Customer system using SSH, and how to connect to an Exadata Cloud at Customer database using Oracle Net Services (SQL*Net).

Connecting to a Compute Node with SSH

You can connect to the compute nodes in an Exadata Cloud at Customer system by using a Secure Shell (SSH) connection. Most UNIX-style systems (including Linux, Solaris, BSD, and

OS X) include an SSH client. For Windows, you can download a free SSH client called PuTTY from <http://www.putty.org>.

PREREQUISITES

To access a compute node in an Exadata Cloud at Customer system by using SSH, you need the following:

- An SSH private key file that corresponds to a public key that is registered in the system. When you create a VM cluster on your Exadata Cloud at Customer system, you must specify the public key portion of one or more SSH key pairs. You can also register extra keys separately after you create the VM cluster. The public keys are stored in the `authorized_keys` file at `~/.ssh/authorized_keys`. Separate `authorized_keys` files are located under the home directories of the `oracle` and `opc` OS users.
- The host name or IP address for the compute node that you want to access. See [To check the status of a VM cluster compute node](#).

To connect from a UNIX-style system

Use the following SSH command to access a compute node:

```
$ ssh -i private-key user@node
```

In the preceding command:

- *private-key* is the full path and name of the file that contains the SSH private key that corresponds to a public key that is registered in the system.
- *user* is the operating system user you want to connect as:
 - Connect as `oracle` to perform operations as the Oracle Database software owner; this user does not have `root` user access to the compute node.
 - Connect as `opc` to perform operations that require `root` access to the compute node, such as patching. This user can use the `sudo -s` command to gain `root` user access to the compute node.
- *node* is the host name or IP address for the compute node that you want to access.

To connect from a Windows system using PuTTY

1. Run the PuTTY program (`putty.exe`).
The PuTTY Configuration window is displayed, showing the Session panel.
2. In the **Host Name (or IP address)** field, enter the host name or IP address of the compute node that you want to access.
3. Confirm that the **Connection type** option is set to **SSH**.
4. In the Category tree, expand **Connection** if necessary and then click **Data**.
The Data panel is displayed
5. In the **Auto-login username** field, enter the operating system user you want to connect as:
 - Connect as `oracle` to perform operations as the Oracle Database software owner; this user does not have `root` user access to the compute node.
 - Connect as `opc` to perform operations that require `root` access to the compute node, such as patching. This user can use the `sudo -s` command to gain `root` user access to the compute node.
6. Confirm that the **When username is not specified** option is set to **Prompt**.
7. In the Category tree, expand **SSH** and then click **Auth**.
The Auth panel is displayed.
8. Click the **Browse** button next to the **Private key file for authentication** field. Then, navigate to and open the private key file that corresponds to a public key that is registered in the system.
9. In the Category tree, click **Session**.
The Session panel is displayed.
10. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
11. Click **Open** to open the connection.
The PuTTY Configuration window closes and the PuTTY terminal window displays.

To access a database after you connect to the compute node

After you connect to a compute node, you can use the following series of commands to identify a database and connect to it.

1. Log in as the `oracle` user. For example:

```
$ ssh -i keyfile oracle@node01
[oracle@node01 ~]$
```

2. Use the `srvctl` utility located under the Oracle Grid Infrastructure home directory to list the databases on the system. For example:

```
[oracle@node01 ~]$ /u01/app/12.2.0.1/grid/bin/srvctl config database -v
nc122 /u02/app/oracle/product/12.2.0/dbhome_6 12.2.0.1.0
s12c /u02/app/oracle/product/12.2.0/dbhome_2 12.2.0.1.0
```

3. Identify the database instances for the database that you want to access. For example:

```
[oracle@node01 ~]$ /u01/app/12.2.0.1/grid/bin/srvctl status database -d s12c
Instance s12c1 is running on node node01
Instance s12c2 is running on node node02
```

4. Configure the environment settings for the database that you want to access. For example:

```
[oracle@node01 ~]$ . oraenv
ORACLE_SID = [oracle] ? s12c
The Oracle base has been set to /u02/app/oracle
[oracle@node01 ~]$ export ORACLE_SID=s12c1
```

5. You can use the `srvctl` command to display more detailed information about the database. For example:

```
[oracle@node01 ~]$ srvctl config database -d s12c
Database unique name: s12c
Database name:
Oracle home: /u02/app/oracle/product/12.2.0/dbhome_2
Oracle user: oracle
Spfile: +DATAC4/s12c/spfiles12c.ora
Password file: +DATAC4/s12c/PASSWORD/passwd
Domain: example.com
```

CHAPTER 11 Database

```
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: DATA4
Mount point paths:
Services:
Type: RAC
Start concurrency:
Stop concurrency:
OSDBA group: dba
OSOPER group: racoper
Database instances: s12c1,s12c2
Configured nodes: node01,node02
CSS critical: no
CPU count: 0
Memory target: 0
Maximum memory: 0
Default network number for database services:
Database is administrator managed
```

6. You can access the database by using SQL*Plus. For example:

```
[oracle@node01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production ...

Copyright (c) 1982, 2016, Oracle. All rights reserved.

Connected to:
Oracle Database 12c EE Extreme Perf Release 12.2.0.1.0 - 64bit Production

SQL>
```

Connecting to a Database with Oracle Net Services

Oracle Database Exadata Cloud at Customer supports remote database access by using Oracle Net Services.

Because Exadata Cloud at Customer uses Oracle Grid Infrastructure, you can make Oracle Net Services connections by using Single Client Access Name (SCAN). SCAN is a feature that provides a consistent mechanism for clients to access the Oracle databases running in a cluster.

By default, the SCAN is associated with three virtual IP addresses (VIPs). Each SCAN VIP is also associated with a SCAN listener that provides a connection endpoint for Oracle Database connections using Oracle Net Services. To maximize availability, Oracle Grid Infrastructure distributes the SCAN VIPs and SCAN listeners across the available cluster nodes. In addition, the SCAN VIPs and SCAN listeners are automatically migrated to a surviving node if there is a node shutdown or failure. The aim is to ensure that Oracle Database clients always have a reliable set of connection endpoints that can service all of the databases running in the cluster.

The SCAN listeners are in addition to the Oracle Net Listeners that run on every node in the cluster, which are also known as the node listeners. When an Oracle Net Services connection comes through SCAN, the SCAN listener routes the connection to one of the node listeners and plays no further part in the connection. A combination of factors, including listener availability, database instance placement, and workload distribution, determines which node listener receives each connection.



Note

This section outlines the basic requirements for connecting to your Exadata Cloud at Customer databases by using Oracle Net Services. For more guidance on achieving continuous service during planned maintenance, node failure, or Oracle Database instance failure, see the [Continuous Availability](#) white paper.

PREREQUISITES

To connect to an Exadata Cloud at Customer database by using Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a compute node that hosts the database that you want to access.
- The database identifier, either the database SID or service name.

To connect using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches. You can:

- Use a connect descriptor that references all of the SCAN VIPs.

This approach requires you to supply all of the SCAN VIP addresses and allows Oracle Net Services to connect to an available SCAN listener.

You can use the following template to define a Net Services alias, which is typically used to provide a convenient name for the connect descriptor:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=
    (ADDRESS= (PROTOCOL=tcp) (HOST=SCAN-VIP-1) (PORT=1521))
    (ADDRESS= (PROTOCOL=tcp) (HOST=SCAN-VIP-2) (PORT=1521))
    (ADDRESS= (PROTOCOL=tcp) (HOST=SCAN-VIP-3) (PORT=1521)))
  (CONNECT_DATA= (sid-or-service-entry)))
```

In the preceding definition:

- *alias-name* is the name you use to identify the alias.
- *SCAN-VIP-[1-3]* are the IP addresses for the SCAN VIPs.
- *sid-or-service-entry* identifies the database SID or service name using one of the following formats:
 - SID=*sid-name*; for example, SID=S12C1.
 - SERVICE_NAME=*service-name*; for example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

**Note**

By default, Oracle Net Services randomly selects one of the addresses in the address list to balance the load between the SCAN listeners.

- Use a connect descriptor that references a custom SCAN name.
Using this approach, you define a custom SCAN name in your DNS, which resolves to the three SCAN VIPs.
You can use the following template to define a Net Services alias that references the custom SCAN name:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp) (HOST=scan-name) (PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
scan-name:1521/sid-or-service-entry
```

For example:

```
exalscan.example.com:1521/S12C1
```

or

```
exalscan.example.com:1521/PDB1.example.oraclecloudatcust.com
```

To connect using a node listener

You can create an Oracle Net Services connection by using a connect descriptor that bypasses the SCAN listeners and routes your connection directly to a node listener.

By using this method, you give up the high-availability and load-balancing provided by SCAN. However, this method may be desirable if you want to direct connections to a specific node or

CHAPTER 11 Database

network interface. For example, you might want to ensure that connections from a program that performs bulk data loading use the backup network.

Using this approach, you direct your connection using the hostname or IP address of the node.

You can use the following template to define a Net Services alias that directly references the node:

```
alias-name = (DESCRIPTION=
(CONNECT_TIMEOUT=timeout)
(ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp) (HOST=node) (PORT=1521)))
(CONNECT_DATA=(sid-or-service-entry)))
```

In the preceding definition:

- *alias-name* is the name you use to identify the alias.
- *timeout* specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The (CONNECT_TIMEOUT=*timeout*) parameter is optional.
- *node* is the hostname or IP address for the compute node that you want to use.
- *sid-or-service-entry* identifies the database SID or service name using one of the following formats:
 - SID=*sid-name*; for example, SID=S12C1.
 - SERVICE_NAME=*service-name*; for example, SERVICE_NAME=PDB1.example.oraclecloudatcust.com.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
node:1521/sid-or-service-entry
```

For example:

```
exalnode01.example.com:1521/S12C1
```

or

```
exalnode01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

Patching and Updating an Exadata Cloud at Customer System

This topic describes the responsibilities and procedures for patching and updating various components in Exadata Cloud at Customer.



Note

For more guidance on achieving continuous service during patching operations, see the [Continuous Availability](#) white paper.

Patching Performed by Oracle

Oracle performs patches and updates to all of the Oracle-managed system components. This includes the physical compute nodes (Dom0), network switches, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and the Exadata Storage Servers.

In all but rare exceptional circumstances, you receive advance communication about these updates to help you plan for them. If there are corresponding recommended updates for your compute node virtual machines (VMs), then Oracle provides notification about them.

Wherever possible, scheduled updates are performed in a manner that preserves service availability throughout the update process. However, there may be some noticeable impact on performance and throughput while individual system components are unavailable during the update process.

For example, Dom0 patching typically requires a reboot. In such cases, wherever possible, the compute nodes are rebooted in a rolling manner, one at a time, to ensure that the service remains available throughout the process. However, each compute node is unavailable for a short time while it reboots, and the overall service capacity diminishes accordingly. Also, if your applications cannot tolerate the reboots, you may need to take mitigating action. For example, you may need to shut down an application while Dom0 patching occurs.

Managing Oracle Database and Oracle Grid Infrastructure Patches

You are responsible for routine patching of the Oracle Database and Oracle Grid Infrastructure software. On Exadata Cloud at Customer, routine patching of the Oracle Database and Oracle Grid Infrastructure software is facilitated by using the `dbaascli` utility. The `dbaascli` utility provides a simple means for applying routine patches, which Oracle periodically loads on to the Cloud Control Plane servers.

The `dbaascli` utility is part of the cloud-specific tooling bundle that is included with Exadata Cloud at Customer. Therefore, before performing the following procedures, ensure that you have the latest version of the cloud-specific tooling on all of the compute nodes in the VM cluster. For more information, see [Cloud Tooling Updates](#).

To list available patches

You can produce a list of available patches by using the `dbaascli` command as follows:

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Execute the `dbaascli patch db list` command:

```
# dbaascli patch db list --oh hostname:oracle_home
```

In the preceding command, `--oh` specifies a compute node and Oracle Home directory for which you want to list the available patches. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.

For example:

```
# dbaascli patch db list --oh hostname1:/u02/app/oracle/product/12.1.0.2/dbhome_1
```

**Note**

The list of available patches is determined by interrogating the database to establish the patches that have already been applied. When a patch is applied, the corresponding database entry is made as part of the SQL patching operation, which is executed at the end of the patch workflow. Therefore, the list of available patches may include partially applied patches along with patches that are currently being applied.

To check prerequisites before applying a patch

You can perform the prerequisites-checking operation using the `dbaascli` command as follows:

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Execute the `dbaascli patch db prereq` command:

- On a specific instance:

```
# dbaascli patch db prereq --patchid patchid --instance1 hostname:oracle_home [--dbnames dbname[,dbname2 ...]]
```

- By specifying only database names:

```
# dbaascli patch db prereq --patchid patchid --dbnames dbname[,dbname2 ...] [-all dbs]
```

In the preceding commands:

- `patchid` identifies the patch to be pre-checked. For details about how to find the available patch identifiers, see [To list available patches](#).

- `--instance1` specifies a compute node and Oracle Home directory that is subject to the pre-check operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.
- `--dbnames` specifies the database names for the databases that are the target of the pre-check operation.
- `-alldbs` specifies that you want to pre-check all of the databases that share the same Oracle Database binaries (Oracle Home) as the specified databases.

For example:

```
# dbaascli patch db prereq 12345678 --instance1
hostname1:/u02/app/oracle/product/12.1.0.2/dbhome_1
```

To apply a patch

You can apply a patch by using the `dbaascli` command.

The patching operation:

- Can be used to patch some or all of your compute nodes using one command.
- Coordinates multi-node patching in a rolling manner.
- Can execute patch-related SQL after patching all the compute nodes in the cluster.

You can perform a patching operation using the `dbaascli` command as follows:

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Execute the `dbaascli patch db apply` command:

- On a specific instance:

```
# dbaascli patch db apply --patchid patchid --instance1 hostname:oracle_home [--dbnames
dbname[,dbname2 ...] [--run_datasql 1]
```

- By specifying only database names:

```
# dbaascli patch db apply --patchid patchid --dbnames dbname[,dbname2 ...] [--run_datasql  
1] [-alldbs]
```

In the preceding commands:

- *patchid* identifies the patch to be applied. For details about how to find the available patch identifiers, see [To list available patches](#).
- `--instance1` specifies a compute node and Oracle Home directory that is subject to the patching operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory. If you use this argument to specify a shared Oracle Home directory and you do not specify the `--dbnames` argument, then all of the databases that share the specified Oracle Home are patched. After the operation, the Oracle Home directory location remains unchanged; however, the patch level information embedded in the Oracle Home name is adjusted to reflect the patching operation.
- `--dbnames` specifies the database names for the databases that are the target of the patching operation.

If you use this argument to patch a database that uses a shared Oracle Home and you do not specify the `-alldbs` option, then a new Oracle Home containing the patched Oracle Database binaries is created and the database is moved to the new Oracle Home.

- `-alldbs` patches all of the databases that share the same Oracle Database binaries (Oracle Home) as the databases specified in the `--dbnames` argument. After the operation, the Oracle Home directory location remains unchanged; however, the patch level information embedded in the Oracle Home name is adjusted to reflect the patching operation.
- `--run_datasql 1` instructs the command to execute patch-related SQL commands.



Note

- Patch-related SQL should only be executed after all of the compute nodes are patched. Take care not to specify this argument if you are patching a node and further nodes remain to be patched.
- This argument can only be specified along with a patching operation on a compute node. If you have patched all of your nodes and you did not specify this argument, you need to manually execute the SQL commands associated with the patch, which typically involves running the `catbundle.sql` script for Oracle Database 11g or the `datapatch` utility for Oracle Database 12c, or later. Refer to the patch documentation for full details.

For example:

```
# dbaascli patch db apply 23456789 --instance1 hostname1:/u02/app/oracle/product/12.1.0.2/dbhome_1 --run_datasql 1
```

To list applied patches

You can use the `opatch` utility to list the patches that have been applied to an Oracle Database or Grid Infrastructure installation.

To produce a list of applied patches for an Oracle Database installation:

1. Connect to a compute node as the `oracle` user.
For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).
2. Set the `ORACLE_HOME` variable to the location of the Oracle Database installation you want to examine. For example:

```
$ export ORACLE_HOME=/u02/app/oracle/product/12.1.0.2/dbhome_1
```

3. Execute the `opatch` command with the `lspatches` option:

```
$ $ORACLE_HOME/OPatch/opatch lspatches
```

To produce a list of applied patches for Oracle Grid Infrastructure:

1. Connect to a compute node as the `opc` user.
2. Become the `grid` user:

```
$ sudo -s  
# su - grid
```

3. Execute the `opatch` command with the `lspatches` option:

```
$ $ORACLE_HOME/OPatch/opatch lspatches
```

To roll back a patch

You can roll back a patch or failed patch attempt by using the `dbaascli` command.

The patch rollback operation:

- Can be used to roll back a patch on some or all of your compute nodes using one command.
- Coordinates multi-node operations in a rolling manner.
- Can execute rollback-related SQL after rolling back the patch on all the compute nodes in the cluster.

You can perform a patch rollback operation using the `dbaascli` command as follows:

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Execute the `dbaascli` command with the `-rollback_async` action:

- On specific instances:

```
# dbaascli patch db switchback --patchid patchid --instance1 hostname:oracle_home [--dbnames dbname[,dbname2 ...]] [--run_datasql 1]
```

- By specifying only database names:

```
# dbaascli patch db switchback --patchid patchid --dbnames dbname[,dbname2 ...] [--run_datasql 1] [-alldbs]
```

In the preceding commands:

- *patchid* identifies the patch to be rolled back.
- `--instance1` specifies a compute node and Oracle Home directory that is subject to the rollback operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory. If you use this argument to specify a shared Oracle Home directory and you do not specify the `--dbnames` argument, then all of the databases that share the specified Oracle Home are rolled back.
- `--dbnames` specifies the database names for the databases that are the target of the rollback operation.
- `-alldbs` specifies that you want to roll back all of the databases that share the same Oracle Database binaries (Oracle Home) as the databases specified in the `--dbnames` argument.
- `--run_datasql 1` instructs the command to execute rollback-related SQL commands.



Note

- Rollback-related SQL should only be executed after all of the compute nodes are rolled back. Take care not to specify this argument if you are rolling back a node and further nodes remain to be rolled back.
- This argument can only be specified along with a rollback operation on a compute node. Therefore, if you have rolled back all of your nodes and you did not specify this argument, you need to manually execute the SQL commands associated with the rollback operation. Refer to the patch documentation for full details.

For example:

```
# dbaascli patch db switchback 34567890 --instance1  
hostname1:/u02/app/oracle/product/12.1.0.2/dbhome_1 --run_dataSQL 1
```

MANUALLY PATCHING ORACLE DATABASE AND ORACLE GRID INFRASTRUCTURE SOFTWARE

In general, Oracle recommends that you use the facilities provided by Exadata Cloud at Customer to perform routine patching of Oracle Database and Oracle Grid Infrastructure software. However, you may need to manually patch the Oracle Database or Oracle Grid Infrastructure software in the following circumstances:

- **Oracle Java Virtual Machine (OJVM) Patching:** Because they cannot be applied in a rolling fashion, patches for the Oracle Database OJVM component are not included in the routine patch sets for Exadata Cloud at Customer. If you need to apply patches to the OJVM component of Oracle Database, you must do so manually. See [Oracle JavaVM](#)

[Component Database PSU and RU \(OJVM PSU and OJVM RU\) Patches.](#)

- **Daylight Savings Time (DST) Patching:** Because they cannot be applied in a rolling fashion, patches for the Oracle Database DST definitions are not included in the routine patch sets for Exadata Cloud at Customer. If you need to apply patches to the Oracle Database DST definitions, you must do so manually. See [Updated DST Transitions and New Time Zones in Oracle RDBMS and OJVM Time Zone File Patches.](#)
- **Non-routine or One-off Patching:** If you encounter a problem that requires a patch which is not included in any routine patch set, work with Oracle Support Services to identify and apply the appropriate patch.

For general information about patching Oracle Database, see "Patch Set Updates and Requirements for Upgrading Oracle Database" in the *Oracle Database Upgrade Guide* for Release [19,18](#), [12.2](#), [12.1](#), or [11.2](#).

Updating the Compute Node Operating System

You are responsible for managing patches and updates to the operating system environment on the compute node VMs. This section outlines the standard Exadata tools and techniques that you can use to update the operating system components on the Exadata Cloud at Customer compute nodes. For further information, see [Updating Exadata Database Servers](#) in the *Oracle Exadata Database Machine Maintenance Guide*.

Preparing for an OS update

- Before you begin an update, review *Exadata Cloud Service Software Versions* ([Doc ID 2333222.1](#)) to determine the latest software to use.
- You are able to apply Exadata software release updates to the compute nodes at your convenience. For feature release updates only, Oracle recommends that you lodge a service request with Oracle Support Services to ensure that Oracle is aware of your plans and is primed to assist if there are any difficulties.

A feature release update is an update that changes any of the first four digits in the Exadata software release identifier. For example, upgrading from Exadata software release 12.1.2.2.0 to release 12.1.2.3.0 would be a feature release update. However,

upgrading from Exadata software release 12.1.2.3.0 to release 12.1.2.3.4 would not be considered a feature release update. You can determine the current Exadata software release by executing the `imageinfo` command on any compute node.

- Some steps in the update process require you to specify a YUM repository. The YUM repository URL is:

```
http://yum.oracle.com/repo/EngineeredSystems/exadata/dbserver/latest-version/base/x86_64.
```

In the preceding URL, *latest-version* is the YUM repository version. You can examine the output from the following `curl` command to determine the latest version of the YUM repository:

```
curl -s -X GET http://yum.oracle.com/repo/EngineeredSystems/exadata/dbserver/index.html
```

- To apply OS updates, the network hosting your Exadata Cloud at Customer system must be configured to allow access to the YUM repository.

To update the OS on all compute nodes of an Exadata Cloud at Customer system

You update the operating system on the compute node virtual machines by using the `patchmgr` tool. This utility manages the entire update of one or more compute nodes remotely, including the pre-reboot, reboot, and post-reboot steps.

You can run the utility from one of your Exadata Cloud at Customer compute nodes or a non-Exadata server running Oracle Linux. The server on which you run the utility is known as the *driving system*. You cannot use the driving system to update itself. Therefore, if the driving system is one of the compute nodes in a VM cluster that you are updating then you must run the `patchmgr` utility more than once.

The following scenarios describe typical ways of performing the updates:

- **Non-Exadata Driving System**

The simplest way to run the update the system is to use a separate Oracle Linux server to update all compute nodes in one operation.

- **Exadata Compute Node Driving System**

You can use one compute node to drive the updates for the rest of the compute nodes in the VM cluster. Then, you can use one of the updated nodes to drive the update on the original driving system. For example, consider updating a half rack system with four compute nodes; `node1`, `node2`, `node3`, and `node4`. You could first use `node1` to drive the updates of `node2`, `node3`, and `node4`. Then, you could use `node2` to drive the update of `node1`.

The driving system requires `root` user SSH access to each compute node being updated.

The following procedure is based on an example that assumes the following:

- The system has two compute nodes, `node1` and `node2`.
- The target Exadata software version is `18.1.4.0.0.180125.3`.
- Each node is used as the driving system to update the other node.

1. Gather the environment details.

- a. Using SSH, connect to `node1` as `root` and run the following command to determine the current Exadata software version:

```
[root@node1 ~]# imageinfo -ver
12.2.1.1.4.171128
```

- b. Switch to the grid user, and identify all nodes in the cluster.

```
[root@node1 ~]# su - grid
[grid@node1 ~]$ olsnodes
node1
node2
```

2. Configure the driving system.

- a. Switch back to the `root` user on `node1` and check whether an SSH key pair (`id_rsa` and `id_rsa.pub`) exists. If not, then generate it.

```
[root@node1 ~]# ls /root/.ssh/id_rsa*
ls: cannot access /root/.ssh/id_rsa*: No such file or directory
[root@node1 ~]# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
```

```

Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
93:47:b0:83:75:f2:3e:e6:23:b3:0a:06:ed:00:20:a5 root@node1.example.com
The key's randomart image is:
+--[ RSA 2048]-----+
|O..      + .      |
|O.       o *      |
|E        . o o     |
|. .      =        |
| o .     S =      |
| +       = .      |
|  +      o o      |
| . .     + .      |
|        ...       |
+-----+

```

- b. Distribute the public key to the target nodes, and verify this step. In the example, the only target node is `node2`.

```

[root@node1 ~]# scp -i ~root/.ssh/id_rsa.pub opc@node2:/tmp/id_rsa.node1.pub

[root@node2 ~]# ls -al /tmp/id_rsa.node1.pub
-rw-r--r-- 1 opc opc 442 Feb 28 03:33 /tmp/id_rsa.node1.pub
[root@node2 ~]# date
Wed Feb 28 03:33:45 UTC 2018

```

- c. On the target node (`node2` in the example), add the root public key of `node1` to the root `authorized_keys` file.

```

[root@node2 ~]# cat /tmp/id_rsa.node1.pub >> ~root/.ssh/authorized_keys

```

- d. Download `patchmgr` into `/root/patch` on the driving system (`node1` in this example).

You can download the `patchmgr` bundle from Oracle Support by using Patch ID [21634633](#).

For further information, see also *dbnodeupdate.sh* and *dbserver.patch.zip: Updating Exadata Database Server Software using the DBNodeUpdate Utility and patchmgr* ([Doc ID 1553103.1](#)).

- e. Unzip the `patchmgr` bundle.

The name of your ZIP file may differ depending on the version that you downloaded.

```
[root@node1 ~]# cd /root/patch
[root@node1 patch]# unzip p21634633_181400_Linux-x86-64.zip
Archive:  p21634633_181400_Linux-x86-64.zip   creating:  dbserver_patch_5.180228.2/
creating:  dbserver_patch_5.180228.2/ibdiagtools/
inflating: dbserver_patch_5.180228.2/ibdiagtools/cable_check.pl
inflating: dbserver_patch_5.180228.2/ibdiagtools/setup-ssh
inflating: dbserver_patch_5.180228.2/ibdiagtools/VERSION_FILE
extracting: dbserver_patch_5.180228.2/ibdiagtools/xmonib.sh
inflating: dbserver_patch_5.180228.2/ibdiagtools/monitord
inflating: dbserver_patch_5.180228.2/ibdiagtools/checkbadlinks.pl
creating:  dbserver_patch_5.180228.2/ibdiagtools/topologies/
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/VerifyTopologyUtility.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/verifylib.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Node.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Rack.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Group.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Switch.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/topology-zfs
inflating: dbserver_patch_5.180228.2/ibdiagtools/dcli
creating:  dbserver_patch_5.180228.2/ibdiagtools/netcheck/
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteScriptGenerator.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/CommonUtils.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/SolarisAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/LinuxAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteLauncher.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteConfig.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/spawnProc.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/runDiagnostics.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/OSAdapter.pm
inflating: dbserver_patch_5.180228.2/ibdiagtools/SampleOutputs.txt
inflating: dbserver_patch_5.180228.2/ibdiagtools/infinicheck
inflating: dbserver_patch_5.180228.2/ibdiagtools/ibping_test
```

CHAPTER 11 Database

```
inflating: dbserver_patch_5.180228.2/ibdiagtools/tar_ibdiagtools
inflating: dbserver_patch_5.180228.2/ibdiagtools/verify-topology
inflating: dbserver_patch_5.180228.2/installfw_exadata_ssh
creating: dbserver_patch_5.180228.2/linux.db.rpms/
inflating: dbserver_patch_5.180228.2/md5sum_files.lst
inflating: dbserver_patch_5.180228.2/patchmgr
inflating: dbserver_patch_5.180228.2/xcp
inflating: dbserver_patch_5.180228.2/ExadataSendNotification.pm
inflating: dbserver_patch_5.180228.2/ExadataImageNotification.pl
inflating: dbserver_patch_5.180228.2/kernelupgrade_oldbios.sh
inflating: dbserver_patch_5.180228.2/cellboot_usb_pci_path
inflating: dbserver_patch_5.180228.2/exadata.img.env
inflating: dbserver_patch_5.180228.2/README.txt
inflating: dbserver_patch_5.180228.2/exadataLogger.pm
inflating: dbserver_patch_5.180228.2/patch_bug_26678971
inflating: dbserver_patch_5.180228.2/dcli
inflating: dbserver_patch_5.180228.2/patchReport.py
extracting: dbserver_patch_5.180228.2/dbnodeupdate.zip
creating: dbserver_patch_5.180228.2/plugins/
inflating: dbserver_patch_5.180228.2/plugins/010-check_17854520.sh
inflating: dbserver_patch_5.180228.2/plugins/020-check_22468216.sh
inflating: dbserver_patch_5.180228.2/plugins/040-check_22896791.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_bash
inflating: dbserver_patch_5.180228.2/plugins/050-check_22651315.sh
inflating: dbserver_patch_5.180228.2/plugins/005-check_22909764.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_perl
inflating: dbserver_patch_5.180228.2/plugins/030-check_24625612.sh
inflating: dbserver_patch_5.180228.2/patchmgr_functions
inflating: dbserver_patch_5.180228.2/exadata.img.hw
inflating: dbserver_patch_5.180228.2/libxcp.so.1
inflating: dbserver_patch_5.180228.2/imageLogger
inflating: dbserver_patch_5.180228.2/ExaXMLNode.pm
inflating: dbserver_patch_5.180228.2/fwverify
```

- f. In the directory that contains the `patchmgr` utility, create the `dbserver_group` file, which contains the list of compute nodes to update. Include the nodes listed after running the `olsnodes` command in step 1, except for the driving system. In this example, `dbserver_group` only contains `node2`.

CHAPTER 11 Database

```
[root@node1 patch]# cd /root/patch/dbserver_patch_5.180228
[root@node1 dbserver_patch_5.180228]# cat dbs_group
node2
```

3. Run a patching precheck operation.

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -precheck -yum_repo yum-
repository -target_version target-version -nomodify_at_prereq
```



Important

Run the precheck operation with the `-nomodify_at_prereq` option to prevent any changes to the system that could impact the backup you take in the next step. Otherwise, the backup might not be able to roll the system back to its original state, should it be necessary.

The output should look similar to the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -precheck -yum_repo
http://yum.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_
version 18.1.4.0.0.180125.3 -nomodify_at_prereq

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****
2018-02-28 21:22:45 +0000      :Working: DO: Initiate precheck on 1 node(s)
```

CHAPTER 11 Database

```
2018-02-28 21:24:57 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:26:15 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:26:47 +0000      :Working: DO: dbnodeupdate.sh running a precheck on node(s).
2018-02-28 21:28:23 +0000      :SUCCESS: DONE: Initiate precheck on node(s).
```

4. Back up the current system.

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -backup -yum_repo yum-
repository -target_version target-version -allow_active_network_mounts
```



Important

Ensure that you take the backup at this point, before any modifications are made to the system.

The output should look similar to the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -backup -yum_repo
http://yum.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_
version 18.1.4.0.0.180125.3 -allow_active_network_mounts

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on 1 node(s).
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on node(s)
2018-02-28 21:29:01 +0000      :Working: DO: Check free space and verify SSH equivalence for
```

```

the root user to node2
2018-02-28 21:30:18 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:30:51 +0000      :Working: DO: dbnodeupdate.sh running a backup on node(s).
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on node(s).
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on 1 node(s).

```

5. Remove all custom RPMs from the target compute nodes. Custom RPMs are reported in precheck results. They include RPMs that were manually installed after the system was provisioned.



Note

- If you are updating the system from version 12.1.2.3.4.170111, and the precheck results include `krb5-workstation-1.10.3-57.el6.x86_64`, remove it. This item is considered a custom RPM for this version.
- Do **not** remove `exadata-sun-vm-computenode-exact` **or** `oracle-ofed-release-guest`. These two RPMs are handled automatically during the update process.

6. Perform the update. Use the `nohup` command to ensure that the update process is not interrupted.

```

[root@node1 dbserver_patch_5.180228]# nohup ./patchmgr -dbnodes dbs_group -upgrade -nobackup -
yum_repo yum-repository -target_version target-version -allow_active_network_mounts &

```

The output should look similar to the following example:

```

[root@node1 dbserver_patch_5.180228]# nohup ./patchmgr -dbnodes dbs_group -upgrade -nobackup -
yum_repo http://yum.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -
target_version 18.1.4.0.0.180125.3 -allow_active_network_mounts &

```

CHAPTER 11 Database

```
*****
*****
NOTE   patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
NOTE   Database nodes will reboot during the update process.
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****
2018-02-28 21:36:26 +0000      :Working: DO: Initiate prepare steps on node(s).
2018-02-28 21:36:26 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:37:44 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:38:43 +0000      :SUCCESS: DONE: Initiate prepare steps on node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on 1 node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on node(s)
2018-02-28 21:38:49 +0000      :Working: DO: Get information about any required OS upgrades
from node(s).
2018-02-28 21:38:59 +0000      :SUCCESS: DONE: Get information about any required OS upgrades
from node(s).
2018-02-28 21:38:59 +0000      :Working: DO: dbnodeupdate.sh running an update step on all
nodes.
2018-02-28 21:48:41 +0000      :INFO    : node2 is ready to reboot.
2018-02-28 21:48:41 +0000      :SUCCESS: DONE: dbnodeupdate.sh running an update step on all
nodes.
2018-02-28 21:48:41 +0000      :Working: DO: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :SUCCESS: DONE: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :Working: DO: Waiting to ensure node2 is down before reboot.
2018-02-28 21:56:18 +0000      :Working: DO: Initiate prepare steps on node(s).
2018-02-28 21:56:19 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:57:37 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:57:42 +0000      :SEEMS ALREADY UP TO DATE: node2
```

CHAPTER 11 Database

```
2018-02-28 21:57:43 +0000 :SUCCESS: DONE: Initiate update on node(s)
```

7. After the update operation completes, verify the version of the Exadata software on the compute node that was updated.

```
[root@node2 ~]# imageinfo -ver  
18.1.4.0.0.180125.3
```

8. Repeat steps 2 through 7 of this procedure using the updated compute node as the driving system to update the remaining compute node. In this example update, you would now use `node2` to update `node1`.
9. As `root` On each compute node, run the `uptrack-install` command to install the available `ksplice` updates.

```
[root@node1 ~]# uptrack-install --all -y
```

```
[root@node2 ~]# uptrack-install --all -y
```

INSTALLING ADDITIONAL OS PACKAGES

You are permitted to install and update OS packages on Exadata Cloud at Customer as long as you do not modify the kernel or InfiniBand-specific packages. However, Oracle technical support, including installation, testing, certification and error resolution, does not apply to any non-Oracle software that you install.

Also, adding or updating packages may introduce problems when applying an Exadata software update because the additional software may add new dependencies that may interrupt an Exadata update. For this reason, minimal customization is recommended.

If you install additional packages, it is recommended that you have scripts to automate the removal and reinstallation of those packages. After an Exadata update, verify that the additional packages are still compatible and are still needed, before reinstalling them.

See also [Installing, Updating, and Managing Non-Exadata Software](#) in the *Oracle Exadata Database Machine Maintenance Guide*.

Cloud Tooling Updates

You are responsible for updating the cloud-specific tooling included on the Exadata Cloud at Customer compute nodes. You can update the cloud-specific tooling by downloading and

applying a software package containing the updated tools as described in this section.

To check the installed cloud tooling release and check for updates

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Use the following command to display information about the installed cloud tooling and to list the available updates:

```
# dbaascli patch tools list
```

The command output displays:

- The version of the cloud tooling that is installed on the compute node.
- The list of available updates.
- Notification of the cloud tooling version that is installed on the other compute nodes in the VM cluster.

To update the cloud tooling

1. Connect to a compute node as the `opc` user and start a command shell as the `root` user.

For detailed instructions, see [Connecting to an Exadata Cloud at Customer System](#).

2. Download and apply the cloud tooling update:

- To update to the latest available cloud tooling, use the following command:

```
# dbaascli patch tools apply --patchid LATEST
```

- To update to a specific cloud tooling release, use the following command:

```
# dbaascli patch tools apply --patchid patchid
```

In the preceding command, *patchid* is a cloud tooling patch identifier, as reported in the output of the `dbaascli patch tools list` command.

The cloud tooling update is applied to all nodes in the VM cluster.

Using EM Express and Database Control on Exadata Cloud at Customer

This topic explains how to access Enterprise Manager Database Express (EM Express) and Enterprise Manager Database Control (Database Control), which are web-based tools for monitoring and managing Oracle Database.

Accessing Enterprise Manager Database Express

Enterprise Manager Database Express is available on Exadata Cloud at Customer databases created using Oracle Database 12c Release 1 (12.1) or later.

How you access EM Express depends on whether you want to manage a CDB or PDB:

- **To manage the CDB.** When a database is created, Exadata Cloud at Customer automatically sets a port for EM Express access to the CDB. You do not need to perform any manual configuration steps. Each database is allocated a unique port number. The allocations use ports in a range starting with 5500, 5501, 5502, and so on.
- **To manage a PDB.** With Oracle Database 12c Release 2 or later, EM Express can be configured to access the CDB and all PDBs on a single port, which is known as the global port. The global port lets you use EM Express to connect to all of the PDBs in the CDB using the HTTPS port for the CDB. You do not need to perform any manual configuration steps. Each database is allocated a unique port number. The allocations use ports in a range starting with 5500, 5501, 5502, and so on.

For a version 12.1 database, you must manually set a port for each PDB you want to manage using EM Express.

To confirm the port that is in use for a specific database, connect to the database as a database administrator and execute the query shown in the following example:

```
SQL> select dbms_xdb_config.getHttpsPort() from dual;
```

```
DBMS_XDB_CONFIG.GETHTTSPORT()
```

SETTING THE PORT FOR EM EXPRESS TO MANAGE A PDB (ORACLE DATABASE 12.1 ONLY)

In Oracle Database 12c Release 1, a unique HTTPS port must be configured for the root container (CDB) and each PDB that you manage using EM Express.

To configure a HTTPS port so that you can manage a PDB with EM Express:

1. Invoke SQL*Plus and log in to the PDB as the `sys` user with `sysdba` privileges.
2. Execute the `DBMS_XDB_CONFIG.SETHTTPS` procedure.

```
SQL> exec dbms_xdb_config.sethttpsport(port-number)
```

ACCESSING EM EXPRESS

You can access EM Express by directing your browser to the URL:

```
https://node-ip-address:port/em
```

where *node-ip-address* is the client network IP address of the compute node hosting EM Express, and *port* is the EM Express port used by the database.

If your browser displays a warning that your connection is not secure or not private, use the browser's advanced option to ignore the warning and continue. You get this warning because Exadata Cloud at Customer uses a self-signed certificate to provide HTTPS (secure HTTP) connectivity, and such certificates are considered suspicious by many web browsers.

Accessing Enterprise Manager 11g Database Control

Enterprise Manager 11g Database Control is available on Exadata Cloud at Customer databases created using Oracle Database 11g Release 2. Each database is allocated a unique port number in a range starting with 1158, 1159, 1160, and so on.

You can confirm the Database Control port for a database by searching for `REPOSITORY_URL` in the `$ORACLE_HOME/host_sid/sysman/config/emd.properties` file. In the preceding file name, *host* is the host name of the compute node hosting Database Control, and *sid* is the Oracle Database system identifier (SID).

You can access Database Control by directing your browser to the URL:

CHAPTER 11 Database

```
https://node-ip-address:port/em
```

where *node-ip-address* is the client network IP address of the compute node hosting Database Control, and *port* is the Database Control port used by the database.

If your browser displays a warning that your connection is not secure or not private, use the browser's advanced option to ignore the warning and continue. You get this warning because Exadata Cloud at Customer uses a self-signed certificate to provide HTTPS (secure HTTP) connectivity, and such certificates are considered suspicious by many web browsers.

Using the dbaascli Utility on Exadata Cloud at Customer

This topic describes the facilities provided by the `dbaascli` utility on Exadata Cloud at Customer.

You can use the `dbaascli` utility to perform various database life-cycle and administration operations on Exadata Cloud at Customer such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more. The capabilities of the `dbaascli` utility are in addition to, and separate from, the Oracle Cloud Infrastructure Console, API, or CLI.

To use the utility, you must be connected to an Exadata Cloud at Customer compute node. See [Connecting to a Compute Node with SSH](#).

Many `dbaascli` commands can be run as the `oracle` user, but some commands require `root` administrator privileges. The following outlines the specific requirements for each command.

Command Reference

`dbaascli database bounce`

You can use this command to shut down and restart the specified database:

```
$ dbaascli database bounce --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to bounce.

The command performs a database shutdown in immediate mode. The database is then restarted and opened. In Oracle Database 12c, or later, all of the PDBs are also opened.

Run the command as the `oracle` user.

dbaascli database changepassword

You can use this command to change the password of a database user:

```
# dbaascli database changepassword --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to act on.

Enter the database user name and new password when prompted.

Run the command as the `root` user.

dbaascli database move

You can use this command to move a database to another Oracle Home directory location:

```
# dbaascli database move --dbname dbname --ohome oraclehome
```

In the command:

- *dbname* specifies the name of the database that you want to move.
- *oraclehome* specifies the path to an existing Oracle Home directory location, which you want the specified database to use.

Before performing a move operation, ensure that all of the database instances associated with the database are up and running.

Run the command as the `root` user.

dbaascli database start

You can use this command to start the specified database:

CHAPTER 11 Database

```
$ dbaascli database start --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to start.

The command starts and opens the database. In Oracle Database 12c, or later, all of the PDBs are also opened.

Run the command as the `oracle` user.

dbaascli database status

You can use this command to check the status of the specified database:

```
$ dbaascli database status --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to check.

Output from the command includes the open mode of the database, the software release and edition of the database, and release version of other software components.

Run the command as the `oracle` user.

dbaascli database stop

You can use this command to stop the specified database:

```
$ dbaascli database stop --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to stop.

The command performs a database shutdown in immediate mode. No new connections or new transactions are permitted. Active transactions are rolled back and all connected users are disconnected.

Run the command as the `oracle` user.

dbaascli database update

You can use this command to perform the following database configuration changes.

Run the command as the `root` user.

- To modify the globally unique database name (`DB_UNIQUE_NAME`), run the following command:

```
# dbaascli database update --dbname dbname --db_unique_name dbname_uniquename [--precheck]
```

In the command:

- *dbname* specifies the name of the database that you want to move.
- *uniquename* specifies the user configurable portion of the new globally unique database name.

The command modifies the `DB_UNIQUE_NAME` database parameter and related configuration entries that reference it, including entries in the Oracle Cluster Registry (OCR) and database server parameter file (SPFILE). File locations that reference the globally unique database name are also updated, including the location of the data files and keystore.

The value for the `--db_unique_name` option must commence with the *dbname* value followed immediately by an underscore character. If this convention is not observed, then the command fails with an error .

Before performing the update operation, you can use the `--precheck` option to run a series of prerequisite checks to ensure that the update can proceed. No changes are made when using the `--precheck` option.

- To reconfigure the online redo log files, run the following command:

```
# dbaascli database update --dbname dbname --redosize redosize [--groups numgroups] [--precheck]
```

In the command:

- *dbname* specifies the name of the database that you want to move.
- *redosize* specifies the size of each online redo log file in megabytes. The valid range is between 1000m and 16000m.
- *numgroups* optionally specifies the number of online redo log groups to create. The default value is 4.

Before performing the update operation, you can use the `--precheck` option to run a series of prerequisite checks to ensure that the update can proceed. No changes are made when using the `--precheck` option.

dbaascli dbhome info

You can use this command to view information about Oracle Home directory locations:

```
# dbaascli dbhome info
```

When prompted:

- Press `Enter` to view information about all Oracle Homes registered in the VM cluster.
- Specify an Oracle Home name to view information about that Oracle Home.

Run the command as the `root` user.

dbaascli dbhome purge

You can use this command to delete an unused Oracle Home directory location:

```
# dbaascli dbhome purge
```

When prompted, enter:

- `1` if you want to specify the Oracle Home name for the location being purged.
- `2` if you want to specify the Oracle Home directory path for the location being purged.

When next prompted, enter the Oracle Home name or directory path for the location being purged.

If your entries are valid and the Oracle Home is not associated with a database, then the Oracle binaries are removed from the Oracle Home directory location and the associated metadata is removed from the system.

Run the command as the `root` user.

dbaascli dbimage list

You can use this command to display information about Oracle Database software images that are downloaded to your Exadata Cloud at Customer environment:

```
# dbaascli dbimage list
```

The command displays a list of software images that are downloaded to your Exadata Cloud at Customer environment, including version and bundle patch information.

Run the command as the `root` user.

dbaascli listener bounce

You can use this command to stop and restart the Oracle Net listener that is associated with the specified database:

```
$ dbaascli listener bounce --dbname dbname
```

In the command, *dbname* specifies the name of the database whose listener you want to bounce.

Run the command as the `oracle` user.

dbaascli listener start

You can use this command to start the Oracle Net listener that is associated with the specified database:

```
$ dbaascli listener start --dbname dbname
```

In the command, *dbname* specifies the name of the database whose listener you want to start.

Run the command as the `oracle` user.

dbaascli listener status

You can use this command to check the status of the Oracle Net listener that is associated with

the specified database:

```
$ dbaascli listener start --dbname dbname
```

In the command, *dbname* specifies the name of the database whose listener you want to check.

The command displays status information about the listener, including a summary of the listener configuration settings, listening protocol addresses, and a summary of services that are registered with the listener.

Run the command as the `oracle` user.

dbaascli listener stop

You can use this command to stop the Oracle Net listener that is associated with the specified database:

```
$ dbaascli listener stop --dbname dbname
```

In the command, *dbname* specifies the name of the database whose listener you want to stop.

Run the command as the `oracle` user.

dbaascli patch db apply

You can use this command to apply an Oracle Database or Oracle Grid Infrastructure patch.

- To apply a patch to a specific instance, use the following command:

```
# dbaascli patch db apply --patchid patchid --instance1 hostname:oracle_home [--dbnames dbname  
[, dbname2 ...]] [--run_datasp 1]
```

- To apply a patch by specifying only database names, use the following command:

```
# dbaascli patch db apply --patchid patchid --dbnames dbname [, dbname2 ...] [--run_datasp 1] [-  
alldbs]
```

In the preceding commands:

- `patchid` identifies the patch to be applied.
- `--instance1` specifies a compute node and Oracle Home directory that is subject to the patching operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.
If you use this argument to specify a shared Oracle Home directory and you do not specify the `--dbnames` argument, then all of the databases that share the specified Oracle Home are patched. After the operation, the Oracle Home directory location remains unchanged; however, the patch level information embedded in the Oracle Home name is adjusted to reflect the patching operation.
- `--dbnames` specifies the database names for the databases that are the target of the patching operation.
If you use this argument to patch a database that uses a shared Oracle Home and you do not specify the `-alldbs` option, then a new Oracle Home containing the patched Oracle Database binaries is created and the database is moved to the new Oracle Home.
- `-alldbs` patches all of the databases that share the same Oracle Database binaries (Oracle Home) as the databases specified in the `--dbnames` argument.
After the operation, the Oracle Home directory location remains unchanged; however, the patch level information embedded in the Oracle Home name is adjusted to reflect the patching operation.
- `--run_dataSQL 1` instructs the command to execute patch-related SQL commands.



Note

- Patch-related SQL should only be executed after all of the compute nodes are patched. Take care not to specify this argument if you are patching a node and further nodes remain to be patched.
- This argument can only be specified along with a patching operation on a compute node. If you have patched all of your nodes and you did not specify this argument, you need to manually execute the SQL commands associated with the patch, which typically involves running the `catbundle.sql` script for Oracle Database 11g or the `datapatch` utility for Oracle Database 12c, or later. Refer to the patch documentation for full details.

Run the command as the `root` user.

`dbaascli patch db list`

You can use this command to check whether any Oracle Database or Oracle Grid Infrastructure patches are available:

```
# dbaascli patch db list --oh hostname:oracle_home
```

In the preceding command, `--oh` specifies a compute node and Oracle Home directory for which you want to list the available patches. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.



Note

The list of available patches is determined by interrogating the database to establish the patches that have already been applied. When a patch is applied, the corresponding database entry is made as part of the SQL patching operation, which is executed at the end of the patch workflow. Therefore, the list of available patches may include partially applied patches along with patches that are currently being applied.

Run the command as the `root` user.

`dbaascli patch db prereq`

You can use this command to check the prerequisites for an Oracle Database or Oracle Grid Infrastructure patch.

- To check patch prerequisites on a specific instance, use the following command:

```
# dbaascli patch db prereq --patchid patchid --instance1 hostname:oracle_home [--dbnames dbname [, dbname2 ...]]
```

- To check patch prerequisites by specifying only database names, use the following command:

```
# dbaascli patch db prereq --patchid patchid --dbnames dbname [, dbname2 ...] [-all dbs]
```

In the preceding commands:

- *patchid* identifies the patch to be pre-checked.
- `--instance1` specifies a compute node and Oracle Home directory that is subject to the pre-check operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.

- `--dbnames` specifies the database names for the databases that are the target of the pre-check operation.
- `-alldbs` specifies that you want to pre-check all of the databases that share the same Oracle Database binaries (Oracle Home) as the specified databases.

Run the command as the `root` user.

`dbaascli patch db switchback`

You can use this command to roll back an Oracle Database or Oracle Grid Infrastructure patch.

- To roll back a patch on specific instances, use the following command:

```
# dbaascli patch db switchback --patchid patchid --instance1 hostname:oracle_home [--dbnames dbname[,dbname2 ...]] [--run_datasp1 1]
```

- To roll back a patch by specifying only database names, use the following command:

```
# dbaascli patch db switchback --patchid patchid --dbnames dbname[,dbname2 ...] [--run_datasp1 1] [-alldbs]
```

In the preceding commands:

- *patchid* identifies the patch to be rolled back.
- `--instance1` specifies a compute node and Oracle Home directory that is subject to the rollback operation. In this context, an Oracle Home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory. If you use this argument to specify a shared Oracle Home directory and you do not specify the `--dbnames` argument, then all of the databases that share the specified Oracle Home are rolled back.
- `--dbnames` specifies the database names for the databases that are the target of the rollback operation.

- `-alldbs` specifies that you want to roll back all of the databases that share the same Oracle Database binaries (Oracle Home) as the databases specified in the `--dbnames` argument.
- `--run_datasql 1` instructs the command to execute rollback-related SQL commands.



Note

- Rollback-related SQL should only be executed after all of the compute nodes are rolled back. Take care not to specify this argument if you are rolling back a node and further nodes remain to be rolled back.
- This argument can only be specified along with a rollback operation on a compute node. Therefore, if you have rolled back all of your nodes and you did not specify this argument, you need to manually execute the SQL commands associated with the rollback operation. Refer to the patch documentation for full details.

Run the command as the `root` user.

`dbaascli patch tools apply`

You can use this command to download and apply a cloud tooling update:

- To update to the latest available cloud tooling, use the following command:

```
# dbaascli patch tools apply --patchid LATEST
```

- To update to a specific cloud tooling release, use the following command:

CHAPTER 11 Database

```
# dbaascli patch tools apply --patchid patchid
```

In the preceding command, *patchid* is a cloud tooling patch identifier, as reported in the output of the [dbaascli patch db list](#) command.

The cloud tooling update is applied to all nodes in the VM cluster.

Run the command as the `root` user.

dbaascli patch tools list

You can use this command to display information about the installed cloud tooling and to list the available updates:

```
# dbaascli patch tools list
```

The command output displays:

- The version of the cloud tooling that is installed on the compute node.
- The list of available updates.
- Notification of the cloud tooling version that is installed on the other compute nodes in the VM cluster.

Run the command as the `root` user.

dbaascli pdb checkdb

You can use this command to display information about a container database (CDB):

```
$ dbaascli pdb checkdb --dbname dbname
```

In the command, *dbname* specifies the name of the CDB for which you want display information.

The information returned by this command includes the number of instances and the CPU count that are associated with the CDB.

Run the command as the `oracle` user.

dbaascli pdb checknode

You can use this command to display information about pluggable databases (PDBs) that are associated with a specific container database (CDB) and a specific compute node:

```
$ dbaascli pdb checknode --node nodenum --dbname dbname
```

In the command:

- *nodenum* specifies the node number for a compute node in the Exadata Cloud at Customer environment. You can display a list of compute nodes and corresponding node numbers by using the `olsnodes` command.
- *dbname* specifies the name of the container database that hosts the PDB.

The command displays status information for all PDBs that are associated with the specified compute node and CDB, including the open mode for each PDB.

Run the command as the `oracle` user.

dbaascli pdb checkpdb

You can use this command to display information about a pluggable database (PDB):

```
$ dbaascli pdb checkpdb --pdbname pdbname --dbname dbname
```

In the command:

- *pdbname* specifies the name of the PDB that you want to check.
- *dbname* specifies the name of the container database that hosts the PDB.

The command displays status information for the specified PDB, including the open mode and restricted status.

Run the command as the `oracle` user.

dbaascli pdb close

You can use this command to close a pluggable database (PDB):

CHAPTER 11 Database

```
$ dbaascli pdb close --pdbname pdbname --dbname dbname
```

In the command:

- *pdbname* specifies the name of the PDB that you want to close.
- *dbname* specifies the name of the container database that hosts the PDB.

Upon successful completion, the PDB is closed on all of the container database instances.

Run the command as the `oracle` user.

dbaascli pdb connect_info

You can use this command to retrieve network connection information for a pluggable database (PDB):

```
$ dbaascli pdb connect_info --pdbname pdbname --dbname dbname
```

In the command:

- *pdbname* specifies the name of the PDB for which you want to retrieve connection information.
- *dbname* specifies the name of the container database that hosts the PDB.

The command outputs a zip file that contains `tnsnames.ora`, `sqlnet.ora`, and `ojdbc8` properties for the PDB.

Run the command as the `oracle` user.

dbaascli pdb connect_string

You can use this command to display Oracle Net connect string information for a pluggable database (PDB):

```
$ dbaascli pdb connect_string --pdbname pdbname --dbname dbname
```

In the command:

CHAPTER 11 Database

- *pdbname* specifies the name of the PDB for which you want to display connect string information.
- *dbname* specifies the name of the container database that hosts the PDB.

Run the command as the `oracle` user.

dbaascli pdb create

You can use this command to create a new pluggable database (PDB):

```
$ dbaascli pdb create --pdbname pdbname --dbname dbname [--maxsize maxsize] [--maxcpu maxcpu]
```

In the command:

- *pdbname* specifies the name of the new PDB that you want to create.
- *dbname* specifies the name of the container database that hosts the new PDB.
- *maxsize* optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE` PDB storage clause in the `CREATE PLUGGABLE DATABASE SQL` command. You can impose a limit by specifying an integer followed by a size unit (`K`, `M`, `G`, or `T`), or you can specify `UNLIMITED` to explicitly enforce no limit.
- *maxcpu* optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB.

During the PDB creation process, you are prompted to specify the administration password for the new PDB.

Run the command as the `oracle` user.

dbaascli pdb delete

You can use this command to delete a pluggable database (PDB):

```
$ dbaascli pdb delete --pdbname pdbname --dbname dbname
```

In the command:

- *pdbname* specifies the name of the PDB that you want to delete.
- *dbname* specifies the name of the container database that hosts the PDB.

Run the command as the `oracle` user.

`dbaascli pdb info`

You can use this command to display more detailed information about a pluggable database (PDB):

```
$ dbaascli pdb info [--pdbname pdbname] --dbname dbname [--detailed]
```

In the command:

- *pdbname* optionally specifies the name of the PDB for which you want to display information. If this option is not specified, then the command displays information about all of the PDBs in the specified container database.
- *dbname* specifies the name of the container database that hosts the PDB.

The command displays information such as the CPU count and storage usage that is associated with a PDB. You can add the optional `--detailed` argument to display extra information, including the list of compute nodes where a PDB is open in read/write mode.

Run the command as the `oracle` user.

`dbaascli pdb local_clone`

You can use this command to create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB):

```
$ dbaascli pdb local_clone --pdbname sourcepdbname --target_pdbname targetpdbname --dbname dbname
```

In the command:

- *sourcepdbname* specifies the name of the PDB that you want to clone.
- *targetpdbname* specifies the name of the new PDB that you want to create.
- *dbname* specifies the name of the container database that hosts the PDBs.

The newly cloned PDB inherits administration passwords from the source PDB.

Run the command as the `oracle` user.

dbaascli pdb open

You can use this command to open a pluggable database (PDB):

```
$ dbaascli pdb open --pdbname pdbname --dbname dbname
```

In the command:

- *pdbname* specifies the name of the PDB that you want to open.
- *dbname* specifies the name of the container database that hosts the PDB.

Upon successful completion, the PDB is opened on all of the container database instances.

Run the command as the `oracle` user.

dbaascli pdb remote_clone

You can use this command to create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB):

```
$ dbaascli pdb remote_clone --pdbname sourcepdbname --source_db sourcedbname --source_db_scan  
sourcedbscan --dbname dbname
```

In the command:

- *sourcepdbname* specifies the name of the source PDB that you want to clone.
- *sourcedbname* specifies the unique name of the CDB that hosts the source PDB.

CHAPTER 11 Database

- *sourcedbscan* specifies the Single Client Access Name (SCAN) that is used to connect to the source database.
- *dbname* specifies the name of the CDB that hosts the newly cloned PDB.

When promoted, you must supply the `sys` user password for the source PDB.

The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: *dbname_sourcepdbname*.

Run the command as the `oracle` user.

dbaascli pdb rename

You can use this command to change the name of a pluggable database (PDB):

```
$ dbaascli pdb rename --pdbname oldname --newname newname --dbname dbname
```

In the command:

- *oldname* specifies the old name of the PDB that you want to rename.
- *newname* specifies the new name of the PDB that you want to rename.
- *dbname* specifies the name of the container database that hosts the PDB.

Run the command as the `oracle` user.

dbaascli pdb resize

You can use this command to modify the size limits for a pluggable database (PDB):

```
$ dbaascli pdb resize --pdbname pdbname --dbname dbname [--maxsize maxsize] [--maxcpu maxcpu]
```

In the command:

- *pdbname* specifies the name of the new PDB that you want to modify.
- *dbname* specifies the name of the container database that hosts the PDB.

- *maxsize* optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE` PDB storage clause in the `CREATE PLUGGABLE DATABASE SQL` command. You can impose a limit by specifying an integer followed by a size unit (`K`, `M`, `G`, or `T`), or you can specify `UNLIMITED` to explicitly enforce no limit.
- *maxcpu* optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB.

When you run the command, you must specify at least one of optional attributes, `--maxsize` or `--maxcpu`. You can specify both optional attributes in a single command.

Run the command as the `oracle` user.

`dbaascli pdb start_service`

You can use this command to start the Oracle Database service that is associated with a pluggable database (PDB):

```
$ dbaascli pdb start_service --pdbname pdname --dbname dbname
```

In the command:

- *pdname* specifies the name of the PDB that is associated with the database service that you want to start.
- *dbname* specifies the name of the container database that hosts the PDB.

Run the command as the `oracle` user.

`dbaascli tde rotate masterkey`

You can use this command to change (rotate) the master encryption key for the specified database:

```
$ dbaascli tde rotate masterkey --dbname dbname
```

CHAPTER 11 Database

In the command, *dbname* specifies the name of the database that you want to act on.

Enter the keystore password when prompted. The keystore password is initially set to the administration password that you specified when you created the database.

Run the command as the `oracle` user.

`dbaascli tde status`

You can use this command to display information about the keystore for the specified database:

```
$ dbaascli tde status --dbname dbname
```

In the command, *dbname* specifies the name of the database that you want to check.

Output from the command includes the type of keystore and the status of the keystore.

Run the command as the `oracle` user.

Monitoring and Managing Exadata Storage Servers on Exadata Cloud at Customer

This topic explains how you can use the ExaCLI utility on Exadata Cloud at Customer to perform monitoring and management functions on the Exadata Storage Servers.

ExaCLI is a command-line administration tool that runs on Exadata compute nodes and storage server nodes and enables you to manage other nodes remotely. See [Using the ExaCLI Utility](#).

On Exadata Cloud at Customer, ExaCLI is configured on the compute nodes to support essential Exadata Storage Server monitoring and management functions, which are performed without direct administrative access to the Exadata Storage Servers. The ExaCLI utility is located on each Exadata Cloud at Customer compute node at `/usr/local/sbin/exacli`. To use the ExaCLI utility you require:

- The ExaCLI command, which specifies the operation that you want to perform. For Exadata Storage Server targets, ExaCLI supports the same command syntax as CellCLI. See [CellCLI Command Reference](#). However, ExaCLI on Exadata Cloud at Customer supports only a specific set of monitoring and management operations. For a full list, see [Supported ExaCLI Commands on Exadata Cloud at Customer](#).

- The username and password to connect to the Exadata Storage Server. On Exadata Cloud at Customer, the preconfigured user is `cloud_user_`*clustername*, where *clustername* is the name of the Oracle Grid Infrastructure cluster that is being used. You can determine the Oracle Grid Infrastructure cluster name by running the following command as the `grid` user on any cluster node:

```
$ crsctl get cluster name
```

The password for `cloud_user_`*clustername* is initially set to a random value, which you can view by running the following command as the `opc` user on any cluster node:

```
$ /opt/exacloud/get_cs_data.py
```

- The network address of the Exadata Storage Server that is the target of the operation. You can determine the IP addresses for your Exadata Storage Servers by examining the `/etc/oracle/cell/network-config/cellip.ora` file on any Exadata Cloud at Customer compute node.

Supported ExaCLI Commands on Exadata Cloud at Customer

You can use ExaCLI on Exadata Cloud at Customer to view information about the following Exadata Storage Server objects by using the [LIST](#) command:

- `activerequest`
- `alertdefinition`
- `alerthistory`
- `cell`
- `celldisk`
- `database`
- `flashcache`

CHAPTER 11 Database

- flashcachecontent
- flashlog
- griddisk
- ibport
- iormprofile
- lun
- metriccurrent
- metricdefinition
- metrichistory
- offloadgroup
- physicaldisk
- pluggabledatabase

You can use ExaCLI on Exadata Cloud at Customer to act on the following Exadata Storage Server objects by using the [CREATE](#), [ALTER](#), [DROP](#), and [LIST](#) commands:

- diagpack
- iormplan
- quarantine

You can use ExaCLI on Exadata Cloud at Customer to alter the following Exadata Storage Server cell attributes by using the [ALTER CELL](#) command:

- dbPerfDataSuppress
- diagHistoryDays
- events
- metricCollection
- metricHistoryDays
- offloadGroupEvents
- securityCert

- `securityPrivKey`
- `securityPrivKeyPW`
- `securityPubKey`
- `traceLevel`

Exadata DB Systems

Exadata DB systems allow you to leverage the power of Exadata within the Oracle Cloud Infrastructure. An Exadata DB system consists of a base system, quarter rack, half rack, or full rack of compute nodes and storage servers, tied together by a high-speed, low-latency InfiniBand network and intelligent Exadata software. You can configure automatic backups, optimize for different workloads, and scale up the system to meet increased demands.



Note

Exadata DB systems launched on or after March 14, 2019 run Oracle Linux 7 (OL7). Previously launched systems are running Oracle Linux 6 (OL6). See [OS Updates](#) for important information about updating existing Exadata DB system operating systems.

Supported Database Edition and Versions

Exadata DB systems require Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all the database enterprise management packs and all the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (RAC).

Exadata DB systems support the following software releases:

- Oracle Database 19c (19.0)
- Oracle Database 18c (18.0)

- Oracle Database 12c Release 2 (12.2)
- Oracle Database 12c Release 1 (12.1)
- Oracle Database 11g Release 2 (11.2)



Note

If you plan to run Oracle Database 19c on your Exadata DB system, you must specify version 19c when you create the DB system. Earlier database versions are supported on a 19c Exadata DB system and can be created at anytime. Exadata DB systems created with earlier Oracle Database versions will not automatically support Oracle Database 19c. The DB system must be upgraded manually.

Subscription Types

The only subscription type available for Exadata DB systems is the Monthly Flex purchase model under Universal Credit Pricing. See <https://www.oracle.com/cloud/bring-your-own-license/faq/universal-credit-pricing.html> for more information.

Metering Frequency

For each Exadata DB system you provision, you are billed for the infrastructure for the first month, and then by the hour after that. Each OCPU you add to the system is billed by the hour from the time you add it.

Scaling an Exadata DB System

Two kinds of scaling operations are supported for an Exadata DB system:

- Scaling within an Exadata DB system lets you modify compute node processing power within the system.
- Scaling across Exadata DB system configurations lets you move to a different configuration, for example, from a quarter rack to a half rack.

Scaling Within an Exadata System

If an Exadata DB system requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system. For a base system or a quarter rack, you can scale in multiples of 2 across the 2 database compute nodes. For a half rack, you can scale in multiples of 4 across the 4 database compute nodes. For a full rack, you can scale in multiples of 8 across the 8 database compute nodes.

For a non-metered Exadata DB system, you can temporarily modify the compute node processing power (bursting) or add compute node processing power on a more permanent basis. For a metered Exadata DB system, you can simply modify the number of enabled CPU cores.

You can provision an X7 Exadata DB system or a base system with zero CPU cores, or scale the DB system down to zero cores after you provision it. With zero cores, you are billed only for the infrastructure until you scale up the system. For detailed information about pricing, see <https://www.oracle.com/database/exadata-cloud-service-pricing.html>.

For information on CPU cores per configuration, see [System Configuration](#). To learn how to scale a system, see [To scale an Exadata DB system](#).

Scaling Across Exadata DB System Configurations

Scaling across Exadata DB system configurations enables you to move to a different system configuration. This is useful when a database deployment requires:

- Processing power that is beyond the capacity of the current system configuration.
- Storage capacity that is beyond the capacity of the current system configuration.
- A performance boost that can be delivered by increasing the number of available compute nodes.

- A performance boost that can be delivered by increasing the number of available Exadata Storage Servers.

Scaling from a base system or a quarter rack to a half rack, or from a half rack to a full rack, requires that the data associated with your database deployment is backed up and restored on a different Exadata DB system, which requires planning and coordination between you and Oracle. To start the process, submit a service request to Oracle.

System Configuration

Exadata DB systems are offered in base system, quarter rack, half rack or full rack configurations, and each configuration consists of compute nodes and storage servers. The compute nodes are each configured with a Virtual Machine (VM). You have root privilege for the compute node VMs, so you can load and run additional software on them. However, you do not have administrative access to the Exadata infrastructure components, including the physical compute node hardware, network switches, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, or the Exadata Storage Servers, which are all administered by Oracle.

You have full administrative privileges for your databases, and you can connect to your databases by using Oracle Net Services from outside the Oracle Cloud Infrastructure. You are responsible for database administration tasks such as creating tablespaces and managing database users. You can also customize the default automated maintenance set up, and you control the recovery process in the event of a database failure.

The subsections that follow provide the details for each shape's configuration.

Exadata X7 Shapes

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter2.92	Exadata.Half2.184	Exadata.Full2.368
Number of Compute Nodes	2	4	8

CHAPTER 11 Database

Property	Quarter Rack	Half Rack	Full Rack
Total Minimum Number of Enabled CPU Cores	0	0	0
Total Maximum Number of Enabled CPU Cores	92	184	368
Total RAM Capacity	1440 GB	2880 GB	5760 GB
Number of Exadata Storage Servers	3	6	12
Total Raw Flash Storage Capacity	76.8 TB	153.6 TB	307.2 TB
Total Usable Storage Capacity	106 TB	212 TB	424 TB

Exadata X7 shapes provide 1 TB of user disk space for database homes.

Exadata X6 Shapes

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter1.84	Exadata.Half1.168	Exadata.Full1.336
Number of Compute Nodes	2	4	8
Total Minimum (Default) Number of Enabled CPU Cores	22	44	88

CHAPTER 11 Database

Property	Quarter Rack	Half Rack	Full Rack
Total Maximum Number of Enabled CPU Cores	84	168	336
Total RAM Capacity	1440 GB	2880 GB	5760 GB
Number of Exadata Storage Servers	3	6	12
Total Raw Flash Storage Capacity	38.4 TB	76.8 TB	153.6 TB
Total Usable Storage Capacity	84 TB	168 TB	336 TB

Exadata X6 shapes provide 200 GB of user disk space for database homes.

Exadata Base System

The system configuration of an Exadata base system is similar to a quarter rack with some differences in capacity.

Property	Value
Shape Name	Exadata.Base.48
Number of Compute Nodes	2
Total Minimum Number of Enabled CPU Cores	0
Total Maximum Number of Enabled CPU Cores	48
Total RAM Capacity	720 GB
Number of Exadata Storage Servers	3

Property	Value
Total Raw Flash Storage Capacity	38.4 TB
Total Usable Storage Capacity	74.8 TB

Storage Configuration

When you launch an Exadata DB system, the storage space inside the Exadata storage servers is configured for use by Oracle Automatic Storage Management (ASM). By default, the following ASM disk groups are created:

- The DATA disk group is intended for the storage of Oracle Database data files.
- The RECO disk group is primarily used for storing the Fast Recovery Area (FRA), which is an area of storage where Oracle Database can create and manage various files related to backup and recovery, such as RMAN backups and archived redo log files.
- The DBFS and ACFS disk groups are system disk groups that support various operational purposes. The DBFS disk group is primarily used to store the shared clusterware files (Oracle Cluster Registry and voting disks), while the ACFS disk groups are primarily used to store Oracle Database binaries. Compared to the DATA and RECO disk groups, the system disk groups are so small that they are typically ignored when discussing the overall storage capacity. You should not store Oracle Database data files or backups inside the system disk groups.

The disk group names contain a short identifier string that is associated with your Exadata Database machine environment. For example, the identifier could be C2, in which case the DATA disk group would be named DATA2, the RECO disk group would be named RECO2, and so on.

In addition, you can create a SPARSE disk group. A SPARSE disk group is required to support Exadata snapshots. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily. Snapshot clones are often used for development, testing, or other purposes that require a transient database.

Note that you cannot change the disk group layout after service creation.

Impact of Configuration Settings on Storage

If you choose to perform database backups to the Exadata storage, or to create a sparse disk group, or to do both, your choices profoundly affect how storage space in the Exadata storage servers is allocated to the ASM and sparse disk groups.

The table that follows shows the approximate percentages of storage allocated for DATA, RECO, and SPARSE disk groups for each possible configuration.

Configuration Settings	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Database backups on Exadata storage: No Sparse disk group: No	80 %	20 %	0 %
Database backups on Exadata storage: Yes Sparse disk group: No	40 %	60 %	0 %
Database backups on Exadata storage: No Sparse disk group: Yes	60 %	20 %	20 %
Database backups on Exadata storage: Yes Sparse disk group: Yes	35 %	50 %	15 %

Best Practices for Exadata DB Systems

Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Exadata DB system:

- Wherever possible, use the Oracle-supplied cloud interfaces such as the Oracle Cloud Infrastructure Console, API, or CLI, or cloud-specific tools such as `dbaascli` and `dbaasapi` to perform lifecycle management and administrative operations on your Exadata DB system. For example, use the `exadbcpatchmulti` command to apply Oracle Database patches instead of manually running `opatch`. In addition, if an operation can be performed by using the Console as well as a command line utility, Oracle recommends that you use the Console. For example, use the Console instead of using `dbaasapi` to create databases.
- Do not change the compute node OS users or manually manipulate SSH key settings associated with your Exadata DB system.
- Apply *only* patches that are available through the Database service. Do *not* apply patches from any other source unless you are directed to do so by Oracle Support.
- Apply the quarterly patches regularly, every quarter if possible.
- Do not change the ports for Oracle Net Listener.

Network Setup for Exadata DB Systems

Before you set up an Exadata DB system, you must set up a virtual cloud network (VCN) and other [Networking service components](#). This topic describes the recommended configuration for the VCN and several related requirements for the Exadata DB system.

VCN and Subnets

To launch an Exadata DB system, you must have:

- A [VCN](#) in the region where you want the DB system
- At least two subnets in the VCN. The two subnets are:
 - Client subnet
 - Backup subnet

In general, Oracle recommends using regional subnets, which span all availability domains in the region. If you instead use AD-specific subnets, both the client and backup subnets must be in the same availability domain. The important thing to know for your DB system is that the

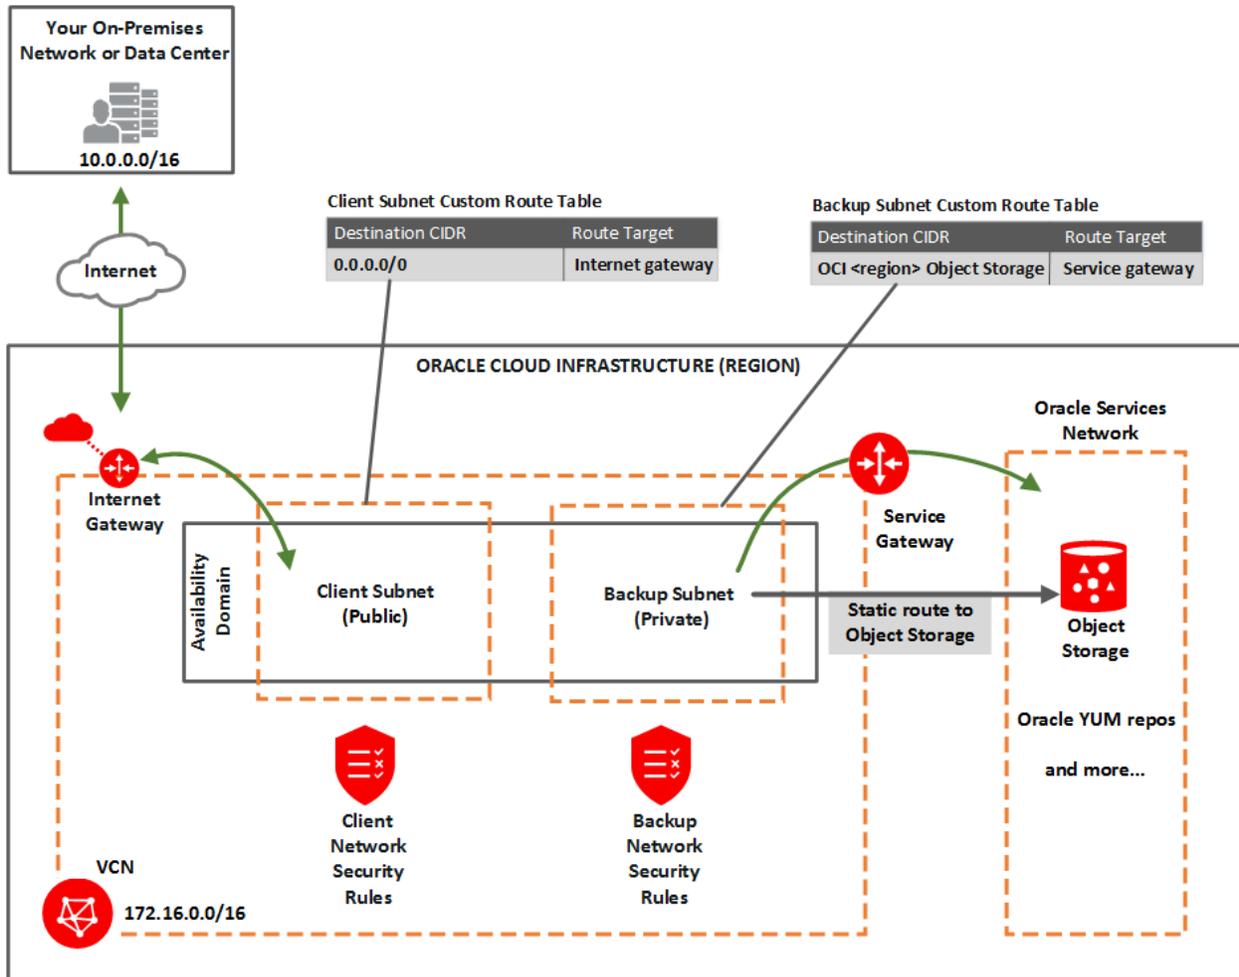
resources you create in the two subnets must be in the same availability domain. For more information, see [About Regional Subnets](#).

You will create custom route tables for each subnet. You will also create security rules to control traffic to and from the client network and backup network of the Exadata compute nodes. More information follows about those items.

OPTION 1: PUBLIC CLIENT SUBNET WITH INTERNET GATEWAY

This option can be useful when doing a proof-of-concept or development work. You can use this setup in production if you want to use an internet gateway with the VCN, or if you have services that run only on a public network and need access to the database. See the following diagram and description.

CHAPTER 11 Database



You set up:

- [Subnets](#):
 - *Public* client subnet (*public* means that the resources in the subnet can have public IP addresses at your discretion).

- Private backup subnet (*private* means that the resources in the subnet cannot have public IP addresses and therefore cannot receive incoming connections from the internet).
- Gateways for the VCN:
 - [Internet gateway](#) (for use by the client subnet).
 - [Service gateway](#) (for use by the backup subnet). Also see [Option 1: Service Gateway Access Only to Object Storage](#).
- [Route tables](#):
 - Custom route table for the public client subnet, with a route for 0.0.0.0/0, and target = the internet gateway.
 - Separate custom route table for the private backup subnet, with a route rule for the [service CIDR label](#) called **OCI <region> Object Storage**, and target = the service gateway. Also see [Option 1: Service Gateway Access Only to Object Storage](#).
- [Security rules](#) to enable the desired traffic to and from the Exadata nodes. See [Security Rules for the Exadata System](#).
- [Static route](#) on the DB system's compute nodes (to enable access to Object Storage by way of the backup subnet).



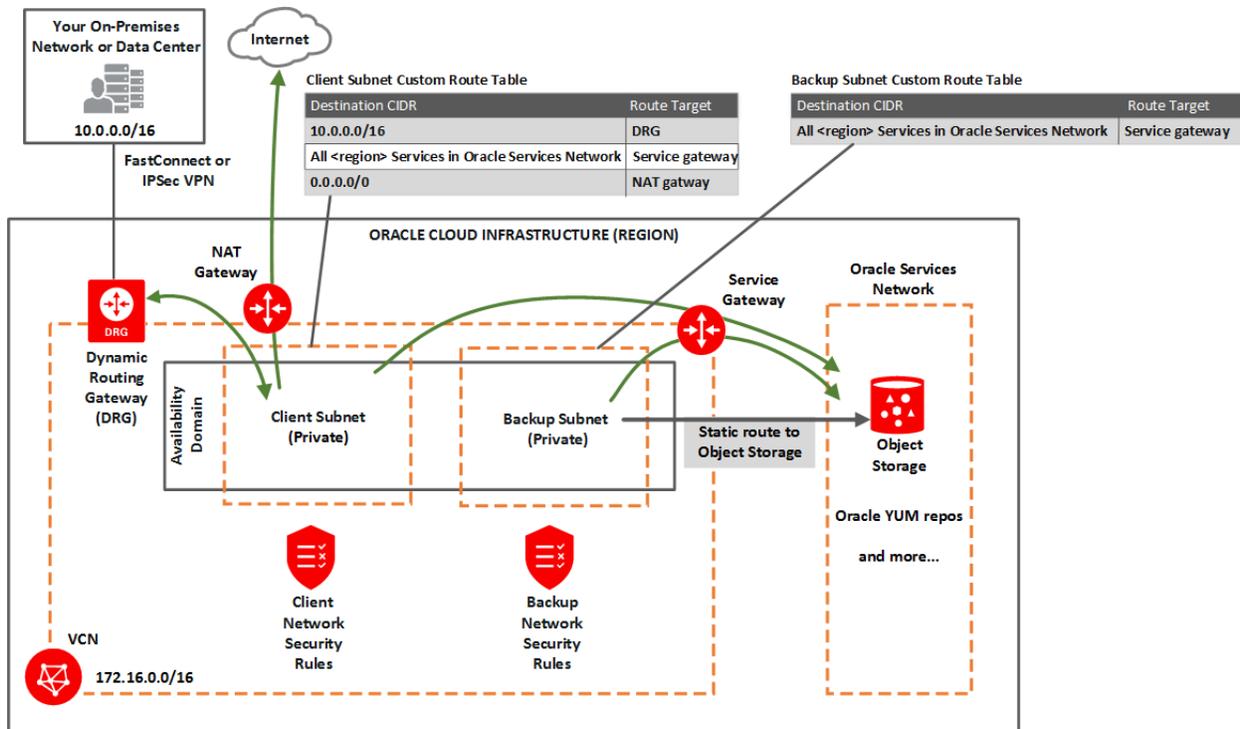
Important

See this [known issue](#) for information about configuring route rules with *service gateway* as the target on route tables associated with public subnets.

OPTION 2: PRIVATE SUBNETS

Oracle recommends this option for a production system. Both subnets are private and cannot be reached from the internet. See the following diagram and description.

CHAPTER 11 Database



You set up:

- [Subnets](#):
 - *Private* client subnet.
 - Private backup subnet.
- Gateways for the VCN:
 - [Dynamic routing gateway \(DRG\)](#), with a [FastConnect](#) or [IPSec VPN](#) to your on-premises network (for use by the client subnet).
 - [Service gateway](#) (for use by the backup subnet to reach Object Storage, and for use by the client subnet to reach the Oracle YUM repo for OS updates). Also see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#).

- [NAT gateway](#) (for use by the client subnet to reach public endpoints not supported by the service gateway).
- [Route tables](#):
 - Custom route table for the private client subnet, with two rules:
 - A rule for the on-premises network's CIDR, and target = DRG.
 - A rule for the [service CIDR label](#) called **All <region> Services in Oracle Services Network**, and target = the service gateway. The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. The rule enables the client subnet to reach the regional Oracle YUM repo for OS updates. Also see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#).
 - A rule for 0.0.0.0/0, and target = NAT gateway.
 - Separate custom route table for the private backup subnet, with one rule:
 - The same rule as for the client subnet: for the [service CIDR label](#) called **All <region> Services in Oracle Services Network**, and target = the service gateway. This rule enables the backup subnet to reach the regional Object Storage for backups.
- [Security rules](#) to enable the desired traffic to and from the Exadata nodes. See [Security Rules for the Exadata System](#).
- [Static route](#) on the DB system's compute nodes (to enable access to Object Storage by way of the backup subnet).

REQUIREMENTS FOR IP ADDRESS SPACE

If you're setting up Exadata DB systems (and thus VCNs) in more than one region, make sure the IP address space of the VCNs does not overlap. This is important if you want to set up disaster recovery with Oracle Data Guard.

The two subnets you create for the Exadata DB system must not overlap with 192.168.128.0/20.

The following table lists the *minimum* required subnet sizes, depending on the Exadata rack size. For the client subnet, each node requires two IP addresses, and in addition, three

addresses are reserved for Single Client Access Names ([SCANs](#)). For the backup subnet, each node requires one address.



Tip

The Networking service [reserves three IP addresses in each subnet](#). Allocating a larger space for the subnet than the minimum required (for example, at least /25 instead of /28) can reduce the relative impact of those reserved addresses on the subnet's available space.

Rack Size	Client Subnet: # Required IP Addresses	Client Subnet: Minimum Size	Backup Subnet: # Required IP Addresses	Backup Subnet: Minimum Size
Base System or Quarter Rack	$(2 \text{ addresses} * 2 \text{ nodes}) + 3 \text{ for } \text{SCANs} + 3 \text{ reserved in subnet} = 10$	/28 (16 IP addresses)	$(1 \text{ address} * 2 \text{ nodes}) + 3 \text{ reserved in subnet} = 5$	/29 (8 IP addresses)
Half Rack	$(2 * 4 \text{ nodes}) + 3 + 3 = 14$	/28 (16 IP addresses)	$(1 * 4 \text{ nodes}) + 3 = 7$	/29 (8 IP addresses)
Full Rack	$(2 * 8 \text{ nodes}) + 3 + 3 = 22$	/27 (32 IP addresses)	$(1 * 8 \text{ nodes}) + 3 = 11$	/28 (16 IP addresses)

VCN CREATION WIZARD: NOT FOR PRODUCTION

The Networking section of the Console includes a handy wizard that creates a VCN along with related resources. It can be useful if you just want to try launching an instance. However, the wizard automatically chooses the address ranges and creates public subnets and an internet gateway. You may not want this for your production network, so Oracle recommends you create the VCN and other resources individually yourself instead of using the wizard.

DNS: SHORT NAMES FOR THE VCN, SUBNETS, AND DB SYSTEM

For the nodes to communicate, the VCN must use the [Internet and VCN Resolver](#). It enables hostname assignment to the nodes, and DNS resolution of those hostnames by resources in the VCN. It enables round robin resolution of the database's [SCANs](#). It also enables resolution of important service endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata DB system. The Internet and VCN Resolver is the VCN's default choice for DNS in the VCN. For more information, see [DNS in Your Virtual Cloud Network](#) and also [DHCP Options](#).

When you create the VCN, subnets, and Exadata, you must carefully set the following identifiers, which are related to DNS in the VCN:

- VCN domain label
- Subnet domain label
- Hostname prefix for the Exadata DB system

These values make up the node's fully qualified domain name (FQDN):

```
<hostname_prefix>-#####.<subnet_domain_label>.<vcn_domain_label>.oraclevcn.com
```

For example:

```
exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com
```

In this example, you assign `exacs` as the hostname prefix when you create the Exadata DB system. The Database service automatically appends a hyphen and a five-letter string with the node number at the end. For example:

- Node 1: `exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 2: `exacs-abcde2.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 3: `exacs-abcde3.clientpvtad1.acmevcniad.oraclevcn.com`
- And so on

Requirements for the hostname prefix:

- Maximum 12 characters
- Cannot be the string *localhost*

Requirements for the VCN and subnet domain labels:

- Recommended maximum: 14 characters each. The actual underlying requirement is a total of 28 characters *across both domain labels* (excluding the period between the labels). For example, both of these are acceptable: `subnetad1.verylongvcnphx` or `verylongsubnetad1.vcnphx`. For simplicity, the recommendation is 14 characters each.
- No hyphens or underscores.
- Recommended: include the region name in the VCN's domain label, and include the availability domain name in the subnet's domain label.

`<12_chars_max>-#####.<14_chars_max>.<14_chars_max>.oraclevcn.com`

In general, the FQDN has a maximum total limit of 63 characters.

The preceding maximums are not enforced when you create the VCN and subnets. However, if the labels exceed the maximum, the Exadata deployment fails.

DNS: BETWEEN ON-PREMISES NETWORK AND VCN

To enable the use of hostnames when on-premises hosts and VCN resources communicate with each other, you have two options:

- Set up an instance in the VCN to be a custom DNS server. For an example of an implementation of this scenario with the Oracle Terraform provider, see [Hybrid DNS Configuration](#).
- Manage hostname resolution yourself manually.

Node Access to Object Storage: Static Route

Access to Oracle Cloud Infrastructure Object Storage is required for backing up databases, patching, and updating the cloud tooling on an Exadata DB system. Regardless of how you set up the VCN with that access (for example, with a service gateway), you must configure a static route to Object Storage on each of the compute nodes in the cluster. This is required

because, by default, all traffic in an Exadata DB system is routed through the data network. You need the traffic destined for Object Storage to be routed instead through the backup interface (BONDETH1).



Important

You must configure a static route for Object Storage access *on each compute node* in an Exadata DB system. Otherwise, attempts to back up databases, patch, or update tooling on the system might fail.

Object Storage IP allocations

Oracle Cloud Infrastructure Object Storage uses the CIDR block IP range 134.70.0.0/17 for all regions. This range was introduced in April and May of 2018.

As of June 1, 2018, Object Storage no longer supports the following discontinued IP ranges. Oracle recommends that you remove these older IP addresses from your access-control lists, firewall rules, and other rules after you have adopted the new IP ranges.

The **discontinued** IP ranges are:

- Germany Central (Frankfurt): 130.61.0.0/16
- UK South (London): 132.145.0.0/16
- US East (Ashburn): 129.213.0.0/16
- US West (Phoenix): 129.146.0.0/16

To configure a static route for Object Storage access

1. SSH to a compute node in the Exadata DB system.

```
ssh -i <private_key_path> opc@<node_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the

root user's profile.

```
login as: opc
[opc@dbsys ~]$ sudo su -
```

3. Identify the gateway configured for the BONDETH1 interface.

```
[root@dbsys ~]# grep GATEWAY /etc/sysconfig/network-scripts/ifcfg-bondeth1 |awk -F"=" '{print $2}'
10.0.4.1
```

4. Add the following static rule for BONDETH1 to the `/etc/sysconfig/network-scripts/route-bondeth1` file:

```
10.0.X.0/XX dev bondeth1 table 211
default via <gateway> dev bondeth1 table 211
134.70.0.0/17 via <gateway_from_previous_step> dev bondeth1
```

5. Restart the interface.

```
[root@dbsys ~]# ifdown bondeth1; ifup bondeth1;
```

The file changes from the previous step take effect immediately after the `ifdown` and `ifup` commands run.

6. Repeat the preceding steps on *each* compute node in the Exadata DB system.

Service Gateway for the VCN

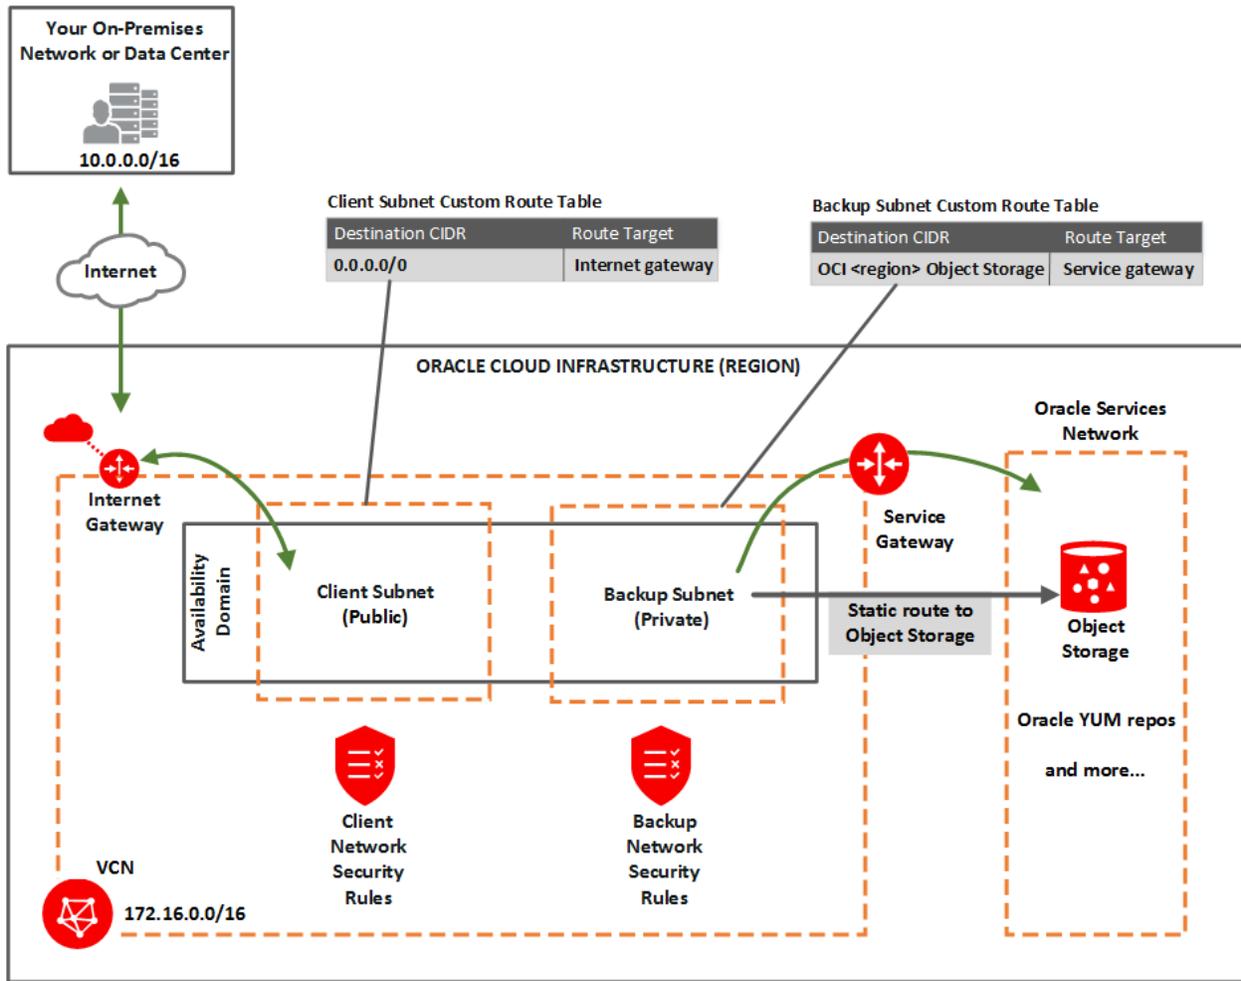
Your VCN needs access to both Object Storage for backups and Oracle YUM repos for OS updates.

Depending on whether you use [option 1](#) or [option 2](#) described previously, you use the service gateway in different ways. See the next two sections.

Option 1: Service Gateway Access Only to Object Storage

You configure the *backup subnet* to use the [service gateway](#) for access only to Object Storage.

As a reminder, here's the diagram for option 1:



In general, you must:

- Perform the [tasks for setting up a service gateway on a VCN](#), and specifically enable the service CIDR label called **OCI <region> Object Storage**.

- In the task for updating routing, add a route rule to the *backup* subnet's custom route table. For the destination service, use **OCI <region> Object Storage** and target = the service gateway.
- In the task for updating security rules in the subnet, perform the task on the *backup* network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to **OCI <region> Object Storage**. See [Rule Required Specifically for the Backup Network](#).

Option 2: Service Gateway Access to Both Object Storage and YUM Repos

You configure *both the client subnet and backup subnet* to use the [service gateway](#) for access to the [Oracle Services Network](#), which includes both Object Storage and the Oracle YUM repos.

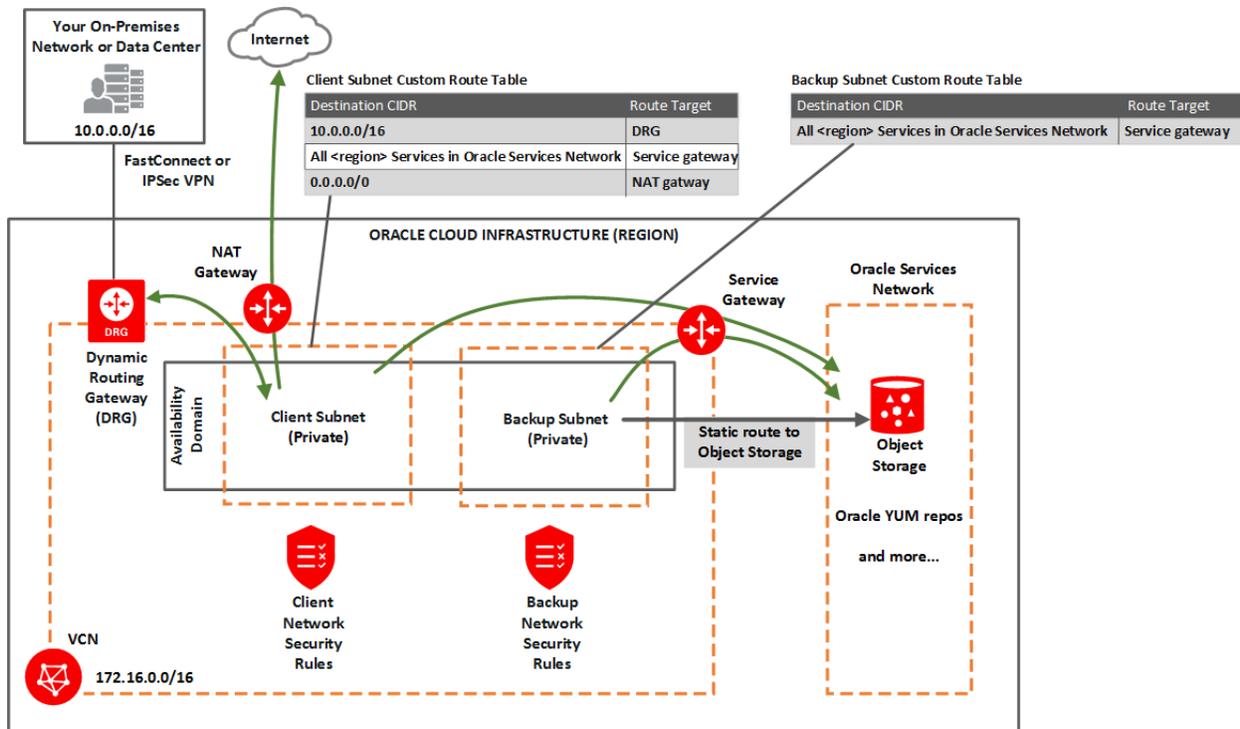


Important

See this [known issue](#) for information about accessing Oracle YUM services through the service gateway.

As a reminder, here's the diagram for option 2:

CHAPTER 11 Database



In general, you must:

- Perform the [tasks for setting up a service gateway on a VCN](#), and specifically enable the service CIDR label called **All <region> Services in Oracle Services Network**.
- In the task for updating routing in each subnet, add a rule to each subnet's custom route table. For the destination service, use **All <region> Services in Oracle Services Network** and target = the service gateway.
- In the task for updating security rules for the subnet, perform the task on the *backup* network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to **OCI <region> Object Storage**. See [Rule Required Specifically for the Backup Network](#). Note that the client subnet already has a broad egress rule that covers access to the YUM repos.

Here are a few additional details about using the service gateway for option 2:

- Both the client subnet and backup subnet use the service gateway, but to access different services. You cannot enable both the **OCI <region> Object Storage** service CIDR label and the **All <region> Services in Oracle Services Network** for the service gateway. To cover the needs of both subnets, you must enable **All <region> Services in Oracle Services Network** for the service gateway. The VCN can have only a single service gateway.
- Any route rule that targets a given service gateway must use an enabled service CIDR label and not a CIDR block as the destination for the rule. That means for option 2, the route tables for both subnets must use **All <region> Services in Oracle Services Network** for their service gateway rules.
- Unlike route rules, security rules can use either *any* service CIDR label (whether the VCN has a service gateway or not) or a CIDR block as the source or destination CIDR for the rule. Therefore, although the backup subnet has a route rule that uses **All <region> Services in Oracle Services Network**, the subnet can have a security rule that uses **OCI <region> Object Storage**. See [Rule Required Specifically for the Backup Network](#).

Security Rules for the Exadata System

This section lists the [security rules](#) to use with your Exadata system. Security rules control the types of traffic allowed for the client network and backup network of the Exadata's compute nodes. The rules are divided into three sections.

There are different ways to implement these rules. For more information, see [Ways to Implement the Security Rules](#).

RULES REQUIRED FOR BOTH THE CLIENT NETWORK AND BACKUP NETWORK

This section has several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules: the following rules are included by default in the [default security list](#).

General ingress rule 1: Allows SSH traffic from anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

This rule enables hosts in the VCN to receive Path MTU Discovery fragmentation messages. Without access to these messages, hosts in the VCN can have problems communicating with hosts outside the VCN.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** 4

General ingress rule 3: Allows connectivity error messages within the VCN

This rule enables the hosts in the VCN to receive connectivity error messages from each other.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your VCN's CIDR

- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** All

General egress rule 1: Allows all egress traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

RULES REQUIRED SPECIFICALLY FOR THE CLIENT NETWORK

The following security rules are important for the client network.

Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet

The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 6200

Client ingress rule 2: Allows SQL*NET traffic from within the client subnet

This rule is for SQL*NET traffic and is required in these cases:

- If you need to enable client connections to the database
- If you plan to use Oracle Data Guard
- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521



Important

The two preceding client ingress rules only cover connections initiated from within the client subnet. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.

The next four rules (two ingress, two egress) allow TCP and ICMP traffic inside the client network and enable the nodes to communicate with each other. If TCP connectivity fails across the nodes, the Exadata DB system fails to provision.

Client ingress rule 3: Allows all TCP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP

- **Source Port Range:** All
- **Destination Port Range:** All

Client ingress rule 4: Allows all ICMP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** ICMP
- **Type and Code:** All

Client egress rule 1: Allows all TCP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** All

Client egress rule 2: Allows all ICMP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** Client subnet's CIDR
- **IP Protocol:** ICMP
- **Type and Code:** All

The next egress rule is important because it allows connections to the Oracle YUM repos. It is redundant with the general egress rule in [Rules Required for Both the Client Network and Backup Network](#) (and in the [default security list](#)). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

Client egress rule 3: Allows all egress traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

RULE REQUIRED SPECIFICALLY FOR THE BACKUP NETWORK

The following security rule is important for the backup network because it enables the DB system to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them). It is redundant with the general egress rule in [Rules Required for Both the Client Network and Backup Network](#) (and in the [default security list](#)). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

Backup egress rule: Allows access to Object Storage

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
 - The service CIDR label called **OCI <region> Object Storage**
 - If the client network does not have access to the Oracle YUM repos, use the service CIDR label called **All <region> Services in Oracle Services Network**
- **IP Protocol:** TCP

- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)

Ways to Implement the Security Rules

The Networking service offers two ways to implement security rules within your VCN:

- [Network security groups](#)
- [Security lists](#)

For a comparison of the two methods, see [Comparison of Security Lists and Network Security Groups](#).

If you use network security groups

If you choose to use [network security groups](#) (NSGs), here is the recommended process:

1. Create an NSG for the client network. Add the following security rules to that NSG:
 - The rules listed in [Rules Required for Both the Client Network and Backup Network](#)
 - The rules listed in [Rules Required Specifically for the Client Network](#)
2. Create a separate NSG for the backup network. Add the following security rules to that NSG:
 - The rules listed in [Rules Required for Both the Client Network and Backup Network](#)
 - The rules listed in [Rule Required Specifically for the Backup Network](#)
3. When the database administrator [creates the Exadata DB system](#), they must choose several networking components (for example, which VCN and subnets to use):
 - When they choose the client subnet, they can also choose which NSG or NSGs to use. Make sure they choose the client network's NSG.
 - When they choose the backup subnet, they can also choose which NSG or NSGs to use. Make sure they choose the backup network's NSG.

You could instead create a separate NSG for the general rules. Then when the database administrator chooses which NSGs to use for the client network, make sure they choose both the general NSG and the client network NSG. Similarly for the backup network, they choose both the general NSG and the backup network NSG.

If you use security lists

If you choose to use [security lists](#), here is the recommended process:

1. Configure the client subnet to use the required security rules:
 - a. Create a custom security list for the client subnet and add the rules listed in [Rules Required Specifically for the Client Network](#).
 - b. Associate the following two security lists with the client subnet:
 - VCN's [default security list](#) with all its default rules. This automatically comes with the VCN. By default it contains the rules in [Rules Required for Both the Client Network and Backup Network](#).
 - The new custom security list you created for the client subnet.
2. Configure the backup subnet to use the required security rules:
 - a. Create a custom security list for the backup subnet and add the rules listed in [Rule Required Specifically for the Backup Network](#).
 - b. Associate the following two security lists with the backup subnet:
 - VCN's [default security list](#) with all its default rules. This automatically comes with the VCN. By default it contains the rules in [Rules Required for Both the Client Network and Backup Network](#).
 - The new custom security list you created for the backup subnet.

Later when the database administrator creates the Exadata DB system, they must choose several networking components. When they select the client subnet and backup subnet that you've already created and configured, the security rules are automatically enforced for the nodes created in those subnets.



Warning

Do not remove the default egress rule from the default security list. If you do, make sure to instead include the following replacement egress rule in the client subnet's security list:

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

Creating Exadata DB Systems

This topic explains how to launch an Exadata DB system. It also describes how to configure required access to the Oracle Cloud Infrastructure Object Storage service and set up DNS.

When you launch an Exadata DB system using the Console or the API, the system is provisioned to support Oracle databases. The service creates an initial database based on the options you provide and some default options described later in this topic.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Prerequisites

- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the DB System via SSH. A sample public key, abbreviated for readability, is shown below.

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA...lo/gKMLVM2xzclxJr/Hc26biw3TXWGEakrK1OQ== rsa-key-20160304
```

For more information, see [Managing Key Pairs on Linux Instances](#).

- A correctly configured virtual cloud network (VCN) to launch the DB system in. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary for the DB system. For more information, see [Network Setup for Exadata DB Systems](#).

Default Options for the Initial Database

To simplify launching a DB system in the Console and when using the API, the following default options are used for the initial database.

- **Console Enabled:** False
- **Create Container Database:** False for version 11.2.0.4 databases. Otherwise, true.
- **Create Instance Only (for standby and migration):** False
- **Database Home ID:** Creates a database home

- **Database Language:** AMERICAN
- **Database Sizing Template:** odb2
- **Database Storage:** Automatic Storage Management (ASM)
- **Database Territory:** AMERICA
- **Database Unique Name:** The user-specified database name and a system-generated suffix, for example, dbtst_phx1cs.
- **PDB Admin Name:** pdbuser (Not applicable for version 11.2.0.4 databases.)

For a list of the database options that you can set in the Console, see [To create an Exadata DB system](#).

Using the Console

To create an Exadata DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Click **Create DB System**.
3. On the **Create DB System** page, provide the basic information for the DB system:
 - **Select a compartment:** By default, the DB system launches in your current compartment and you can use the network resources in that compartment.
 - **Name your DB system:** A friendly, display name for the DB system. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system.
 - **Select an availability domain:** The availability domain in which the DB system resides.
 - **Select a shape type:** The shape type you select sets the default shape and filters the shape options in the next field.
 - **Select a shape:** The shape determines the type of DB system and the resources allocated to the system. To specify a shape other than the default, click **Change Shape**, and select an available shape from the list.

Exadata shapes

Exadata X7 shapes:

- **Exadata.Quarter2.92:** Provides a 2-node Exadata DB system with up to 92 CPU cores, and 106 TB of usable storage.
- **Exadata.Half2.184:** Provides a 4-node Exadata DB system with up to 184 CPU cores, and 212 TB of usable storage.
- **Exadata.Full2.368:** Provides an 8-node Exadata DB system with up to 368 CPU cores, and 424 TB of usable storage.

Exadata X6 shapes:

- **Exadata.Quarter1.84:** Provides a 2-node Exadata DB system with 22 enabled CPU cores, with up to 62 additional CPU cores, and 84 TB of usable storage.
- **Exadata.Half1.168:** Provides a 4-node Exadata DB system with 44 enabled CPU cores, with up to 124 additional CPU cores, and 168 TB of usable storage.
- **Exadata.Full1.336:** Provides an 8-node Exadata DB system with 88 enabled CPU cores, with up to 248 additional CPU cores, and 336 TB of usable storage.

Exadata base system:

Exadata.Base.48: Provides a 2-node Exadata DB system with up to 48 CPU cores, and 74.8 TB of usable storage.

All Exadata shapes provide 720 GB RAM per node and unlimited I/O, and support only Enterprise Edition - Extreme Performance. For more details about Exadata shapes, see [System Configuration](#).

- **Configure the DB system:** Specify the following:
 - **Total node count:** The number of nodes in the DB system. The number depends on the shape you select.

- **Oracle Database software edition:** The database edition supported by the DB system. Exadata DB systems only support Enterprise Edition - Extreme Performance.
- **CPU core count:** The number of CPU cores for the DB system. The text below the field indicates the acceptable values for that shape. The core count is evenly divided across the nodes.
You can increase the CPU cores to accommodate increased demand after you launch the DB system.
For an Exadata X7 DB system or a base system, you can specify zero (0) cores when you launch the system. This will provision the system and immediately stop it. See [Scaling Within an Exadata System](#) for information about CPU core scaling and the impact on billing.
- **Configure storage:** Specify the following:
 - **Cluster Name:** (*Optional*) A unique cluster name for a multi-node DB system. The name must begin with a letter and contain only letters (a-z and A-Z), numbers (0-9) and hyphens (-). The cluster name can be no longer than 11 characters and is not case sensitive.
 - **Storage Allocation:** The configuration settings that determine the percentage of storage assigned to DATA, RECO, and optionally, SPARSE disk:
 - **Database Backups on Exadata Storage:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata DB system environment. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.
 - **Create Sparse Disk Group:** Select this configuration option if you intend to use snapshot functionality within your Exadata DB system environment. If you select this option, the SPARSE disk group is

created, which enables you to use Exadata DB system snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and Exadata DB system snapshot functionality will not be available on any database deployments that are created in the environment.



Important

Creating a sparse disk group impacts the storage available for the ASM disk groups (DATA and RECO) and you cannot change the storage allocation configuration after you provision your DB system. For information about the percentage of storage that will be assigned to DATA, RECO, and SPARSE disk based on your configuration, see [Impact of Configuration Settings on Storage](#). Similar information will display under the options in the Console dialog.

- **Add public SSH keys:** The public key portion of each key pair you want to use for SSH access to the DB system. You can browse or drag and drop .pub files, or paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.
- **Choose a license type:** The type of license you want to use for the DB system. Your choice affects metering for billing.
 - **License Included** means the cost of the cloud service includes a license for the Database service.

- **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
4. Specify the network information:
- **Virtual Cloud Network:** The VCN in which to launch the DB system. Click **Change Compartment** to select a VCN in a different compartment.
 - **Client subnet:** The subnet to which the Exadata DB system should attach. Click **Change Compartment** to select a subnet in a different compartment.
Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet will cause the private interconnect to malfunction.
 - **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from Oracle Cloud Infrastructure Object Storage, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable.
Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.
If you plan to back up databases to Object Storage, see the network prerequisites in [Managing Exadata Database Backups](#).
 - **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress [security rules](#) to your DB system. A maximum of five NSGs can be specified. For more information, see [Network Security Groups](#) and [Network Setup for Exadata DB Systems](#).
Note that if you choose a subnet with a [security list](#), the security rules for the DB system will be a union of the rules in the security list and the NSGs.

To use network security groups

- a. Check the **Use Network Security Groups to Control Client Traffic** check box. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.
 - b. Specify the NSG to use with the client network. You might need to use more than one NSG. If you're not sure, contact your network administrator.
 - c. To use additional NSGs with the client network, click **+ Another Network Security Group**.
 - d. Check the **Use Network Security Groups to Control Backup Traffic** check box.
 - e. Specify the NSG to use with the backup network. You might need to use more than one NSG. If you're not sure, contact your network administrator.
 - f. To use additional NSGs with the backup network, click **+ Another Network Security Group**.
- **Hostname prefix:** Your choice of host name for the Exadata DB system. The host name must begin with an alphabetic character, and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata DB system is 12.



Important

The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host domain name:** The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not

permitted.

If you plan to store database backups in Object Storage, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.

- **Host and domain URL:** Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
5. Click **Show Advanced Options** to specify advanced options for the DB system:
 - **Disk redundancy:** Exadata DB systems support only high redundancy (3-way mirroring).
 - **Time zone:** The default time zone for the DB system is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system. For more information, see [DB System Time Zone](#).



Tip

If you want to set a time zone other than UTC or the browser-detected time zone and if you do not see the time zone you are interested in, try selecting "Miscellaneous" as the time zone prefix.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. After you completed the network configuration and any advanced options, click **Next**.

7. Provide information for the initial database:

- **Database name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database version:** The version of the initial database created on the DB system when it is launched. After the DB system is active, you can create additional databases on it. You can mix database versions on the DB system.



Note

If you plan to run Oracle Database 19c on your Exadata DB system, you must specify version 19c when you create the DB system. Earlier database versions are supported on a 19c Exadata DB system and can be created at anytime. Exadata DB systems created with earlier Oracle Database versions will not automatically support Oracle Database 19c. The DB system must be upgraded manually.

- **PDB name:** *Not applicable to version 11.2.0.4.* The name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of 8 alphanumeric characters. The only special character permitted is the underscore (_).
- **Create administrator credentials:** A database administrator `SYS` user will be created with the password you supply.
 - **Username:** `SYS`
 - **Password:** Supply the password for this user. The password must meet the following criteria:

A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

- **Confirm password:** Re-enter the SYS password you specified.
- **Select workload type:** Choose the workload type that best suits your application:
 - **Online Transactional Processing (OLTP)** configures the database for a transactional workload, with a bias towards high volumes of random data access.
 - **Decision Support System (DSS)** configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations.
- **Configure database backups:** Specify the settings for backing up the database to Object Storage:
 - **Enable automatic backups:** Check the check box to enable automatic incremental backups for this database.
 - **Backup retention period:** *(Optional)* If you enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The default selection is 30 days.
 - **Backup Scheduling:** If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, the six-hour default window of 00:00 to 06:00 (in the time zone of the DB system's region) is used for your database. See [Backup Scheduling](#) for more information.
- Click **Show Advanced Options** to specify advanced options for the initial database:

- **Character set:** The character set for the database. The default is AL32UTF8.
 - **National character set:** The national character set for the database. The default is AL16UTF16.
8. Click **Create DB System**.
- The DB system appears in the list with a status of Provisioning. The DB system's icon changes from yellow to green (or red to indicate errors).
- After the DB system's icon turns green, with a status of Available, you can click the highlighted DB system name to see details about the DB system. Note the IP addresses. You'll need the private or public IP address, depending on network configuration, to connect to the DB system.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to create DB system components.

DB systems:

- [GetDbSystem](#)
- [LaunchDbSystem](#)
- [ListDbSystems](#)

Database homes:

- [CreateDbHome](#)
- [GetDbHome](#)
- [ListDbHomes](#)

Shapes and database versions:

- [ListDbSystemShapes](#)
- [ListDbVersions](#)

Configuring a Static Route for Accessing the Object Store

All the traffic in an Exadata DB system is, by default, routed through the data network. To route backup traffic to the backup interface (BONDETH1), you need to configure a static route on *each* of the compute nodes in the cluster. For instructions, see [Node Access to Object Storage: Static Route](#).

Setting Up DNS for a DB System

DNS lets you use hostnames instead of IP addresses to communicate with a DB system. You can use the *Internet and VCN Resolver* (the DNS capability built into the VCN) as described in [DNS in Your Virtual Cloud Network](#). Oracle recommends using a VCN Resolver for DNS name resolution for the client subnet. It automatically resolves the Swift endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata DB system.

Maintaining an Exadata DB System

Maintaining a secure Exadata DB system in the best working order requires you to perform the following tasks regularly:

- Patching the grid infrastructure and Database software on the compute nodes. See [Patching an Exadata DB System](#) for information and instructions.
- Updating the operating system and the tooling on the compute nodes. See [Updating an Exadata DB System](#) for information and instructions.

In addition to the maintenance tasks you perform, Oracle manages the patching and updating of all other infrastructure components, including the physical compute nodes (Dom0), network switches, power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and the Exadata storage servers.

These Oracle updates occur on a quarterly basis, typically in January, April, July, and October. Occasionally, Oracle might need to update your system apart from the regular quarterly

updates to apply time-sensitive changes such as security patches. While you cannot opt out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal to help you plan for them. For further information about the update policy and details such as the duration and impact on your system's availability and performance, see [Oracle Database Cloud Exadata Service Supported Software Versions and Planning for Updates](#).

Managing Exadata DB Systems

You can start, stop, terminate, scale, manage licenses for, and check the status of, an Exadata DB system by using the Oracle Cloud Infrastructure Console or the API.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Using the Console

To check the status of an Exadata DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. In the list of DB systems, find the system you're interested in and check its icon. The color of the icon and the text next to it indicates the status of the system.
 - **Provisioning:** Yellow icon. Resources are being reserved for the DB system, the system is booting, and the initial database is being created. Provisioning can take several minutes. The system is not ready to use yet.
 - **Available:** Green icon. The DB system was successfully provisioned. A few minutes after the system enters this state, you can SSH to it and begin using it.
 - **Terminating:** Gray icon. The DB system is being deleted by the terminate action in the Console or API.
 - **Terminated:** Gray icon. The DB system has been deleted and is no longer available.
 - **Failed:** Red icon. An error condition prevented the provisioning or continued operation of the DB system.

To view the status of a database node, under Resources, click **Nodes** to see the list of nodes. In addition to the states listed for a DB system, a node's status can be one of the following:

- **Starting:** Yellow icon. The database node is being powered on by the start or reboot action in the Console or API.
- **Stopping:** Yellow icon. The database node is being powered off by the stop or reboot action in the Console or API.
- **Stopped:** Yellow icon. The database node was powered off by the stop action in the Console or API.

You can also check the status of DB systems and database nodes using the [ListDbSystems](#) or [ListDbNodes](#) API operations, which return the `lifecycleState` attribute.

To start, stop, or reboot an Exadata DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. In the list of DB systems, find the DB system you want to stop or start, and then click its name to display details about it.
4. In the list of nodes, click the Actions icon (three dots) for a node and then click one of the following actions:
 - **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
 - **Stop:** Shuts down the node. After the node is powered off, the **Start** action is enabled.
 - **Reboot:** Shuts down the node, and then restarts it.



Note

- For billing purposes, the **Stop** state has no effect on the resources you consume. Billing continues for nodes that you stop, and related resources continue to apply against any relevant quotas. You must **Terminate** a DB system to remove its resources from billing and quotas.
- After you restart or reboot a node, the floating IP address might take several minutes to be updated and display in the Console.

To scale an Exadata DB system

If an Exadata DB system requires more compute node processing power, you can scale up (burst) the number of enabled CPU cores (OCPU) in the system.

You can also scale an X7 Exadata DB system or a base system down to zero CPU cores to temporarily stop the system and be charged only for the hardware infrastructure. For more information about scaling down, see [Scaling Within an Exadata System](#).

CPU cores must be scaled symmetrically across all nodes in the DB system. Use multiples of two for a base system or quarter rack, multiples of four for a half rack, and multiples of eight for a full rack. The total number of CPU cores in a rack must not exceed the maximum limit for that shape.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.

3. In the list of DB systems, find the system you want to scale and click its highlighted name.
The system details are displayed.
4. Click **Scale CPU Cores** and then change the number in the **CPU Core Count** field. The text below the field indicates the acceptable values, based on the shape used when the DB system was launched.
5. Click **Update**.



Note

If you scale an X7 Exadata DB system or a base system down to zero CPU cores, the floating IP address of the nodes might take several minutes to be updated and display in the Console.

To move an Exadata DB system to another compartment



Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see [Details for the Database Service](#).

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.

3. In the list of DB systems, find the system you want to move and click its highlighted name.

The system details are displayed.

4. Click **Move Resource**.
5. Select the new compartment.
6. Click **Move Resource**.

For information about dependent resources for Database resources, see [Moving Database Resources to a Different Compartment](#).

To terminate an Exadata DB system

Terminating a DB system permanently deletes it and any databases running on it.



Note

The database data is local to the DB system and is lost when the system is terminated. Oracle recommends that you back up any data in the DB system before terminating it.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. For the DB system you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
4. Confirm when prompted.
The DB system's icon indicates Terminating.

At this point, you cannot connect to the system and any open connections are terminated.

To edit the network security groups (NSGs) for your client or backup network

Your client and backup networks can each use up to five network security groups (NSGs). Note that if you choose a subnet with a [security list](#), the security rules for the DB system will be a union of the rules in the security list and the NSGs. For more information, see [Network Security Groups](#) and [Network Setup for Exadata DB Systems](#).

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. In the list of DB systems, find the system you want to manage and click its highlighted name.
The system details are displayed.
4. In the **Network** details, click the **Edit** link to the right of the **Client Network Security Groups** or **Backup Network Security Groups** field.
5. In the **Edit Network Security Groups** dialog, click **+ Another Network Security Group** to add an NSG to the network.
To change an assigned NSG, click the drop-down menu displaying the NSG name, then select a different NSG.
To remove an NSG from the network, click the **X** icon to the right of the displayed NSG name.
6. Click **Save**.

To manage your BYOL database licenses

If you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.

3. In the list of DB systems, find the system you want to scale and click its highlighted name.
The system details are displayed.
4. Click **Scale CPU Cores**, and then change the number.

To manage tags for your DB systems and database resources

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system or database resource you're interested in, and click the name.
4. Click the **Tags** tab to view or edit the existing tags. Or click **Apply Tag(s)** to add new tags.

For more information, see [Resource Tags](#).

To view a work request for your DB systems and database resources

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system or database resource you're interested in, and click the name.
4. In the **Resources** section, click **Work Requests**. The status of all work requests appears on the page.
5. To see the log messages, error messages, and resources that are associated with a specific work request, click the operation name. Then, select an option in the **More information** section.
For associated resources, you can click the Actions icon (three dots) next to a resource to copy the resource's OCID.

For more information, see [Work Requests](#).

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage DB system components.

DB systems:

- [GetDbSystem](#)
- [ListDbSystems](#)
- [ChangeDbSystemCompartment](#)
- [TerminateDbSystem](#)
- [UpdateDbSystem](#)

Nodes:

- [DbNodeAction](#): Use this operation to power cycle a node in the DB system.
- [ListDbNodes](#)
- [GetDbNode](#)

Managing Exadata DB System I/O Resources

This topic explains the I/O Resource Management (IORM) feature and how to enable it, modify the IORM settings, and disable it by using the Console or the API.

About IORM

The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata DB system.

On an Exadata DB system, all databases share dedicated storage servers which include flash storage. By default, the databases are given equal priority with respect to these resources. The Exadata storage management software uses a first come, first served approach for query

processing. If a database executes a major query that overloads I/O resources, overall system performance can be slowed down.

IORM allows you to assign priorities to your databases to ensure critical queries are processed first when workloads exceed their resource allocations. You assign priorities by creating directives that specify the number of shares for each database. The number of shares corresponds to a percentage of resources given to that database when I/O resources are stressed.

Directives work together with an overall optimization objective you set for managing the resources. The following objectives are available:

- **Auto** - Recommended. IORM determines the optimization objective and continuously and dynamically determines the optimal settings, based on the workloads observed, and resource plans enabled.
- **Balanced** - For critical OLTP and DSS workloads. This setting balances low disk latency and high throughput. This setting limits disk utilization of large I/Os to a lesser extent than low latency to achieve a balance between good latency and good throughput.
- **High throughput** - For critical DSS workloads that require high throughput.
- **Low latency** - For critical OLTP workloads. This setting provides the lowest possible latency by significantly limiting disk utilization.

For more information about IORM, see [Exadata System Software User's Guide](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

Using the Console

To enable IORM on your Exadata DB system

Enabling IORM includes specifying an optimization objective and configuring your resource plan directives.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, find the Exadata DB system for which you want to enable IORM, and click its highlighted name.
The system details are displayed, showing the IORM status as "Disabled."
4. Click **Enable IORM**.
It might take a minute for the Enable I/O Resource Management dialog to retrieve the DB system information.
5. Select the objective to apply to the resource plan:
 - **Auto** - (Recommended) Dynamically changes the objective based on the resource plan and observed workloads.
 - **Balanced** - Weighs high throughput and low latency evenly.
 - **High throughput** - Provides the best throughput for DSS workloads.
 - **Low latency** - Provides the best latency for critical OLTP workloads.
6. Configure the resource plan default directive by setting the number of shares. This number of shares is assigned to each database not associated with a specific directive.
7. In the Resource Plan Directives section, add a directive for each database you want to assign a greater or lesser number of shares than the default directive.
To add a directive, click **+ Additional Directive**, then specify the database and the number of shares for that database.

8. When you are done adding directives, click **Enable**.

While the IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your DB system. After a successful update, the IORM status shows as "Enabled."

To modify the IORM configuration on your Exadata DB system

Use this procedure to change your IORM settings or to disable IORM.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. In the list of DB systems, find the Exadata DB system for which you want to modify the IORM configuration, and click its highlighted name.
The system details are displayed, showing the IORM status as "Enabled."
4. Click **Update IORM**.
5. In the Update I/O Resource Management dialog, take one of the following actions:
 - Change your settings - Specify a new objective and adjust your directives, as applicable, and then click **Update**.
 - Disable IORM - Click **Disable IORM**. Disabling IORM removes all your resource plan directives and restores a basic objective for I/O resource management.

While the new IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your DB system. After a successful update, the IORM status shows as "Enabled" or "Disabled," depending on the action you took.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage the I/O resources of your Exadata DB system.

- [ListDbSystems](#)
- [GetDbSystem](#)
- [GetExadataIormConfig](#)
- [UpdateExadataIormConfig](#)

Connecting to an Exadata DB System

This topic explains how to connect to an Exadata DB System using SSH or SQL Developer. How you connect depends on how your cloud network is set up. You can find information on various networking scenarios in [Overview of Networking](#), but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

Prerequisites

For SSH access to a compute node in an Exadata DB System, you'll need the following:

- The full path to the file that contains the private key associated with the public key used when the system was launched.
- The public or private IP address of the DB System. Use the private IP address to connect to the DB system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the DB System's public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Oracle Cloud Infrastructure Console on the **Database** page.

Connecting to a Compute Node with SSH

You can connect to the compute nodes in an Exadata DB System by using a Secure Shell (SSH) connection. Most UNIX-style systems (including Linux, Solaris, BSD, and OS X) include an

SSH client by default. For Windows, you can download a free SSH client called PuTTY from <http://www.putty.org>.

To connect from a UNIX-style system

Use the following SSH command to access a compute node:

```
$ ssh -i <private key> opc@<DB System IP address>
```

<private key> is the full path and name of the file that contains the private key associated with the Exadata DB System you want to access.

Use the private or public IP address depending on your network configuration. For more information, see [Prerequisites](#).

To connect from a Windows system

1. Open `putty.exe`.
2. In the **Category** pane, select **Session** and enter the following fields:
 - **Host Name (or IP address):** `opc@<ip_address>`
Use the compute node's private or public IP address depending on your network configuration. For more information, see [Prerequisites](#).
 - **Connection type:** SSH
 - **Port:** 22
3. In the **Category** pane, expand **Connection**, expand **SSH**, and then click **Auth**, and browse to select your private key.
4. Optionally, return to the **Session** category screen and save this session information for reuse later.
5. Click **Open** to start the session.

To access a database after you connect to the compute node

1. Log in as opc and then sudo to the oracle user.

```
login as: opc  
[opc@<host_name> ~]$ sudo su - oracle
```

2. Source the database's .env file to set the environment.

```
[oracle@<host_name>]# . <database_name>.env
```

In the following example, the host name is "ed1db01" and the database name is "cdb01".

```
[oracle@ed1db01]# . cdb01.env
```

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open port 1521 for the Oracle default listener by updating the security list used for the DB System. This method provides more durable access to the database. For more information, see [Updating the Security List](#).

After you've created an SSH tunnel or opened port 1521 as described above, you can connect to a Exadata DB System using SCAN IP addresses or public IP addresses, depending on how your network is set up and where you are connecting from. You can find the IP addresses in the Console, in the **Database** details page.

To connect using SCAN IP addresses

You can connect to the database using the SCAN IP addresses if your client is on-premises and you are connecting using a FastConnect or VPN connection. You have the following options:

- Use the private SCAN IP addresses, as shown in the following `tnsnames.ora` example:

```
testdb=
  (DESCRIPTION =
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP) (HOST = <scanIP1>) (PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP) (HOST = <scanIP2>) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

- Define an external SCAN name in your on-premises DNS server. Your application can resolve this external SCAN name to the DB System's private SCAN IP addresses, and then the application can use a connection string that includes the external SCAN name. In the following `tnsnames.ora` example, `extscanname.example.com` is defined in the on-premises DNS server.

```
testdb =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = <extscanname.example.com>) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

To connect using public IP addresses

You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway. However, there are important implications to consider:

- When the client uses the public IP address, the client bypasses the SCAN listener and reaches the node listener, so server side load balancing is not available.

- When the client uses the public IP address, it cannot take advantage of the VIP failover feature. If a node becomes unavailable, new connection attempts to the node will hang until a TCP/IP timeout occurs. You can set client side sqlnet parameters to limit the TCP/IP timeout.

The following `tnsnames.ora` example shows a connection string that includes the `CONNECT_TIMEOUT` parameter to avoid TCP/IP timeouts.

```
test=
  (DESCRIPTION =
    (CONNECT_TIMEOUT=60)
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP) (HOST = <publicIP1>) (PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP) (HOST = <publicIP2>) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

Managing Exadata DB System Software Images

You can control the version of Oracle binaries that is installed when you provision a new database on an Exadata DB system by maintaining the software images on the system. Oracle provides a library of cloud software images that you can view and download onto your Exadata DB system by using the `dbaascli` utility.

When you create a new Exadata DB system database with a new Oracle Home (Database Home) directory location, the Oracle Database binaries are sourced from a software image that is stored on your Exadata DB system. Over time, the software images on your Exadata DB system become outdated if they are not maintained. Using an outdated software image makes it necessary for you to apply patches to newly installed binaries to bring them up to date. Oracle recommends that you maintain your Exadata DB system environment with up-to-date software images to avoid this extra patching step which can be time-consuming and error prone.

Viewing Information About Available Software Images

You can view information about Oracle Database software images that are available to download to your Exadata DB system by using the `cswlib list` subcommand of the `dbaascli` utility.

To view information about available software images

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to a Compute Node with SSH](#).

2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the `dbaascli` command with the `cswlib list` subcommand:

```
# dbaascli cswlib list
```

The command displays a list of available software images, including version and bundle patch information that you can use to download the software image.

4. Exit the root-user command shell:

```
# exit
$
```

Downloading Software Images

You can download available software images onto your Exadata DB system by using the `cswlib download` subcommand of the `dbaascli` utility.

To download a software image

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to a Compute Node with SSH](#).
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the `dbaascli` command with the `cswlib download` subcommand:

```
# dbaascli cswlib download [--version <software_version>] [--bp <software_bundle_patch>]
```

The command displays a list of software images that are downloaded to your Exadata Cloud Service environment, including version and bundle patch information.

The optional parameters are:

- **version:** specifies an Oracle Database software version. For example, 19000, 18000, or 12201.
- **bp:** identifies a bundle patch release. For example, APR2019, JAN2019, or OCT2018.

If you do not include the optional parameters, the `dbaascli cswlib download` command downloads the latest available software image for all available Oracle Database software versions.

4. Exit the root-user command shell:

```
# exit
$
```

Updating an Exadata DB System

This topic covers how to update the operating system and the tooling on the database server nodes (also known as "compute nodes") of an Exadata DB system. Review all of the information carefully before you begin the updates.

OS Updates



Important

Starting on March 14, 2019, Exadata DB system images run Oracle Linux 7 (OL7). Previously launched systems are running Oracle Linux 6 (OL6). The underlying infrastructure of existing Exadata DB systems will be patched in the April 2019 time frame. This patch will allow you to upgrade the DB system operating system to OL7.

After you receive notification that your infrastructure patch update is complete, follow the instructions in [Patching an Exadata DB System](#) to patch the Oracle Grid Infrastructure and the databases, and then update the OS. Review the minimum software requirements and other details in [How to update the Exadata System Software \(DomU\) to 19c on the Exadata Cloud Service in OCI \(Doc ID 2521053.1\)](#) before you perform the OS update tasks.

You update the operating systems of Exadata compute nodes by using the `patchmgr` tool. This utility manages the entire update of one or more compute nodes remotely, including running pre-reboot, reboot, and post-reboot steps. You can run the utility from either an Exadata compute node or a non-Exadata server running Oracle Linux. The server on which you run the utility is known as the "driving system." You cannot use the driving system to update itself. Therefore, if the driving system is one of the Exadata compute nodes on a system you are updating, you must run a separate operation on a different driving system to update that server.

The following two scenarios describe typical ways of performing the updates:

Scenario 1: Non-Exadata Driving System

The simplest way to run the update the Exadata system is to use a separate Oracle Linux server to update all Exadata compute nodes in the system.

Scenario 2: Exadata Node Driving System

You can use one Exadata compute node to drive the updates for the rest of the compute nodes in the system, and then use one of the updated nodes to drive the update on the original Exadata driver node.

For example: You are updating a half rack Exadata system, which has four compute nodes - node1, node2, node3, and node4. First, use node1 to drive the updates of node2, node3, and node4. Then, use node2 to drive the update of node1.

The driving system requires root user `SSH` access to each compute node the utility will update.

PREPARING FOR THE OS UPDATES



Warning

Do not install NetworkManager on the DB system. Installing this package and rebooting the system results in severe loss of access to the system.

- Before you begin your updates, review *Exadata Cloud Service Software Versions* ([Doc ID 2333222.1](#)) to determine the latest software version and target version to use.
- Some steps in the update process require you to specify a YUM repository. The YUM repository URL is:

```
http://yum-<region_key>.oracle.com/repo/EngineeredSystems/exadata/dbserver/<latest_
version>/base/x86_64.
```

[Region keys](#) are three-letter abbreviations, for example PHX.

You can run the following `curl` command to determine the latest version of the YUM repository for your DB system region:

```
curl -s -X GET http://yum-<region_
key>.oracle.com/repo/EngineeredSystems/exadata/dbserver/index.html |egrep "18.1."
```

This example returns the most current version of the YUM repository for the US West (Phoenix):

```
curl -s -X GET http://yum-phx.oracle.com/repo/EngineeredSystems/exadata/dbserver/index.html
|egrep "18.1."
<a href="18.1.4.0.0/">18.1.4.0.0/</a> 01-Mar-2018 03:36 -
```

- To apply OS updates, the DB system's VCN must be configured to allow access to the YUM repository. For more information, see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos.](#)

To update the OS on all compute nodes of an Exadata DB system

This example procedure assumes the following:

- The system has two compute nodes, `node1` and `node2`.
- The target version is 18.1.4.0.0.180125.3.
- Each of the two nodes is used as the driving system for the update on the other one.

1. Gather the environment details.

- a. SSH to `node1` as `root` and run the following command to determine the version of Exadata:

```
[root@node1]# imageinfo -ver
12.2.1.1.4.171128
```

- b. Switch to the `grid` user, and identify all computes in the cluster.

```
[root@node1]# su - grid
[grid@node1]$ olsnodes
node1
node1
```

2. Configure the driving system.

- a. Switch back to the `root` user on `node1`, check whether a root ssh key pair (`id_rsa` and `id_rsa.pub`) already exists. If not, then generate it.

```
[root@node1 .ssh]# ls /root/.ssh/id_rsa*
ls: cannot access /root/.ssh/id_rsa*: No such file or directory
[root@node1 .ssh]# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
93:47:b0:83:75:f2:3e:e6:23:b3:0a:06:ed:00:20:a5
root@node1.fraad1client.exadataclientne.oraclevcn.com
The key's randomart image is:
+--[ RSA 2048]-----+
|O..    + .    |
|O.     o *    |
|E      . o o   |
| . .    =     |
| o .    S =   |
| +      = .   |
| +      o o   |
| . .    + .   |
|      ...    |
+-----+

```

- b. Distribute the public key to the target nodes, and verify this step. In this example, the only target node is node2.

```
[root@node1 .ssh]# scp -i ~opc/.ssh/id_rsa ~root/.ssh/id_rsa.pub opc@node2:/tmp/id_
rsa.node1.pub
id_rsa.pub

[root@node2 ~]# ls -al /tmp/id_rsa.node1.pub
-rw-r--r-- 1 opc opc 442 Feb 28 03:33 /tmp/id_rsa.node1.pub
[root@node2 ~]# date
Wed Feb 28 03:33:45 UTC 2018

```

- c. On the target node (node2, in this example), add the root public key of node1 to the root `authorized_keys` file.

```
[root@node2 ~]# cat /tmp/id_rsa.node1.pub >> ~root/.ssh/authorized_keys

```

- d. Download `dbserver_patch.zip` as `p21634633_12*_Linux-x86-64.zip` onto the driving system (`node1`, in this example), and unzip it. See *dbnodeupdate.sh* and *dbserver_patch.zip: Updating Exadata Database Server Software using the DBNodeUpdate Utility and patchmgr* ([Doc ID 1553103.1](#)) for information about the files in this `.zip`.

```
[root@node1 patch]# mkdir /root/patch
[root@node1 patch]# cd /root/patch
[root@node1 patch]# unzip p21634633_181400_Linux-x86-64.zip
Archive:  p21634633_181400_Linux-x86-64.zip   creating: dbserver_patch_5.180228.2/
  creating: dbserver_patch_5.180228.2/ibdiagtools/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/cable_check.pl
  inflating: dbserver_patch_5.180228.2/ibdiagtools/setup-ssh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/VERSION_FILE
  extracting: dbserver_patch_5.180228.2/ibdiagtools/xmonib.sh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/monitord
  inflating: dbserver_patch_5.180228.2/ibdiagtools/checkbadlinks.pl
  creating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/VerifyTopologyUtility.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/verifylib.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Node.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Rack.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Group.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Switch.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topology-zfs
  inflating: dbserver_patch_5.180228.2/ibdiagtools/dcli
  creating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteScriptGenerator.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/CommonUtils.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/SolarisAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/LinuxAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteLauncher.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/remoteConfig.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/spawnProc.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/runDiagnostics.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/OSAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/SampleOutputs.txt
  inflating: dbserver_patch_5.180228.2/ibdiagtools/infinicheck
  inflating: dbserver_patch_5.180228.2/ibdiagtools/ibping_test
```

```
inflating: dbserver_patch_5.180228.2/ibdiagtools/tar_ibdiagtools
inflating: dbserver_patch_5.180228.2/ibdiagtools/verify-topology
inflating: dbserver_patch_5.180228.2/installfw_exadata_ssh
  creating: dbserver_patch_5.180228.2/linux.db.rpms/
inflating: dbserver_patch_5.180228.2/md5sum_files.lst
inflating: dbserver_patch_5.180228.2/patchmgr
inflating: dbserver_patch_5.180228.2/xcp
inflating: dbserver_patch_5.180228.2/ExadataSendNotification.pm
inflating: dbserver_patch_5.180228.2/ExadataImageNotification.pl
inflating: dbserver_patch_5.180228.2/kernelupgrade_oldbios.sh
inflating: dbserver_patch_5.180228.2/cellboot_usb_pci_path
inflating: dbserver_patch_5.180228.2/exadata.img.env
inflating: dbserver_patch_5.180228.2/README.txt
inflating: dbserver_patch_5.180228.2/exadataLogger.pm
inflating: dbserver_patch_5.180228.2/patch_bug_26678971
inflating: dbserver_patch_5.180228.2/dcli
inflating: dbserver_patch_5.180228.2/patchReport.py
extracting: dbserver_patch_5.180228.2/dbnodeupdate.zip
  creating: dbserver_patch_5.180228.2/plugins/
inflating: dbserver_patch_5.180228.2/plugins/010-check_17854520.sh
inflating: dbserver_patch_5.180228.2/plugins/020-check_22468216.sh
inflating: dbserver_patch_5.180228.2/plugins/040-check_22896791.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_bash
inflating: dbserver_patch_5.180228.2/plugins/050-check_22651315.sh
inflating: dbserver_patch_5.180228.2/plugins/005-check_22909764.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_perl
inflating: dbserver_patch_5.180228.2/plugins/030-check_24625612.sh
inflating: dbserver_patch_5.180228.2/patchmgr_functions
inflating: dbserver_patch_5.180228.2/exadata.img.hw
inflating: dbserver_patch_5.180228.2/libxcp.so.1
inflating: dbserver_patch_5.180228.2/imageLogger
inflating: dbserver_patch_5.180228.2/ExaXMLNode.pm
inflating: dbserver_patch_5.180228.2/fwverify
```

- e. Create the `dbserver_group` file that contains the list of compute nodes to update. Include the nodes listed after running the `olsnodes` command in step 1 except for the driving system node. In this example, `dbserver_group` should include only `node2`.

CHAPTER 11 Database

```
[root@node1 patch]# cd /root/patch/dbserver_patch_5.180228
[root@node1 dbserver_patch_5.180228]# cat dbs_group
node2
```

3. Run a patching precheck operation.

```
patchmgr -dbnodes dbs_group -precheck -yum_repo <yum_repository> -target_version <target_version>
-nomodify_at_prereq
```



Important

You must run the precheck operation with the `-nomodify_at_prereq` option to prevent any changes to the system that could impact the backup you take in the next step. Otherwise, the backup might not be able to roll back the system to its original state, should that be necessary.

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -precheck -yum_repo
http://yum-phx.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_
version 18.1.4.0.0.180125.3 -nomodify_at_prereq

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****
2018-02-28 21:22:45 +0000          :Working: DO: Initiate precheck on 1 node(s)
```

```
2018-02-28 21:24:57 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:26:15 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:26:47 +0000      :Working: DO: dbnodeupdate.sh running a precheck on node(s).
2018-02-28 21:28:23 +0000      :SUCCESS: DONE: Initiate precheck on node(s).
```

4. Back up the current system.

```
patchmgr -dbnodes dbs_group -backup -yum_repo <yum_repository> -target_version <target_version>
-allow_active_network_mounts
```



Important

This is the proper stage to take the backup, before any modifications are made to the system.

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -backup -yum_repo
http://yum-phx.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_
version 18.1.4.0.0.180125.3 -allow_active_network_mounts

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on 1 node(s).
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on node(s)
2018-02-28 21:29:01 +0000      :Working: DO: Check free space and verify SSH equivalence for
```

```

the root user to node2
2018-02-28 21:30:18 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:30:51 +0000      :Working: DO: dbnodeupdate.sh running a backup on node(s) .
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on node(s) .
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on 1 node(s) .
    
```

5. Remove all custom RPMs from the target compute nodes that will be updated. Custom RPMs are reported in precheck results. They include RPMs that were manually installed after the system was provisioned.



Note

- If you are updating the system from version 12.1.2.3.4.170111, and the precheck results include `krb5-workstation-1.10.3-57.el6.x86_64`, remove it. (This item is considered a custom RPM for this version.)
- Do **not** remove `exadata-sun-vm-computenode-exact` or `oracle-ofed-release-guest`. These two RPMs are handled automatically during the update process.

6. Run the `nohup` command to perform the update.

```

nohup patchmgr -dbnodes dbs_group -upgrade -nobackup -yum_repo <yum_repository> -target_version
<target_version> -allow_active_network_mounts &
    
```

The output should look like the following example:

```

[root@node1 dbserver_patch_5.180228]# nohup ./patchmgr -dbnodes dbs_group -upgrade -nobackup -
yum_repo http://yum-phx.oracle.com/repo/EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_
64 -target_version 18.1.4.0.0.180125.3 -allow_active_network_mounts &

*****
*****
    
```

CHAPTER 11 Database

```
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the latest release of
dbserver.patch.zip)
NOTE
NOTE    Database nodes will reboot during the update process.
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.

*****
*****

2018-02-28 21:36:26 +0000      :Working: DO: Initiate prepare steps on node(s).
2018-02-28 21:36:26 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:37:44 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:38:43 +0000      :SUCCESS: DONE: Initiate prepare steps on node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on 1 node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on node(s)
2018-02-28 21:38:49 +0000      :Working: DO: Get information about any required OS upgrades
from node(s).
2018-02-28 21:38:59 +0000      :SUCCESS: DONE: Get information about any required OS upgrades
from node(s).
2018-02-28 21:38:59 +0000      :Working: DO: dbnodeupdate.sh running an update step on all
nodes.
2018-02-28 21:48:41 +0000      :INFO    : node2 is ready to reboot.
2018-02-28 21:48:41 +0000      :SUCCESS: DONE: dbnodeupdate.sh running an update step on all
nodes.
2018-02-28 21:48:41 +0000      :Working: DO: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :SUCCESS: DONE: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :Working: DO: Waiting to ensure node2 is down before reboot.
2018-02-28 21:56:18 +0000      :Working: DO: Initiate prepare steps on node(s).
2018-02-28 21:56:19 +0000      :Working: DO: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:57:37 +0000      :SUCCESS: DONE: Check free space and verify SSH equivalence for
the root user to node2
2018-02-28 21:57:42 +0000      :SEEMS ALREADY UP TO DATE: node2
2018-02-28 21:57:43 +0000      :SUCCESS: DONE: Initiate update on node(s)
```

7. After the update operation completes, verify the version of the kernel on the compute node that was updated.

```
[root@node2 ~]# imageinfo -ver  
18.1.4.0.0.180125.3
```

8. If the driving system is a compute node that needs to be updated (as in this example), repeat steps 2 through 7 of this procedure using an updated compute node as the driving system to update the remaining compute node. In this example update, you would use `node2` to update `node1`.
9. On each compute node, run the `uptrack-install` command as root to install the available `ksplice` updates.

```
uptrack-install --all -y
```

Updating Tooling on an Exadata DB System

You can update the cloud-specific tooling included on an Exadata DB system compute node by downloading and applying an RPM file containing the latest version of the tools.



Note

Oracle highly recommends that you maintain the same version of cloud tooling across your Exadata DB system environment. Perform the following procedure on every compute node in the Exadata DB system.

PREREQUISITE

The compute nodes in the Exadata DB system must be configured to access the Oracle Cloud Infrastructure Object Storage service. For more information, see [Node Access to Object Storage: Static Route](#).

UPDATING THE CLOUD TOOLING ON EACH COMPUTE NODE MANUALLY

The method for updating the tooling depends on the tooling release that is currently installed on the compute node.

To check the installed tooling release

1. Connect to the compute node as the `opc` user.
2. Start a root-user command shell.

```
$ sudo -s
#
```

3. Use the following command to display information about the installed cloud tooling and note the release label, shown in red in the example that follows.

```
# rpm -qa | grep -i dbaastools_exa
dbaastools_exa-1.0-1+18.1.2.1.0_180511.0801.x86_64
```

In this example, the release version is **18.1.2.1.0_180511.0801**.

To update the tooling if the release label is higher than 17430

You use the `patch tools` subcommand of the `dbaascli` utility to update the cloud tooling.



Important

If you are updating the tooling on an Exadata DB system that includes a Data Guard configuration, you must perform these steps on both the primary database's DB system and on the standby database's DB system.

1. Connect as the `opc` user to the compute node.
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Check whether any cloud tooling updates are available:

```
# dbaascli patch tools list
```

Example output:

```
[root@exacs-node1 ]# dbaascli patch tools list
DBAAS CLI version 19.4.1.0.0
Executing command patch tools list
Checking tools on all nodes
Current Patchid on stb-elbdc1: 19.4.1.0.0_190822.1034
Available Patches
Patchid : 19.4.1.0.0_190827.1034
Patchid : 19.4.1.0.0_190912.0440 (LATEST)
Install tools patch using
dbaascli patch tools apply --patchid 19.4.1.0.0_190912.0440    or
dbaascli patch tools apply --patchid LATEST
All Nodes have the same tools version
```

4. In the command response, locate the patch ID of the cloud tooling update. The patch ID is listed as the "Patchid" value. If multiple patches are listed, choose the latest one.
5. Apply the patch containing the latest cloud tooling update by using one of the following methods:

- Specify the patch ID of the latest patch:

```
# dbaascli patch tools apply --patchid <patch_ID>
```

- Specify the patch ID as LATEST:

```
# dbaascli patch tools apply --patchid LATEST
```

- Run the update process in the background:

```
# dbaascli patch tools apply --patchid LATEST &
```

6. Reset the backup configuration:

```
# /var/opt/oracle/ocde/assistants/bkup/bkup
```

7. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

8. If you are updating cloud tooling on a DB system hosting a Data Guard configuration, repeat the preceding steps on the compute node of the peer (primary or standby database's) DB system.

To update the tooling if the release label is 17430 or lower

1. Download the RPM file using the Swift object storage API endpoint URL for your region.

```
wget <swift_API_endpoint>/v1/exadata/patches/dbaas_patch/shome/dbaastools_exa.rpm
```

The following example downloads the RPM file from the US West (Phoenix).

```
wget https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/exadata/patches/dbaas_patch/shome/dbaastools_exa.rpm
```

See [API Reference and Endpoints](#) for the Swift API endpoint for your region.

2. Apply the RPM file.

```
# rpm -ev dbaastools_exa
# rpm -ivh dbaastools_exa.rpm
```

3. Repeat the previous steps on each compute node in the Exadata DB system.

CONFIGURING AUTOMATIC CLOUD TOOLING UPDATES

You can configure automatic cloud tooling updates for Exadata DB systems. When you configure these updates, an entry is added to the `/etc/crontab` file to regularly check for cloud tooling updates and apply new updates to the compute node when they become available.



Note

These procedures apply only if the release label is higher than 17430.

To check whether automatic cloud tooling updates are enabled for an Exadata DB system

1. Connect to the compute node as the opc user.
2. Start a root-user command shell:

```
$ sudo -s  
#
```

3. Use the following command to check whether automatic tooling updates are enabled:

```
# dbaascli patch tools auto status
```

If the command response includes "INFO: auto rpm update is enabled", then automatic updates are enabled. If the response includes "INFO: auto rpm update is disabled", then automatic updates are disabled.

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit  
$ exit
```

5. If you are checking the status of automatic cloud tooling updates on a DB system hosting a Data Guard configuration, repeat the preceding steps on the compute node of the peer (primary or standby database's) DB system.

To enable automatic cloud tooling updates for an Exadata DB system

1. Connect to the compute node as the opc user.
2. Start a root-user command shell:

```
$ sudo -s  
#
```

3. Use the following command to enable automatic tooling updates:

```
# dbaascli patch tools auto enable
```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

5. If you are enabling automatic cloud tooling updates on a DB system hosting a Data Guard configuration, repeat the preceding steps on the compute node of the peer (primary or standby database's) DB system.

To run a tooling update on demand when automatic cloud tooling updates are enabled

You can perform an update at any time between automatic updates by running the `dbaascli patch tools auto enable` subcommand. This command checks whether there is a newer version of the tooling than the version on the compute node and applies the newer version if it finds one.

1. Connect to the compute node as the `opc` user.
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Use the following command to check for a newer tooling version and apply it:

```
# dbaascli patch tools auto execute
```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

5. If you are performing the on-demand cloud tooling update on a DB system hosting a Data Guard configuration, repeat the preceding steps on the compute node of the peer (primary or standby database's) DB system.

To disable automatic cloud tooling updates for an Exadata DB system

1. Connect to the compute node as the `opc` user.

2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Use the following command to disable automatic tooling updates:

```
# dbaascli patch tools auto disable
```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

5. If you are disabling automatic cloud tooling updates on a DB system hosting a Data Guard configuration, repeat the preceding steps on the compute node of the peer (primary or standby database's) DB system.

Patching an Exadata DB System

This topic explains how to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata DB system. The utility requires root or sudo administration privileges.



Note

You must update the cloud specific tooling on all the compute nodes in your Exadata DB system before performing the following procedures. For more information, see [Updating an Exadata DB System](#).

Prerequisites

Patches are stored in Oracle Cloud Infrastructure Object Storage, so the Exadata DB system requires access to that service. To enable this access, Oracle recommends using a service gateway with the VCN. For more information, see [Network Setup for Exadata DB Systems](#). In

that topic, pay particular attention to:

- [Service Gateway for the VCN](#)
- [Node Access to Object Storage: Static Route](#)
- [Rule Required Specifically for the Backup Network](#)

Managing Patches

To list available patches

You can produce a list of available patches using the `patch db list` subcommand of the `dbaascli` utility:

1. Connect to the compute node as the `opc` user.
For detailed instructions, see [Connecting to an Exadata DB System](#).
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the following command:

```
# dbaascli patch db list --oh=<hostname>:<oracle_home>
```

where:

- `--oh` specifies a compute node and Oracle home directory for which you want to list the available patches. In this context, an Oracle home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.

Example (Oracle Database):

```
# dbaascli patch db list --oh exacs-nodel:/u02/app/oracle/product/18.0.0.0/dbhome_1
DBAAS CLI version 18.2.3.2.0
Executing command patch db list --oh exacs-nodel:/u02/app/oracle/product/18.0.0.0/dbhome_1
INFO : EXACS patching

Available Patches
patchid :29708703 (Database Release Update : 18.7.0.0.190716 (Jul 2019))
```

CHAPTER 11 Database

```
Install database patch using
dbascli patch db apply --patchid 29708703 (Database Release Update : 18.7.0.0.190716 (Jul 2019))
--dbnames <>
```

Example (Oracle Grid Infrastructure):

```
# dbascli patch db list --oh exacs-node1:/u01/app/18.1.0.0/grid
```



Note

The list of available patches is determined by interrogating the database to establish the patches that have already been applied. When a patch is applied, the corresponding database entry is made as part of the SQL patching operation, which is executed at the end of the patch workflow. Therefore, the list of available patches may include partially applied patches along with patches that are currently being applied.

To learn more about the `patch db list` subcommand, including available options, execute the following command:

```
# dbascli patch db list ?
```

4. Exit the root-user command shell.

```
# exit
$
```

To check prerequisites before applying a patch

You can perform the prerequisites-checking operation using the `patch db prereq` subcommand of the `dbascli` utility:

1. Connect to the compute node as the `opc` user.
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the following command:

```
# dbaascli patch db prereq --patchid <patchid> --dbnames <dbname>
```

where:

- `--patchid` identifies the patch to be pre-checked.
- `--dbnames` specifies the name of the database you want to pre-check.

Example (Oracle Database):

```
# dbaascli patch db prereq --patchid 29708703 --dbnames DB18
```

Example (Oracle Grid Infrastructure):

```
# dbaascli patch db prereq --patchid 29708703-GI --dbnames GRID
```

To run the command in the background, append an ampersand (&):

```
# dbaascli patch db prereq --patchid 29708703 --dbnames DB18 &
```

To learn more about the `patch db prereq` subcommand, including available options, execute the following command:

```
# dbaascli patch db prereq ?
```

4. Exit the root-user command shell:

```
# exit
$
```

To apply a patch

You can apply a patch by using the `patch db apply` subcommand of the `dbaascli` utility.

The patching operation:

- Can be used to patch some or all of your compute nodes using one command.
- Coordinates multi-node patching in a rolling manner.
- Can execute patch-related SQL after patching all the compute nodes in the cluster.

To perform the patching operation:

1. Connect to the compute node as the `opc` user.
2. Start a root-user command shell:

```
$ sudo -s  
#
```

3. Execute the following command:

```
# dbaascli patch db apply --patchid <patchid> --dbnames <dbname> --run_datasql 1
```

where:

- `--patchid` identifies the patch to be applied.
- `--dbnames` specifies the name of the database you want to apply the patch to.
- `--run_datasql 1` instructs the command to execute patch-related SQL commands.



Note

- Patch-related SQL should only be executed after all of the compute nodes are patched. Therefore, take care not to specify this argument if you are patching a subset of nodes and further nodes remain to be patched.
- This argument can only be specified in conjunction with a patching operation on a set of compute nodes. Therefore, if you have patched all of your nodes and you did not specify this argument, you will need to manually execute the SQL commands associated with the patch. Refer to the patch documentation for further details.

Example (Oracle Database):

```
# dbaascli patch db apply --patchid 29708703 --dbnames DB18
```

Example (Oracle Grid Infrastructure):

```
# dbaascli patch db apply --patchid 29708703-GI --dbnames GRID
```

To run the command in the background, append an ampersand (&):

```
# dbaascli patch db apply --patchid 29708703 --dbnames DB18 &
```

To learn more about the `patch db apply` subcommand, including available options, execute the following command:

```
# dbaascli patch db apply ?
```

4. Exit the root-user command shell:

```
# exit
$
```

To list applied patches

You can produce a list of applied patches to determine which patches have been applied.

You can use the `opatch` utility to determine the patches that have been applied to an Oracle Database or Grid Infrastructure installation.

To produce a list of applied patches for an Oracle Database installation:

1. Connect to a compute node as the `oracle` user.
2. Go to the Oracle user's home directory:

```
$ cd
```

3. Ensure that you are in the Oracle user's home directory:

```
$ pwd
```

```
/home/oracle
```

4. Source the environment file.
Example (using the environment file for a database named "DB18"):

```
$ . DB18.env
```

5. Execute the `opatch` command with the `lspatches` option:

```
$ opatch lspatches
```

To produce a list of applied patches for Oracle Grid Infrastructure:

1. Connect to a compute node as the `opc` user.
2. Become the `grid` user:

```
$ sudo -s
```

```
# su - grid
```

3. Execute the `opatch` command with the `lspatches` option:

```
$ opatch lspatches
```

To switchback a patch

You can switchback (roll back) a patch by using the `dbaascli` utility.

The patch switchback (roll back) operation:

- Can be used to roll back a patch on some or all of your compute nodes using one command.
- Coordinates multi-node operations in a rolling manner.
- Can execute rollback-related SQL after rolling back the patch on all the compute nodes in the cluster.

To perform a patch switchback (roll back) operation:

1. Connect to the compute node as the `opc` user.
2. Start a root-user command shell:

```
$ sudo -s  
#
```

3. Execute the following command:

```
# dbaascli patch db switchback --patchid <patchid> --instance1 <hostname>:<oracle_home> --dbnames  
<dbname> --run_datasql 1
```

where:

- `--patchid` identifies the patch to be rolled back.
- `--instanceN` specifies a compute node and one or more Oracle home directories that are subject to the rollback operation. In this context, an Oracle home directory may be an Oracle Database home directory or the Oracle Grid Infrastructure home directory.

- `--dbnames` specifies the name of the database you want to apply the switchback operation to.
- `--run_datasql 1` instructs the command to execute rollback-related SQL commands.



Note

- Rollback-related SQL should only be executed after all of the compute nodes are rolled back. Therefore, take care not to specify this argument if you are rolling back a subset of nodes and further nodes remain to be rolled back.
- This argument can only be specified in conjunction with a rollback operation on a set of compute nodes. Therefore, if you have rolled back all of your nodes and you did not specify this argument, you will need to manually execute the SQL commands associated with the rollback operation. Refer to the patch documentation for further details.

Example (Oracle Database):

```
# dbaascli patch db switchback --patchid 29708703 --dbnames DB18
```

Example (Oracle Grid Infrastructure):

```
# dbaascli patch db switchback --patchid 29708703-GI --dbnames GRID
```

To run the command in the background, append an ampersand (&):

```
# dbaascli patch db switchback --patchid 29708703 --dbnames DB18 &
```

To learn more about the `patch db switchback` subcommand, including available options, execute the following command:

```
# dbaascli patch db switchback ?
```

4. Exit the root-user command shell:

```
# exit  
$
```

Monitoring a Database on an Exadata DB System

This topic explains how to access Enterprise Manager Database Express and Enterprise Manager Database Control, which are web-based tools for managing Oracle Database.

Accessing Enterprise Manager Database Express 12c

Enterprise Manager Database Express 12c (EM Express) is available on Exadata DB system database deployments created using Oracle Database 12c Release 1 (12.1) or later.

How you access EM Express depends on whether you want to manage a CDB or PDB.

- **To manage the CDB.** When a database deployment is created, Database automatically sets port 5500 on the deployment's compute nodes for EM Express access to the CDB.
- **To manage a PDB.** For an Oracle Database 12.2 or later deployment, a single port (known as the global port) is automatically set on the deployment's compute nodes. The global port lets you use EM Express to connect to all of the PDBs in the CDB using the HTTPS port for the CDB.
For an Oracle Database 12.1 deployment, you must manually set a port on the deployment's compute nodes for each PDB you want to manage using EM Express.

For both CDBs and PDBs, you must add the port to a security list as described in [Updating the Security List](#).

CHAPTER 11 Database

To confirm the port that is in use for a specific database, connect to the database as a database administrator and execute the query shown in the following example:

```
SQL> select dbms_xdb_config.getHttpsPort() from dual;
```

```
DBMS_XDB_CONFIG.GETHTTPSPT()
-----
                             5502
```

SETTING THE PORT FOR EM EXPRESS TO MANAGE A PDB (ORACLE DATABASE 12.1 ONLY)

In Oracle Database 12c Release 1, a unique HTTPS port must be configured for the root container (CDB) and each PDB that you manage using EM Express.

To configure a HTTPS port so that you can manage a PDB with EM Express:

1. Invoke SQL*Plus and log in to the PDB as the SYS user with SYSDBA privileges.
2. Execute the `DBMS_XDB_CONFIG.SETHTTPSPT` procedure.

```
SQL> exec dbms_xdb_config.sethttpspt(port-number)
```

ACCESSING EM EXPRESS

Before you access EM Express, add the port to the security list. See [Updating the Security List](#).

After you update the security list, you can access EM Express by directing your browser to the URL `https://<node-ip-address>:<port>/em`, where `node-ip-address` is the public IP address of the compute node hosting EM Express, and `port` is the EM Express port used by the database.

Accessing Enterprise Manager 11g Database Control

Enterprise Manager 11g Database Control (Database Control) is available on Exadata DB system database deployments created using Oracle Database 11g Release 2. Database Control is allocated a unique port number for each database deployment. By default, access to Database Control is provided using port 1158 for the first deployment. Subsequent deployments are allocated ports in a range starting with 5500, 5501, 5502, and so on.

You can confirm the Database Control port for a database by searching for `REPOSITORY_URL` in the `$ORACLE_HOME/host_sid/sysman/config/emd.properties` file.

Before you access Database Control, add the port for the database to the security list associated with the Exadata DB system's client subnet. For more information, see [Updating the Security List](#).

After you update the security list, you can access Database Control by directing your browser to the URL `https://<node-ip-address>:<port>/em`, where `node-ip-address` is the public IP address of the compute node hosting Database Control, and `port` is the Database Control port used by the database.

Updating the Security List

Before you can access EM Express or Database Control, you must add the port for the database to the security list associated with the Exadata DB system's data (client) subnet. To update an existing security list, complete the following steps using the Console:

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Locate the DB system in the list.
4. Note the DB system's **Client Subnet** name and click its **Virtual Cloud Network**.
5. Locate the subnet in the list, and then click its security list under **Security Lists**.
6. Click **Edit All Rules** and add an ingress rule with source type=CIDR, source CIDR= *<source CIDR>*, protocol=TCP, and port= *<port number or port range>*.
The source CIDR should be the CIDR block that includes the ports you open for the client connection.

For detailed information about creating or updating a security list, see [Security Lists](#).

Managing Exadata Databases

When you launch an Exadata DB system, an initial database is created in that system. You can create additional databases in that DB system at any time by using the Console or the Oracle Cloud Infrastructure API.

When you add a database to an Exadata DB system, the database versions you can select from depend on the current patch level of that DB system. You might have to patch your DB system to add later database versions. For information about patching the DB system, see [Patching an Exadata DB System](#).

Each new database is created in a separate database home.



Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

You can also add and remove databases, and perform other management tasks on a database by using command line utilities. For information and instructions on how to use these utilities, see [Managing Exadata Databases Manually](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).



Note

See [Known Issues](#) for information about using the Oracle Cloud Infrastructure Console, API, or CLI to manage Exadata DB systems if your system was provisioned before June 15, 2018.

Using the Console

To create a new database in an existing Exadata DB system



Note

If IORM is enabled on the DB system, the default directive will apply to the new database and system performance might be impacted. Oracle recommends that you review the IORM settings and make applicable adjustments to the configuration after the new database is provisioned.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. In the list of DB systems, find the Exadata DB system in which you want to create the database, and then click its name to display details about it.
4. Click **Create Database**.

5. In the **Create Database** dialog, enter the following:
- **Database name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
 - **Database version:** The version of the database. You can mix database versions on the Exadata DB system.
 - **PDB name** (*Optional*) For version 12.1.0.2 and later, you can specify the name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of 8 alphanumeric characters. The only special character permitted is the underscore (_).
 - **Create administrator credentials:** A database administrator `SYS` user will be created with the password you supply.
 - **Username:** `SYS`
 - **Password:** Supply the password for this user. The password must meet the following criteria:

A strong password for `SYS`, `SYSTEM`, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (`SYS`, `SYSTEM`, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
 - **Confirm password:** Re-enter the `SYS` password you specified.
 - **Select workload type:** Choose the workload type that best suits your application:
 - **Online Transactional Processing (OLTP)** configures the database for a transactional workload, with a bias towards high volumes of random data access.

- **Decision Support System (DSS)** configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations.
 - **Configure database backups:** Specify the settings for backing up the database to Object Storage:
 - **Enable automatic backup:** Check the check box to enable automatic incremental backups for this database.
 - **Backup retention period:** If you enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The default selection is 30 days.
 - **Backup Scheduling:** If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, the six-hour default window of 00:00 to 06:00 (in the time zone of the DB system's region) is used for your database. See [Backup Scheduling](#) for more information.
6. Click **Show Advanced Options** to specify advanced options for the database:
- **Character set:** The character set for the database. The default is AL32UTF8.
 - **National character set:** The national character set for the database. The default is AL16UTF16.
 - Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Database**.

When the database creation is complete, the status changes from Provisioning to Available.

To terminate a database

Oracle recommends that you create a final backup before you terminate any production (non-

test) database. See [Managing Exadata Database Backups by Using bkup_api](#) to learn how to back up an Exadata database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. In the list of DB Systems, find the DB System that contains the database you want to terminate, and then click its name to display details about it.
4. In the list of databases, find the database you want to terminate, and then click its name to display details about it.
5. Click **Actions**, and then click **Terminate**.
6. In the confirmation dialog, indicate whether you want to back up the database before terminating it, and type the name of the database to confirm the termination.
7. Click **Terminate Database**.
The database's status indicates Terminating.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage databases.

Database homes:

- [ListDbHomes](#)
- [GetDbHome](#)
- [CreateDbHome](#)
- [DeleteDbHome](#)

Databases:

- [ListDatabases](#)
- [GetDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Changing the Database Passwords

The password that you specify in the Database Admin Password field when you create a new Exadata DB system or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB Admin credentials. Use the following procedures if you need to change passwords for an existing database.

Note that if you are enabling Data Guard for a database, the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

To change the SYS password for an Exadata DB system database

1. Log onto the DB system host as `opc`.
2. Run the following command:

```
sudo dbaascli database changepassword --dbname <database_name>
```

To change the TDE wallet password for an Exadata DB system database

1. Log onto the DB system host as `opc`.
2. Run the following command:

```
sudo dbaascli tde changepassword --dbname <database_name>
```

Managing Exadata Databases Manually

Exadata DB systems include these command line tools for performing various tasks to manage individual databases:

- `dbaasapi` - For adding and removing databases from the Exadata DB system. See [Using dbaasapi](#).
- `dbaascli` - For a variety of life-cycle and administration operations such as:
 - Starting and stopping a database
 - Starting and stopping the Oracle Net listener
 - Viewing information about Oracle Homes
 - Moving a database to another Oracle Home
 - Deleting an unused Oracle Home
 - Performing database configuration changes
 - Managing Oracle Database software images
 - Managing pluggable databases (PDBs)
 - Performing database recovery
 - Rotating the master encryption key

For details about how to use this CLI, see [The dbaascli Utility](#).

Using dbaasapi

You can use the `dbaasapi` command line utility to create and delete databases on an Exadata DB system. The utility operates like a REST API. It reads a JSON request body and produces a JSON response body in an output file.

The utility is located in the `/var/opt/oracle/dbaasapi/` directory on the compute nodes and must be run as the root user.

To learn how to add or remove Exadata databases by using the Oracle Cloud Infrastructure Console or API instead, see [Managing Exadata Databases](#).



Note

- Databases that you create by using `dbaasapi` do *not* appear in your Exadata DB system's list of databases that you see in the Console.
- You must update the cloud-specific tooling on all the compute nodes in your Exadata DB system before performing the following procedures. For more information, see [Updating an Exadata DB System](#).
- Only one `dbaasapi` operation can execute at a given time. Oracle recommends that you check the status of an operation to ensure it completed before you run another operation.

Prerequisites

If you plan to create a database and store its backups in the Oracle Cloud Infrastructure Object Storage, refer to the prerequisites in [Managing Exadata Database Backups](#), and ensure that the system meets the networking requirements for backing up to Object Storage. Review the [Create Database Parameters](#) and gather the information you'll need to supply in the input file you create for the `dbaasapi` operation.



Warning

Oracle recommends that you avoid specifying parameter values that include confidential information when you use the `dbaasapi` commands.

Creating a Database

The following procedure creates directory called `dbinput`, a sample input file called `myinput.json`, and a sample output file called `createdb.out`.

1. SSH to a compute node in the Exadata DB system.

```
ssh -i <private_key_path> opc@<node_ip_address>
```

2. Log in as `opc` and then `sudo` to the root user.

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Make a directory for the input file and change to the directory.

```
[root@dbsys ~]# mkdir -p /home/oracle/dbinput  
# cd /home/oracle/dbinput
```

4. Create the input file in the directory. The following sample file will create a database configured to store backups in an existing bucket in Object Storage. For parameter descriptions, see [Create Database Parameters](#).

```
{  
  "object": "db",  
  "action": "start",  
  "operation": "createdb",  
  "params": {  
    "nodelist": "",  
    "dbname": "exadb",  
    "edition": "EE_EP",  
    "version": "12.1.0.2",  
    "adminPassword": "<password>",  
    "sid": "exadb",  
    "pdbName": "PDB1",  
    "charset": "AL32UTF8",  
    "ncharset": "AL16UTF16",  
    "backupDestination": "OSS",  
    "cloudStorageContainer": "https://swiftobjectstorage.<region_<br>name>.oraclecloud.com/v1/mycompany/DBBackups",  
    "cloudStorageUser": "<name@example.com>",
```

```
    "cloudStoragePwd":      "<auth_token>"
  },
  "outputfile": "/home/oracle/createdb.out",
  "FLAGS": ""
}
```

5. Run the utility and specify the input file.

```
[root@dbsys ~]# /var/opt/oracle/dbaasapi/dbaasapi -i myinput.json
```

6. Check the output file and note the ID.

```
[root@dbsys ~]# cat /home/oracle/createdb.out
{
  "msg" : "",
  "object" : "db",
  "status" : "Starting",
  "errmsg" : "",
  "outputfile" : "/home/oracle/createdb.out",
  "action" : "start",
  "id" : "170",
  "operation" : "createdb",
  "logfile" : "/var/opt/oracle/log/gsa1/dbaasapi/db/createdb/1.log"
}
```

7. Create a JSON file to check the database creation status. Note the action of "status". Replace the ID and the dbname with the values from the previous steps.

```
{
  "object": "db",
  "action": "status",
  "operation": "createdb",
  "id": 170,
  "params": {
    "dbname": "exadb"
  },
  "outputfile": "/home/oracle/createdb.out",
  "FLAGS": ""
}
```

8. Run the utility with the status file as input and then check the utility output. Rerun the status action regularly until the response indicates that the operation succeeded or failed.

CHAPTER 11 Database

```
[root@dbsys ~]# /var/opt/oracle/dbaasapi/dbaasapi -i db_status.json

[root@dbsys ~]# cat /home/oracle/createdb.out

{
  "msg" : "Sync sqlnet file...[done]\\n##Done executing tde\\nWARN: Could not register elogger_
parameters: elogger.pm::_init: /var/opt/oracle/dbaas_acfs/events does not exist\\n##Invoking
assistant bkup\\nUsing cmd : /var/opt/oracle/ocde/assistants/bkup/bkup -out
/var/opt/oracle/ocde/res/bkup.out -sid=\"exadb1\" -reco_grp=\"RECO1\" -
hostname=\"edldb01.data.customer1.oraclevcn.com\" -oracle_
home=\"/u02/app/oracle/product/12.1.0/dbhome_5\" -dbname=\"exadb\" -dbtype=\"exarac\" -
exabm=\"yes\" -edition=\"enterprise\" -bkup_cfg_files=\"no\" -acfs_vol_
dir=\"/var/opt/oracle/dbaas_acfs\" -bkup_oss_url=\"bkup_oss_url\" -bkup_oss_user=\"bkup_oss_
user\" -version=\"12102\" -oracle_base=\"/u02/app/oracle\" -firststrun=\"no\" -action=\"config\" -
bkup_oss=\"no\" -bkup_disk=\"no\" -data_grp=\"DATA1\" -action=config \\n\\n##Done executing
bkup\\nWARN: Could not register elogger_parameters: elogger.pm::_init: /var/opt/oracle/dbaas_
acfs/events does not existRemoved all entries from creg file : /var/opt/oracle/creg/exadb.ini
matching passwd or decrypt_key\\n\\n#### Completed OCDE Successfully ####\\nWARN: Could not
register elogger_parameters: elogger.pm::_init: /var/opt/oracle/dbaas_acfs/events does not
exist",
  "object" : "db",
  "status" : "Success",
  "errmsg" : "",
  "outputfile" : "/home/oracle/createdb_exadb.out",
  "action" : "start",
  "id" : "170",
  "operation" : "createdb",
  "logfile" : "/var/opt/oracle/log/exadb/dbaasapi/db/createdb/170.log"
}
```

CREATE DATABASE PARAMETERS

Use the following parameters to create a database.

Parameter	Description
object	The value "db".
action	The value "start".

CHAPTER 11 Database

Parameter	Description
operation	The value "createdb".
nodelist	The value "" (an empty string). The database will be created across all nodes in the cluster.
dbname	The database name, in quotes.
edition	The value "EE_EP". (Only Enterprise Edition - Extreme Performance is supported .)
version	The database version as 18.0.0.0, 12.2.0.1, 12.1.0.2, or 11.2.0.4, in quotes.
adminPassword	The administrator (SYS and SYSTEM) password to use for the new database, in quotes. The password must be nine to thirty characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be <code>_</code> , <code>#</code> , or <code>-</code> .
sid	The SID of the database, in quotes.
pdbName	The name of the pluggable database, in quotes.

Parameter	Description
charset	<p>The database character set, in quotes.</p> <p>Allowed values</p> <p>AL32UTF8, AR8ADOS710, AR8ADOS720, AR8APTEC715, AR8ARABICMACS, AR8ASMO8X, AR8ISO8859P6, AR8MSWIN1256, AR8MUSSAD768, AR8NAFITHA711, AR8NAFITHA721, AR8SAKHR706, AR8SAKHR707, AZ8ISO8859P9E, BG8MSWIN, BG8PC437S, BLT8CP921, BLT8ISO8859P13, BLT8MSWIN1257, BLT8PC775, BN8BSCII, CDN8PC863, CEL8ISO8859P14, CL8ISO8859P5, CL8ISOIR111, CL8KOI8R, CL8KOI8U, CL8MACCYRILLICS, CL8MSWIN1251, EE8ISO8859P2, EE8MACCES, EE8MACCROATIANS, EE8MSWIN1250, EE8PC852, EL8DEC, EL8ISO8859P7, EL8MACGREEKS, EL8MSWIN1253, EL8PC437S, EL8PC851, EL8PC869, ET8MSWIN923, HU8ABMOD, HU8CWI2, IN8ISCII, IS8PC861, IW8ISO8859P8, IW8MACHEBREWS, IW8MSWIN1255, IW8PC1507, JA16EUC, JA16EUCTILDE, JA16SJIS, JA16SJISTILDE, JA16VMS, KO16KSCCS, KO16MSWIN949, LA8ISO6937, LA8PASSPORT, LT8MSWIN921, LT8PC772, LT8PC774, LV8PC1117, LV8PC8LR, LV8RST104090, N8PC865, NE8ISO8859P10, NEE8ISO8859P4, RU8BESTA, RU8PC855, RU8PC866, SE8ISO8859P3, TH8MACTHAIS, TH8TISASCII, TR8DEC, TR8MACTURKISHS, TR8MSWIN1254, TR8PC857, US7ASCII, US8PC437, UTF8, VN8MSWIN1258, VN8VN3, WE8DEC, WE8DG, WE8ISO8859P15, WE8ISO8859P9, WE8MACROMAN8S, WE8MSWIN1252, WE8NCR4970, WE8NEXTSTEP, WE8PC850, WE8PC858, WE8PC860, WE8ROMAN8, ZHS16CGB231280, ZHS16GBK, ZHT16BIG5, ZHT16CCDC, ZHT16DBT, ZHT16HKSCS, ZHT16MSWIN950, ZHT32EUC, ZHT32SOPS, ZHT32TRIS</p>

Parameter	Description
ncharset	The database national character set. The value AL16UTF16 or UTF8, in quotes.
backupDestination	<p>The database backup destination, in quotes. You can configure the following backup destinations.</p> <p>NONE No backup destination is configured.</p> <p>DISK Configure database backups to the local disk Fast Recovery Area.</p> <p>OSS Configure database backups to an existing bucket in the Oracle Cloud Infrastructure Object Storage service. You must specify all the <code>cloudStorage</code> parameters.</p> <p>BOTH Configure database backups to both local disk and an existing bucket in Object Storage. You must specify all the <code>cloudStorage</code> parameters.</p> <p>For example:</p> <pre>"backupDestination": "BOTH"</pre>

Parameter	Description
cloudStorageContainer= <i><swift_url></i>	<p>Required if you specify a backup destination of <code>OSS</code> or <code>BOTH</code>. The Object Storage URL, your Oracle Cloud Infrastructure tenant, and an existing bucket in the object store to use as the backup destination, in the following format:</p> <pre>https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>/<bucket></pre> <p>See Regions and Availability Domains to look up the region name string.</p> <p>For example:</p> <pre>"cloudStorageContainer":"https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<company_name>/DBBackups"</pre>
cloudStorageUser= <i><user_name></i>	<p>Required if you specify a backup destination of <code>OSS</code> or <code>BOTH</code>. The user name for the Oracle Cloud Infrastructure user account, for example:</p> <pre>"cloudStorageUser":"name@company.com"</pre> <p>This is the user name you use to sign in to the Console. The user name must be a member of the Administrators group, as described in Prerequisites.</p>

Parameter	Description
cloudStoragePwd= <i><auth_token></i>	Required if you specify a backup destination of <code>OSS</code> or <code>BOTH</code> . The auth token generated by using the Console or IAM API, in quotes, for example: "cloudStoragePwd": " <i><auth_token></i> " For more information, see Managing User Credentials . This is not the password for the Oracle Cloud Infrastructure user.
outputfile	The absolute path for the output of the request, for example, "outputfile": "/home/oracle/createdb.out".
FLAGS	The value "" (an empty string).

Deleting a Database

Oracle recommends that you create a final backup before you delete any production (non-test) database. See [Managing Exadata Database Backups by Using bkup_api](#) to learn how to back up an Exadata database.

1. SSH to a compute node in the Exadata DB system.

```
ssh -i <private_key_path> opc@<node_ip_address>
```

2. Log in as opc and then sudo to the root user.

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Make a directory for the input file and change to the directory.

```
[root@dbsys ~]# mkdir -p /home/oracle/dbinput
```

```
# cd /home/oracle/dbinput
```

4. Create the input file in the directory and specify the database name to delete and an output file. For more information, see [Delete Database Parameters](#) .

```
{
  "object": "db",
  "action": "start",
  "operation": "deletedb",
  "params": {
    "dbname": "exadb"
  },
  "outputfile": "/home/oracle/delete_exadb.out",
  "FLAGS": ""
}
```

5. Run the utility and specify the input file.

```
[root@dbsys ~]# /var/opt/oracle/dbaasapi/dbaasapi -i myinput.json
```

6. Check the output file and note the ID.

```
[root@ed1db01 ~]# cat /home/oracle/delete_exadb.out
{
  "msg" : "",
  "object" : "db",
  "status" : "Starting",
  "errmsg" : "",
  "outputfile" : "/home/oracle/deletedb.out",
  "action" : "start",
  "id" : "17",
  "operation" : "deletedb",
  "logfile" : "/var/opt/oracle/log/exadb/dbaasapi/db/deletedb/17.log"
}
```

7. Create a JSON file to check the database deletion status. Note the action of "status" in the sample file below. Replace the ID and the dbname with the values from the previous steps.

```
{
  "object": "db",
  "action": "status",
  "operation": "deletedb",
```

```

"id": 17,
"params": {
  "dbname": "exadb"
},
"outputfile": "/home/oracle/deletedb.out",
"FLAGS": ""
}

```

8. Run the utility with the status file as input and then check the utility output. Rerun the status action regularly until the response indicates that the operation succeeded.

```

[root@dbsys ~]# /var/opt/oracle/dbaasapi/dbaasapi -i db_status.json

[root@dbsys ~]# cat /home/oracle/deletedb.out

{
  "msg" : "Using cmd : su - root -c \"/var/opt/oracle/ocde/assistants/dg/dgcc -dbname exadb -
action delete\" \\n\\n##Done executing dg\\nWARN: Could not register elogger_parameters:
elogger.pm::_init: /var/opt/oracle/dbaas_acfs/events does not exist\\n##Invoking assistant
bkup\\nUsing cmd : /var/opt/oracle/ocde/assistants/bkup/bkup -out
/var/opt/oracle/ocde/res/bkup.out -bkup_oss_url=\"bkup_oss_url\" -bkup_daily_time=\"0:13\" -bkup_
oss_user=\"bkup_oss_user\" -dbname=\"exadb\" -dbtype=\"exarac\" -exabm=\"yes\" -firstrun=\"no\" -
action=\"delete\" -bkup_cfg_files=\"no\" -bkup_oss=\"no\" -bkup_disk=\"no\" -action=delete
\\n\\n##Done executing bkup\\nWARN: Could not register elogger_parameters: elogger.pm::_init:
/var/opt/oracle/dbaas_acfs/events does not exist\\n##Invoking assistant dbda\\nUsing cmd :
/var/opt/oracle/ocde/assistants/dbda/dbda -out /var/opt/oracle/ocde/res/dbda.out -em=\"no\" -pga_
target=\"2000\" -dbtype=\"exarac\" -sga_target=\"2800\" -action=\"delete\" -build=\"no\" -
nid=\"no\" -dbname=\"exadb\" -action=delete \\n",
  "object" : "db",
  "status" : "InProgress",
  "errmsg" : "",
  "outputfile" : "/home/oracle/deletedb.out",
  "action" : "start",
  "id" : "17",
  "operation" : "deletedb",
  "logfile" : "/var/opt/oracle/log/exadb/dbaasapi/db/deletedb/17.log"
}

```

DELETE DATABASE PARAMETERS

Use the following parameters to delete a database.

Parameter	Description
object	The value "db".
action	The value "start".
operation	The value "deletedb".
dbname	The database name, in quotes.
outputfile	The absolute path for the output of the request, for example, "/home/oracle/deletedb.out".
FLAGS	The value "" (an empty string).

Managing Exadata Database Backups

This topic explains how to work with Exadata database backups managed by Oracle Cloud Infrastructure. You do this by using the Console or the API. (For unmanaged backups, see [Managing Exadata Database Backups by Using bkup_api.](#))



Important

If you previously used `bkup_api` to configure backups and then you switch to using the Console or the API for backups:

- A new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to work.
- `bkup_api` uses cron jobs to schedule backups. These jobs are not automatically removed when you switch to using managed backups.

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Prerequisites

- Review the information and instructions in [Configuring a Static Route for Accessing the Object Store](#) and ensure that you configure the static route for the backup subnet on each compute node in the Exadata DB system.
- Your DB system must have connectivity to the applicable Swift endpoint for Object Storage. See <https://www.oracle.com/cloud/storage/object-storage-faq.html> for information about the Swift endpoints to use.



Important

To avoid backup failures, ensure that the database's archiving mode is set to `ARCHIVELOG` (the default).

Using the Console

You can use the Console to enable automatic incremental backups, create full backups on demand, and view the list of managed backups for a database. The Console also allows you to delete full backups.



Note

The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `bkup_api`).

All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.

The database and DB system must be in an “Available” state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching operations) while a backup operation is in progress. If an automatic backup operation fails, the Database service retries the operation during the next day’s backup window. If an on-demand full backup fails, you can try the operation again when the DB system and database availability are restored.

AUTOMATIC INCREMENTAL BACKUPS

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

BACKUP RETENTION

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

BACKUP SCHEDULING

The automatic backup process starts at any time during your daily backup window. You can optionally specify a 2-hour scheduling window for your database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting on

an even-numbered hour (for example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM). Backups jobs do not necessarily complete within the scheduling window

The default backup window of 00:00 to 06:00 in the time zone of the DB system's region is assigned to your database if you do not specify a window. Note that the default backup scheduling window is six hours long, while the windows you specify are two hours long. See [note](#) for backup window time zone information.



Note

- **Data Guard** - You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association. However, automatic backups for that database will not be created until it assumes the primary role.
- **Retention Period Changes** - If you shorten your database's automatic backup retention period in the future, existing backups falling outside the updated retention period are deleted by the system.
- **Object Storage Costs** - Automatic backups incur Object Storage usage costs.

ON-DEMAND FULL BACKUPS

You can create a full backup of your database at any time.

STANDALONE BACKUPS

When you terminate a DB system or a database, all of its resources are deleted, along with any automatic backups. Full backups remain in Object Storage as standalone backups. You can use a standalone backup to create a new database.

To configure automatic backups for a database

When you launch a DB system, you can optionally enable automatic backups for the initial database. Use this procedure to enable or disable automatic backups after the database is created.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.
A list of databases is displayed.
4. Find the database for which you want to enable or disable automatic backups, and click its name to display database details. The details indicate whether automatic backups are enabled.
5. Click **Configure Automatic Backups**.
6. In the **Configure Automatic Backups** dialog, check or uncheck **Enable Automatic Backup**, as applicable. If you are enabling automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The default selection is 30 days.
7. Click **Save Changes**.

To create an on-demand full backup of a database

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.
A list of databases is displayed.

4. Find the database for which you want to create an on-demand full backup and click its name to display database details.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. Click **Create Backup**.

To delete full backups from Object Storage



Note

You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in Object Storage for 30 days, after which time they are automatically deleted.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system where the database is located and click the DB system name to display details.
A list of databases is displayed.
4. Find the database you are interested in and click its name to display database details.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Delete**.
7. Confirm when prompted.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage database backups:

- [ListBackups](#)
- [GetBackup](#)
- [CreateBackup](#)
- [DeleteBackup](#)
- [UpdateDatabase](#) - To enable and disable automatic backups.

For the complete list of APIs for the Database service, see [Database Service API](#).

WHAT'S NEXT?

See [Recovering an Exadata Database from Object Storage](#).

Managing Exadata Database Backups by Using bkup_api

You can use Exadata's backup utility, `bkup_api`, to back up databases on an Exadata DB system to an existing bucket in the Oracle Cloud Infrastructure Object Storage service and to the local disk Fast Recovery Area.

For backups managed by Oracle Cloud Infrastructure, see [Managing Exadata Database Backups](#).

This topic explains how to:

- Create a backup configuration file that indicates the backup destination, when the backup should run, and how long backups are retained. If the backup destination is Object Storage, the file also contains the credentials to access the service.
- Associate the backup configuration file with a database. The database will be backed up as scheduled, or you can create an on-demand backup.



Note

You must update the cloud-specific tooling on all the compute nodes in your Exadata DB system before performing the following procedures. For more information, see [Updating an Exadata DB System](#).

Prerequisites

- The Exadata DB system requires access to the Oracle Cloud Infrastructure Object Storage service. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see [Network Setup for Exadata DB Systems](#). In that topic, pay particular attention to:
 - [Service Gateway for the VCN](#)
 - [Node Access to Object Storage: Static Route](#)
 - [Rule Required Specifically for the Backup Network](#)
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket. For more information, see [Managing Buckets](#).
- An [auth token](#) generated by Oracle Cloud Infrastructure. You can use the Console or the IAM API to generate the password. For more information, see [Working with Auth Tokens](#).
- The user name specified in the backup configuration file must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to *all* of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

```
Allow group <group_name> to manage objects in compartment <compartment_name> where  
target.bucket.name = '<bucket_name>'
```

```
Allow group <group_name> to read buckets in compartment <compartment_name>
```

For more information about adding a user to a group, see [Managing Groups](#). For more information about policies, see [Getting Started with Policies](#).

Default Backup Configuration

The backup configuration follows a set of Oracle best-practice guidelines:

- Full (level 0) backup of the database followed by rolling incremental (level 1) backups on a seven-day cycle (a 30-day cycle for the Object Storage destination).
- Full backup of selected system files.
- Automatic backups daily at a specific time set during the database deployment creation process.

Retention period:

- Both Object Storage and local storage: 30 days, with the 7 most recent days' backups available on local storage.
- Object Storage only: 30 days.
- Local storage only: Seven days.

Encryption:

- Both Object Storage and local storage: All backups to cloud storage are encrypted.
- Object Storage only: All backups to cloud storage are encrypted.

Managing Backups

To create a backup configuration file



Important

The following procedure must be performed on the first



compute node in the Exadata DB system. To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

1. SSH to the first compute node in the Exadata DB system.

```
ssh -i <private_key_path> opc@<node_1_ip_address>
```

2. Log in as `opc` and then `sudo` to the root user.

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Create a new backup configuration file in `/var/opt/oracle/ocde/assistants/bkup/` as shown in the sample configuration file below. This example uses the file name `bkup.cfg`, but you can provide your own file name. The following file schedules a backup to both local storage and an existing bucket in Object Storage. The parameters are described below this procedure.

```
[root@dbsys ~]# cd /var/opt/oracle/ocde/assistants/bkup/
```

```
vi bkup.cfg
```

```
bkup_disk=yes
```

```
bkup_oss=yes
```

```
bkup_oss_url=https://swiftobjectstorage.<region>.oraclecloud.com/v1/companyabc/DBBackups
```

```
bkup_oss_user=jsmith@example.com
```

```
bkup_oss_passwd=<password>
```

```
bkup_oss_recovery_window=7
```

```
bkup_daily_time=06:45
```

4. Change the permissions of the file.

```
[root@dbsys bkup]# chmod 600 bkup.cfg
```

CHAPTER 11 Database

5. Use the following command to install the backup configuration, configure the credentials, schedule the backup, and associate the configuration with a database name.

```
[root@dbsys bkup]# ./bkup -cfg bkup.cfg -dbname=<database_name>
```

The backup is scheduled via cron and can be viewed at `/etc/crontab`.

When the scheduled backup runs, you can check its progress with the following command.

```
[root@dbsys bkup]# /var/opt/oracle/bkup_api/bkup_api bkup_status
```

The backup configuration file parameters are described in the following table:

Parameter	Description
<code>bkup_disk=[yes no]</code>	Whether to back up locally to disk (Fast Recovery Area).
<code>bkup_oss=[yes no]</code>	Whether to back up to Object Storage. If yes, you must also provide the parameters <code>bkup_oss_url</code> , <code>bkup_oss_user</code> , <code>bkup_oss_passwd</code> , and <code>bkup_oss_recovery_window</code> .

Parameter	Description
bkup_oss_url= <i><swift_url></i>	<p>Required if bkup_oss=yes.</p> <p>The Object Storage URL including the tenant and bucket you want to use. The URL is:</p> <pre>https:// swiftobjectstorage .<region_ name> . oraclecloud . com /v1/<tenant>/<bucket></pre> <p>where <i><tenant></i> is the <i>lowercase</i> tenant name (even if it contains uppercase characters) that you specify when signing in to the Console and <i><bucket></i> is the name of the existing bucket you want to use for backups.</p>

Parameter	Description
<code>bkup_oss_user= <oci_user_name></code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>The user name for the Oracle Cloud Infrastructure user account, for example <code>jsmith@<example>.com</code>. The user must be a member of the Administrators group, as described in Prerequisites.</p> <p>This is the user name you use to sign in to the Console.</p>
<code>bkup_oss_passwd= <auth_token></code>	<p>Required if <code>bkup_oss=yes</code>.</p> <p>The auth token generated by using the Console or IAM API, as described in Prerequisites.</p> <p>This is not the password for the Oracle Cloud Infrastructure user.</p>

Parameter	Description
<code>bkup_oss_recovery_window=<i>n</i></code>	Required if <code>bkup_oss=yes</code> . The number of days for which backups and archived redo logs are maintained in the Object Storage bucket. Specify 1 to 30 days.
<code>bkup_daily_time=<i>hh:mm</i></code>	The time at which the daily backup is scheduled, specified in hours and minutes (hh:mm), in 24-hour format.

To create an on-demand backup

You can use the `bkup_api` utility to create an on-demand backup of a database.

1. SSH to the first compute node in the Exadata DB system.

```
ssh -i <private_key_path> opc@<node_1_ip_address>
```

To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

2. Log in as `opc` and then `sudo` to the root user.

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. You can let the backup follow the current retention policy, or you can create a long-term

backup that persists until you delete it:

- To create a backup that follows the current retention policy, enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_start --dbname=<database_name>
```

- To create a long-term backup, enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_start --keep --dbname=<database_name>
```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit  
$ exit
```

By default, the backup is given a timestamp-based tag. To specify a custom backup tag, add the `--tag` option to the `bkup_api` command; for example, to create a long-term backup with the tag "monthly", enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_start --keep --tag=monthly
```

After you enter a `bkup_api bkup_start` command, the `bkup_api` utility starts the backup process, which runs in the background. To check the progress of the backup process, enter the following command:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_status --dbname=<database_name>
```

To remove the backup configuration

A backup configuration can contain the credentials to access the Object Storage bucket. For this reason, you might want to remove the file after successfully configuring the backup.

```
[root@dbsys bkup]# rm bkup.cfg
```

To delete a local backup

To delete a backup of a database deployment on the Exadata DB system, use the `bkup_api` utility.

CHAPTER 11 Database

1. Connect to the first compute node in your Exadata DB system as the `opc` user.
To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

2. Start a root-user command shell:

```
$ sudo -s#
```

3. List the available backups:

```
# >/var/opt/oracle/bkup_api/bkup_api recover_list --dbname=<database_name>
```

where `dbname` is the database name for the database that you want to act on.

A list of available backups is displayed.

4. Delete the backup you want:

```
# /var/opt/oracle/bkup_api/bkup_api bkup_delete --bkup=<backup-tag> --dbname=<database_name>
```

where `backup-tag` is the tag of the backup you want to delete.

5. Exit the root-user command shell:

```
# exit$
```

To delete a backup in Object Storage

Use the `RMAN delete backup` command to delete a backup from the Object Store.

WHAT NEXT?

If you used Object Storage as a backup destination, you can display the backup files in your bucket in the Console on the **Storage** page, by selecting **Object Storage**.

You can manually restore a database backup by using the RMAN utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide* for Release [18.1](#), [12.2](#), [12.1](#), or [11.2](#).

Recovering an Exadata Database from Object Storage

This topic explains how to recover an Exadata database from a backup stored in Object Storage by using the Console or the API. The Object Storage service is a secure, scalable, on-demand storage solution in Oracle Cloud Infrastructure. For information on backing up your Exadata DB system to Object Storage, see [Managing Exadata Database Backups](#).

Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Using the Console

You can use the Console to restore the database from a backup in the Object Storage that was created by using the Console or the API. You can restore to the last known good state of the database, or you can specify a point in time or an existing System Change Number (SCN).



Note

The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `bkup_api`).

RESTORING AN EXISTING DATABASE

To restore a database

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.

2. Choose your **Compartment**.

A list of DB systems is displayed.

3. Find the DB system where the database is located, and click the system name to display details about it.

A list of databases is displayed.

4. Find the database you want to restore, and click its name to display details about it.

5. Click **Restore**.

6. Select one of the following options, and click **Restore Database**:

- **Restore to the latest:** Restores the database to the last known good state with the least possible data loss.
- **Restore to the timestamp:** Restores the database to the timestamp specified.
- **Restore to System Change Number (SCN):** Restores the database using the SCN specified. This SCN must be valid.



Tip

You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

7. Confirm when prompted.

If the restore operation fails, the database will be in a "Restore Failed" state. You can try restoring again using a different restore option. However, Oracle recommends that you review the `RMAN` logs on the host and fix any issues before reattempting to restore the database. These log files can be found in subdirectories of the `/var/opt/oracle/log` directory.

To restore a database using a specific backup from Object Storage

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.
A list of databases is displayed.
4. Find the database you want to restore, and click its name to display details about it.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Restore**.
7. Confirm when prompted.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to recover a database:

- [ListBackups](#)
- [GetBackup](#)
- [RestoreDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Recovering an Exadata Database by Using RMAN

If you backed up your Exadata database by using `bkup_api`, you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility. For information about

using `RMAN`, see the *Oracle Database Backup and Recovery User's Guide* for Release [18.1](#), [12.2](#), [12.1](#), or [11.2](#).

To restore an Exadata database from a managed backup, see [Recovering an Exadata Database from Object Storage](#).

Managing Oracle Homes Manually

This topic describes how to manage Oracle Homes using the `dbaascli` utility. Oracle Homes are also called Database Homes.

An Oracle Home is a directory location on the compute nodes that contains Oracle Database binaries. Exadata DB systems enable multiple database deployments to share a set of Oracle Database binaries in a shared Oracle Home directory location.

Viewing Information About Oracle Homes

You can view information about Oracle Home directory locations by using the `dbhome info` subcommand of the `dbaascli` utility as follows.

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to an Exadata DB System](#).

2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the `dbaascli` command with the `dbhome info` subcommand:

```
# dbaascli dbhome info
```

4. When prompted, press **Enter** to view information about all Oracle Homes registered in your Exadata DB system, or specify an Oracle Home name to view information only about that Oracle Home.

5. Exit the root-user command shell:

```
# exit
$
```

Moving a Database to Another Oracle Home

Moving a database to another Oracle Home enables you to consolidate existing Oracle Homes and manage the storage that they consume. You can move a database to another Oracle Home by using the `database move` subcommand of the `dbaascli` utility as follows.

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to an Exadata DB System](#).

2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Ensure that all database instances associated with the database deployment are up and running.

```
# dbaascli database status --dbname <dbname>
```

In the preceding command, `<dbname>` specifies the name of the database that you want to check.

Restart any database instances that are not running and open.

4. Execute the `dbaascli` command with the `database move` subcommand:

```
# dbaascli database move --dbname <dbname> --ohome <oracle_home>
```

In the preceding command:

- `<dbname>` — specifies the name of the database that you want to move.
- `<oracle_home>` — specifies the path to an existing Oracle Home directory location, which you want the specified database to use.

When performing a move operation to an Oracle Home with a different patch level, if the database is part of an Exadata DB system Data Guard implementation, then ensure that you move the standby database to the new patchset before you move the primary database.

5. Exit the root-user command shell:

```
# exit
$
```

Creating an Oracle Home

You can create an Oracle Home directory location and software installation, without creating a database, by using the `dbhome create` subcommand of the `dbaascli` utility as follows.

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to an Exadata DB System](#).

2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Run the `dbaascli` command with the `dbhome create` subcommand:

```
# dbaascli dbhome create --version <software_version>
```

In the preceding command, `<software_version>` specifies an Oracle Database software version. For example, 19000, 18000, 12201, 12102, or 11204. The latest available bundle patch for the specified software version is automatically used.

To see information about Oracle Database software images that are available in your Exadata DB system, including software version and bundle patch details, use [the `dbaascli dbimage list` command](#).

When prompted, type `yes` to confirm that the installation is based on a local software image.

4. Exit the root-user command shell:

```
# exit
$
```

Deleting an Oracle Home

If an Oracle Home directory does not support any databases, you can delete it by using the `dbhome purge` subcommand of the `dbaascli` utility as follows.

1. Connect to a compute node as the `opc` user.
For detailed instructions, see [Connecting to an Exadata DB System](#).
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the `dbaascli` command with the `dbhome purge` subcommand:

```
# dbaascli dbhome purge
```

4. When prompted, enter:
 - 1 — if you want to specify the Oracle Home name for the location being purged.
 - 2 — if you want to specify the Oracle Home directory path for the location being purged.
5. When next prompted, enter the Oracle Home name or directory path for the location being purged.

If your entries are valid and the Oracle Home is not associated with a database, then the Oracle binaries are removed from the Oracle Home directory location and the associated metadata is removed from the system.

6. Exit the root-user command shell:

```
# exit
$
```

Using Oracle Data Guard with Exadata DB Systems

This topic explains how to use the Console or the API to manage Data Guard associations in your Exadata DB system. When you use the Console or the API to enable Data Guard for an Oracle Cloud Infrastructure Exadata DB system database:

- The standby database is a physical standby.
- The peer databases (primary and standby) are in the same compartment, they are the same shape, and their database versions are identical.
- You are limited to one standby database for each primary database.

To configure a Data Guard system across regions or between on-premises and Oracle Cloud Infrastructure DB systems, or to configure your database with multiple standbys, you must access the database host directly and set up Data Guard manually.

For complete information on Oracle Data Guard, see the [Data Guard Concepts and Administration](#) documentation on the [Oracle Document Portal](#).

Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

Prerequisites

An Exadata DB system Data Guard implementation requires two Exadata DB systems, one containing the primary database and one containing the standby database. When you enable Data Guard for an Exadata DB system database, the DB system with the database to be used as the standby must already exist before you enable Data Guard.

NETWORK REQUIREMENTS

Ensure that your environment meets the following network requirements:

- Peer DB systems in the Data Guard association can use different subnets but they must use the same VCN, and port 1521 must be open.
- **Important!** Properly configure the security list ingress and egress rules for the subnets of both DB systems in the Data Guard association to allow TCP traffic to flow between the applicable ports. Ensure that the rules you create are stateful (the default). For example, if the subnet of the primary DB System uses the source CIDR 10.0.0.0/24 and the subnet of the standby DB system uses the source CIDR 10.0.1.0/24, create rules as shown in the following example.



Note

The egress rules in the example show how to enable TCP traffic only for port 1521, which is a minimum requirement for Data Guard to work. If TCP traffic is already enabled on all of your outgoing ports (0.0.0.0/0), then you need not explicitly add these specific egress rules.

Security List for Primary DB System's Subnet

Ingress Rules:

```
Stateless: No
Source: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Egress Rules:

```
Stateless: No
Destination: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Security List for Standby DB System's Subnet

Ingress Rules:

```
Stateless: No
Source: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
```

CHAPTER 11 Database

```
Allows: TCP traffic for ports: 1521
```

Egress Rules:

```
Stateless: No  
Destination: 10.0.0.0/24  
IP Protocol: TCP  
Source Port Range: All  
Destination Port Range: 1521  
Allows: TCP traffic for ports: 1521
```

For information about creating and editing rules, see [Security Lists](#).

PASSWORD REQUIREMENTS

For Data Guard operations to work, the SYS password and the TDE wallet password of the primary and standby databases must all be the same. If you change any one of these passwords, you must update the rest of the passwords to match. See [Changing the Database Passwords](#) to learn how to change the SYS password or the TDE wallet password.

If you make any change to the TDE wallet (such as adding a master key for a new PDB or changing the wallet password), you must copy the wallet from the primary to the standby so that Data Guard can continue to operate. For Oracle Database versions earlier than 12.2, if you change the SYS password on one of the peers, you need to manually sync the password file between the DB systems.

Availability Domain Considerations for Data Guard

Oracle recommends that the DB system of the standby database be in a different availability domain from the DB system of the primary database to improve availability and disaster recovery.

Working with Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. The Oracle Cloud Infrastructure Database Data Guard implementation requires two databases, one in a primary role and one in a standby role. The two databases compose a Data Guard association. Most of your applications access the primary database. The standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switch or fail over the standby database to the primary role.

SWITCHOVER

A switchover reverses the primary and standby database roles. Each database continues to participate in the Data Guard association in its new role. A switchover ensures no data loss. Performing planned maintenance on a DB system with a Data Guard association is typically done by switching the primary to the standby role, performing maintenance on the standby, and then switching it back to the primary role.

FAILOVER

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable. A failover might result in some data loss when you use **Maximum Performance** protection mode.

REINSTATE

Reinstates a database into the standby role in a Data Guard association. You can use the `reinstat` command to return a failed database into service after correcting the cause of failure.



Note

You can't terminate a primary database that has a Data Guard association with a peer (standby) database. Delete the standby database first. Alternatively, you can switch over the primary database to the standby role, and then terminate it.

You can't terminate a Exadata DB system that includes Data Guard enabled databases. You must first remove the Data Guard association by terminating the standby database.

Using the Console

The Console allows you to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a *switchover* or a *failover* operation, and *reinststate* a failed database.

When you enable Data Guard, a separate Data Guard association is created for the primary and the standby database.

To enable Data Guard on an Exadata DB system

If you don't already have Exadata DB systems with the databases that will assume the primary and standby roles, create them as described in [To create an Exadata DB system](#). A new DB system includes an initial database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the Exadata DB system with the database for which you want to enable Data Guard.
3. Click the name of the Exadata DB system that contains the database you want to assume the primary role, and then click the name of that database.



Tip

If Data Guard is already enabled, a shield icon appears next to the database name.

4. Under **Resources**, click **Data Guard Associations**.
5. Click **Enable Data Guard**.
6. In the **Enable Data Guard** dialog box, configure your Data Guard association.
 - **Protection Mode:** (Informational) The protection mode used for this Data Guard association. The Console supports only **Maximum Performance**.
 - **Availability Domain:** The availability domain of the peer DB system.

- **Peer DB System:** Select the DB system that contains the peer (standby) database. The peer DB system must be in the same compartment, and must be the same shape.
- **Transport Type:** (Informational) The redo transport type used for this Data Guard association. The Console supports only **Async**.
- **Database Admin Password:** Enter the primary database admin password. The same password is used for the standby database.



Important

The admin password and the TDE password must be the same. If they are not, follow the instructions in [Changing the Database Passwords](#) to align them.

7. Click **Enable**.

When the association is created, a shield icon appears next to the name of this database and its peer, and their respective roles (primary or standby) are displayed.

To perform a database switchover

You initiate a switchover operation by using the Data Guard association of the primary database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the Exadata DB system with the primary database you want to switch over.
3. Click the DB system name, and then click the name of the primary database.
4. Under **Resources**, click **Data Guard Associations**.

5. For the Data Guard association on which you want to perform a switchover, click the Actions icon (three dots), and then click **Switchover**.
6. In the **Switchover Database** dialog box, enter the database admin password, and then click **OK**.
This database should now assume the role of the standby, and the standby should assume the role of the primary in the Data Guard association.

To perform a database failover

You initiate a failover operation by using the Data Guard association of the standby database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the Exadata DB system with the primary database's peer standby you want to fail over to.
3. Click the DB system name, and then click the name of the standby database.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to perform a failover, click **Failover**.
6. In the **Failover Database** dialog box, enter the database admin password, and then click **OK**.
This database should now assume the role of the primary, and the old primary's role should display as **Disabled Standby**.

To reinstate a database

After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the Exadata DB system with the failed database you want to reinstate.
3. Click the DB system name, and then click the database name.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to reinstate this database, click the Actions icon (three dots), and then click **Reinstate**.
6. In the **Reinstate Database** dialog box, enter the database admin password, and then click **OK**.

This database should now be reinstated as the standby in the Data Guard association.

To terminate a Data Guard association on an Exadata DB system

On an Exadata DB system, you remove a Data Guard association by terminating the standby database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the Exadata DB system that includes the standby database you want to terminate.
3. Click the DB system name.
4. For the standby database you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
5. In the **Terminate Database** dialog box, enter the name of the database, and then click **OK**.

Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Data Guard associations on an Exadata DB system:

- [CreateDataGuardAssociation](#)
- [GetDataGuardAssociation](#)
- [ListDataGuardAssociations](#)
- [SwitchoverDataGuardAssociation](#)
- [FailoverDataGuardAssociation](#)
- [ReinstateDataGuardAssociation](#)
- [DeleteDbHome](#) - To terminate an Exadata DB system Data Guard association, you delete the standby database.

For the complete list of APIs for the Database service, see [Database Service API](#).

Configuring Oracle Database Features for Exadata DB Systems

This topic describes how to configure Oracle Multitenant, tablespace encryption, and Huge Pages for use with your Exadata DB systems.

Using Oracle Multitenant on an Exadata DB system

When you create an Exadata DB system that uses Oracle Database 12c or later, an Oracle Multitenant environment is created.

The multitenant architecture enables an Oracle database to function as a multitenant container database (CDB) that includes zero, one, or many pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appears to an Oracle Net Services client as a non-CDB. All Oracle databases using versions earlier than Oracle Database 12c are non-CDBs.

To use Oracle Transparent Data Encryption (TDE) in a pluggable database (PDB), you must create and activate a master encryption key for the PDB.

In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers.

You must export and import the master encryption key for any encrypted PDBs you plug into your Exadata DB system CDB.

If your source PDB is encrypted, you must export the master encryption key and then import it.

You can export and import all of the TDE master encryption keys that belong to the PDB by exporting and importing the TDE master encryption keys from within a PDB. Export and import of TDE master encryption keys support the PDB unplug and plug operations. During a PDB unplug and plug, all of the TDE master encryption keys that belong to a PDB, as well as the metadata, are involved.

See "Exporting and Importing TDE Master Encryption Keys for a PDB" in *Oracle Database Advanced Security Guide* for Release [19](#), [18](#), [12.2](#) or [12.1](#).

See "ADMINISTER KEY MANAGEMENT" in *Oracle Database SQL Language Reference* for Release [19](#), [18](#), [12.2](#) or [12.1](#).

To determine if you need to create and activate an encryption key for the PDB

1. Invoke SQL*Plus and log in to the database as the SYS user with SYSDBA privileges.
2. Set the container to the PDB:

```
SQL> ALTER SESSION SET CONTAINER = pdb;
```

3. Query V\$ENCRYPTION_WALLET as follows:

```
SQL> SELECT wr1_parameter, status, wallet_type FROM v$encryption_wallet;
```

If the STATUS column contains a value of OPEN_NO_MASTER_KEY, you need to create and activate the master encryption key.

To create and activate the master encryption key in a PDB

1. Set the container to the PDB:
- ```
SQL> ALTER SESSION SET CONTAINER = pdb;
```
2. Create and activate a master encryption key in the PDB by executing the following command:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEY USING TAG 'tag' FORCE KEYSTORE IDENTIFIED BY keystore-
password WITH BACKUP USING 'backup_identifier';
```

In the previous command:

- `keystore-password` is the keystore password. By default, the keystore password is set to the value of the administration password that is specified when the database is created.
- The optional `USING TAG 'tag'` clause can be used to associate a tag with the new master encryption key.
- The `WITH BACKUP` clause, and the optional `USING 'backup_identifier'` clause, can be used to create a backup of the keystore before the new master encryption key is created.

See also `ADMINISTER KEY MANAGEMENT` in *Oracle Database SQL Language Reference for Release [19](#), [18](#) or [12.2](#)*.

**Note**

To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.

If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:

- a. Close the keystore.
- b. Open the password-based keystore.
- c. Create and activate a master encryption key in the PDB by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.
- d. Update the auto-login keystore by using `ADMINISTER KEY MANAGEMENT` with the `CREATE AUTO_LOGIN KEYSTORE FROM KEYSTORE` option.

3. Query `V$ENCRYPTION_WALLET` again to verify that the `STATUS` column is set to `OPEN`:

```
SQL> SELECT wr1_parameter, status, wallet_type FROM v$encryption_wallet;
```

4. Query `V$INSTANCE` and take note of the value in the `HOST_NAME` column, which identifies the database server that contains the newly updated keystore files:

```
SQL> SELECT host_name FROM v$instance;
```

5. Copy the updated keystore files to all of the other database servers.

To distribute the updated keystore, you must perform the following actions on each database server that does not contain the updated keystore files:

- a. Connect to the root container and query `V$ENCRYPTION_WALLET`. Take note of the keystore location contained in the `WRL_PARAMETER` column:

```
SQL> SELECT wrl_parameter, status FROM v$encryption_wallet;
```

- b. Copy the updated keystore files.

You must copy all of the updated keystore files from a database server that is already updated. Use the keystore location observed in the `WRL_PARAMETER` column of `V$ENCRYPTION_WALLET`.

Open the updated keystore:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEYSTORE open FORCE KEYSTORE IDENTIFIED BY keystore-password
CONTAINER=all;
```



### Note

To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.

If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:

- a. Close the keystore before copying the updated keystore files.
- b. Copy the updated keystore files.
- c. Open the updated keystore by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.

6. Query `GV$ENCRYPTION_WALLET` to verify that the `STATUS` column is set to `OPEN` across all of the database instances:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM gv$encryption_wallet;
```

### To export and import a master encryption key

1. Export the master encryption key.
  - a. Invoke `SQL*Plus` and log in to the PDB.
  - b. Execute the following command:

## CHAPTER 11 Database

---

```
SQL> ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET "secret" TO 'filename'
IDENTIFIED BY keystore-password;
```

2. Import the master encryption key.
  - a. Invoke SQL\*Plus and log in to the PDB.
  - b. Execute the following command:

```
SQL> ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS WITH SECRET "secret" FROM 'filename'
IDENTIFIED BY keystore-password;
```

### Managing Tablespace Encryption

By default, all new tablespaces that you create in an Exadata database are encrypted.

However, the tablespaces that are initially created when the database is created may not be encrypted by default.

- For databases that use Oracle Database 12c Release 2 or later, only the `USERS` tablespaces initially created when the database was created are encrypted. No other tablespaces are encrypted including the non-`USERS` tablespaces in:
  - The root container (`CDB$ROOT`).
  - The seed pluggable database (`PDB$SEED`).
  - The first PDB, which is created when the database is created.
- For databases that use Oracle Database 12c Release 1 or Oracle Database 11g, none of the tablespaces initially created when the database was created are encrypted.

For further information about the implementation of tablespace encryption in Exadata, along with how it impacts various deployment scenarios, see [Oracle Database Tablespace Encryption Behavior in Oracle Cloud](#).

### CREATING ENCRYPTED TABLESPACES

User-created tablespaces are encrypted by default.

By default, any new tablespaces created by using the `SQL CREATE TABLESPACE` command are encrypted with the AES128 encryption algorithm. You do not need to include the `USING 'encrypt_algorithm'` clause to use the default encryption.

You can specify another supported algorithm by including the `USING 'encrypt_algorithm'` clause in the `CREATE TABLESPACE` command. Supported algorithms are AES256, AES192, AES128, and 3DES168.

### MANAGING TABLESPACE ENCRYPTION

You can manage the software keystore (known as an Oracle wallet in Oracle Database 11g), the master encryption key, and control whether encryption is enabled by default.

#### *MANAGING THE MASTER ENCRYPTION KEY*

Tablespace encryption uses a two-tiered, key-based architecture to transparently encrypt (and decrypt) tablespaces. The master encryption key is stored in an external security module (software keystore). This master encryption key is used to encrypt the tablespace encryption key, which in turn is used to encrypt and decrypt data in the tablespace.

When a database is created on an Exadata DB system, a local software keystore is created. The keystore is local to the compute nodes and is protected by the administration password specified during the database creation process. The auto-login software keystore is automatically opened when the database is started.

You can change (rotate) the master encryption key by using the `ADMINISTER KEY MANAGEMENT` SQL statement. For example:

```
SQL> ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY USING TAG 'tag'
IDENTIFIED BY password WITH BACKUP USING 'backup';

keystore altered.
```

See "Managing the TDE Master Encryption Key" in *Oracle Database Advanced Security Guide* for Release [19](#), [18](#), [12.2](#) or [12.1](#) or "Setting and Resetting the Master Encryption Key" in *Oracle Database Advanced Security Administrator's Guide* for Release [11.2](#).

#### *CONTROLLING DEFAULT TABLESPACE ENCRYPTION*

The `ENCRYPT_NEW_TABLESPACES` initialization parameter controls the default encryption of new tablespaces. In Exadata databases, this parameter is set to `CLOUD_ONLY` by default.

Values of this parameter are as follows.

| Value      | Description                                                                                                                                                                                                                                                                                                                                 |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ALWAYS     | During creation, tablespaces are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the <code>ENCRYPTION</code> clause.                                                                                                                                                                         |
| CLOUD_ONLY | Tablespaces created in an Exadata database are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the <code>ENCRYPTION</code> clause. For non-cloud databases, tablespaces are only encrypted if the <code>ENCRYPTION</code> clause is specified. <code>ENCRYPTION</code> is the default value. |
| DDL        | During creation, tablespaces are not transparently encrypted by default, and are only encrypted if the <code>ENCRYPTION</code> clause is specified.                                                                                                                                                                                         |



### Note

With Oracle Database 12c Release 2 (12.2), or later, you can no longer create an unencrypted tablespace in an Exadata database. An error message is returned if you set `ENCRYPT_NEW_TABLESPACES` to `DDL` and issue a `CREATE TABLESPACE` command without specifying an `ENCRYPTION` clause.

## Managing Huge Pages

Huge Pages provide considerable performance benefits for Oracle Database on systems with large amounts of memory. Oracle Database on an Exadata DB system hosted in Oracle Cloud Infrastructure provides configuration settings that make use of Huge Pages by default; however, you can make manual adjustments to optimize the configuration of Huge Pages.

Huge Pages is a feature integrated into the Linux kernel 2.6. Enabling Huge Pages makes it possible for the operating system to support large memory pages. Using Huge Pages can improve system performance by reducing the amount of system CPU and memory resources required to manage Linux page tables, which store the mapping between virtual and physical

memory addresses. For Oracle Databases, using Huge Pages can drastically reduce the number of page table entries associated with the System Global Area (SGA).

On Exadata DB systems hosted in Oracle Cloud Infrastructure, a standard page is 4 KB, while a Huge Page is 2 MB by default. Therefore, an Oracle Database on an Exadata DB system with a 50 GB SGA requires 13,107,200 standard pages to house the SGA, compared with only 25,600 Huge Pages. The result is much smaller page tables, which require less memory to store and fewer CPU resources to access and manage.

### ADJUSTING THE CONFIGURATION OF HUGE PAGES

The configuration of Huge Pages for Oracle Database is a two-step process:

- At the operating system level, the overall amount of memory allocated to Huge Pages is controlled by the `vm.nr_hugepages` entry in the `/etc/sysctl.conf` file. This setting is made on each compute node in the environment and it is strongly recommended that the setting is consistent across all of the compute nodes. To alter the Huge Page allocation, you can execute the following command on each compute node as the root user:

```
sysctl -w vm.nr_hugepages=value
```

where `value` is the number of Huge Pages that you want to allocate.

On Exadata DB systems hosted in Oracle Cloud Infrastructure, each Huge Page is 2 MB by default. Therefore, to allocate 50 GB of memory to Huge Pages you can execute the following command:

```
sysctl -w vm.nr_hugepages=25600
```

- At the Oracle Database level, the use of Huge Pages is controlled by the `USE_LARGE_PAGES` instance parameter setting. This setting applies to each database instance in a clustered database. Oracle strongly recommends a consistent setting across all of the database instances associated with a database. The following options are available:
  - `TRUE` — specifies that the database instance can use Huge Pages if they are available. For all versions of Oracle Database after 11.2.0.3, Oracle allocates as much of the SGA as it can, using Huge Pages. When the Huge Page allocation is exhausted, standard memory pages are used.

- `FALSE` — specifies that the database instance does not use Huge Pages. This setting is generally not recommended if Huge Pages are available.
- `ONLY` — specifies that the database instance must use Huge Pages. With this setting, the database instance fails to start if the entire SGA cannot be accommodated in Huge Pages.

If you make any adjustments at either the operating system or Oracle Database level, ensure that the overall configuration works.

For more information, see the *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for Release [19](#), [18](#), [12.1](#), or [11.2](#) for a general overview of Huge Pages and more information about configuring Huge Pages. Also, see `USE_LARGE_PAGES` in the *Oracle Database Reference* for Release [12.2](#), [12.1](#), or [11.2](#).

### DB System Time Zone



#### Note

This topic applies only to Exadata DB systems.

The Time Zone field in the Console and in the API allows you to launch an Exadata DB system with a time zone other than UTC (the default). Although UTC is the recommended time zone to use, having a common time zone for your database clients and application hosts can simplify management and troubleshooting for the database administrator.

The time zone that you specify when you create the DB system applies to the host and to the Oracle Grid Infrastructure, and controls the time zone of the database log files. The time zone of the database itself is not affected, however, the database's time zone affects only the timestamp datatype. By default, it is set to UTC and although you can change it manually, Oracle recommends that you keep it as UTC to avoid data conversion and improve performance when data is transferred among databases. This is especially important for distributed databases, replication, and exporting and importing.

### Time Zone Options

Whether you use the Console or the API, the time zone options you can select from are represented in the named region format, for example, *America/Los\_Angeles*. The Console allows you to select UTC, the time zone detected in your browser (if your browser supports time zone detection), or an alternate time zone.

To specify an alternate time zone (the **Select Another Time Zone** option), you first select a value in the Time Zone Prefix field to narrow the list of time zones from which to select in the Time Zone Suffix field. In the *America/Los\_Angeles* example, *America* is the time zone prefix and *Los\_Angeles* is the time zone suffix. The items you see in the Time Zone Prefix and Time Zone Suffix fields roughly correlate with the time zones supported in both the `Java.util.TimeZone` class and on the Linux operating system. If you do not see the time zone you are looking for, try selecting "Miscellaneous" as the time zone prefix.



#### Tip

If you are using the API and would like to see a list of supported time zones, you can examine the time zone options in the Console. These options appear in the **Launch DB System** dialog when you show advanced options after you select an Exadata DB system shape.

### Changing Time Zones After Provisioning

Follow these steps if you need to change the time zone of the Exadata DB system, Oracle Grid Infrastructure, or database, after you launch the DB system:

#### To change the time zone of the DB system host

1. Log on to the host system as `root`.
2. Stop the CRS stack on all of the Exadata compute nodes.

```
#Grid_Home/bin/crsctl stop crs
```

3. Run the following commands to check the current time zone and to change it to the time zone you choose:

```
$ cat /etc/sysconfig/clock
ZONE="America/New_York"
$ cp -p /etc/sysconfig/clock /etc/sysconfig/clock.20160629

$ vi /etc/sysconfig/clock
ZONE="Europe/Berlin"

$ date
Wed Jun 29 10:35:17 EDT 2016
$ ln -sf /usr/share/zoneinfo/Europe/Berlin /etc/localtime
$ date
Wed Jun 29 16:35:27 CEST 2016
```

In this example, the time zone was changed from *America/New\_York* to *Europe/Berlin*.



### Tip

To see a list of valid time zones on the host, you can run the `ls -l /usr/share/zoneinfo` command.

4. (Optional) Verify that `/opt/oracle.cellos/cell.conf` indicates the correct time zone. Using our example, the time zone entry in this file would be `<Timezone>Europe/Berlin</Timezone>`.
5. Restart the CRS stack on all of the Exadata compute nodes.

```
#Grid_Home/bin/crsctl start crs
```

## To change the time zone of the Oracle Grid Infrastructure

The time zone of the Oracle Grid Infrastructure determines the time zone of the database log files. You can change this time zone by updating the `TZ` property in the `GRID_HOME/crs/install/s_crsconfig_<node_name>_env.txt` configuration file.

1. Ensure that you are logged onto the host as `root` and that the CRS stack is stopped on all of the Exadata compute nodes. See [To change the time zone of the DB system host](#).
2. Inspect the current time zone value in the `GRID_HOME/crs/install/s_crsconfig_<node_name>_env.txt` file.

```
$ cat /u01/app/12.1.0.2/grid/crs/install/s_crsconfig_node1_env.txt
#####
#This file can be used to set values for the NLS_LANG and TZ environment
#variables and to set resource limits for Oracle Clusterware and
#Database processes.
#1. The NLS_LANG environment variable determines the language and
character set used for messages. For example, a new value can be
configured by setting NLS_LANG=JAPANESE_JAPAN.UTF8
#2. The Time zone setting can be changed by setting the TZ entry to
the appropriate time zone name. For example, TZ=America/New_York
#3. Resource limits for stack size, open files and number of processes
can be specified by modifying the appropriate entries.
#
#Do not modify this file except as documented above or under the
#direction of Oracle Support Services.
#####
TZ=UTC
NLS_LANG=AMERICAN_AMERICA.WE8ISO8859P1
CRS_LIMIT_STACK=2048
CRS_LIMIT_OPENFILE=65536
CRS_LIMIT_NPROC=16384
TNS_ADMIN=
```

In this example, the time zone is set to UTC.

3. Modify the time zone value, as applicable. Perform this task for all nodes in the cluster.
4. Restart the CRS stack on all of the Exadata compute nodes.

```
#Grid_Home/bin/crsctl start crs
```

For more information about changing the time zone of the Grid Infrastructure, see [How To Change Timezone for Grid Infrastructure \(Doc ID 1209444.1\)](#).

### To change the time zone of a database

Use the `ALTER DATABASE SET TIME_ZONE` command to change the time zone of a database. This command takes either a named region such as `America/Los_Angeles` or an absolute offset from UTC.

This example sets the time zone to UTC:

```
ALTER DATABASE SET TIME_ZONE = '+00:00';
```

You must restart the database for the change to take effect. For more information, see [Setting the Database Time Zone](#).

## Bare Metal and Virtual Machine DB Systems

Oracle Cloud Infrastructure offers 1-node DB systems on either bare metal or virtual machines, and 2-node RAC DB systems on virtual machines. If you need to quickly spin up a DB system for development or testing purposes, a special [fast provisioning](#) 1-node VM system is available.

You can manage these systems by using the Console, the API, the Oracle Cloud Infrastructure CLI, the Database CLI (DBCLI), Enterprise Manager, Enterprise Manager Express, or SQL Developer.



#### Note

This documentation is intended for Oracle database administrators and assumes familiarity with Oracle databases and tools. If you need additional information, see the product documentation available at <http://docs.oracle.com/en/database/>.

## Supported Database Editions and Versions

All 1- node RAC DB systems support the following Oracle Database editions:

## CHAPTER 11 Database

---

- Standard Edition
- Enterprise Edition
- Enterprise Edition - High Performance
- Enterprise Edition - Extreme Performance

2-node RAC DB systems require Oracle Enterprise Edition - Extreme Performance.

For standard provisioning of DB systems (using [Oracle Automatic Storage Management](#) (ASM) as your storage management software), the supported database versions are:

- Oracle Database 19c (19.0) *(available for virtual machine DB systems only)*
- Oracle Database 18c (18.0)
- Oracle Database 12c Release 2 (12.2)
- Oracle Database 12c Release 1 (12.1)
- Oracle Database 11g Release 2 (11.2)

For [fast provisioning](#) of 1-node Virtual Machine DB systems (using [Logical Volume Manager](#) as your storage management software), the supported database versions are:

- Oracle Database 19c (19.0) *(available for virtual machine DB systems only)*
- Oracle Database 18c (18.0)



### Tip

Your DB system's operating system will periodically need to be updated, just as your Oracle Database software will need to be updated. Before attempting an OS update, be sure to read the information in [Updating a DB System](#) and back up your DB system's databases.

### Availability of Older Database Versions for Virtual Machine DB Systems

For virtual machine DB systems, Oracle Cloud Infrastructure also supports the creation of DB systems using older database versions. For each shape, the latest version and the two prior versions of the release are available at provisioning.



#### Warning

If you need to launch your DB system with an older database version, see [Critical Patch Updates](#) for information on known security issues with your chosen database version. You will also need to analyze and patch known security issues for the operating system included with the older database version. See [Securing Database](#) for information on security best practices for databases in Oracle Cloud Infrastructure.

### Bare Metal DB Systems

Bare metal DB systems consist of a single bare metal server running Oracle Linux 6.8, with locally attached NVMe storage. If the node fails, you can simply launch another system and restore the databases from current backups.

When you launch a bare metal DB system, you select a single Oracle Database Edition that applies to all the databases on that DB system. The selected edition cannot be changed. Each DB system can have multiple database homes, which can be different versions. Each database home can have only one database, which is the same version as the database home.

#### Shapes for Bare Metal DB Systems

When you launch a DB system, you choose a *shape*, which determines the resources allocated to the DB system. The available shapes for a bare metal DB system are:

- **BM.DenseIO2.52:** Provides a 1-node DB system (one bare metal server), with up to 52 CPU cores, 768 GB memory, and eight 6.4 TB locally attached NVMe drives (51.2 TB total) to the DB system.
- **BM.DenseIO1.36:** *Limited availability.* Provides a 1-node DB system (one bare metal server), with up to 36 CPU cores, 512 GB memory, and nine 3.2 TB locally attached NVMe drives (28.8 TB total) to the DB system.

*Note:* BM.DenseIO1.36 is available only to monthly universal credit customers existing on or before November 9th, 2018. This shape is available only in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.

### Storage Considerations

The shape you choose for a bare metal DB system determines its total raw storage, but other options, like 2- or 3-way mirroring and the space allocated for data files, affect the amount of usable storage on the system. The following table shows how various configurations affect the usable storage for bare metal DB systems.

| Shape                                      | Raw Storage  | Usable Storage with Normal Redundancy (2-way Mirroring) | Usable Storage with High Redundancy (3-way Mirroring) |
|--------------------------------------------|--------------|---------------------------------------------------------|-------------------------------------------------------|
| BM.DenseIO2.52                             | 51.2 TB NVMe | DATA 16 TB<br>RECO 4 TB                                 | DATA 9 TB<br>RECO 2.3 TB                              |
| BM.DenseIO1.36<br><a href="#">see note</a> | 28.8 TB NVMe | DATA 9.4 TB<br>RECO 1.7 TB                              | DATA 5.4 TB<br>RECO 1 TB                              |

**Note:** BM.DenseIO1.36 availability is limited to monthly universal credit customers existing on or before November 9th, 2018, in the us-phoenix-1, us-ashburn-1, and eu-frankfurt-1 regions.

### Virtual Machine DB Systems

There are two types of DB systems on virtual machines:

- A 1-node virtual machine DB system consists of one virtual machine.
- A 2-node virtual machine DB system consists of two virtual machines.

When you launch a virtual machine DB system, you select the Oracle Database Edition that applies to the database on that DB system. The selected edition cannot be changed. Unlike a bare metal DB system, a virtual machine DB system can have only a single database home, which in turn can have only a single database. The database can be a container database (CDB) with multiple pluggable databases (PDBs), if the edition is High Performance or Extreme Performance. The database will be the same version as the database home.

Virtual machine DB systems also differ from bare metal DB systems in the following ways:

- A virtual machine DB system database uses Oracle Cloud Infrastructure block storage instead of local storage. You specify a storage size when you launch the DB system, and you can scale up the storage as needed at any time.
- The number of CPU cores on an existing virtual machine DB system cannot be changed.

#### **Fast Provisioning Option for 1-node Virtual Machine DB Systems**

For 1-node virtual machine DB systems, Oracle Cloud Infrastructure provides have a "fast provisioning" option that allows you to create your DB system using [Logical Volume Manager](#) as your storage management software. The alternative ("standard provisioning") is to provision with [Oracle Automatic Storage Management](#) (ASM).

**Note**

- When using the fast provisioning option, the number and size of the block volumes specified during provisioning determines the maximum total storage available through scaling. See [Storage Scaling Considerations for Virtual Machine Databases Using Fast Provisioning](#) for details.
- Multi-node Virtual Machine DB systems require Oracle Automatic Storage Management and cannot be created using the fast-provisioning option.

**Fault Domain Considerations for 2-node Virtual Machine DB Systems**

When you provision a 2-node RAC DB systems, the system assigns each node to a different fault domain by default. Using the **Advanced Options** link in the provisioning dialog, you can select the fault domain(s) to be used for your 2-node RAC DB systems and the system will assign the nodes to your selected fault domains. Oracle recommends that you place each node of a 2-node RAC DB system in a different fault domain. For more information on fault domains, see [Fault Domains](#).

**Shapes for Virtual Machine DB Systems**

When you launch a DB system, you choose a *shape*, which determines the resources allocated to the DB system.

The following table shows the available shapes for a virtual machine DB system on X7.

| Shape          | CPU Cores | Memory |
|----------------|-----------|--------|
| VM.Standard2.1 | 1         | 15 GB  |
| VM.Standard2.2 | 2         | 30 GB  |

## CHAPTER 11 Database

| Shape           | CPU Cores | Memory |
|-----------------|-----------|--------|
| VM.Standard2.4  | 4         | 60 GB  |
| VM.Standard2.8  | 8         | 120 GB |
| VM.Standard2.16 | 16        | 240 GB |
| VM.Standard2.24 | 24        | 320 GB |

The following table shows the available shapes for a virtual machine DB system on X5. [see note](#)

| Shape                               | CPU Cores | Memory |
|-------------------------------------|-----------|--------|
| VM.Standard1.1 <sup>see note</sup>  | 1         | 7 GB   |
| VM.Standard1.2 <sup>see note</sup>  | 2         | 14 GB  |
| VM.Standard1.4 <sup>see note</sup>  | 4         | 28 GB  |
| VM.Standard1.8 <sup>see note</sup>  | 8         | 56 GB  |
| VM.Standard1.16 <sup>see note</sup> | 16        | 112 GB |

**Note:** X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9th, 2018, in the us-phoenix-1, us-ashburn-1, and eu-frankfurt-1 regions.

### Storage Options for Virtual Machine DB Systems

Virtual machine DB systems use Oracle Cloud Infrastructure block storage. The following table shows details of the storage options for a virtual machine DB system. Total storage includes available storage plus recovery logs.

## CHAPTER 11 Database

---

| Available Storage (GB) | Total Storage (GB) |
|------------------------|--------------------|
| 256                    | 712                |
| 512                    | 968                |
| 1024                   | 1480               |
| 2048                   | 2656               |
| 4096                   | 5116               |
| 6144                   | 7572               |
| 8192                   | 10032              |
| 10240                  | 12488              |
| 12288                  | 14944              |
| 14336                  | 17404              |
| 16384                  | 19860              |
| 18432                  | 22320              |
| 20480                  | 24776              |
| 22528                  | 27232              |
| 24576                  | 29692              |
| 26624                  | 32148              |
| 28672                  | 34608              |
| 30720                  | 37064              |
| 32768                  | 39520              |

| Available Storage (GB) | Total Storage (GB) |
|------------------------|--------------------|
| 34816                  | 41980              |
| 36864                  | 44436              |
| 38912                  | 46896              |
| 40960                  | 49352              |

For 2-node RAC virtual machine DB systems, storage capacity is shared between the nodes.

### Database Backups

See [Backing Up a Database](#) for information about the backup options you have for your cloud databases. See [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#) for information about managed automatic backups in Oracle Cloud Infrastructure.

### Network Setup for DB Systems



#### Note

This topic is not applicable to Exadata DB systems. For information on the network setup for an Exadata DB system, see [Network Setup for Exadata DB Systems](#).

Before you set up a bare metal or virtual machine DB system, you must set up a virtual cloud network (VCN) and other Networking service components. This topic describes the recommended configuration for the VCN.

### VCN and Subnets

To launch a DB system, you must have:

- A [VCN](#) in the region where you want the DB system
- At least one subnet in the VCN (either a public subnet or a private subnet)

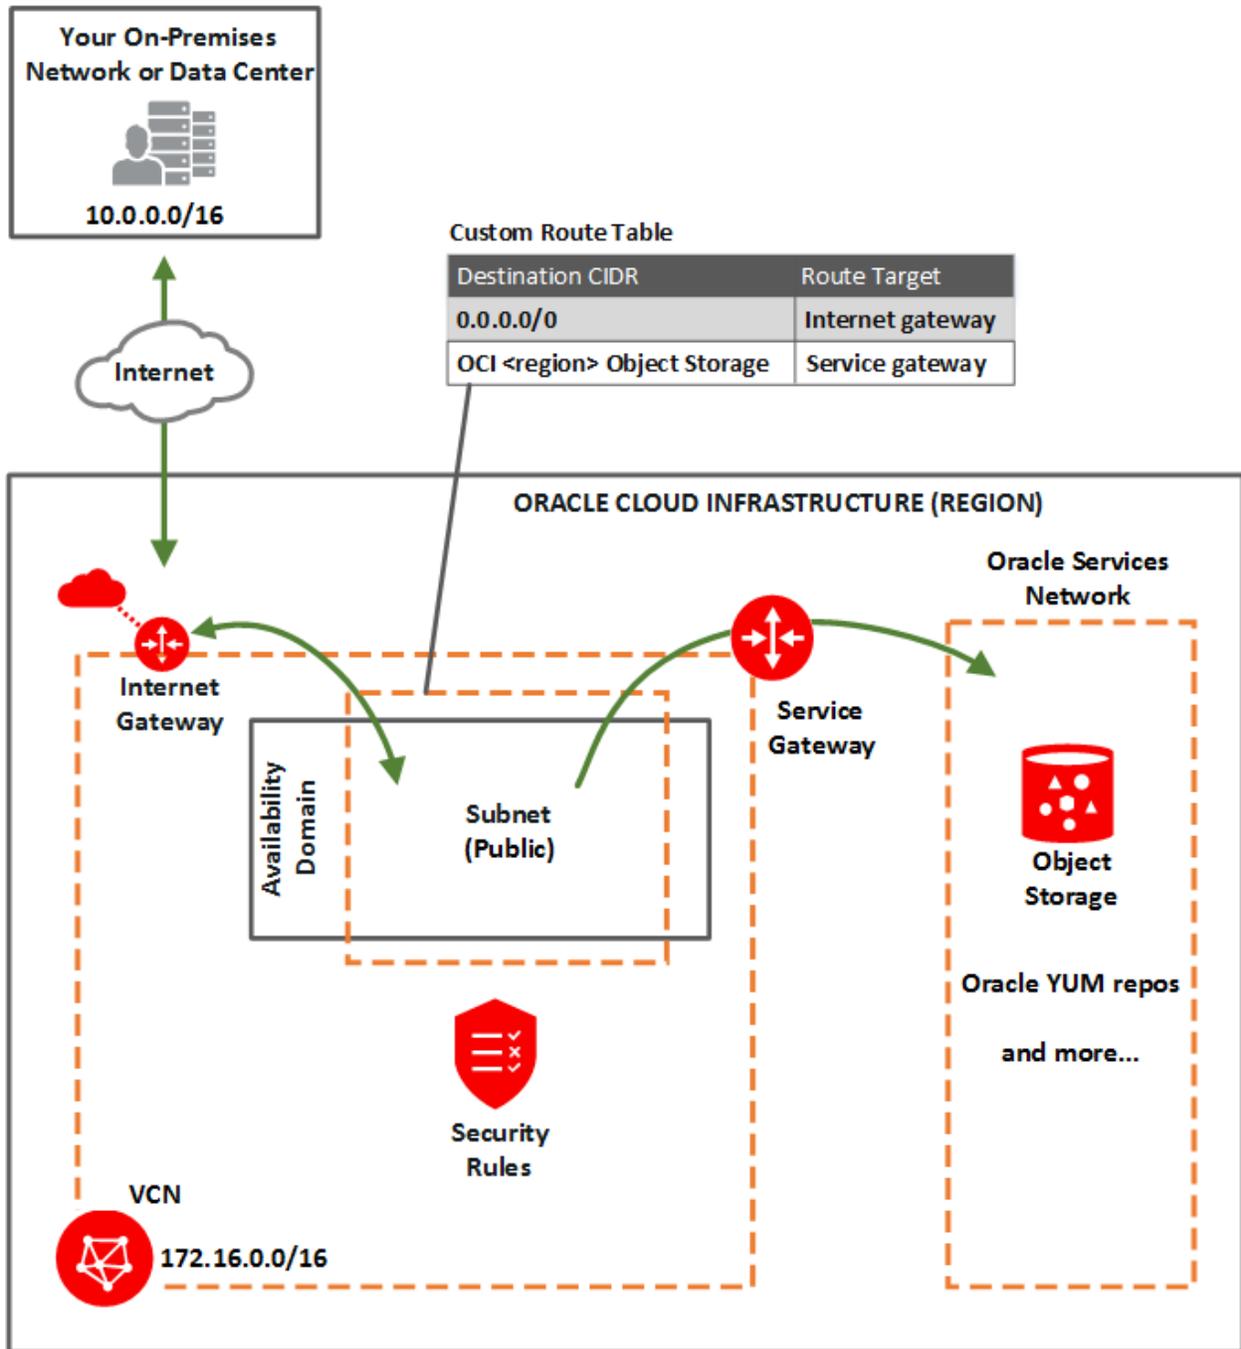
In general, Oracle recommends using regional subnets, which span all availability domains in the region. For a bare metal or virtual machine DB system, either a regional subnet or AD-specific subnet works. For more information, see [About Regional Subnets](#).

You will create a custom route table. You will also create security rules to control traffic to and from the DB system's compute nodes. More information follows about that.

Certain details of the VCN and subnet configuration depend on your choice for DNS resolution within the VCN. For more information, see [DNS for the DB System](#).

### **OPTION 1: PUBLIC SUBNET WITH INTERNET GATEWAY**

This option can be useful when doing a proof-of-concept or development work. You can use this setup in production if you want to use an internet gateway with the VCN, or if you have services that run only on a public network and need access to the database. See the following diagram and description.



You set up:

- [Public subnet](#).
- [Internet gateway](#).
- [Service gateway](#) to reach Object Storage for database backups and patching. Also see [Option 1: Service Gateway Access Only to Object Storage](#).
- [Route table](#): A custom route table for the subnet, with two rules:
  - A rule for 0.0.0.0/0, and target = internet gateway.
  - A rule for the [service CIDR label](#) called **OCI <region> Object Storage**, and target = the service gateway. Also see [Option 1: Service Gateway Access Only to Object Storage](#).
- [Security rules](#) to enable the desired traffic to and from the DB system nodes. See [Security Rules for the DB System](#).

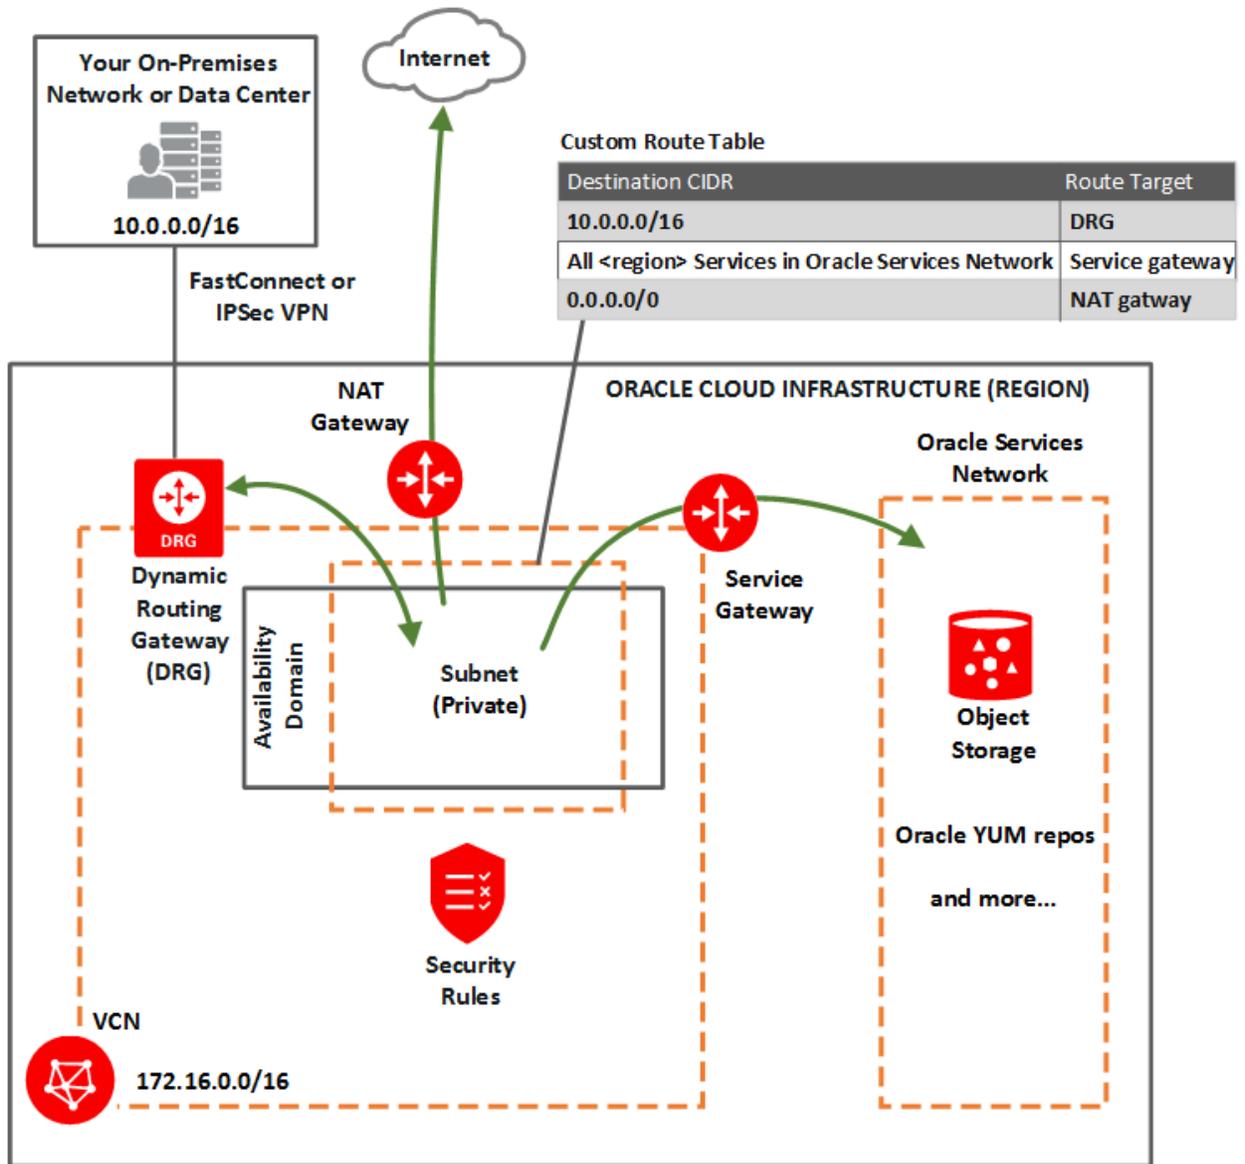


### Important

See this [known issue](#) for information about configuring route rules with *service gateway* as the target on route tables associated with public subnets.

### OPTION 2: PRIVATE SUBNET

Oracle recommends this option for a production system. The subnet is private and cannot be reached from the internet. See the following diagram and description.



You set up:

- [Private subnet](#).
- Gateways for the VCN:
  - [Dynamic routing gateway \(DRG\)](#), with a [FastConnect](#) or [IPSec VPN](#) to your on-premises network.
  - [Service gateway](#) to reach Object Storage for database backups and patching, and to reach Oracle YUM repos for OS updates. Also see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#).
  - [NAT gateway](#) (to reach public endpoints not supported by the service gateway).
- [Route table](#): A custom route table for the subnet, with these rules:
  - A route for the on-premises network's CIDR, and target = DRG.
  - A rule for the [service CIDR label](#) called **All <region> Services in Oracle Services Network**, and target = the service gateway. Also see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#).
  - If you want to access the Oracle YUM repos through the NAT gateway, add a route rule for the [regional YUM repo's public IP address](#), and target = the NAT gateway. If you just use the next rule only, the traffic to the YUM repo would still be routed to the service gateway, because the service gateway route is more specific than 0.0.0.0/0.
  - A rule for 0.0.0.0/0, and target = NAT gateway.
- [Security rules](#) to enable the desired traffic to and from the DB system nodes. See [Security Rules for the DB System](#).

### REQUIREMENTS FOR IP ADDRESS SPACE

If you are setting up DB systems (and thus VCNs) in more than one region, make sure the IP address space of the VCNs does not overlap.

The subnet you create for a bare metal or virtual machine DB system cannot overlap with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance.

The following table lists the *minimum* required subnet size.

**Tip**

The Networking service [reserves three IP addresses in each subnet](#). Allocating a larger space for the subnet than the minimum required (for example, at least /25 instead of /28) can reduce the relative impact of those reserved addresses on the subnet's available space.

| DB System Type                       | # Required IP Addresses                                                           | Minimum Subnet Size   |
|--------------------------------------|-----------------------------------------------------------------------------------|-----------------------|
| 1-node bare metal or virtual machine | 1 + 3 reserved in subnet = 4                                                      | /30 (4 IP addresses)  |
| 2-node RAC virtual machine           | (2 addresses * 2 nodes) + 3 for <a href="#">SCANs</a> + 3 reserved in subnet = 10 | /28 (16 IP addresses) |

**VCN CREATION WIZARD: NOT FOR PRODUCTION**

The Networking section of the Console includes a handy wizard that creates a VCN along with related resources. It can be useful if you just want to try launching an instance. However, the wizard automatically chooses the address ranges and creates public subnets and an internet gateway. You may not want this for your production network, so Oracle recommends you create the VCN and other resources individually yourself instead of using the wizard.

**DNS for the DB System**

There are two choices for DNS and hostname resolution for the DB system:

- Recommended: Use the default DNS functionality in the VCN (called the *Internet and VCN Resolver*)
- Use a custom DNS resolver of your choice

The following table shows which choices are supported with each type of DB system, and the endpoints that need to be resolved for the DB system to function.

| DB System Type                       | Supported DNS Choices                                                                                                                              | Endpoints to Be Resolved                                                                                                                                                                                                                         |
|--------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1-node bare metal or virtual machine | <ul style="list-style-type: none"> <li>• Recommended: Default (Internet and VCN Resolver)</li> <li>• Custom DNS resolver of your choice</li> </ul> | <ul style="list-style-type: none"> <li>• Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints)</li> <li>• Oracle YUM repo endpoints</li> </ul>                                                               |
| 2-node RAC virtual machine           | <ul style="list-style-type: none"> <li>• Default (Internet and VCN Resolver)</li> </ul>                                                            | <ul style="list-style-type: none"> <li>• Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints)</li> <li>• Oracle YUM repo endpoints</li> <li>• <a href="#">Single Client Access Names (SCANs)</a></li> </ul> |

The following sections give more details about the DNS choices.

**DEFAULT (INTERNET AND VCN RESOLVER)**

See the preceding table for the types of DB systems that support the Internet and VCN Resolver.

Oracle recommends using the Internet and VCN Resolver for DNS. It's the default, built-in DNS functionality that comes with each VCN. It enables hosts in a VCN to resolve these items:

- Hostnames of other hosts in the same VCN
- Hostnames that are publicly published on the Internet

For general information about the Internet and VCN Resolver, see [DNS in Your Virtual Cloud Network](#).

For a DB system, the Internet and VCN Resolver handles resolution of all necessary endpoints: Object Storage endpoints (includes both the Object Storage endpoints and Swift endpoints), YUM repos, and SCANs (SCANs are used only with 2-node RAC systems).

By default, each VCN is configured to use the Internet and VCN Resolver. If you plan to use a custom DNS resolver, you must configure the VCN in a different way. For more information, see [Custom DNS Resolver](#).

### To use the Internet and VCN Resolver with your DB System

As part of the overall network setup, perform these tasks:

1. Create the VCN with the required DNS settings:
  - When [creating the VCN](#), select the check box for **Use DNS Hostnames in this VCN**.
  - Specify a DNS label for the VCN. See the restrictions in [Hostname restrictions for using the Internet and VCN Resolver](#).
  - Notice that you cannot change these VCN DNS settings after you create the VCN.
2. Create each subnet with the required DNS settings:
  - When [creating a subnet](#) in the VCN, select the check box for **Use DNS Hostnames in this Subnet**.
  - Specify a DNS label for the subnet. See the restrictions in [Hostname restrictions for using the Internet and VCN Resolver](#).
  - Notice that you cannot change these subnet DNS settings after you create the subnet.
3. Use the [default set of DHCP options](#) that come with the VCN:
  - When creating each subnet, configure it to use the VCN's default set of DHCP options.
  - By default, the default set of DHCP options is configured to use the Internet and VCN Resolver.
4. Create the DB system with a hostname prefix:

- Later, when creating the DB system, specify a value in the **Hostname Prefix** field. See the restrictions in [Hostname restrictions for using the Internet and VCN Resolver](#).
- Notice that the DB system's **Host Domain Name** value is automatically assigned based on the VCN and subnet DNS labels.

The resulting DB system has a fully qualified domain name (FQDN) based on the hostname prefix, VCN label, and subnet label you specify.

### Hostname restrictions for using the Internet and VCN Resolver

When you create the VCN, subnet, and DB system, you must carefully set the following identifiers, which are related to DNS in the VCN:

- VCN DNS label
- Subnet DNS label
- Hostname prefix for the DB system

These values make up the node's fully qualified domain name (FQDN):

```
<hostname_prefix><RAC_node_#>.<subnet_dns_label>.<vcn_dns_label>.oraclevcn.com
```

For RAC systems only, the Database service automatically appends a node number after the hostname prefix.

For example:

- **Node 1:** `dbsys1.ad1.acmevcniad.oraclevcn.com`
- **Node 2:** `dbsys2.ad1.acmevcniad.oraclevcn.com`

Requirement for the DB system's hostname prefix:

- Maximum 16 characters, otherwise the DB system deployment will fail.
- Cannot be the string `localhost`.

Requirements for the VCN and subnet DNS labels:

- Recommended maximum: 15 characters.
- No hyphens or underscores.
- Recommended: Include the region name in the VCN's name, and include the availability domain name in the subnet's name.
- The FQDN has a maximum total limit of 63 characters, so set the VCN and subnet DNS labels short enough to meet that requirement. Here is a safe general rule:  
`<16_chars_max>#. <15_chars_max>. <15_chars_max>.oraclevcn.com`
- The recommended maximums are not enforced when you create the VCN and subnets. However, the DB system deployment fails if the FQDN has more than 63 characters.

### CUSTOM DNS RESOLVER

See the preceding table for the types of DB systems that support the use of a custom DNS resolver.

A custom DNS resolver is a DNS server that you set up in your on-premises network and maintain yourself. It must resolve the endpoints required by the DB system.

By default, the VCN is configured to use the Internet and VCN Resolver. Therefore, if you instead want to use a custom DNS resolver, you must configure the VCN and DHCP options in a different way. See the following process.

### To use a custom DNS resolver with your DB system

As part of the overall network setup, perform these tasks:

1. Create the VCN with the recommended DNS settings:
  - When [creating the VCN](#), Oracle recommends that you select the check box for **Use DNS Hostnames in this VCN** and then specify a DNS label for the VCN. See the restrictions listed in [Hostname restrictions when using a custom DNS resolver](#).
  - Notice that you cannot change the preceding VCN DNS settings after you create the VCN. They are optional for a custom DNS server, but required if you use the

Internet and VCN Resolver. Therefore, Oracle recommends that you configure them now in case you later want to use the Internet and VCN Resolver.

2. Create each subnet with the recommended DNS settings:
  - When [creating a subnet](#) in the VCN, Oracle recommends that you select the check box for **Use DNS Hostnames in this Subnet** and then specify a DNS label for the subnet. See the restrictions listed in [Hostname restrictions when using a custom DNS resolver](#).
  - Notice that you cannot change the preceding subnet DNS settings after you create the subnet. They are optional for a custom DNS server, but required if you use the Internet and VCN Resolver. Therefore, Oracle recommends that you configure them now in case you later want to use the Internet and VCN Resolver.
3. Edit the [default set of DHCP options](#) to use a custom resolver:
  - When creating each subnet, configure it to use the VCN's default set of DHCP options.
  - [Edit the default set of DHCP options](#) so that **DNS Type** is set to **Custom Resolver**. Provide the IP address for at least one DNS server (maximum three). Optionally provide a single search domain (which will automatically be added to the host's `/etc/resolv.conf` file).
4. Create the DB system with required DNS entries:
  - Later, when creating the DB system, specify a **Hostname Prefix**.
  - For the **Host Domain Name**: If you selected the check box for **Use DNS Hostnames** in the preceding steps, the **Host Domain Name** is automatically generated from the VCN and subnet DNS labels. Otherwise, you must provide a value for the **Host Domain Name**. See the restrictions listed in [Hostname restrictions when using a custom DNS resolver](#).
  - Notice that when launching the DB system, the Database service automatically assigns an IP address from the VCN's CIDR block and resolves the address locally based on the host's `/etc/hosts` file. Your custom DNS resolver does not need to resolve the hostname in advance for the DB system launch to succeed.

### Hostname restrictions when using a custom DNS resolver

Requirement for the DB system's hostname prefix:

- Maximum 16 characters, otherwise the DB system deployment will fail.
- Cannot be the string *localhost*.

Requirements for the VCN and subnet DNS labels:

- You can provide a value for the DNS labels only if you select the check box for **Use DNS Hostnames** when creating the VCN and subnets. The resulting FQDN for the DB system follows this format:

*<hostname\_prefix>.<subnet\_DNS\_label>.<VCN\_DNS\_label>.oraclevcn.com*

- Recommended maximum for each DNS label: 15 characters.
- No hyphens or underscores.
- Recommended: Include the region name in the VCN's name, and include the availability domain name in the subnet's name.
- The FQDN has a maximum total limit of 63 characters, so set the VCN and subnet DNS labels short enough to meet that requirement. Here is a safe general rule:  
*<16\_chars\_max>.<15\_chars\_max>.<15\_chars\_max>.oraclevcn.com*
- The recommended maximums are not enforced when you create the VCN and subnets. However, the DB system deployment fails if the FQDN has more than 63 characters.

Requirements for the DB system's host domain name:

- You can provide a value in the **Host Domain Name** field only if you did not select the check box for **Use DNS Hostnames** when creating the VCN and subnets.
- No hyphens or underscores.
- Ensure that the value results in an FQDN that is no longer than 63 characters. Otherwise the DB system deployment will fail.

### **DNS: BETWEEN ON-PREMISES NETWORK AND VCN**

If you are using the Internet and VCN Resolver and want to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other, you can set up an instance in the VCN to be a custom DNS server. For an example of an implementation of this scenario with the Oracle Terraform provider, see [Hybrid DNS Configuration](#).

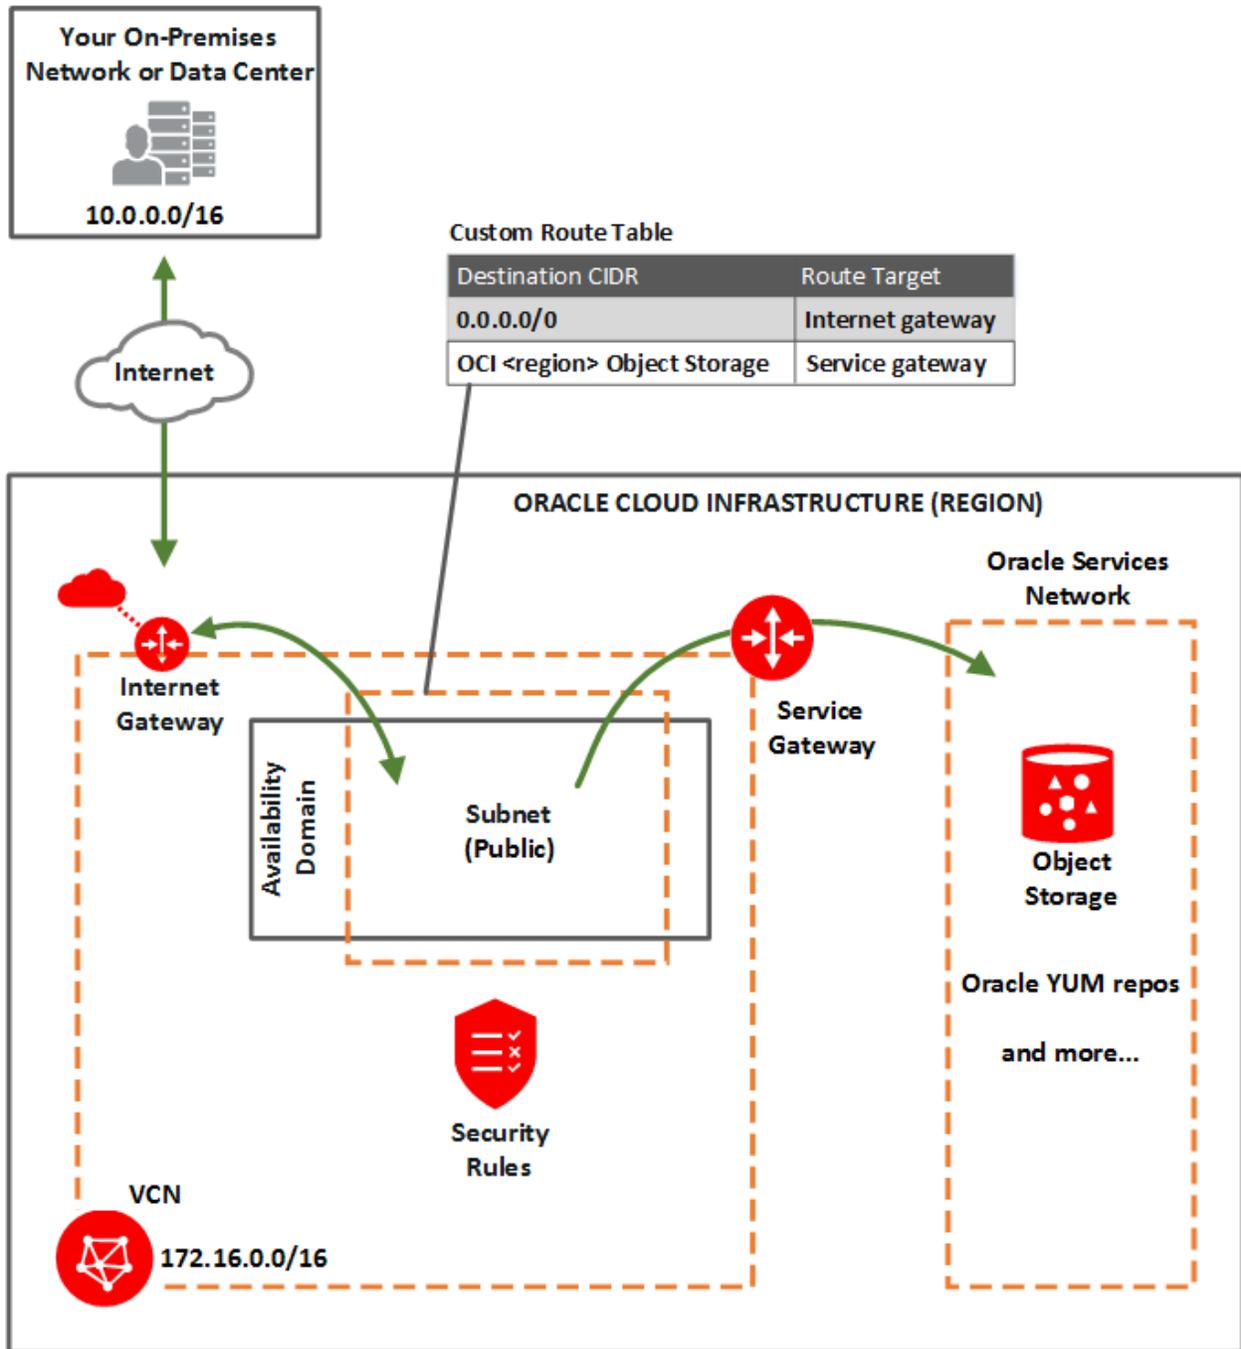
### **Service Gateway for the VCN**

Your VCN needs access to both Object Storage (for backing up databases, patching, and updating the cloud tooling on a DB system) and Oracle YUM repos for OS updates.

Depending on whether you use [option 1](#) or [option 2](#) described previously, you use the service gateway in different ways. See the next two sections.

#### **Option 1: Service Gateway Access Only to Object Storage**

You configure the subnet to use the [service gateway](#) for access only to Object Storage. As a reminder, here's the diagram for option 1:



In general, you must:

- Perform the [tasks for setting up a service gateway on a VCN](#), and specifically enable the service CIDR label called **OCI <region> Object Storage**.
- In the task for updating routing, add a route rule to the subnet's custom route table. For the destination service, use **OCI <region> Object Storage** and target = the service gateway.
- In the task for updating security rules for the subnet, perform the task on the DB system's custom network security group (NSG) or security list. Here you set up a security rule with the destination service set to **OCI <region> Object Storage**. See [Custom Security Rules](#).

### Option 2: Service Gateway Access to Both Object Storage and YUM Repos

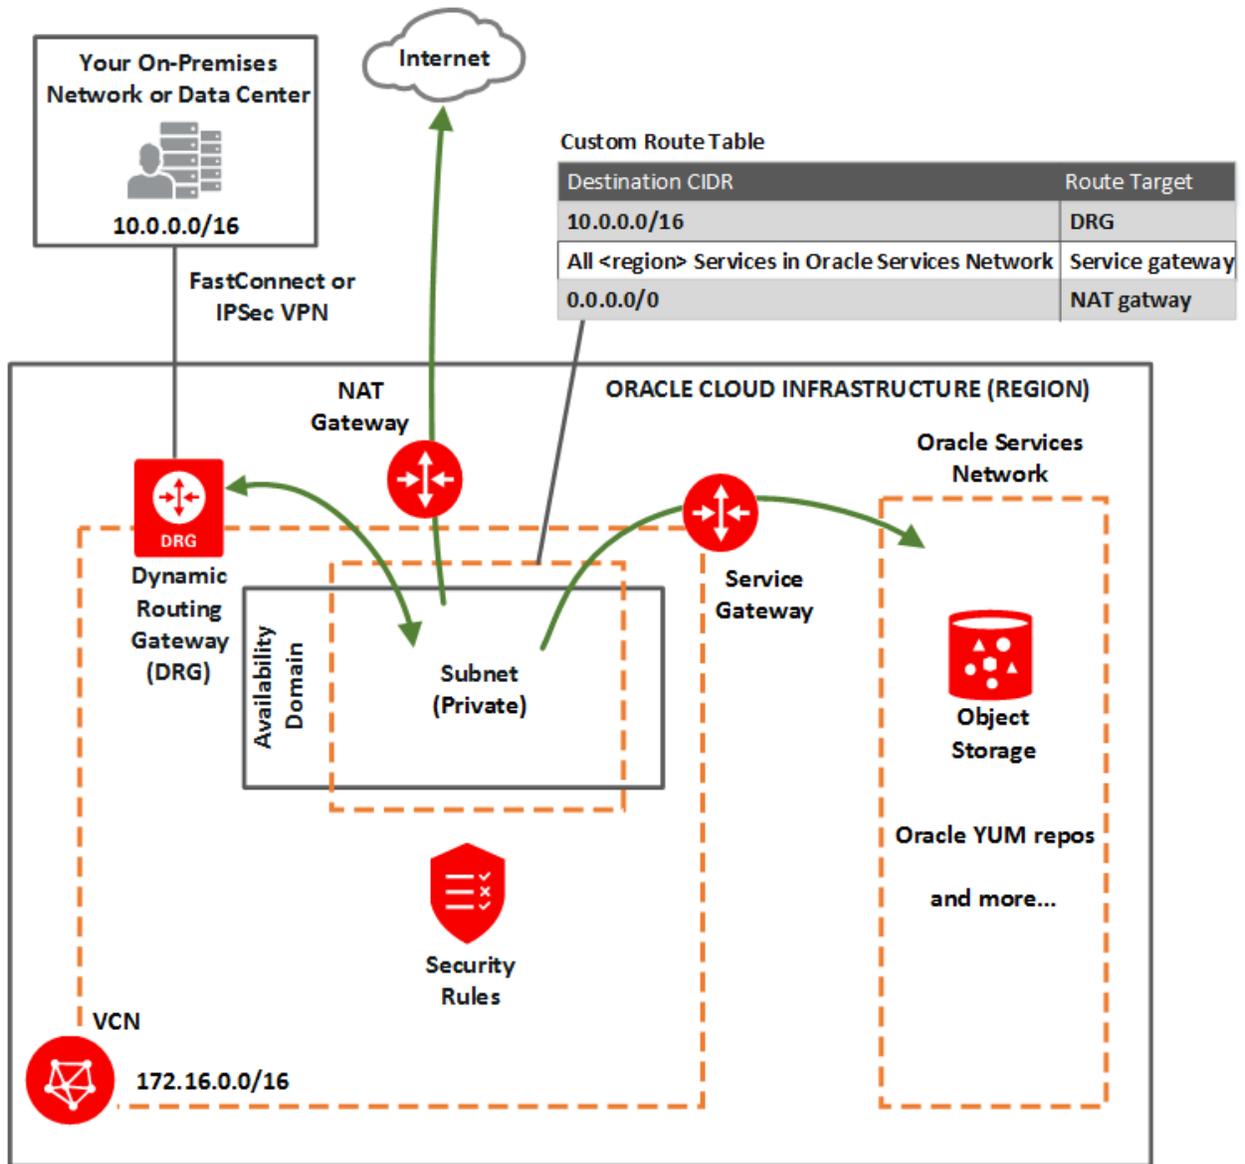
You configure the subnet to use the [service gateway](#) for access to the [Oracle Services Network](#), which includes both Object Storage and the Oracle YUM repos.



#### Important

See this [known issue](#) for information about accessing Oracle YUM services through the service gateway.

As a reminder, here's the diagram for option 2:



In general, you must:

- Perform the [tasks for setting up a service gateway on a VCN](#), and specifically enable the service CIDR label called **All <region> Services in Oracle Services Network**.
- In the task for updating routing in the subnet, add a rule to the subnet's custom route table. For the destination service, use **All <region> Services in Oracle Services Network** and target = the service gateway.
- In the task for updating security rules for the subnet, perform the task on the subnet's custom network security group (NSG) or security list. Here you set up a security rule with the destination service set to **All <region> Services in Oracle Services Network**. See [Custom Security Rules](#).

### Security Rules for the DB System

This section lists the [security rules](#) to use with your DB system. Security rules control the types of traffic allowed in and out of the DB system's compute nodes. The rules are divided into two sections.

There are different ways to implement these rules. For more information, see [Ways to Implement the Security Rules](#).



#### Important

Your instances running Oracle-provided DB system images also have firewall rules that control access to the instance. Make sure that both the instance's security rules and firewall rules are set correctly. Also see [Opening Ports on the DB System](#).

### GENERAL RULES REQUIRED FOR BASIC CONNECTIVITY

This section has several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules: the following rules are included by default in the [default security list](#).

### General ingress rule 1: Allows SSH traffic from anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22

### General ingress rule 2: Allows Path MTU Discovery fragmentation messages

This rule enables hosts in the VCN to receive Path MTU Discovery fragmentation messages. Without access to these messages, hosts in the VCN can have problems communicating with hosts outside the VCN.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** 4

### General ingress rule 3: Allows connectivity error messages within the VCN

This rule enables the hosts in the VCN to receive connectivity error messages from each other.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Your VCN's CIDR

- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** All

### General egress rule 1: Allows all egress traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

### CUSTOM SECURITY RULES

The following rules are necessary for the DB system's functionality.

### Custom ingress rule 1: Allows ONS and FAN traffic from within the VCN

This rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** VCN's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 6200

### Custom ingress rule 2: Allows SQL\*NET traffic from within the VCN

This rule is for SQL\*NET traffic and is required only if you need to enable client connections to

the database.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** VCN's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521



### Important

The preceding custom ingress rules 1 and 2 only cover connections initiated from within the VCN. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.

### Custom egress rule 1: Allows outbound SSH access

This rule enables SSH access between nodes in a 2-node DB system. It is redundant with the general egress rule in [General Rules Required for Basic Connectivity](#) (and in the [default security list](#)). It is optional but recommended in case the general rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP

- **Source Port Range:** All
- **Destination Port Range:** 22

### Custom egress rule 2: Allows access to Object Storage and YUM repos

This rule enables the DB system to communicate with Object Storage alone (for [option 1](#)), or with the Oracle Services Network, which includes both Object Storage and the Oracle YUM repos (for [option 2](#)). It is redundant with the general egress rule in [General Rules Required for Basic Connectivity](#) (and in the [default security list](#)). It is optional but recommended in case the general rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
  - For [option 1](#), use the service CIDR label called **OCI <region> Object Storage**
  - For [option 2](#), use the service CIDR label called **All <region> Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)

### Ways to Implement the Security Rules

The Networking service offers two ways to implement security rules within your VCN:

- [Network security groups](#)
- [Security lists](#)

For a comparison of the two methods, see [Comparison of Security Lists and Network Security Groups](#).

### If you use network security groups

If you choose to use [network security groups](#) (NSGs), here is the recommended process:

1. Create a network security group for DB systems. Add the following security rules to that NSG:
  - The rules listed in [General Rules Required for Basic Connectivity](#)
  - The rules listed in [Custom Security Rules](#)
2. When the database administrator [creates the DB system](#), they must choose several networking components (for example, which VCN and subnet to use). They can also choose which NSG or NSGs to use. Make sure they choose the NSG you created.

You could instead create one NSG for the general rules and a separate NSG for the custom rules. Then when the database administrator chooses which NSGs to use for the DB system, make sure they choose both NSGs.

### If you use security lists

If you choose to use [security lists](#), here is the recommended process:

1. Configure the subnet to use the required security rules:
  - a. Create a custom security list for the subnet and add the rules listed in [Custom Security Rules](#).
  - b. Associate the following two security lists with the subnet:
    - VCN's [default security list](#) with all its default rules. This automatically comes with the VCN.
    - The new custom security list you created for the subnet
2. Later when the database administrator creates the DB system, they must choose several networking components. When they select the subnet that you have already created and configured, the security rules are automatically enforced for the compute nodes created in the subnet.



### Warning

**Do not remove the default egress rule from the default security list.** If you do, instead make sure to include the following replacement egress rule in the subnet's custom security list:

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

## Managing Bare Metal and Virtual Machine DB Systems

This topic explains how to launch, start, stop, terminate, scale, manage licenses for, and check the status of a bare metal and virtual machine DB system, and set up DNS for a 1-node or 2-node RAC DB system.

When you launch a DB system using the Console, the API, or the CLI, the system is provisioned to support Oracle databases and an initial database is created based on the options you provide and some default options described later in this topic.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

### Prerequisites

You'll need the following items to launch any DB system.

- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the DB System via SSH. A sample public key, abbreviated for readability, is shown below.

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA...lo/gKMLVM2xzclxJr/Hc26biw3TXWGEakrK1OQ== rsa-key-20160304
```

For more information, see [Managing Key Pairs on Linux Instances](#).

- A correctly configured virtual cloud network (VCN) to launch the DB system in. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary for the DB system. For more information, see [Network Setup for DB Systems](#).
- If you plan to back up your DB system to Object Storage or to use the managed patching feature, Oracle recommends using a service gateway to enable access to Object Storage.
- For a 2-node RAC DB system, ensure that port 22 is open for both ingress and egress on the subnet, and that the security rules you create are stateful (the default), otherwise, the DB system might fail to provision successfully.

### Default Options for the Initial Database

To simplify launching a DB system in the Console and when using the API, the following default options are used for the initial database and for any additional databases that you create. (Several advanced options such as Time Zone can be set when you can use the `dbcli` command line interface to create databases.)

- **Console Enabled:** False
- **Create Container Database:** False for version 11.2.0.4 databases. Otherwise, true.
- **Create Instance Only (for standby and migration):** False
- **Database Home ID:** Creates a new database home
- **Database Language:** AMERICAN
- **Database Sizing Template:** odb2
- **Database Storage:** ACFS for version 11.2.0.4 databases. Otherwise, ASM.
- **Database Territory:** AMERICA
- **Database Unique Name:** The user-specified database name and a system-generated suffix, for example, `dbtst_phx1cs`.
- **PDB Admin Name:** `pdbuser` (Not applicable for version 11.2.0.4 databases.)

For a list of the database options that you can set, see [To launch a DB system](#).

### Using the Console

#### To launch a DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. Click **Launch DB System**.
4. In the **Launch DB System** dialog, enter the following:

*DB SYSTEM INFORMATION*

- **Compartment:** By default, the DB system launches in your current compartment and you can use the network resources in that compartment. Click the **click here** link in the dialog box if you want to enable compartment selection for the DB system, network, and subnet resources.
- **Display Name:** A friendly, display name for the DB system. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system.
- **Availability Domain:** The availability domain in which the DB system resides.
- **Shape Type:** The type of shape to use to launch the DB system. The shape type filters the list of available shapes to select from.
- **Shape:** The shape to use to launch the DB system. The shape determines the type of DB system and the resources allocated to the system.

### Bare metal shapes

- **BM.DenseIO2.52:** Provides a 1-node DB system (one bare metal server), with up to 52 CPU cores, 768 GB memory, and eight 6.4 TB locally attached NVMe drives (51.2 TB total) to the DB system.
- **BM.DenseIO1.36:** *Limited availability.* Provides a 1-node DB system (one bare metal server), with up to 36 CPU cores, 512 GB memory, and nine 3.2 TB locally attached NVMe drives (28.8 TB total) to the DB system.  
*Note:* BM.DenseIO1.36 is available only to monthly universal credit customers existing on or before November 9th, 2018. This shape is available only in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.

### Virtual machine shapes

Virtual machine X7 shapes:

- **VM.Standard2.1:** Provides a 1-node DB system with 1 core.
- **VM.Standard2.2:** Provides a 1- or 2-node DB system with 2 cores.
- **VM.Standard2.4:** Provides a 1- or 2-node DB system with 4 cores.
- **VM.Standard2.8:** Provides a 1- or 2-node DB system with 8 cores.
- **VM.Standard2.16:** Provides a 1- or 2-node DB system with 16 cores.
- **VM.Standard2.24:** Provides a 1- or 2-node DB system with 24 cores.

Virtual machine X5 shapes:

- **VM.Standard1.1:** Provides a 1-node DB system with 1 core.
- **VM.Standard1.2:** Provides a 1- or 2-node DB system with 2 cores.
- **VM.Standard1.4:** Provides a 1- or 2-node DB system with 4 cores.
- **VM.Standard1.8:** Provides a 1- or 2-node DB system with 8 cores.
- **VM.Standard1.16:** Provides a 1- or 2-node DB system with 16 cores.



### Note

- X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9th, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.
- VM.Standard1.1 and VM.Standard2.1 shapes cannot be used for 2-node RAC clusters.

- **Cluster Name:** A unique cluster name for a multi-node DB system. The name must begin with a letter and contain only letters (a-z and A-Z), numbers (0-9) and

hyphens (-). The cluster name can be no longer than 11 characters and is not case sensitive.

- **Total Node Count:** *Virtual machine DB systems only.* The number of nodes in the DB system. The number depends on the shape you select. You can specify 1 or 2 nodes for virtual machine DB systems, except for VM.Standard2.1 and VM.Standard1.1, which are single-node DB systems.
- **Oracle Database Software Edition:** The database edition supported by the DB system. For bare metal systems, you can mix supported database releases on the DB system to include older database versions, but not editions. The database edition cannot be changed and applies to all the databases in this DB system. Virtual machine systems support only one database.
- **Available Storage Size:** *Virtual machine DB systems only.* The amount of Block Storage you wish to allocate to the virtual machine DB system.
- **Cluster Name:** A unique cluster name for a multi-node DB system. The name must begin with a letter and contain only letters (a-z and A-Z), numbers (0-9) and hyphens (-). The cluster name can be no longer than 11 characters and is not case sensitive.
- **CPU Core Count:** The number of CPU cores for the DB system. Displays only if you select a shape that allows you to configure the number of cores. The text below the field indicates the acceptable values for that shape. For a multi-node DB system, the core count is evenly divided across the nodes.  
Except for virtual machine DB systems, you can increase the CPU cores to accommodate increased demand after you launch the DB system.
- **License Type:** The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this Oracle Cloud Infrastructure Database service resource will include both the Oracle Database software licenses and the service.

- **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this Oracle Cloud Infrastructure Database service resource. See [Bring Your Own License](#) for more information.
- **SSH Public Key:** The public key portion of the key pair you want to use for SSH access to the DB system. To provide multiple keys, paste each key on a new line. Make sure each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
- **Data Storage Percentage:** *For bare metal DB systems only.* The percentage (40% or 80%) assigned to DATA storage (user data and database files). The remaining percentage is assigned to RECO storage (database redo logs, archive logs, and recovery manager backups).
- **Available Storage Size:** *For virtual machine DB systems only.* The amount of block storage to allocate to the virtual machine DB system. Available storage can be scaled up or down as needed after provisioning your DB system.
- **Advanced Options:** *For bare metal DB systems only.*
  - **Disk Redundancy:** *For bare metal systems only.* The type of redundancy configured for the DB system.
    - **Normal** is 2-way mirroring, recommended for test and development systems.
    - **High** is 3-way mirroring, recommended for production systems.

### NETWORK INFORMATION

- **Virtual Cloud Network Compartment:** The compartment containing the network in which to launch the DB system.
- **Virtual Cloud Network:** The VCN in which to launch the DB system.
- **Subnet Compartment:** The compartment containing a subnet within the cloud network to attach the DB system to.
- **Client Subnet:** The subnet to which the bare metal or virtual machine DB system

should attach.

*For 1- and 2-node RAC DB systems:* Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet will cause the private interconnect to malfunction.

- **Hostname Prefix:** Your choice of host name for the bare metal or virtual machine DB system. The host name must begin with an alphabetic character, and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for bare metal and virtual machine DB systems is 16.



### Important

The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host Domain Name:** The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.
- **Host and Domain URL:** Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.

### DATABASE INFORMATION

- **Database Name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database Version:** The version of the initial database created on the DB system

when it is launched. After the DB system is active, you can create additional databases on it. You can mix database versions on the DB system, but not editions.

If you are launching a DB system with a virtual machine shape, you have option of selecting an older database version. Check **Display all database versions** to include older database versions in the drop-down list of database version choices. See [Availability of Older Database Versions for Virtual Machine DB Systems](#) for more information.



### Note

When you display all database versions in the drop-down list, the latest database version of each release is represented twice in the list, as follows:

- Once using four numeric segments and the notation "(latest)". For example:  
12.2.0.1 (latest)
  - Once using five numeric segments, without the "(latest)" notation. For example:  
12.2.0.1.180417
- **PDB Name:** *Not applicable to version 11.2.0.4.* The name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of 8 alphanumeric characters. The only special character permitted is the underscore ( \_).
  - **Database Admin Password:**  
A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2

numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.

- **Confirm Database Admin Password:** Re-enter the Database Admin Password you specified.
- **Automatic Backup:** (Optional) If you enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The default selection is 30 days.
- **Database Workload:**

Select the workload type that best suits your application.

  - **Online Transactional Processing (OLTP)** configures the database for a transactional workload, with a bias towards high volumes of random data access.
  - **Decision Support System (DSS)** configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations.
- **Character Set:** The character set for the database. The default is AL32UTF8.

**Advanced Options:**

  - **Character Set:** The character set for the database. The default is AL32UTF8.
  - **National Character Set:** The national character set for the database. The default is AL16UTF16.
  - **Fault Domain:** The fault domain(s) in which the DB system resides. You can choose which fault domain to use for your DB system. For 2-node RAC DB systems, you can specify which two fault domains are to be used. Oracle recommends that you place each node of a 2-node RAC DB system in a different fault domain. For more information on fault domains, see [Fault Domains](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a

resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Launch DB System**.

The DB system appears in the list with a status of Provisioning. The DB system's icon changes from yellow to green (or red to indicate errors).

6. Wait for the DB system's icon to turn green, with a status of Available, and then click the highlighted DB system name.

Details about the DB system are displayed.

7. Note the IP addresses; you'll need the private or public IP address, depending on network configuration, to connect to the DB system.

### To launch a new DB system from a backup

Before you begin, note the following:

- When you launch a new DB system from a backup, the availability domain will be the same as where the backup is hosted.
- The shape you specify must be the same type as the database from which the backup was taken. For example, if you are using a backup of a 1-node database, then the DB system you select as your target must also be a 1-node DB system.
- The Oracle database software edition you specify must be an equal or greater edition than that of the backed up database.
- If you specify a virtual machine DB system shape, the Available Storage Size will default to the data size of the backup, rounded up to the closest storage size option. However, you can specify a larger storage size.
- If you are creating a database from an automatic backup, you may choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0

backup. For more information on automatic backups, see [Automatic Incremental Backups](#)

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Navigate to the backup you wish to use to create a new database. You can select a backup from a database details page, or select a backup that appears in your compartment's list of standalone backups.

### To navigate to a database's list of backups

- a. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
- b. Find the database associated with the backup you wish to use, and click its name to display details about it.  
A list of backups is displayed in the default view of the database details. You can also access the list of backups for a database by clicking on **Backups** under **Resources**.

### To navigate to the list of standalone backups for your current compartment

- a. Click **Standalone Backups** under **Bare Metal, VM, and Exadata**.
  - b. In the list of standalone backups, find the backup you want to use to create the database.
4. Click the Actions icon (three dots) for the backup you are interested in, and then click **Create Database**.

5. In the **Create Database from Backup** dialog, enter the following:

### *DB SYSTEM INFORMATION*

- **DB System Information:** Select **Launch New DB System**.
- **Display Name:** A friendly, display name for the DB system. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system.
- **Shape:** The shape to use to launch the DB system. The shape determines the type of DB system and the resources allocated to the system.  
The selected shape must support the same number of nodes as the DB system from which the backup was created.

### Bare metal shapes

- **BM.DenseIO2.52:** Provides a 1-node DB system (one bare metal server), with up to 52 CPU cores, 768 GB memory, and eight 6.4 TB locally attached NVMe drives (51.2 TB total) to the DB system.
- **BM.DenseIO1.36:** *Limited availability.* Provides a 1-node DB system (one bare metal server), with up to 36 CPU cores, 512 GB memory, and nine 3.2 TB locally attached NVMe drives (28.8 TB total) to the DB system.  
*Note:* BM.DenseIO1.36 is available only to monthly universal credit customers existing on or before November 9th, 2018. This shape is available only in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.

### Virtual machine shapes)

Virtual machine X7 shapes:

- **VM.Standard2.1:** Provides a 1-node DB system with 1 core.
- **VM.Standard2.2:** Provides a 1- or 2-node DB system with 2 cores.

- **VM.Standard2.4:** Provides a 1- or 2-node DB system with 4 cores.
- **VM.Standard2.8:** Provides a 1- or 2-node DB system with 8 cores.
- **VM.Standard2.16:** Provides a 1- or 2-node DB system with 16 cores.
- **VM.Standard2.24:** Provides a 1- or 2-node DB system with 24 cores.

Virtual machine X5 shapes:

- **VM.Standard1.1:** Provides a 1-node DB system with 1 core.
- **VM.Standard1.2:** Provides a 1- or 2-node DB system with 2 cores.
- **VM.Standard1.4:** Provides a 1- or 2-node DB system with 4 cores.
- **VM.Standard1.8:** Provides a 1- or 2-node DB system with 8 cores.
- **VM.Standard1.16:** Provides a 1- or 2-node DB system with 16 cores.



### Note

- X5-based shapes availability is limited to monthly universal credit customers existing on or before November 9th, 2018, in the US West (Phoenix), US East (Ashburn), and Germany Central (Frankfurt) regions.
- VM.Standard1.1 and VM.Standard2.1 shapes cannot be used for 2-node RAC clusters.

- **Total Node Count:** *Virtual machine DB systems only.* The number of nodes in the DB system. The number depends on the shape you select. You can specify 1 or 2 nodes for virtual machine DB systems, except for VM.Standard2.1 and VM.Standard1.1, which are single-node DB systems.
- **Oracle Database Software Edition:** The database edition supported by the DB

system. For bare metal systems, you can mix supported database releases on the DB system to include older database versions, but not editions. The database edition cannot be changed and applies to all the databases in this DB system. Virtual machine systems support only one database.

- **Available Storage Size:** *Virtual machine DB systems only.* The amount of Block Storage you wish to allocate to the virtual machine DB system.
- **Cluster Name:** A unique cluster name for a multi-node DB system. The name must begin with a letter and contain only letters (a-z and A-Z), numbers (0-9) and hyphens (-). The cluster name can be no longer than 11 characters and is not case sensitive.
- **License Type:** The type of license you want to use for the DB system. Your choice affects metering for billing.
  - **License included** means the cost of this Oracle Cloud Infrastructure Database service resource will include both the Oracle Database software licenses and the service.
  - **Bring Your Own License (BYOL)** means you will use your organization's Oracle Database software licenses for this Oracle Cloud Infrastructure Database service resource. See [Bring Your Own License](#) for more information.
- **SSH Public Key:** The public key portion of the key pair you want to use for SSH access to the DB system. To provide multiple keys, paste each key on a new line. Make sure each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
- **Data Storage Percentage:** *For bare metal DB systems only.* The percentage (40% or 80%) assigned to DATA storage (user data and database files). The remaining percentage is assigned to RECO storage (database redo logs, archive logs, and recovery manager backups).
- **Available Storage Size:** *For virtual machine DB systems only.* The amount of block storage to allocate to the virtual machine DB system. Available storage can be scaled up or down as needed after provisioning your DB system.

If you are creating a DB system from a backup, the minimum value for available storage is determined by the size of the backup.

- **Advanced Options:**

- **Disk Redundancy:** *For bare metal systems only.* The type of redundancy configured for the DB system.
  - **Normal** is 2-way mirroring, recommended for test and development systems.
  - **High** is 3-way mirroring, recommended for production systems.
- **Fault Domain:** The fault domain(s) in which the DB system resides. You can choose which fault domain to use for your DB system. For 2-node RAC DB systems, you can specify which two fault domains are to be used. Oracle recommends that you place each node of a 2-node RAC DB system in a different fault domain. For more information on fault domains, see [Fault Domains](#).

### NETWORK INFORMATION

- **Virtual Cloud Network:** The VCN in which to launch the DB system.
- **Client Subnet:** The subnet to which the bare metal or virtual machine DB system should attach.

*For 1- and 2-node RAC DB systems:* Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet will cause the private interconnect to malfunction.

- **Hostname Prefix:** Your choice of host name for the bare metal or virtual machine DB system. The host name must begin with an alphabetic character, and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for bare metal and virtual machine DB systems is 16.



### Important

The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

#### *DATABASE INFORMATION*

- **Database Name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
  - **Database Admin Password:**  
A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
  - **Confirm Database Admin Password:** Re-enter the Database Admin Password you specified.
  - **Password for Transparent Data Encryption (TDE) Wallet or RMAN Encryption:**  
Enter either the TDE wallet password or the RMAN encryption password for the backup, whichever is applicable. The TDE wallet password is the SYS password provided when the database was created by using the Oracle Cloud Infrastructure Console, API, or CLI. The RMAN encryption password is typically required instead if the password was subsequently changed manually.
6. Click **Create Database**.
  7. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging,

see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

### To check the status of a DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the system you're interested in and check its icon. The color of the icon and the text below it indicates the status of the system.
  - **Provisioning:** Yellow icon. Resources are being reserved for the DB system, the system is booting, and the initial database is being created. Provisioning can take several minutes. The system is not ready to use yet.
  - **Available:** Green icon. The DB system was successfully provisioned. A few minutes after the system enters this state, you can SSH to it and begin using it.
  - **Terminating:** Gray icon. The DB system is being deleted by the terminate action in the Console or API.
  - **Terminated:** Gray icon. The DB system has been deleted and is no longer available.
  - **Failed:** Red icon. An error condition prevented the provisioning or continued operation of the DB system.

To view the status of a database node, under Resources, click **Nodes** to see the list of nodes. In addition to the states listed for a DB system, a node's status can be one of the following:

- **Starting:** Yellow icon. The database node is being powered on by the start or reboot action in the Console or API.
- **Stopping:** Yellow icon. The database node is being powered off by the stop or reboot action in the Console or API.

- **Stopped:** Yellow icon. The database node was powered off by the stop action in the Console or API.

You can also check the status of DB systems and database nodes by using the [ListDbSystems](#) or [ListDbNodes](#) API operations, which return the `lifecycleState` attribute.

To start, stop, or reboot a DB system



### Tip

Oracle recommends that you run a Network Time Protocol (NTP) daemon to keep system clocks stable during rebooting. If you need information about an NTP daemon, see [Setting Up "NTP \(Network Time Protocol\) Server" in RHEL/CentOS 7](#).

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the DB system you want to stop or start, and then click its name to display details about it.
4. In the list of nodes, click the Actions icon (three dots) for a node and then click one of the following actions:
  - **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
  - **Stop:** Shuts down the node. After the node is powered off, the **Start** action is enabled.
  - **Reboot:** Shuts down the node, and then restarts it.



### Note

- Resource billing differs between bare metal and virtual machine DB systems as follows:
  - **Bare metal DB systems** - The **Stop** state has no effect on the resources you consume. Billing continues for nodes that you stop, and related resources continue to apply against any relevant quotas. You must **Terminate** a DB system to remove its resources from billing and quotas.
  - **Virtual machine DB systems** - Stopping a node stops billing for all OCPUs associated with that node. Billing resumes if you restart the node.
- After you restart or reboot a node, the floating IP address might take several minutes to be updated and display in the Console.

### To scale the CPU cores for a bare metal DB system

If a multi-node DB system requires more compute node processing power, you can scale up (burst) the number of enabled CPU cores in the system.



### Note

You cannot change the number of CPU cores for a virtual machine DB system.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the system you want to scale and click its highlighted name.  
The system details are displayed.
4. Click **Scale Up/Down** and then change the number in **Total CPU Core Count**. The text below the field indicates the acceptable values, based on the shape used when the DB system was launched.
5. Click **Scale Up/Down DB System**.

### To scale up the storage for a virtual machine DB system

If a virtual machine DB system requires more block storage, you can increase the storage at any time without impacting the system.



### Note

This procedure does not apply to bare metal DB systems.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.

A list of DB systems is displayed.

3. In the list of DB systems, find the system you want to scale up and click its highlighted name.

The system details are displayed.

4. Click **Scale Storage Up**, and then select the new storage size from the drop-down list.
5. Click **Scale Storage Up**.

### To terminate a DB system

Terminating a DB system permanently deletes it and any databases running on it.



#### Note

The database data is local to the DB system and will be lost when the system is terminated. Oracle recommends that you back up any data in the DB system prior to terminating it.

Terminating a DB system removes all automatic incremental backups of all databases in the DB system from Oracle Cloud Infrastructure Object Storage. Full backups remain in Object Storage as standalone backups which you can use to create a new database. See [Recovering a Database from Object Storage](#).

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. For the DB system you want to terminate, click the Actions icon (three dots) and then click **Terminate**.

4. Confirm when prompted. .  
The DB system's icon indicates Terminating.

At this point, you cannot connect to the system and any open connections will be terminated.

### To manage your BYOL database licenses

If you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the system you want to scale and click its highlighted name.  
The system details are displayed.
4. Click **Scale Up/Down OCPU**, and then change the number.

### To manage tags for your DB systems and database resources

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. Find the DB system or database resource you're interested in, and click the name.
4. Click the **Tags** tab to view or edit the existing tags. Or click **Apply Tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage DB system components.

DB systems:

- [ListDbSystems](#)
- [GetDbSystem](#)
- [LaunchDbSystem](#)
- [TerminateDbSystem](#)

Database homes:

- [ListDbHomes](#)
- [GetDbHome](#)
- [CreateDbHome](#)
- [DeleteDbHome](#)

Databases:

- [ListDatabases](#)
- [GetDatabase](#)

Nodes:

- [DbNodeAction](#): Use this operation to power cycle a node in the DB system.
- [ListDbNodes](#)
- [GetDbNode](#)

Shapes and database versions:

- [ListDbSystemShapes](#)
- [ListDbVersions](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

### Setting up DNS for a DB System

DNS lets you use host names instead of IP addresses to communicate with a DB system. You can use the *Internet and VCN Resolver* (the DNS capability built into the VCN) as described in [DNS in Your Virtual Cloud Network](#).

Alternatively, you can use your choice of DNS server. You associate the host name and domain name to the public or private IP address of the DB system. You can find the host and domain names and IP addresses for the DB system in the Oracle Cloud Infrastructure Console on the **Database** page.

To associate the host name to the DB system's public or private IP address, contact your DNS administrator and request a custom DNS record for the DB system's IP address. For example, if your domain is example.com and you want to use clouddb1 as the host name, you would request a DNS record that associates clouddb1.example.com to your DB system's IP address.

If you provide the public IP address to your DNS administrator as described above, you should also associate a custom domain name to the DB system's public IP address:

1. Register your domain name through a third-party domain registration vendor, such as register.com.
2. Resolve your domain name to the DB system's public IP address, using the third-party domain registration vendor console. For more information, refer to the third-party domain registration documentation.

### Connecting to a DB System

This topic explains how to connect to an active DB system. How you connect depends on the client tool or protocol you use, the purpose of the connection, and how your cloud network is set up. You can find information on various networking scenarios in [Overview of Networking](#),

but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

### Prerequisites

This section describes prerequisites you'll need to perform various tasks in this topic.

- To use the Console or the API to get the default administration service connection strings, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. See [Authentication and Authorization](#) for more information on user authorizations for the Oracle Cloud Infrastructure Database service.
- To connect to the database, you'll need the public or private IP address of the DB system. Use the private IP address to connect to the DB system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the DB System's public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Oracle Cloud Infrastructure Console on the **Database** page.
- For Secure Shell (SSH) access to the DB system, you'll need the full path to the file that contains the private key associated with the public key used when the DB system was launched.

If you have problems connecting, see [Troubleshooting Connection Issues](#).

### Database Services and Connection Strings

Database services allow you to control client access to a database instance depending on the functionality needed. For example, you might need to access the database for administration purposes only or you might need to connect an application to the database. Connection strings are specific to a database service.

When you provision a DB system, a default database administration service is automatically created. For 12c and later Oracle Databases, this service is for administrating the database at the CDB level. Because this service provides limited functionality, it is not suitable for connecting an application. Oracle recommends that you create a default application service for the initial database after you create your DB system. For 12c and later Oracle Databases, application services connect at the PDB level. Here are some important functions an application service can provide:

- Workload identification
- Load balancing
- Application continuity and Transaction Guard
- Fast Application Notification
- Resource assignment based on the service name

For details about these and other High Availability capabilities, see [Client Failover Best Practices for Highly Available Oracle Databases](#).

### CREATING AN APPLICATION SERVICE

You use the `srvctl` utility to create an application service. Before you can connect to the service, you must start it.

### To create an application service for a PDB or an 11g Oracle database

1. Log in to the DB system host as `opc`.
2. Switch to the oracle user, and set your environment to the Oracle Database you want to administer.

```
$ sudo su - oracle
$. oraenv
ORACLE_SID = [oracle] ? <database_name>
The Oracle base has been set to /u01/app/oracle
```

3. Create the application service for the database. Include the `pdb` option only if you are creating an application service for a PDB.

## CHAPTER 11 Database

---

```
$ srvctl add service
-db <DB_unique_name>
-pdb <PDB_name>
-service <app_service_name>
-role PRIMARY
-notification TRUE
-session_state dynamic
-failovertype transaction
-failovermethod basic
-commit_outcome TRUE
-failoverretry 30
-failoverdelay 10
-replay_init_time 900
-clbgoal SHORT
-rlbgoal SERVICE_TIME
-preferred <rac_node1>,<rac_node2>
-retention 3600
```

Note that the preferred option is required only for multi-node databases to specify the hostname of the node in the RAC.

#### 4. Start the application service.

```
$ srvctl start service -db <DB_unique_name> -s <app_service_name>
```

For more information about services for a PDB, see [Managing Services for PDBs](#).

### DATABASE CONNECTION STRINGS

You must use the appropriate connection string to access a database administration or application service. You can use the Console or the API to get the string for connecting to the default administration service from within a VCN. For 12c and later Oracle Databases, this service is for administrating the database at the CDB level. The string is provided in both the Easy Connect and in the full connect descriptor (long) format. Use the long format for the connection if hostname resolution is not available. You can also use the long format to create an alias in the tnsnames.ora file.

For accessing a database service within the VCN, the connection string for a Real Application Cluster (RAC) DB system uses the Single Client Access Name (SCAN) while the connection string for single instance DB system uses the hostname instead.

The private SCAN name is a Round Robin DNS entry created when you launch a 2-node RAC DB system. The private SCAN name is resolvable only within the VCN. If the client and the database are in the same VCN, the connection mechanism is the same as an on-premises RAC database; all the features provided by VIPs and SCAN VIPs, such as server side load balancing and VIP failover, are available.



### Note

If you manually change the `DB_UNIQUE_NAME`, `DB_DOMAIN`, or listener port on the DB system, the connection strings you see in the Console or API will not reflect your changes. Ensure that you use the actual values of these parameters when you make a connection.

### To get the connection strings for the default administration service

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.
3. Find the DB system you're interested in, and click the name.
4. Click **DB Connection**.
5. Click the applicable link to view or copy the connection string.

You can derive the connection strings for other database services by replacing part of the default application service connection string with the applicable values.

### To derive the connection string for a PDB administration service or an application service

1. Follow the procedure to get the Easy Connect string for the default administration service. That string should have the following format:

```
<hostname|SCAN>:1521/<DB_unique_name>.<DB_domain>
```

### 2. Make the appropriate substitution:

- For the PDB administration service, replace `DB_unique_name` with the PDB name.

```
<hostname|SCAN>:1521/<PDB_name>.<DB_domain>
```

- For an application service, replace `DB_UNIQUE_NAME` with the name of the application service.

```
<hostname|SCAN>:1521/<app_service_name>.<DB_domain>
```

## Connecting to a Database Service by Using SQL\*Net

This section describes how to connect to a database service from a computer that has a SQL\*Net client installed. Port 1521 must be open to support the SQL\*Net protocol.

### CONNECTING FROM WITHIN THE VCN

For security reasons, Oracle recommends that you connect to your database services from within the VCN. You can use this method whether you are connecting to an administration service or to an application service.

To connect using SQL\*Plus, you run the following command using the applicable connection string:

```
sqlplus system/<password>@<connection_string>
```

Consider the following:

- If your system is not using the VCN Resolver, ensure that the DB system's hostname (for single-node systems) or SCAN name (for multi-node systems) can be resolved. See [DNS in Your Virtual Cloud Network](#) for information about DNS name resolution.
- For connecting to the administration service of a PDB, ensure that the PDB is open or the service will not be available.
- For connecting to an application service, ensure that the service is started. For Fast Application Notification to work, ensure that port 6200 can be reached. See [Client Failover Best Practices for Highly Available Oracle Databases](#) for information about Fast Application Notification.

### CONNECTING FROM THE INTERNET

Although Oracle does not recommend connecting to your database from the Internet, you can connect to a database service by using a public IP address if port 1521 is open to the public for ingress.

To use this method, you run the following command using the public IP address instead of the hostname or SCAN in the connection string:

```
sqlplus system/<password>@<public_IP>:1521/<service_name>.<DB_domain>
```

Consider the following:

- SCANS and hostnames are not resolvable on the Internet, therefore load balancing and failover for multi-node DB systems, which rely on these names, cannot work.
- For multi-node DB systems, which normally use SCANS, you must specify the IP address of one of the RAC hosts to access the database.



#### Important

Do not use this method to connect to the database from within the VCN. Doing so negatively impacts performance because traffic to the database is routed out of the VCN and back in through the public IP address.

### EXAMPLE: CONNECTING IN SQL DEVELOPER USING SQL\*NET

Prerequisites:

- Ensure that port 1521 is open for the Oracle default listener. (You can do this by checking the DB system's security list.)
- If port 1521 is open only to hosts in the VCN, then you must run your SQL Developer client from a machine that has direct access to the VCN. If you are connecting to the database from the Internet instead, then the public IP address of your computer must be granted access to port 1521 in the security list. (Alternatively, the security list can

grant full access to port 1521, however, this is not recommended for security reasons.) You must use the public IP address of the host because connecting from the Internet does not support SCAN name resolution.

### To connect from within the VCN

After the prerequisites are met, start SQL Developer and create a connection by supplying the following connection details:

- **Username:** sys as sysdba
- **Password:** The **Database Admin Password** that was specified in the **Launch DB System** dialog in the Console.
- **Hostname:** The hostname as it appears in the Easy Connect format of the connection string. (See [Database Connection Strings](#) for help with getting the connection string and identifying the hostname.)
- **Port:** 1521
- **Service name:** The concatenated name of the service and host domain name, for example, db1\_phx1tv.example.com. You can identify this value as the last part of the Easy Connect string, *<service\_name>.<DB\_domain>*.

### To connect from The Internet by using public IP addresses

You can use the node's public IP address to connect to the database if the database is in a VCN that has an internet gateway. However, there are important implications to consider:

- When the client uses the public IP address, the client bypasses the SCAN listener and reaches the node listener, so server side load balancing is not available.
- When the client uses the public IP address, it cannot take advantage of the VIP failover feature. If a node becomes unavailable, new connection attempts to the node will hang until a TCP/IP timeout occurs. You can set client side SQL\*Net parameters to limit the TCP/IP timeout.

The following `tnsnames.ora` example shows a connection string that includes the `CONNECT_TIMEOUT` parameter. This parameter controls the TCP timeout for connecting to a node.

```
test=
 (DESCRIPTION =
 (CONNECT_TIMEOUT=3)
 (ADDRESS_LIST=
 (ADDRESS = (PROTOCOL = TCP) (HOST = <public_IP1>) (PORT = 1521))
 (ADDRESS = (PROTOCOL = TCP) (HOST = <public_IP2>) (PORT = 1521))
)
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
)
)
```

For more information about using the `CONNECT_TIMEOUT` parameter, see [Client Failover Best Practices for Highly Available Oracle Databases](#).

### Connecting to a Database with a Public IP by Using SSH Tunneling

You can access the services of DB system databases with public IP addresses by using SSH tunneling. The main advantage of this method is that port 1521 does not need to be opened to the public internet. However, just like accessing the database with a public IP using a SQL\*Net client, load balancing and failover for multi-node DB systems cannot work because they rely on SCANS and hostnames.

Oracle SQL Developer and Oracle SQLcl and are two tools that facilitate the use of tunneling for Oracle Database access.

To open a tunnel, and then connect to a database service by using SQLcl, you run commands like the following:

```
SQL> sstunnel opc@<public_IP> -i <private_key> -L <local_port>:<private_IP>:1521
Using port:22
SSH Tunnel connected
SQL> connect system/<password>@localhost:<local_port>/<service_name>.<DB_domain>
```

See [Oracle SQL Developer](#) and [Oracle SQLcl](#) for information about these tools.

### Connecting to a Database by Using SSH and the Bequeath Protocol

This method allows you to connect to the database without using the network listener. It should be used to connect only for administration purposes.

When connecting to a multi-node DB system, you'll SSH to each individual node in the cluster.

#### To connect from a UNIX-style system

Use the following SSH command to access the DB system:

```
$ ssh -i <private_key> opc@<DB_system_IP_address>
```

<private\_key> is the full path and name of the file that contains the private key associated with the DB system you want to access.

Use the DB system's private or public IP address depending on your network configuration. For more information, see [Prerequisites](#).

#### To connect from a Windows system

1. Open `putty.exe`.
2. In the **Category** pane, select **Session** and enter the following fields:
  - **Host Name (or IP address):** `opc@<DB_system_IP_address>`  
Use the DB system's private or public IP address depending on your network configuration. For more information, see [Prerequisites](#).
  - **Connection type:** SSH
  - **Port:** 22
3. In the **Category** pane, expand **Connection**, expand **SSH**, and then click **Auth**, and browse to select your private key.
4. Optionally, return to the **Session** category screen and save this session information for reuse later.
5. Click **Open** to start the session.

### To access a database after you connect

1. Log in as opc and then sudo to the grid user.

```
login as: opc

[opc@ed1db01 ~]$ sudo su - grid
```

2. List all the databases on the system.

```
root@ed1db01]# srvctl config database -v

cdbm01 /u02/app/oracle/product/12.1.0/dbhome_2 12.1.0.2.0
exadb /u02/app/oracle/product/11.2.0/dbhome_2 11.2.0.4.0
mmdb /u02/app/oracle/product/12.1.0/dbhome_3 12.1.0.2.0
```

3. Connect as the oracle user and get the details about one of the databases by using the `srvctl` command.

```
[root@ed1db01 ~]# su - oracle
[oracle@ed1db01 ~]$. oraenv
ORACLE_SID = [oracle] ? cdbm01
The Oracle base has been set to /u02/app/oracle
[oracle@ed1db01 ~]$ srvctl config database -d cdbm01
Database unique name: cdbm01 <<== DB unique name
Database name:
Oracle home: /u02/app/oracle/product/12.1.0/dbhome_2
Oracle user: oracle
Spfile: +DATA1/cdbm01/spfilecdbm01.ora
Password file: +DATA1/cdbm01/PASSWORD/passwd
Domain: data.customer1.oraclevcn.com
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: DATA1,RECO1
Mount point paths:
Services:
Type: RAC
Start concurrency:
```

## CHAPTER 11 Database

---

```
Stop concurrency:
OSDBA group: dba
OSOPER group: racoper
Database instances: cdbm011,cdbm012 <== SID
Configured nodes: ed1db01,ed1db02
Database is administrator managed
```

4. Set the ORACLE\_SID and ORACLE\_UNIQUE\_NAME using the values from the previous step.

```
[oracle@ed1db01 ~]$ export ORACLE_UNIQUE_NAME=cdbm01
[oracle@ed1db01 ~]$ export ORACLE_SID=cdbm011
[oracle@ed1db01 ~]$ sqlplus / as sysdba

SQL*Plus: Release 12.1.0.2.0 Production on Wed Apr 19 04:10:12 2017

Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c EE Extreme Perf Release 12.1.0.2.0 - 64bit Production
With the Partitioning, Real Application Clusters, Automatic Storage Management, Oracle Label
Security,
OLAP, Advanced Analytics and Real Application Testing options
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [GetDatabase](#) API operation to get the default administration service connection strings.

### Troubleshooting Connection Issues

The following issues might occur when connecting to a DB system or database.

### **ORA-28365: WALLET IS NOT OPEN ERROR**

For a 1-node DB system or 2-node RAC DB system, regardless of how you connect to the DB system, *before* you use OS authentication to connect to a database (for example, `sqlplus / as sysdba`) be sure to set the `ORACLE_UNQNAME` variable. Otherwise, commands that require the TDE wallet will result in the error `ORA-28365: wallet is not open`.

Note that this is not an issue when using a TNS connection because `ORACLE_UNQNAME` is automatically set in the database CRS resource.

### **SSH ACCESS STOPS WORKING**

If the DB system's root volume becomes full, you might lose the ability to SSH to the system (the SSH command will fail with permission denied errors). Before you copy a large amount of data to the root volume, for example, to migrate a database, use the `dbcli create-dbstorage` command to set up storage on the system's NVMe drives and then copy the database files to that storage. For more information, see [Setting Up Storage on the DB System](#).

### **WHAT NEXT?**

Before you begin updating your DB system, review the information in [Updating a DB System](#).

For information about setting up an Enterprise Manager console to monitor your databases, see [Monitoring a Database](#).

## Updating a DB System



### **Note**

This topic is not applicable to Exadata DB systems. For information on how to update an Exadata DB system, see [Updating an Exadata DB System](#)

This topic includes information and instructions on how to update the OS of a bare metal or virtual machine DB system.



### Warning

- Review all of the information before you begin updating the system. Updating the operating system through methods not described on this page can cause permanent loss of access.
- Always back up your databases prior to updating your DB system's operating system.

### Bash Profile Updates

Do not add interactive commands such as `oraenv`, or commands that might return an error or warning message, to the `.bash_profile` file for the `grid` or `oracle` users. Adding such commands can prevent Database service operations from functioning properly.

### Essential Firewall Rules

For a 1-node DB system or 2-node RAC DB system, do not remove or modify the following firewall rules in `/etc/sysconfig/iptables`:

- The firewall rules for ports 1521, 7070, and 7060 allow the Database service to manage the DB system. Removing or modifying them can result in the Database Service no longer operating properly.
- The firewall rules for 169.254.0.2:3260 and 169.254.0.3:80 prevent non-root users from escalating privileges and tampering with the system's boot volume and boot process. Removing or modifying these rules can allow non-root users to modify the system's boot volume.

### OS Updates

Before you update the OS, review the following important guidelines and information:

- Back up your DB system's databases prior to attempting an OS update.
- Do not remove packages from a DB system. However, you might have to remove custom RPMs (packages that were installed after the system was provisioned) for the update to complete successfully.



### Warning

Do not install NetworkManager on the DB system. Installing this package and rebooting the system results in severe loss of access to the system.

- Oracle recommends that you test any updates thoroughly before updating a production system.
- The image used to launch a DB system is updated regularly with the necessary patches. After you launch a DB system, you are responsible for applying the required OS security updates published through the Oracle public YUM server.
- To apply OS updates, the DB system's VCN must be configured to allow access to the YUM repository. For more information, see [Network Setup for DB Systems](#).

### To update an OL7 OS on a DB system host

You can update the OS on 2-node RAC virtual machine DB systems in a rolling fashion.



### Note

Ensure the Oracle Clusterware (CRS) is completely shut down before performing the OS kernel updates.

1. Log on to the DB system host as `opc`, and then `sudo` to the `root` user.

```
login as: opc
[opc@dbsys ~]$ sudo su -
```

2. If your DB system uses an image with the kernel version `4.1.12-124.27.1.el7uek` (used with older images), then change the `bootefi` label before updating the OS.

### To check the kernel version

Run the following command.

```
$ uname -r
```

Example response indicating kernel version `4.1.12-124.27.1.el7uek`:

```
4.1.12-124.27.1.el7uek.x86_64
```

If you have kernel version `4.1.12-124.27.1.el7uek`, then proceed to change the `bootefi` label.

### To change the `bootefi` label (each node)

- a. Edit `/etc/fstab`: Change the label `bootefi` to `BOOTEFI` (uppercase).

Example:

```
LABEL=BOOTEFI /boot/efi vfat defaults 1 2
```

- b. Restart the DB node.
- c. Run the following command to ensure that the required link is created.

```
$ sudo ls -lrt /etc/grub2-efi.cfg
```

Example response indicating that the required link exists:

```
lrwxrwxrwx 1 root root 31 Sep 4 11:49 /etc/grub2-efi.cfg ->
../boot/efi/EFI/redhat/grub.cfg
```

3. Identify the host region by running the following command:

```
curl -s http://169.254.169.254/opc/v1/instance/ |grep region
```

4. With the region you noted from the previous step, determine the region name, and perform the following two steps.

See [Regions and Availability Domains](#) to look up the region name.

- a. Download the repo.

```
wget https://swiftobjectstorage.<region_
name>.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol7repo -O /tmp/oci_
dbaas_ol7repo
```

This example output assumes the region is us-phoenix-1 (PHX).

```
wget https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol7repo -O /tmp/oci_dbaas_
ol7repo
--2019-07-16 10:40:42-- https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol7repo
Resolving swiftobjectstorage.us-phoenix-1.oraclecloud.com... 129.146.13.177,
129.146.13.180, 129.146.12.235, ...
Connecting to swiftobjectstorage.us-phoenix-1.oraclecloud.com|129.146.13.177|:443...
connected.
HTTP request sent, awaiting response... 200 OK
Length: 1394 (1.4K) [binary/octet-stream]
Saving to: `/tmp/oci_dbaas_ol7repo'

100%
[=====
=====
=====] 1,394 --.-K/s in 0s

2019-07-16 10:40:42 (34.5 MB/s) - `/tmp/oci_dbaas_ol7repo' saved [1394/1394]
```

- b. Download the version lock files.

```
wget https://swiftobjectstorage.<region_
name>.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol7.list -O
/tmp/versionlock.list
```

This example output assumes the region is us-phoenix-1 (PHX).

```
wget https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol7.list -O
```

```

/tmp/versionlock.list
--2019-07-16 10:41:38-- https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol7.list
Resolving swiftobjectstorage.us-phoenix-1.oraclecloud.com... 129.146.12.224,
129.146.12.164, 129.146.14.172, ...
Connecting to swiftobjectstorage.us-phoenix-1.oraclecloud.com|129.146.12.224|:443...
connected.
HTTP request sent, awaiting response... 200 OK
Length: 15769 (15K) [binary/octet-stream]
Saving to: `/tmp/versionlock.list'

100%
[=====
=====
=====>] 15,769 --.-K/s in 0.1s

2019-07-16 10:41:39 (123 KB/s) - `/tmp/versionlock.list' saved [15769/15769]

```

5. Copy the repo file to the `/etc/yum.repos.d` directory.

```
cp /tmp/oci_dbaas_ol7repo /etc/yum.repos.d/ol7.repo
```

6. Copy and overwrite the existing version lock file.

```
cp /etc/yum/pluginconf.d/versionlock.list /etc/yum/pluginconf.d/versionlock.list-`date +%Y%m%d`
cp /tmp/versionlock.list /etc/yum/pluginconf.d/versionlock.list
```

The initial version lock file should be empty. However, it is a good practice to back it up in case it is not and you need to refer to it later.

7. Run the update command.

```

| 18 MB 00:00
yum update
Loaded plugins: kernel-update-handler, ulninfo, versionlock
Excluding 250 updates due to versionlock (use "yum versionlock status" to show them)
Resolving Dependencies
--> Running transaction check
--> Package kernel-uek.x86_64 0:4.1.12-124.28.5.el7uek will be installed
--> Package kernel-uek-firmware.noarch 0:4.1.12-124.28.5.el7uek will be installed
--> Package libtalloc.x86_64 0:2.1.10-1.el7 will be updated
--> Package libtalloc.x86_64 0:2.1.13-1.el7 will be an update

```

## CHAPTER 11 Database

```
---> Package pytalloc.x86_64 0:2.1.10-1.e17 will be updated
---> Package pytalloc.x86_64 0:2.1.13-1.e17 will be an update
--> Finished Dependency Resolution

Dependencies Resolved

=====
=====
Package Arch Version
Repository Size
=====
=====
Installing:
kernel-uek x86_64 4.1.12-124.28.5.e17uek
 ol7_UEKR4 44 M
kernel-uek-firmware noarch 4.1.12-124.28.5.e17uek
 ol7_UEKR4 1.0 M
Updating:
libtalloc x86_64 2.1.13-1.e17
 ol7_latest 31 k
pytalloc x86_64 2.1.13-1.e17
 ol7_latest 16 k

Transaction Summary

=====
=====
Install 2 Packages
Upgrade 2 Packages

Total download size: 46 M
Is this ok [y/d/N]: y
Downloading packages:
No Presto metadata available for ol7_UEKR4
No Presto metadata available for ol7_latest
(1/4): kernel-uek-firmware-4.1.12-124.28.5.e17uek.noarch.rpm
 | 1.0 MB 00:00:00
(2/4): libtalloc-2.1.13-1.e17.x86_64.rpm
```

## CHAPTER 11 Database

```
| 31 kB 00:00:00
(3/4): pytalloc-2.1.13-1.el7.x86_64.rpm
| 16 kB 00:00:00
(4/4): kernel-uek-4.1.12-124.28.5.el7uek.x86_64.rpm
44 MB 00:00:01

Total
41 MB/s | 46 MB 00:00:01
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
Warning: RPMDB altered outside of yum.
** Found 7 pre-existing rpmdb problem(s), 'yum check' output follows:
oda-hw-mgmt-19.3.0.0.0_LINUX.X64_190530-1.x86_64 has missing requires of libnfsodm19.so() (64bit)
oda-hw-mgmt-19.3.0.0.0_LINUX.X64_190530-1.x86_64 has missing requires of perl(GridDefParams)
oda-hw-mgmt-19.3.0.0.0_LINUX.X64_190530-1.x86_64 has missing requires of perl(Sys::Syslog)
oda-hw-mgmt-19.3.0.0.0_LINUX.X64_190530-1.x86_64 has missing requires of perl(s_GridSteps)
perl-RPC-XML-0.78-3.el7.noarch has missing requires of perl(DateTime) >= ('0', '0.70', None)
perl-RPC-XML-0.78-3.el7.noarch has missing requires of perl(DateTime::Format::ISO8601) >= ('0',
'0.07', None)
perl-RPC-XML-0.78-3.el7.noarch has missing requires of perl(Module::Load) >= ('0', '0.24', None)
Installing : kernel-uek-firmware-4.1.12-124.28.5.el7uek.noarch
1/6
Updating : libtalloc-2.1.13-1.el7.x86_64
2/6
Updating : pytalloc-2.1.13-1.el7.x86_64
3/6
Installing : kernel-uek-4.1.12-124.28.5.el7uek.x86_64
4/6
Cleanup : pytalloc-2.1.10-1.el7.x86_64
5/6
Cleanup : libtalloc-2.1.10-1.el7.x86_64
6/6
```



### Note

- Ignore the error activating message that results from running the update.
- An update will occur only if a versionlock file has a valid update available to apply to the DB system.

8. Restart the system.

```
$ sudo su -
reboot
```

9. Run the following command to validate the update:

```
uname -r
4.1.12-124.28.5
```

In this example, then new kernel version is 4.1.12-124.28.5.

## To update an OL6 OS on a DB system host

You can update the OS on 2-node RAC virtual machine DB systems in a rolling fashion.



### Note

Ensure the Oracle Clusterware (CRS) is completely shut down before performing the OS kernel updates.

1. Log on to the DB system host as `opc`, and then `sudo` to the `root` user.

```
login as: opc
[opc@dbsys ~]$ sudo su -
```

2. Identify the host region by running the following command:

```
curl -s http://169.254.169.254/opc/v1/instance/ |grep region
```

3. With the region you noted from the previous step, determine the region name, and perform the following two steps.

See [Regions and Availability Domains](#) to look up the region name.

- a. Download the repo.

```
wget https://swiftobjectstorage.<region_
name>.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol6repo -O /tmp/oci_
dbaas_ol6repo
```

This example output assumes the region is us-phoenix-1 (PHX).

```
wget https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol6repo -O /tmp/oci_dbaas_
ol6repo
--2018-03-16 10:40:42-- https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/oci_dbaas_ol6repo
Resolving swiftobjectstorage.us-phoenix-1.oraclecloud.com... 129.146.13.177,
129.146.13.180, 129.146.12.235, ...
Connecting to swiftobjectstorage.us-phoenix-1.oraclecloud.com|129.146.13.177|:443...
connected.
HTTP request sent, awaiting response... 200 OK
Length: 1394 (1.4K) [binary/octet-stream]
Saving to: `/tmp/oci_dbaas_ol6repo'

100%
[=====
=====
=====>] 1,394 --.-K/s in 0s

2018-03-16 10:40:42 (34.5 MB/s) - `/tmp/oci_dbaas_ol6repo' saved [1394/1394]
```

- b. Download the version lock files.

```
wget https://swiftobjectstorage.<region_
name>.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol6.list -O
/tmp/versionlock.list
```

This example output assumes the region is us-phoenix-1 (PHX).

```
wget https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol6.list -O
/tmp/versionlock.list
--2018-03-16 10:41:38-- https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaaspatchstore/DBaaSOSPatches/versionlock_ol6.list
Resolving swiftobjectstorage.us-phoenix-1.oraclecloud.com... 129.146.12.224,
129.146.12.164, 129.146.14.172, ...
Connecting to swiftobjectstorage.us-phoenix-1.oraclecloud.com|129.146.12.224|:443...
connected.
HTTP request sent, awaiting response... 200 OK
Length: 15769 (15K) [binary/octet-stream]
Saving to: `/tmp/versionlock.list'

100%
[=====
=====
=====>] 15,769 --.-K/s in 0.1s

2018-03-16 10:41:39 (123 KB/s) - `/tmp/versionlock.list' saved [15769/15769]
```

#### 4. Enable the repo for your region.

- a. Copy the repo file to the `/etc/yum.repos.d` directory.

```
cp /tmp/oci_dbaas_ol6repo /etc/yum.repos.d/ol6.repo
```

- b. Modify the `ol6.repo` file to enable the repo for your region.

```
vi /etc/yum.repos.d/ol6.repo

[ol6_latest_PHX]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://yum-phx.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1 <= Enabled.

[ol6_UEKR4_PHX]
name=Latest Unbreakable Enterprise Kernel Release 4 for Oracle Linux $releasever
($basearch)
baseurl=http://yum-phx.oracle.com/repo/OracleLinux/OL6/UEKR4/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
```

```
gpgcheck=1
enabled=1 <= Enabled.
```

### 5. Install yum-plugin-versionlock.

```
$ sudo su -
yum repolist
Loaded plugins: kernel-update-handler, security, ulninfo
ol6_UEKR4
| 1.2 kB 00:00
ol6_UEKR4/primary
| 29 MB 00:00
ol6_UEKR4
588/588
ol6_latest
| 1.4 kB 00:00
ol6_latest/primary
| 67 MB 00:00
ol6_latest
39825/39825
repo id repo name
status
ol6_UEKR4 Latest Unbreakable Enterprise Kernel Release 4 for Oracle
Linux 6Server (x86_64) 588
ol6_latest Oracle Linux 6Server Latest (x86_64)
39825
repolist: 40413
[root@jigsosupg ~]# yum install yum-plugin-versionlock
Loaded plugins: kernel-update-handler, security, ulninfo
Setting up Install Process
Resolving Dependencies
--> Running transaction check
--> Package yum-plugin-versionlock.noarch 0:1.1.30-40.0.1.el6 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
=====
Package Arch Version
```

## CHAPTER 11 Database

```
Repository Size
=====
Installing:
yum-plugin-versionlock noarch 1.1.30-40.0.1.el6
 ol6_latest 32 k

Transaction Summary
=====
Install 1 Package(s)

Total download size: 32 k
Installed size: 43 k
Is this ok [y/N]: y
Downloading Packages:
yum-plugin-versionlock-1.1.30-40.0.1.el6.noarch.rpm
 | 32 kB 00:00
warning: rpmts_HdrFromFdno: Header V3 RSA/SHA256 Signature, key ID ec551f03: NOKEY
Retrieving key from file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
Importing GPG key 0xEC551F03:
 Userid : Oracle OSS group (Open Source Software group) <build@oss.oracle.com>
 Package: 6:oraclelinux-release-6Server-8.0.3.x86_64 (@odadom1)
 From : /etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
Is this ok [y/N]: y
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
Warning: RPMDB altered outside of yum.
** Found 4 pre-existing rpmdb problem(s), 'yum check' output follows:
oda-hw-mgmt-12.2.0.1.0_LINUX.X64_170614.TR1221-1.x86_64 has missing requires of
/usr/local/bin/perl
oda-hw-mgmt-12.2.0.1.0_LINUX.X64_170614.TR1221-1.x86_64 has missing requires of libnfsodm12.so()
(64bit)
oda-hw-mgmt-12.2.0.1.0_LINUX.X64_170614.TR1221-1.x86_64 has missing requires of perl
(GridDefParams)
oda-hw-mgmt-12.2.0.1.0_LINUX.X64_170614.TR1221-1.x86_64 has missing requires of perl(s_GridSteps)
```

```
Installing : yum-plugin-versionlock-1.1.30-40.0.1.el6.noarch
 1/1
Verifying : yum-plugin-versionlock-1.1.30-40.0.1.el6.noarch
 1/1

Installed:
 yum-plugin-versionlock.noarch 0:1.1.30-40.0.1.el6

Complete!
```



### Note

Ignore the RPMDB warning messages that refer to oda-hw-mgmt.

#### 6. Copy and overwrite the existing version lock file.

```
cp /etc/yum/pluginconf.d/versionlock.list /etc/yum/pluginconf.d/versionlock.list-`date +%Y%m%d`
cp /tmp/versionlock.list /etc/yum/pluginconf.d/versionlock.list
```

The initial version lock file should be empty. However, it is a good practice to back it up in case it is not and you need to refer to it later.

#### 7. Run the update command.

```
yum update
Loaded plugins: kernel-update-handler, security, ulninfo, versionlock
Setting up Update Process
Resolving Dependencies
--> Running transaction check
--> Package kernel-uek.x86_64 0:4.1.12-112.14.13.el6uek will be installed
--> Package kernel-uek-firmware.noarch 0:4.1.12-112.14.13.el6uek will be installed
--> Package linux-firmware.noarch 0:20160616-44.git43e96ale.0.12.el6 will be updated
--> Package linux-firmware.noarch 0:20171128-56.git17e62881.0.2.el6 will be an update
--> Finished Dependency Resolution

Dependencies Resolved
```

## CHAPTER 11 Database

```
=====
=====
Package Arch Version
Repository Size
=====
Installing:
kernel-uek x86_64 4.1.12-112.14.13.el6uek
 ol6_UEKR4 51 M
kernel-uek-firmware noarch 4.1.12-112.14.13.el6uek
 ol6_UEKR4 2.4 M
Updating:
linux-firmware noarch 20171128-56.git17e62881.0.2.el6
 ol6_UEKR4 74 M

Transaction Summary

=====
Install 2 Package(s)
Upgrade 1 Package(s)

Total download size: 128 M
Is this ok [y/N]:y
Downloading Packages:
(1/3): kernel-uek-4.1.12-112.14.13.el6uek.x86_64.rpm
 | 51 MB 00:00
(2/3): kernel-uek-firmware-4.1.12-112.14.13.el6uek.noarch.rpm
 | 2.4 MB 00:00
(3/3): linux-firmware-20171128-56.git17e62881.0.2.el6.noarch.rpm
 | 74 MB 00:00

Total
 214 MB/s | 128 MB 00:00
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
```

## CHAPTER 11 Database

```
Installing : kernel-uek-firmware-4.1.12-112.14.13.el6uek.noarch
 1/4
Updating : linux-firmware-20171128-56.git17e62881.0.2.el6.noarch
 2/4
Installing : kernel-uek-4.1.12-112.14.13.el6uek.x86_64
 3/4
Cleanup : linux-firmware-20160616-44.git43e96a1e.0.12.el6.noarch
 4/4

ol6_UEKR4/filelists
| 18 MB 00:00
Uploading /boot/vmlinuz-4.1.12-112.14.13.el6uek.x86_64 to http://169.254.0.3/kernel
Uploading /boot/initramfs-4.1.12-112.14.13.el6uek.x86_64.img to http://169.254.0.3/initrd
Uploading /tmp/tmp5HjrRUcmdline to http://169.254.0.3/cmdline

Error activating kernel/initrd/cmdline: 502 - <html>
<head><title>502 Bad Gateway</title></head>
<body bgcolor="white">
<center><h1>502 Bad Gateway</h1></center>
</body>
</html>
```



### Note

- Ignore the error activating message that results from running the update.
- An update will occur only if a versionlock file has a valid update available to apply to the DB system.

### 8. Restart the system.

```
$ sudo su -
reboot
```

### 9. Run the following command to validate the update:

```
uname -r
4.1.12-112.14.13
```

In this example, then new kernel version is 4.1.12-112.14.13.

For information about applying Oracle database patches to a DB system, see [Patching a DB System](#).

### Configuring a DB System

This topic provides information to help you configure your DB system.

#### Network Time Protocol

Oracle recommends that you run a Network Time Protocol (NTP) daemon on your 1-node DB systems to keep system clocks stable during rebooting. If you need information about an NTP daemon, see [Setting Up "NTP \(Network Time Protocol\) Server" in RHEL/CentOS 7](#).

Oracle recommends that you configure NTP on both nodes in a 2-node RAC DB system to synchronize time across the nodes. If you do not configure NTP, then Oracle Clusterware configures and uses the Cluster Time Synchronization Service (CTSS), and the cluster time might be out-of-sync with applications that use NTP for time synchronization.

For information about configuring NTP on a version 12c database, see [Setting Network Time Protocol for Cluster Time Synchronization](#). For a version 11g database, see [Network Time Protocol Setting](#).

#### Transparent Data Encryption

All user-created tablespaces in a DB system database are encrypted by default, using Transparent Data Encryption (TDE).

- For version 12c databases, if you don't want your tablespaces encrypted, you can set the [ENCRYPT\\_NEW\\_TABLESPACES](#) database initialization parameter to DDL.
- On a 1- or 2-node RAC DB system, you can use the [dbcli update-tdekey](#) command to update the master encryption key for a database.

- You must create and activate a master encryption key for any PDBs that you create. After creating or plugging in a new PDB on a 1- or 2-node RAC DB System, use the `dbcli update-tdekey` command to create and activate a master encryption key for the PDB. Otherwise, you might encounter the error `ORA-28374: typed master key not found in wallet` when attempting to create tablespaces in the PDB. In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers. For more information, see "Overview of Managing a Multitenant Environment" in the [Oracle Database Administrator's Guide](#).
- For information about encryption on Exadata DB systems, see [Using Tablespace Encryption in Exadata Cloud Service](#).

For detailed information about database encryption, see the [Oracle Database Security White Papers](#).

### Patching a DB System



#### Note

This topic is not applicable to Exadata DB systems.

This topic explains how to perform patching operations on bare metal and virtual machine DB systems and database homes by using the Console, API, or the database CLI (DBCLI).

Currently, the following patches are available:

Version	DB System Patch	Database Patch
19.0.0.0	October 2019	October 2019, July 2019, April 2019, January 2019
18.0.0.0	October 2019	October 2019, July 2019, April 2019, January 2019
12.2.0.1	October 2019	October 2019, July 2019, April 2019, January 2019

## CHAPTER 11 Database

---

Version	DB System Patch	Database Patch
12.1.0.2	October 2019	October 2019, July 2019, April 2019, January 2019
11.2.0.4	Not applicable	October 2019, July 2019, April 2019, January 2019

For information about operating system updates, see [OS Updates](#).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

### About Patching DB Systems

Because patching a system requires a reboot, plan to run the operations at a time when they will have minimal impact on users. To avoid system interruption, consider implementing a high availability strategy such as Oracle Data Guard. For more information, see [Using Oracle Data Guard with the Database CLI](#).

Oracle recommends that you back up your database and test the patch on a test system before you apply the patch. For information about backing up the databases, see [Backing Up a Database](#).

You must patch a DB system before you patch the databases within that system.

### Prerequisites

The DB system requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see these topics:

- [Network Setup for DB Systems](#): For information about setting up your VCN for the DB system, including the service gateway.
- <https://cloud.oracle.com/infrastructure/storage/object-storage/faq>: For information about the Swift endpoints to use.



#### Important

In addition to the prerequisites listed, ensure that the following conditions are met to avoid patching failures:

- The `/u01` directory on the database host file system has at least 15 GB of free space for the execution of patching processes.
- The Oracle Clusterware is up and running on the DB system.
- All nodes of the DB system are up and running.

See [Patching Failures on Bare Metal and Virtual Machine DB Systems](#) for details on problems that can result from not following these guidelines.

### Using the Console

You can use the Console to view the history of patch operations on a DB system or an individual database, apply patches, and monitor the status of an operation.

Oracle recommends that you use the pre-check action to ensure your DB system or database home has met the requirements for the patch you want to apply.

### PERFORMING PATCH OPERATIONS

#### To perform a patch operation on a DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system on which you want to perform a patch operation, and click its name to display details about it.
4. Under **Resources**, click **Patches**.
5. Review the list of patches.
6. Click the Actions icon (three dots) for the patch you are interested in, and then click one of the following actions:
  - **Pre-check:** Check for any prerequisites to make sure that the patch can be successfully applied.
  - **Apply:** Performs the pre-check, and then applies the patch.
7. Confirm when prompted.
8. In the list of patches, click the patch name to display its patch request and monitor the progress of the patch operation.  
While a patch is being applied, the patch's status displays as **Applying** and the DB system's status displays as **Updating**. If the operation completes successfully, the patch's status changes to **Applied** and the DB system's status changes to **Available**.

#### To perform a patch operation on a database

Before you perform this procedure, ensure that the latest patch was successfully applied to the DB system.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database on which you want to perform the patch operation, and click its name to display details about it.
5. Under **Resources**, click **Patches**.
6. Review the list of patches.
7. Click the Actions icon (three dots) for the patch you are interested in, and then click one of the following actions:
  - **Pre-check:** Check for any prerequisites to make sure that the patch can be successfully applied.
  - **Apply:** Performs the pre-check, and then applies the patch.
8. Confirm when prompted.
9. In the list of patches, click the patch name to display its patch request and monitor the progress of the patch operation.  
While a patch is being applied, the patch's status displays as **Applying** and the database's status displays as **Updating**. If the operation completes successfully, the patch's status changes to **Applied** and the database's status changes to **Available**.

### VIEWING PATCH HISTORY

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry.

Patch history views in the Console do not show patches that were applied by using command line tools like DBCLI or the Opatch utility.

### To view the patch history of a DB system

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system you are interested in, and click the system name to display details about it.
4. Under **Resources**, click **Patch History**.  
The history of patch operations for that DB system is displayed.

### To view the patch history of a database

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database you are interested in, and click its name to display details about it.
5. Under **Resources**, click **Patch History**.  
The history of patch operations for that database is displayed.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage patching DB systems and databases.

DB systems:

- [ListDbSystemPatches](#)
- [ListDbSystemPatchHistoryEntries](#)
- [GetDbSystemPatch](#)
- [GetDbSystemPatchHistoryEntry](#)
- [UpdateDbSystem](#)

Databases:

- [ListDbHomePatches](#)
- [ListDbHomePatchHistoryEntries](#)
- [GetDbHomePatch](#)
- [GetDbHomePatchHistoryEntry](#)
- [UpdateDbHome](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

### Using the Database CLI

This topic explains how to use the command line interface on the DB system to patch a DB system. Patches are available from the Oracle Cloud Infrastructure Object Storage service. You'll use the `dbcli` commands to download and apply patches to some or all of the components in your system.

#### PREREQUISITES

For connecting to the DB system via SSH, you'll need the path to private key associated with the public key used when the DB system was launched.

You also need the public or private IP address of the DB system. Use the private IP address to connect to the DB system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the DB System's public IP address to connect to the system from outside the cloud (with no VPN).

You can find the IP addresses in the Oracle Cloud Infrastructure Console on the **Database** page.

### To update the CLI with the latest commands

Update the CLI to ensure you have the latest patching commands (older DB systems might not include them).

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Update the CLI by using the [cliadm update-dbcli](#) command.

```
[root@dbsys ~]# cliadm update-dbcli
{
 "jobId" : "dc9ce73d-ed71-4473-99cd-9663b9d79bfd",
 "status" : "Created",
 "message" : "Dcs cli will be updated",
 "reports" : [],
 "createTimestamp" : "January 18, 2017 10:19:34 AM PST",
 "resourceList" : [],
 "description" : "dbcli patching",
 "updatedAt" : "January 18, 2017 10:19:34 AM PST"
}
```

4. Wait for the update job to complete successfully. Check the status of the job by using the [dbcli list-jobs](#) command.

```
[root@dbsys ~]# dbcli list-jobs
```

ID	Description	Created
Status		

---

```

dc9ce73d-ed71-4473-99cd-9663b9d79bfd dbcli patching January 18, 2017 10:19:34
AM PST Success
```

### To check for installed and available patches

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Display the installed patch versions by using the [dbcli describe-component](#) command. If the **Available Version** column indicates a version number for a component, you should update the component.

```
[root@dbsys ~]# dbcli describe-component
```

```
System Version
```

```

```

```
12.1.2.10.0
```

Component Name	Installed Version	Available Version
OAK	12.1.2.10.0	up-to-date
GI	12.1.0.2.161018	up-to-date
ORADB12102_HOME1	12.1.0.2.160719	12.1.0.2.161018

4. Display the latest patch versions available in Object Storage by using the [dbcli describe-latestpatch](#) command.

```
[root@dbsys ~]# dbcli describe-latestpatch
```

```
componentType availableVersion
```

```

```

```
gi 12.1.0.2.161018
```

## CHAPTER 11 Database

---

```
db 11.2.0.4.161018
db 12.1.0.2.161018
oak 12.1.2.10.0
```

### To patch server components

You can patch the Grid Infrastructure (GI) and storage management kit (OAK) server components.

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Update the server components by using the [dbcli update-server](#) command.

```
[root@dbsys ~]# dbcli update-server
{
 "jobId" : "9a02d111-e902-4e94-bc6b-9b820ddf6ed8",
 "status" : "Created",
 "reports" : [],
 "createTimestamp" : "January 19, 2017 09:37:11 AM PST",
 "resourceList" : [],
 "description" : "Server Patching",
 "updatedAtTime" : "January 19, 2017 09:37:11 AM PST"
}
```

Note the job ID above.

4. Check the job output by using the [dbcli describe-job](#) command with the job ID.

```
[root@dbsys ~]# dbcli describe-job -i 9a02d111-e902-4e94-bc6b-9b820ddf6ed8
```

```
Job details

```

## CHAPTER 11 Database

```

ID: 9a02d111-e902-4e94-bc6b-9b820ddf6ed8
Description: Server Patching
Status: Running
Created: January 19, 2017 9:37:11 AM PST
Message:

```

Task Name	Status	Start Time	End Time
Create Patching Repository Directories	Success	January 19, 2017 9:37:11 AM PST	January 19, 2017 9:37:11 AM PST
Download latest patch metadata	Success	January 19, 2017 9:37:11 AM PST	January 19, 2017 9:37:11 AM PST
Update System version	Success	January 19, 2017 9:37:11 AM PST	January 19, 2017 9:37:11 AM PST
Update Patching Repository	Success	January 19, 2017 9:37:11 AM PST	January 19, 2017 9:38:35 AM PST
oda-hw-mgmt upgrade	Success	January 19, 2017 9:38:35 AM PST	January 19, 2017 9:38:58 AM PST
Opatch updation	Success	January 19, 2017 9:38:58 AM PST	January 19, 2017 9:38:58 AM PST
Patch conflict check	Success	January 19, 2017 9:38:58 AM PST	January 19, 2017 9:42:06 AM PST
Apply clusterware patch	Success	January 19, 2017 9:42:06 AM PST	January 19, 2017 10:02:32 AM PST
Updating GiHome version	Success	January 19, 2017 10:02:32 AM PST	January 19, 2017 10:02:38 AM PST

5. Verify that the server components were updated successfully by using the [dbcli describe-component](#) command. The **Available Version** column should indicate update-to-date.

### To patch database home components

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the

root user's profile, which will set the PATH to the dbcli directory (/opt/oracle/dcs/bin).

```
login as: opc
[opc@dbsys ~]$ sudo su -
```

3. Get the ID of the database home by using the [dbcli list-dbhomes](#) command.

```
[root@dbsys ~]# dbcli list-dbhomes
ID Name DB Version Home Location

b727bf80-c99e-4846-ac1f-28a81a725df6 OraDB12102_home1 12.1.0.2
/u01/app/orauser/product/12.1.0.2/dbhome_1
```

4. Update the database home components by using the [dbcli update-dbhome](#) command and providing the ID from the previous step.

```
[root@dbsys ~]# dbcli update-dbhome -i b727bf80-c99e-4846-ac1f-28a81a725df6
{
 "jobId" : "31b38f67-f993-4f2e-b7eb-5bccda9901ae",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "January 20, 2017 10:08:48 AM PST",
 "resourceList" : [],
 "description" : "DB Home Patching: Home Id is 52e2e799-946a-4339-964b-c203dee35328",
 "updatedAtTime" : "January 20, 2017 10:08:48 AM PST"
}
```

Note the job ID above.

5. Check the job output by using the [dbcli describe-job](#) command with the job ID.

```
[root@dbsys ~]# dbcli describe-job -i 31b38f67-f993-4f2e-b7eb-5bccda9901ae

Job details

ID: 31b38f67-f993-4f2e-b7eb-5bccda9901ae
Description: DB Home Patching: Home Id is b727bf80-c99e-4846-ac1f-28a81a725df6
Status: Success
Created: January 20, 2017 10:08:48 AM PST
```

## CHAPTER 11 Database

Message:

Task Name	Status	Start Time	End Time
Create Patching Repository Directories	Success	January 20, 2017 10:08:49 AM PST	January 20, 2017 10:08:49 AM PST
Download latest patch metadata	Success	January 20, 2017 10:08:49 AM PST	January 20, 2017 10:08:49 AM PST
Update System version	Success	January 20, 2017 10:08:49 AM PST	January 20, 2017 10:08:49 AM PST
Update Patching Repository	Success	January 20, 2017 10:08:49 AM PST	January 20, 2017 10:08:58 AM PST
Opatch updation	Success	January 20, 2017 10:08:58 AM PST	January 20, 2017 10:08:58 AM PST
Patch conflict check	Success	January 20, 2017 10:08:58 AM PST	January 20, 2017 10:12:00 AM PST
db upgrade	Success	January 20, 2017 10:12:00 AM PST	January 20, 2017 10:22:17 AM PST

6. Verify that the database home components were updated successfully by using the [dbcli describe-component](#) command. The **Available Version** column should indicate update-to-date.

### Applying Interim Patches



#### Note

This topic applies only to database homes in 1-node and 2-node RAC DB systems.

If you are required to apply an interim patch (previously known as a "one-off" patch) to fix a specific defect, follow the procedure in this section. You use the Opatch utility to apply an interim patch to a database home.

In the procedure example, the database home directory is **/u02/app/oracle/product/12.1.0.2/dbhome\_1** and the patch number is **26543344**.

### To apply an interim patch to a database home

1. Obtain the applicable interim patch from My Oracle Support.
2. Review the information in the patch README.txt file. This file might contain additional and/or custom instructions to follow to apply the patch successfully.
3. Use SCP or SFTP to place the patch on your target database.
4. Shut down each database that is running in the database home.

```
srvctl stop database -db <db_name> -stopoption immediate -verbose
```

5. Set the Oracle home environment variable to point to the target Oracle home.

```
sudo su - oracle
export ORACLE_HOME=/u02/app/oracle/product/12.1.0.2/dbhome_1
```

6. Change to the directory where you placed the patch, and unzip the patch.

```
cd <work_dir_where_opatch_is_stored>
unzip p26543344_122010_Linux-x86-64.zip
```

7. Change to the directory with the unzipped patch, and check for conflicts.

```
cd 26543344
$ORACLE_HOME/OPatch/patch prereq CheckConflictAgainstOHWithDetail -ph ./
```

8. Apply the patch.

```
$ORACLE_HOME/OPatch/patch apply
```

9. Verify the patch was applied successfully.

```
$ORACLE_HOME/OPatch/patch lsinventory -detail -oh $ORACLE_HOME
```

10. If the database home contains databases, restart them.

```
$ORACLE_HOME/bin/srvctl start database -db <db_name>
```

Otherwise, run the following command as root user.

```
/u01/app/<db_version>/grid/bin/setasmgidwrap o=/u01/app/oracle/product/<db_version>/dbhome_1/bin/oracle
```

11. If the readme indicates that the patch has a sqlpatch component, run the datapatch command against each database.

Before you run datapatch, ensure that all pluggable databases (PDBs) are open. To open a PDB, you can use SQL\*Plus to execute `ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ WRITE;` against the PDB.

```
$ORACLE_HOME/OPatch/datapatch
```

### Creating a Database



#### Note

This topic is not applicable to virtual machine DB systems.

When you launch a bare metal DB system, an initial database is created in that system. You can create additional databases in that DB system at any time by using the Console or the API. You can create an empty database or reproduce a database by using a backup. Note that if you are creating a database from an automatic backup, you can choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup. For more information on automatic backups, see [Automatic Incremental Backups](#).

The database edition will be the edition of the DB system in which the database is created, and each new database is created in a separate database home.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let database admins manage DB systems](#) lets the specified group do everything with databases and related Database resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Database Service](#).

### Using the Console

To create a new database in an existing DB system



#### Note

The database that you create will be the same edition as the initial database in the DB system.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the DB system in which you want to create the database, and then click its name to display details about it.
4. Click **Create Database**.

5. In the **Create Database** dialog, enter the following:
- **Database name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
  - **Database version:** The version of the database. You can mix database versions on the DB system, but not editions.
  - **PDB name** (Optional) For version 12.1.0.2 and later, you can specify the name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of 8 alphanumeric characters. The only special character permitted is the underscore ( \_).
  - **Create administrator credentials:** A database administrator `SYS` user will be created with the password you supply.
    - **Username:** `SYS`
    - **Password:** Supply the password for this user. The password must meet the following criteria:

A strong password for `SYS`, `SYSTEM`, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (`SYS`, `SYSTEM`, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
    - **Confirm password:** Re-enter the `SYS` password you specified.
  - **Select workload type:** Choose the workload type that best suits your application:
    - **Online Transactional Processing (OLTP)** configures the database for a transactional workload, with a bias towards high volumes of random data access.

- **Decision Support System (DSS)** configures the database for a decision support or data warehouse workload, with a bias towards large data scanning operations.
  - **Configure database backups:** Specify the settings for backing up the database to Object Storage:
    - **Enable automatic backup:** Check the check box to enable automatic incremental backups for this database.
    - **Backup Retention Period:** If you enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The default selection is 30 days.
    - **Backup Scheduling:** If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, the six-hour default window of 00:00 to 06:00 (in the time zone of the DB system's region) is used for your database. See [Backup Scheduling](#) for more information.
6. Click **Show Advanced Options** to specify advanced options for the database:
- **Character set:** The character set for the database. The default is AL32UTF8.
  - **National character set:** The national character set for the database. The default is AL16UTF16.
  - Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Database**.

When the database creation is complete, the status changes from Provisioning to Available.

### To create a database from a backup in an existing DB system

Before you begin, note the following:

- When you create a database from a backup, you can choose a different DB system and compartment. However, the availability domain will be the same as where the source database is hosted.



#### Tip

You can use the [GetBackup](#) API to obtain information about the availability domain of the backup.

- The DB system you specify must support the same type as the system from which the backup was taken. For example, if the backup is from a single-node database, then the target DB system must be a single-node shape.
  - The version of the target DB system must be the same or higher than the version of the backup.
1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
  2. Choose your **Compartment**.  
A list of DB systems is displayed.
  3. Navigate to the backup you wish to use to create a new database. You can select a backup from a database details page, or select a backup that appears in your compartment's list of standalone backups.



### Tip

If you are creating a database from an automatic backup, you may choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup. For more information on automatic backups, see [Automatic Incremental Backups](#).

### To navigate to a database's list of backups

- a. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
- b. Find the database associated with the backup you wish to use, and click its name to display details about it.  
A list of backups is displayed in the default view of the database details. You can also access the list of backups for a database by clicking on **Backups** under **Resources**.

### To navigate to the list of standalone backups for your current compartment

- a. Click **Standalone Backups** under **Bare Metal, VM, and Exadata**.
  - b. In the list of standalone backups, find the backup you want to use to create the database.
4. Click the Actions icon (three dots) for the backup you are interested in, and then click **Create Database**.

5. In the **Create Database from Backup** dialog, enter the following:

- **DB System:** The DB system in which you want to create the database. You must have the **Use Existing DB System** radio button selected to see the drop-down list of DB system choices.



### Note

You cannot create a new database in the same DB system in which the database used to create the backup resides.

- **Database Name:** The name for the database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
- **Database Admin Password:** A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.
- **Confirm Database Admin Password:** Re-enter the Database Admin Password you specified.
- **Password for Transparent Data Encryption (TDE) Wallet or RMAN Encryption:**  
Enter either the TDE wallet password or the RMAN encryption password for the backup, whichever is applicable. The TDE wallet password is the SYS password provided when the database was created by using the Oracle Cloud Infrastructure Console, API, or CLI. The RMAN encryption password is typically required instead if the password was subsequently changed manually.

6. Click **Create Database**.

### To terminate a database

You'll get the chance to back up the database prior to terminating it. This creates a standalone backup that can be used to create a database later. Oracle recommends that you create this final backup for any production (non-test) database.



#### Note

Terminating a database removes all automatic incremental backups of the database from Oracle Cloud Infrastructure Object Storage. However, all full backups that were created on demand, including your final backup, will persist as standalone backups.

You cannot terminate a database that is assuming the primary role in a Data Guard association. To terminate it, you can switch it over to the standby role.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. In the list of DB systems, find the DB system that contains the database you want to terminate, and then click its name to display details about it.
4. In the list of databases, find the database you want to terminate, and then click its name to display details about it.
5. Click **Actions**, and then click **Terminate**.
6. In the confirmation dialog, indicate whether you want to back up the database before terminating it, and type the name of the database to confirm the termination.
7. Click **Terminate Database**.  
The database's status indicates Terminating.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage databases.

Database homes:

- [ListDbHomes](#)
- [GetDbHome](#)
- [CreateDbHome](#)
- [DeleteDbHome](#)

Databases:

- [ListDatabases](#)
- [GetDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

### Monitoring a Database

This topic explains how to set up an:

- Enterprise Manager Express console to monitor a version 12.1.0.2 or later database
- Enterprise Manager Database Control console to monitor a version 11.2.0.4 database

Each console is a web-based database management tool inside the Oracle database. You can use the console to perform basic administrative tasks such as managing user security, memory, and storage, and view performance information.

### Required IAM Policy

Some of the procedures below require permission to create or update security lists. For more information about security list policies, see [Security Lists](#).

### Monitoring a Database with Enterprise Manager Express

On 1- and 2-node RAC DB Systems, by default, the EM Express console is not enabled on version 18.1.0.0, 12.2.0.1, and 12.1.0.2 databases. You can enable it for an existing database as described below, or you can enable it when you create a database by using the [dbcli create-database](#) command with the `-co` parameter.

You must also update the security list and iptables for the DB system as described later in this topic.

When you enable the console, you'll set the port for the console. The procedure below uses port 5500, but each additional console enabled on the same DB system will have a different port.

#### To enable the EM Express console and determine its port number

1. SSH to the DB system, log in as `opc`, `sudo` to the oracle user, and log in to the database as `SYS`.

```
sudo su - oracle
. oraenv
<provide the database SID at the prompt>
sqlplus / as sysdba
```

2. Do one of the following:

- To enable the console and set its port, use the following command.

```
exec DBMS_XDB_CONFIG.SETHTTPSPORT (<port>);
```

For example:

```
SQL> exec DBMS_XDB_CONFIG.SETHTTPSPORT (5500);
PL/SQL procedure successfully completed.
```

- To determine the port for a previously enabled console, use the following command.

```
select dbms_xdb_config.getHttpsPort() from dual;
```

For example:

```
SQL> select dbms_xdb_config.getHttpsPort() from dual;

DBMS_XDB_CONFIG.GETHTTSPSPORT()

 5500
```

3. Return to the operating system by typing `exit` and then confirm that the listener is listening on the port:

```
lsnrctl status | grep HTTP

(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps) (HOST=xxx.us.oracle.com) (PORT=5500)) (Security=(my_wallet_directory=/u01/app/oracle/admin/prod/xdw_wallet)) (Presentation=HTTP) (Session=RAW))
```

4. If you're using a 2-node RAC DB system, see [To set the required permissions on a 2-node RAC DB system](#).
5. Open the console's port as described in [Opening Ports on the DB System](#).
6. Update the security list for the console's port as described in [Updating the Security List for the DB System](#).

## To set the required permissions on a 2-node RAC DB system

If you're using a 2-node RAC DB system, you'll need to add read permissions for the `asmadmin` group on the wallet directory on *both* nodes in the system.

1. SSH to one of the nodes in the DB system, log in as `opc`, `sudo` to the `grid` user.

```
[opc@dbsysHost1 ~]$ sudo su - grid
[grid@dbsysHost1 ~]$. oraenv
ORACLE_SID = [+ASM1] ?
The Oracle base has been set to /u01/app/grid
```

2. Get the location of the wallet directory, shown in red below in the command output.

```
[grid@dbsysHost1 ~]$ lsnrctl status | grep xdb_wallet

(DESCRIPTION=(ADDRESS=(PROTOCOL=tcps) (HOST=dbsysHost1.sub04061528182.dbsysapril6.oraclevcn.com)
```

```
(PORT=5500) (Security=(my_wallet_directory=/u01/app/oracle/admin/dbsys12_phx3wm/xdw_wallet))
(Presentation=HTTP) (Session=RAW)
```

3. Return to the opc user, switch to the oracle user, and change to the wallet directory.

```
[opc@dbsysHost1 ~]$ sudo su - oracle
[oracle@dbsysHost1 ~]$ cd /u01/app/oracle/admin/dbsys12_phx3wm/xdw_wallet
```

4. List the directory contents and note the permissions.

```
[oracle@dbsysHost1 xdw_wallet]$ ls -ltr
total 8
-rw----- 1 oracle asmadmin 3881 Apr 6 16:32 ewallet.p12
-rw----- 1 oracle asmadmin 3926 Apr 6 16:32 cwallet.sso
```

5. Change the permissions:

```
[oracle@dbsysHost1 xdw_wallet]$ chmod 640 /u01/app/oracle/admin/dbsys12_phx3wm/xdw_wallet/*
```

6. Verify that read permissions were added.

```
[oracle@dbsysHost1 xdw_wallet]$ ls -ltr
total 8
-rw-r----- 1 oracle asmadmin 3881 Apr 6 16:32 ewallet.p12
-rw-r----- 1 oracle asmadmin 3926 Apr 6 16:32 cwallet.sso
```

7. **Important!** Repeat the steps above on the other node in the cluster.

### To connect to the EM Express console

After you've enabled the console and opened its port in the security list and iptables, you can connect as follows:

1. From a web browser, connect to the console using the following URL format:

```
https://<ip_address>:<port>/em
```

For example, `https://129.145.0.164:5500/em`

Use the DB system's private or public IP address depending on your network configuration.

## CHAPTER 11 Database

Use the private IP address to connect to the DB system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN.

Use the DB System's public IP address to connect to the system from outside the cloud (with no VPN).

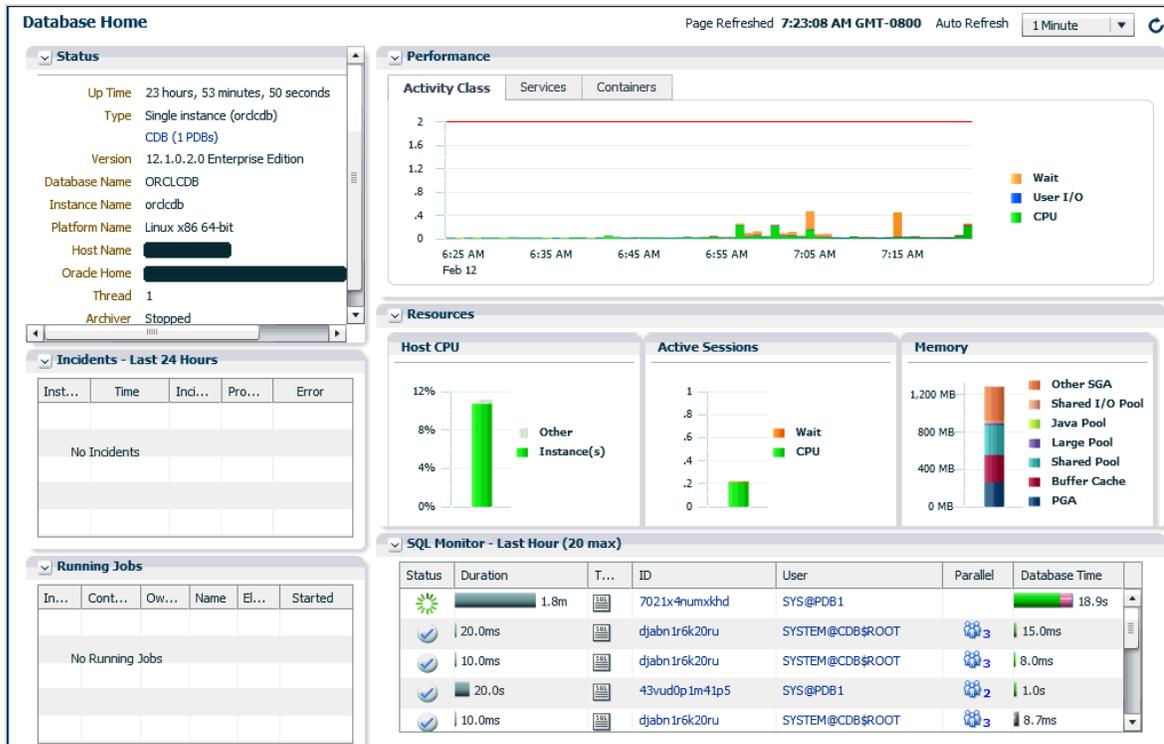
You can find the IP addresses in the Oracle Cloud Infrastructure Console on the **Database** page.

2. A login page is displayed and you can log in with any valid database credentials.



The Database Home page is displayed.

## CHAPTER 11 Database



To learn more about EM Express, see [Introduction to Oracle Enterprise Manager Database Express](#).



### Note

If you're using a 1-node DB system, and you are unable to connect to the EM Express console, see [Database Known Issues](#).

### Monitoring a Database with Enterprise Manager Database Control

By default, the Enterprise Manager Database Control console is not enabled on version 11.2.0.4 databases. You can enable the console:

- when you create a database by using the [dbcli create-database](#) with the `-co` parameter
- for an existing database as described [here](#).

Port 1158 is the default port used for the first console enabled on the DB system, but each additional console enabled on the DB system will have a different port.



#### Note

For a version 11.2.0.4 database on a 2-node RAC DB system, see [To enable the console for a version 11.2.0.4 database on a multi-node DB system](#).

To determine the port for the Enterprise Manager Database Control console

1. SSH to the DB system, log in as `opc`, and `sudo` to the `oracle` user.

```
sudo su - oracle
. oraenv
<provide the database SID at the prompt>
```

2. Use the following command to get the port number.

```
emctl status dbconsole
```

The port is in the URL, as shown in the following example:

```
[oracle@dbsys ~]$ emctl status dbconsole
Oracle Enterprise Manager 11g Database Control Release 11.2.0.4.0
Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.
https://dbprod:1158/em/console/aboutApplication
Oracle Enterprise Manager 11g is running.

```

```
Logs are generated in directory /u01/app/oracle/product/11.2.0.4/dbhome_2/dbprod_db11/sysman/log
```

3. Open the console's port as described in [Opening Ports on the DB System](#).
4. Update the security list for the console's port as described in [Updating the Security List for the DB System](#).

### To connect to the Enterprise Manager Database Control console

After you've enabled the console and opened its port in the security list and iptables, you can connect as follows:

1. From a web browser, connect to the console using the following URL format:

```
https://<ip_address>:<port>/em
```

For example, `https://129.145.0.164:1158/em`

Use the DB system's private or public IP address depending on your network configuration.

Use the private IP address to connect to the DB system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN.

Use the DB System's public IP address to connect to the system from outside the cloud (with no VPN).

You can find the IP addresses in the Oracle Cloud Infrastructure Console on the **Database** page.

2. A login page will be displayed and you can log in with any valid database credentials.

To learn more about Enterprise Manager Database Control, see [Introduction to Oracle Enterprise Manager Database Control](#).

### To enable the console for a version 11.2.0.4 database on a multi-node DB system

A few extra steps are required to enable the console for a version 11.2.0.4 database on a

multi-node DB system.

### Configure SSH Equivalency Between the Two Nodes

You'll create SSH keys on each node and copy the key to the other node, so that each node has the keys for both nodes. The following procedure uses the sample names `node1` and `node2`.

1. SSH to `node1`, log in as `opc`, and `sudo` to the oracle user.

```
sudo su - oracle
```

2. Create a directory called `.ssh`, set its permissions, create an RSA key, and add the public key to the `authorized_keys` file.

```
mkdir .ssh
chmod 755 .ssh
ssh-keygen -t rsa
cat id_rsa.pub > authorized_keys
```

3. Repeat the previous steps on the other node in the cluster.
4. On each node, add the `id_rsa.pub` key for the *other* node to the `authorized_keys` file. When you're done, you should see both keys in `authorized_keys` on each node.
5. On `node1`, create the `known_hosts` file by doing the following:
  - SSH to `node1` and reply yes to the authentication prompt.
  - SSH to `node2` and reply yes to the authentication prompt.
6. On `node2`, create the `known_hosts` file by doing the following:
  - SSH to `node2` and reply yes to the authentication prompt.
  - SSH to `node1` and reply yes to the authentication prompt.
7. On `node1`, verify that SSH equivalency is now configured by using the following Cluster Verification Utility (CVU) command.

```
cluvfy stage -pre crsinst -n all -verbose
```

### Configure the Console

1. On node1, create a file called emca.rsp with the following entries.

```
DB_UNIQUE_NAME=<pdb_unique_name>
SERVICE_NAME=<db_unique_name>.<db_domain>
PORT=<scan listener port>
LISTENER_OH=$GI_HOME
SYS_PWD=<admin password>
DBSNMP_PWD=<admin password>
SYSMAN_PWD=<admin password>
CLUSTER_NAME=<cluster name> <=== to get the cluster name, run: $GI_HOME/bin/cemutlo -n
ASM_OH=$GI_HOME
ASM_SID=+ASM1
ASM_PORT=<asm listener port>
ASM_USER_NAME=ASMSNMP
ASM_USER_PWD=<admin password>
```

2. On node1, run Enterprise Manager Configuration Assistant (EMCA) using the emca.rsp file as input.

```
$ORACLE_HOME/bin/emca -config dbcontrol db -repos create -cluster -silent -respFile <location of
response file above>
```

3. On node2, configure the console so the agent in node1 reports to the console in node1, and the agent in node2 reports to the console in node2.

```
$ORACLE_HOME/bin/emca -reconfig dbcontrol -silent -cluster -EM_NODE <node2 host> -EM_NODE_LIST
<node2 host> -DB_UNIQUE_NAME <db_unique_name>
-SERVICE_NAME <db_unique_name>.<db_domain>
```

4. On each node, verify that console is working properly.

```
$ export ORACLE_UNQNAME=<db_unique_name>

$ emctl status agent
Oracle Enterprise Manager 11g Database Control Release 11.2.0.4.0
Copyright (c) 1996, 2013 Oracle Corporation. All rights reserved.

Agent Version : 10.2.0.4.5
OMS Version : 10.2.0.4.5
Protocol Version : 10.2.0.4.5
Agent Home : /u01/app/oracle/product/11.2.0.4/dbhome_x/<host>_<db_unique_name>
Agent binaries : /u01/app/oracle/product/11.2.0.4/dbhome_x
```

## CHAPTER 11 Database

```
Agent Process ID : 26194
Parent Process ID : 25835
Agent URL : https://<node host>:1831/emd/main
Repository URL : https://<node host>:5501/em/upload/
Started at : 2017-03-15 20:20:34
Started by user : oracle
Last Reload : 2017-03-15 20:27:00
Last successful upload : 2017-03-15 21:06:36
Total Megabytes of XML files uploaded so far : 22.25
Number of XML files pending upload : 0 <=== should be zero
Size of XML files pending upload(MB) : 0.00
Available disk space on upload filesystem : 42.75%
Data channel upload directory : /u01/app/oracle/product/11.2.0.4/dbhome_x/<host>_
<db_unique_name>/sysman/recv
Last successful heartbeat to OMS : 2017-03-15 21:08:45

```

### Update iptables and Security List

1. On each node, edit iptables to open the console's port as described in [Opening Ports on the DB System](#).
2. Update the security list for the console's port as described in [Updating the Security List for the DB System](#).

### Opening Ports on the DB System

Open the following ports as needed on the DB system:

- 6200 - For Oracle Notification Service (ONS).
- 5500 - For EM Express. 5500 is the default port, but each additional EM Express console enabled on the DB system will have a different port. If you're not sure which port to open for a particular console, see [Monitoring a Database with Enterprise Manager Express](#).
- 1158 - For Enterprise Manager Database Control. 1158 is the default port, but each additional console enabled on the DB system will have a different port. If you're not sure which port to open for a particular console, see [Monitoring a Database with Enterprise Manager Database Control](#).

For important information about critical firewall rules, see [Essential Firewall Rules](#).

### To open ports on the DB system

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user.

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Save a copy of iptables as a backup.

```
[root@dbsys ~]# iptables-save > /tmp/iptables.orig
```

(If necessary, you can restore the original file by using the command `iptables-restore < /tmp/iptables.orig`.)

4. Dynamically add a rule to iptables to allow inbound traffic on the console port, as shown in the following sample. Change the port number and comment as needed.

```
[root@dbsys ~]# iptables -I INPUT 8 -p tcp -m state --state NEW -m tcp --dport 5500 -j ACCEPT -m comment --comment "Required for EM Express."
```

5. Make sure the rule was added.

```
[root@dbsys ~]# service iptables status
```

6. Save the updated file to `/etc/sysconfig/iptables`.

```
[root@dbsys ~]# /sbin/service iptables save
```

The change takes effect immediately and will remain in effect when the node is rebooted.

7. Update the DB system's security list as described in [Updating the Security List for the DB System](#).

### Updating the Security List for the DB System

Review the list of ports in [Opening Ports on the DB System](#) and for every port you open in iptables, update the security list used for the DB system, or create a new security list.

Note that port 1521 for the Oracle default listener is included in iptables, but should also be added to the security list.

#### To update an existing security list

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Locate the DB system in the list.
4. Note the DB system's **Subnet** name and click its **Virtual Cloud Network**.
5. Locate the subnet in the list, and then click its security list under **Security Lists**.
6. Click **Edit All Rules** and add an ingress rule with source type = CIDR, source CIDR= *<source CIDR>*, protocol=TCP, and port= *<port number or port range>*.  
The source CIDR should be the CIDR block that includes the ports you open for the client connection.

For detailed information about creating or updating a security list, see [Security Lists](#).

### Backing Up a Database

Backing up your DB System is a key aspect of any Oracle database environment. You can store backups in the cloud or in local storage. Each backup destination has advantages, disadvantages, and requirements that you should consider, as described below.

#### Object Storage (Recommended)

- Backups are stored in the Oracle Cloud Infrastructure Object Storage.
- Durability: High

- Availability: High
- Back Up and Recovery Rate: Medium
- Advantages: High durability, performance, and availability.

### Local Storage

- Backups are stored locally in the Fast Recovery Area of the DB System.
- Durability: Low
- Availability: Medium
- Back Up and Recovery Rate: High
- Advantages: Optimized back up and fast point-in-time recovery.
- Disadvantages: If the DB System becomes unavailable, the backup is also unavailable.

Currently, Oracle Cloud Infrastructure does not provide the ability to attach block storage volumes to a DB System, so you cannot back up to network attached volumes.

For 1- and 2-node RAC DB Systems, see:

- [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#)
- [Backing Up a Database to Local Storage Using the Database CLI](#)

### Backing Up a Database to Oracle Cloud Infrastructure Object Storage



#### Note

This topic is not applicable to Exadata DB systems. For Exadata DB systems, see [Managing Exadata Database Backups](#).

This topic explains how to work with backups managed by Oracle Cloud Infrastructure. You do this by using the Console or the API. (For unmanaged backups, you can use `RMAN` or `dbcli`,

and you must create and manage your own Object Storage buckets for backups. See [Backing Up a Database to Object Storage Using RMAN.](#))



### Warning

If you previously used `RMAN` or `dbcli` to configure backups and then you switch to using the Console or the API for backups, a new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to work.

### REQUIRED IAM POLICY

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### PREREQUISITES

The DB system requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see these topics:

- [Network Setup for DB Systems](#): For information about setting up your VCN for the DB system, including the service gateway.
- <https://cloud.oracle.com/infrastructure/storage/object-storage/faq>: For information about the Swift endpoints to use.



### Important

Note that your database and DB system must be in an “Available” state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching and Data Guard operations) while a backup operation is in progress. If an automatic backup operation fails, the Database service retries the operation during the next day’s backup window. If an on-demand full backup fails, you can try the operation again when the DB system and database availability are restored.

In addition to the prerequisites listed, ensure that the following conditions are met to avoid backup failures:

- The database's archiving mode is set to `ARCHIVELOG` (the default).
- The `/u01` directory on the database host file system has sufficient free space for the execution of backup processes.
- The `.bash_profile` file for the oracle user does not include any interactive commands (such as `oraenv` or one that could generate an error or warning message).
- (For automatic backups) No changes were made to the default `WALLET_LOCATION` entry in the `sqlnet.ora` file.
- No changes were made to `RMAN` backup settings by using standard `RMAN` commands.



See [Backup Failures on Bare Metal and Virtual Machine DB Systems](#) for details on problems that can result from not following these guidelines.

### ORACLE CLOUD INFRASTRUCTURE MANAGED BACKUP FEATURES

The following information applies to managed backups configured using the Oracle Cloud Infrastructure Console or [API](#).

#### *AUTOMATIC INCREMENTAL BACKUPS*

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

#### **Backup Retention**

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

#### **Backup Scheduling**

The automatic backup process starts at any time during your daily backup window. You can optionally specify a 2-hour scheduling window for your database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting on an even-numbered hour (for example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM). Backup jobs do not necessarily complete within the scheduling window

The default backup window of 00:00 to 06:00 in the time zone of the DB system's region is assigned to your database if you do not specify a window. Note that the default backup scheduling window is six hours long, while the windows you specify are two hours long. See [note](#) for backup window time zone information.



### Note

- **Backup Window Time Zone** - Automatic backups enabled for the first time after November 20, 2018 on any database will run between midnight and 6:00 AM in the time zone of the DB system's region. If you have enabled automatic backups on a database before this date, the backup window for the database will continue to be between midnight and 6:00 AM **UTC**. You can create a [My Oracle Support](#) service request to have your automatic backups run in a backup window of your choice.
- **Data Guard** - You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association. However, automatic backups for that database will not be created until it assumes the primary role.
- **Retention Period Changes** - If you shorten your database's automatic backup retention period in the future, existing backups falling outside the updated retention period are deleted by the system.
- **Object Storage Costs** - Automatic backups incur Object Storage usage costs.

### *ON-DEMAND FULL BACKUPS*

You can create a full backup of your database at any time unless your database is assuming the standby role in a Data Guard association.

## CHAPTER 11 Database

---

### STANDALONE BACKUPS

When you terminate a DB system or a database, all of its resources are deleted, along with any automatic backups. Full backups remain in Object Storage as standalone backups. You can use a standalone backup to create a new database.

### USING THE CONSOLE

You can use the Console to enable automatic incremental backups, create full backups on demand, and view the list of managed backups for a database. The Console also allows you to delete full backups.



#### Note

The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `RMAN` or `dbcli`).

All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.

To navigate to the list of standalone backups for your current compartment

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Click **Standalone Backups** under **Bare Metal, VM, and Exadata**.

To configure automatic backups for a database

When you launch a DB system, you can optionally enable automatic backups for the initial database. Use this procedure to configure or disable automatic backups after the database is created.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database for which you want to enable or disable automatic backups, and click its name to display database details. The details indicate whether automatic backups are enabled. When backups are enabled, the details also indicate the chosen backup retention period .
5. Click **Configure Automatic Backups**.
6. In the **Configure Automatic Backups** dialog, check or uncheck **Enable Automatic Backup**, as applicable.  
If you are enabling automatic backups, you can choose to configure the following:
  - **Backup Retention Period:** If you enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, 60 days, or 90 days. The default selection is 30 days.
  - **Backup Scheduling:** If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, the six-hour default window of 00:00 to 06:00 (in the time zone of the DB system's region) is used for your database. See [Backup Scheduling](#) for more information.
7. Click **Save Changes**.

### To create an on-demand full backup of a database

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.

3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database for which you want to create an on-demand full backup and click its name to display database details.
5. Under **Resources**, click **Backups**.  
A list of backups is displayed.
6. Click **Create Backup**.

### To delete full backups from Object Storage



#### Note

You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in Object Storage for 30 days, after which time they are automatically deleted.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located and click the DB system name to display details.  
A list of databases is displayed.
4. Find the database you are interested in and click its name to display database details.
5. Under **Resources**, click **Backups**.  
A list of backups is displayed.

## CHAPTER 11 Database

---

6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Delete**.
7. Confirm when prompted.

### USING THE API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage database backups:

- [ListBackups](#)
- [GetBackup](#)
- [CreateBackup](#)
- [DeleteBackup](#)
- [UpdateDatabase](#) - To enable and disable automatic backups.

For the complete list of APIs for the Database service, see [Database Service API](#).

### WHAT'S NEXT?

See [Recovering a Database from Object Storage](#).

### Backing Up a Database to Object Storage Using RMAN



#### Note

This topic is not applicable to Exadata DB systems. For Exadata DB systems, see [Managing Exadata Database Backups by Using bkup\\_api](#).

This topic explains how to use Recovery Manager (RMAN) to manage backups of your Bare Metal or Virtual Machine DB system database to your own Object Storage. For backups

managed by Oracle Cloud Infrastructure, see [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#).

To back up to the service you'll need to create an Object Storage bucket for the backups, generate a password for the service, install the Oracle Database Cloud Backup Module, and then configure RMAN to send backups to the service. The backup module is a system backup to tape (SBT) interface that's tightly integrated with RMAN, so you can use familiar RMAN commands to perform backup and recovery operations.

You'll notice *Swift* mentioned in the Console and in the endpoint URL for the service. That's because the backup module is typically used to back up to the Oracle Database Backup Cloud Service, which is an OpenStack Swift object store.



### Tip

On a 1-node DB system, you can use the database command line interface (`dbcli`) to back up to Object Storage. This is an alternative to installing the backup module and using RMAN for backups. For more information, see [Objectstoreswift Commands](#). Note that the `dbcli` commands are not available for a 2-node RAC DB system.

### PREREQUISITES

You'll need the following:

- A DB system and a database to back up. For more information, see [Creating Bare Metal and Virtual Machine DB Systems](#).
- The DB system's cloud network (VCN) must be configured with access to Object Storage:
  - For Object Storage access in the same region as the DB system: Oracle recommends using a service gateway. For more information, see [Service Gateway for the VCN](#).

- For Object Storage access in a different region than the DB system: Use an internet gateway. Note that the network traffic between the DB system and Object Storage does not leave the cloud and never reaches the public internet. For more information, see [Internet Gateway](#).
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket. For more information, see [Managing Buckets](#).
- An [auth token](#) generated by Oracle Cloud Infrastructure. You can use the Console or the IAM API to generate the password. For more information, see [Working with Auth Tokens](#).
- The user name (specified when you install and use the backup module) must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to *all* of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

```
Allow group <group_name> to manage objects in compartment <compartment_name> where
target.bucket.name = '<bucket_name>'
```

```
Allow group <group_name> to read buckets in compartment <compartment_name>
```

For more information about adding a user to a group, see [Managing Groups](#). For more information about policies, see [Getting Started with Policies](#).

### INSTALLING THE BACKUP MODULE ON THE DB SYSTEM

1. SSH to the DB system, log in as `opc`, and `sudo` to the oracle user.

```
ssh -i <SSH_key_used_when_launching_the_DB_system> opc@<DB_system_IP_address_or_hostname>
login as: opc
sudo su - oracle
```

2. Change to the directory that contains the backup module `opc_install.jar` file.

```
cd /opt/oracle/oak/pkgrepos/oss/odbc
```

3. Use the following command syntax to install the backup module.

## CHAPTER 11 Database

```
java -jar opc_install.jar -opcId <user_id> -opcPass '<auth_token>' -container <bucket_name> -
walletDir ~/hsbtwallet/ -libDir ~/lib/ -configfile ~/config -host
https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>
```

The parameters are:

Parameter	Description
-opcId	<p>The user name for the Oracle Cloud Infrastructure user account, for example:</p> <pre>-opcId &lt;username&gt;@&lt;example&gt;.com</pre> <p>This is the user name you use to sign in to the Console.</p> <p>The user name must be a member of the Administrators group, as described in <a href="#">Prerequisites</a>.</p> <p>You can also specify the user name in single quotes. This might be necessary if the name contains special characters, for example:<pre>-opcId 'j~smith@&lt;example&gt;.com'</pre><p>Make sure to use straight single quotes and not slanted apostrophes.</p></p>
-opcPass	<p>The <a href="#">auth token</a> generated by using the Console or IAM API, in single quotes, for example:</p> <pre>-opcPass '&lt;password&gt;'</pre> <p>Make sure to use straight single quotes and not slanted apostrophes.</p> <p>For more information, see <a href="#">Managing User Credentials</a>.</p> <p>This is <b>not</b> the password for the Oracle Cloud Infrastructure user.</p>
-container	<p>The name of an existing bucket in Object Storage to use as the backup destination, for example:</p> <pre>-container DBBackups</pre>
-walletDir	<p>The directory where the install tool will create an Oracle Wallet containing the Oracle Cloud Infrastructure user name and <a href="#">auth token</a>.</p> <pre>-walletDir ~/hsbtwallet</pre> creates the wallet in the current user (oracle) home directory.

Parameter	Description
-libDir	The directory where the SBT library is stored. The directory must already exist before you run the command. This parameter causes the latest SBT library to be downloaded.  -libDir ~/lib/ downloads the libopc.so file to the current user's home directory, for example, /home/oracle/lib/libopc.so.
-configfile	The name of the initialization parameter file that will be created by the install tool. This file will be referenced by your RMAN jobs.  -configfile ~/config creates the file in the current user's home directory, for example, /home/oracle/config.
-host	The endpoint URL to which backups are to be sent:  https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>  where <tenant> is the lowercase tenant name (even if it contains uppercase characters) that you specify when signing in to the Console, for example:  https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/companyabc  Do not add a slash after the tenant name.  See <a href="#">Regions and Availability Domains</a> to look up the region name.

### CONFIGURING RMAN

This section describes how to configure RMAN to use the bucket as the default backup destination. The following assumes you are still logged in to the DB system.

1. On the DB system, set the ORACLE\_HOME and ORACLE\_SID environment variables using the oraenv utility.

```
. oraenv
```

2. Connect to the database using RMAN.

## CHAPTER 11 Database

---

```
rman target /
```

3. Configure RMAN to use the SBT device and point to the `config` file that was created when you installed the backup module. A sample command for a version 12 database is shown here.

```
RMAN> CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=/home/oracle/lib/libopc.so,
SBT_PARMS=(OPC_PFILE=/home/oracle/config)';
```

4. Configure RMAN to use SBT\_TAPE by default. The following sample enables the controlfile and spfile autobackup to SBT\_TAPE and configures encryption (recommended). There are other settings that may apply to your installation such as compression, number of backup and recovery channels to use, backup retention policy, archived log deletion policy, and more. See the Oracle Backup and Recovery documentation for your version of Oracle for more information on choosing the appropriate settings.

```
RMAN> CONFIGURE DEFAULT DEVICE TYPE TO SBT_TAPE;
RMAN> CONFIGURE BACKUP OPTIMIZATION ON;
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP ON;
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP FORMAT FOR DEVICE TYPE SBT_TAPE TO '%F';
RMAN> CONFIGURE ENCRYPTION FOR DATABASE ON;
```

Once the RMAN configuration is complete, you can use the same RMAN commands that you regularly use for tape backups.

### BACKING UP THE DATABASE

This section provides examples of commonly used backup commands.

1. Set the database encryption:

```
RMAN> SET ENCRYPTION IDENTIFIED BY "password" ONLY;
```

Note that this setting is not permanent; you must set it for each new RMAN session.

2. Back up the database and archivelogs. Below are some example commands. See the Oracle Backup and Recovery documentation for your version of Oracle for more information about choosing a back up procedure that meets your needs. Be sure to back up regularly to minimize potential data loss and always include a copy of the spfile and controlfile. Note that the example below uses multi-section incremental backups, which

is a feature introduced in 12c. When using 11g, omit the `section size` clause.

```
RMAN> BACKUP INCREMENTAL LEVEL 0 SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;
```

```
RMAN> BACKUP INCREMENTAL LEVEL 1 SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;
```

```
RMAN> BACKUP INCREMENTAL LEVEL 1 CUMULATIVE SECTION SIZE 512M DATABASE PLUS ARCHIVELOG;
```

3. Backup archivelogs frequently to minimize potential data loss, and keep multiple backup copies as a precaution.

```
RMAN> BACKUP ARCHIVELOG ALL NOT BACKED UP 2 TIMES;
```

When the backup job completes, you can display the backup files in your bucket in the Console on the **Storage** page, by selecting **Object Storage**.

### WHAT'S NEXT?

See [Recovering a Database from Object Storage](#).

### Backing Up a Database to Local Storage Using the Database CLI



#### Note

This topic is not applicable to virtual machine DB systems because they have no local storage. For Exadata DB systems, see [Managing Exadata Database Backups](#).

This topic explains how to back up to the local Fast Recovery Area on a bare metal DB system by using the database CLI (dbcli). Some sample dbcli commands are provided below. For complete command syntax, see the [Oracle Database CLI Reference](#).



### Note

Backing up to local storage is fast and provides for fast point-in-time recovery, however, if the DB system becomes unavailable, the backup also becomes unavailable. For information about more durable backup destinations, see [Backing Up a Database](#).

### BACKING UP THE DATABASE TO LOCAL STORAGE

You'll use the dbcli commands to create a backup configuration, associate the backup configuration with the database, initiate the backup operation, and then review the backup job.

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Create a backup configuration by using the [dbcli update-backupconfig](#) command and specify local disk storage as the backup destination.

The following example creates a backup configuration named `prodbackup` and specifies a disk backup destination and a disk recovery window of 5 (backups and archived redo logs will be maintained in local storage for 5 days).

```
[root@dbsys ~]# dbcli create-backupconfig --name prodbackup --backupdestination disk --
recoverywindow 5
{
 "jobId" : "e7050756-0d83-48ce-9336-86592be59827",
 "status" : "Success",
```

```

"message" : null,
"reports" : [{
 "taskId" : "TaskParallel_471",
 "taskName" : "persisting backup config metadata",
 "taskResult" : "Success",
 "startTime" : 1467774813141,
 "endTime" : 1467774813207,
 "status" : "Success",
 "taskDescription" : null,
 "parentTaskId" : "TaskSequential_467",
 "jobId" : "e7050756-0d83-48ce-9336-86592be59827",
 "reportLevel" : "Info",
 "updatedAt" : 1467774813207
}],
"createTimestamp" : 1467774781851,
"description" : "create backup config:prodbackup",
"updatedAt" : 1467774813236
}

```

The example above uses full parameter names for demonstration purposes, but you can abbreviate the parameters like this:

```
dbcli create-backupconfig -n prodbackup -d disk -w 5
```

4. Get the ID of the database you want to back up by using the [dbcli list-databases](#) command.

```
[root@dbsys ~]# dbcli list-databases
```

ID	DB Name	DB Version	CDB	Class	Shape
71ec8335-113a-46e3-b81f-235f4d1b6fde	prod	12.1.0.2	true	OLTP	odbl ACFS

Configured

5. Get the ID of the backup configuration by using the [dbcli list-backupconfigs](#) command.

```
[root@dbackup backup]# /opt/oracle/dcs/bin/dbcli list-backupconfigs
```

ID	Name	DiskRecoveryWindow
	BackupDestination	createTime

```

78a2a5f0-72b1-448f-bd86-cf41b30b64ee prodbackup 5 Disk July 6, 2016 3:13:01
AM UTC
```

6. Associate the backup configuration ID with the database ID by using the [dbcli update-database](#) command.

```
[root@dbsys ~]# dbcli update-database --backupconfigid 78a2a5f0-72b1-448f-bd86-cf41b30b64ee --
dbid 71ec8335-113a-46e3-b81f-235f4d1b6fde
{
 "jobId" : "2b104028-a0a4-4855-b32a-b97a37f5f9c5",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1467775842977,
 "description" : "update database id:71ec8335-113a-46e3-b81f-235f4d1b6fde",
 "updatedAt" : 1467775842978
}
```

You can view details about the update job by using the [dbcli describe-job](#) command and specifying the job ID from the `dbcli update-database` command output, for example:

```
dbcli describe-job --jobid 2b104028-a0a4-4855-b32a-b97a37f5f9c5
```

7. Initiate the database backup by using the [dbcli create-backup](#) command. The backup operation is performed immediately.

The following example creates a backup of the specified database.

```
[root@dbsys ~]# dbcli create-backup --dbid 71ec8335-113a-46e3-b81f-235f4d1b6fde
{
 "createTimestamp": 1467792576854,
 "description": "Backup service creation with db name: prod",
 "jobId": "d6c9edaa-fc80-40a9-bcdd-056430cdc56c",
 "message": null,
 "reports": [],
 "status": "Created",
 "updatedAt": 1467792576855
}
```

Or you can abbreviate the command parameters like this:

```
dbcli create-backup -i 71ec8335-113a-46e3-b81f-235f4d1b6fde
```

You can view details about the back up job by using the [dbcli describe-job](#) command and specifying the job ID from the `dbcli create-backup` command output, for example:

```
dbcli describe-job --jobid d6c9edaa-fc80-40a9-bcdd-056430cdc56c
```

8. **Important!** Manually back up any TDE password-based wallets to your choice of a safe location, preferably not on the DB system. The wallets are required to restore the backup to a new host.
9. Optionally, you can review the backup report. Use the [Oracle Database CLI Reference](#) command to create a report, then use [Oracle Database CLI Reference](#) to get the report ID, and then use [Oracle Database CLI Reference](#) with the report ID to get the report location, as shown in the following example.

```
[root@dbsys ~]# dbcli create-backupreport --dbid 71ec8335-113a-46e3-b81f-235f4d1b6fde --
reporttype summary
```

```
{
 "jobId" : "65ce79fe-4ef4-4d7d-8020-e56a5390026d",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "July 6, 2016 23:06:11 PM UTC",
 "description" : "Creating a report for database 71ec8335-113a-46e3-b81f-235f4d1b6fde",
 "updatedAt" : "July 6, 2016 23:06:11 PM UTC"
}
```

```
[root@dbsys ~]# dbcli list-backupreports
```

ID	Name	ReportType	DbId
createTime	updatedAt		
<i>c0e0a16a-485f-4176-ab73-5b30ccf5c560</i>		summary	71ec8335-113a-46e3-b81f-235f4d1b6fde
July 6, 2016 11:04:05 PM UTC	July 6, 2016 11:04:17 PM UTC		

```
[root@dbsys ~]# dbcli describe-backupreport --id c0e0a16a-485f-4176-ab73-5b30ccf5c560
```

```
Backup Report details
```

```

ID: ed67a2cf-fe63-4755-a5d6-7eda5b669837
Name:
Report Type: summary
Location:
/opt/oracle/dcs/log/LrbvevghazqOGpbatvbpmRZJeVCzaW/rman/bkup/hhUiFnBz/rman_list_backup_
summary/2016-07-06/rman_list_backup_summary_2016-07-06_09-49-20.0832.log
Database ID: 71ec8335-113a-46e3-b81f-235f4d1b6fde
CreatedTime: July 6, 2016 3:49:12 AM UTC
UpdatedTime: July 6, 2016 3:49:24 AM UTC
```

After the backup command completes, the database backup files are available in the Fast Recovery Area on the DB system.

### WHAT'S NEXT?

See [Recovering a Database from a CLI Backup](#).

## Recovering a Database

For information on restoring a database on a bare metal or virtual machine DB system, see the following topics:

- [Recovering a Database from Object Storage](#)
- [Recovering a Database from a CLI Backup](#)
- [Recovering a Database from the Oracle Cloud Infrastructure Classic Object Store](#)

### Recovering a Database from Object Storage



#### Note

This topic is not applicable to Exadata DB systems.

This topic explains how to recover a database from a backup stored in Object Storage. The service is a secure, scalable, on-demand storage solution in Oracle Cloud Infrastructure. For

information on using Object Storage as a backup destination, see [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#).

You can recover a database using the Console, API, or by using RMAN.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### REQUIRED IAM POLICY

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### PREREQUISITES

The DB system requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see these topics:

- [Network Setup for DB Systems](#): For information about setting up your VCN for the DB system, including the service gateway.
- <https://cloud.oracle.com/infrastructure/storage/object-storage/faq>: For information about the Swift endpoints to use.

### USING THE CONSOLE

You can use the Console to restore the database from a backup in the Object Storage that was

created by using the Console or the API. You can restore to the last known good state of the database, or you can specify a point in time or an existing System Change Number (SCN). You can also create a new database by using a standalone backup.



### Note

The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `RMAN` or `dbcli`).

Restoring a database with Data Guard enabled is not supported. You must first remove the Data Guard association by terminating the standby database before you can restore the database.

### RESTORING AN EXISTING DATABASE

#### To restore a database

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database you want to restore, and click its name to display details about it.  
A list of backups is displayed in the default view of the database details. You can also access the list of backups for a database by clicking on **Backups** under **Resources**.
5. Click **Restore**.
6. Select one of the following options, and then click **Restore Database**:

- **Restore to the latest:** Restores the database to the last known good state with the least possible data loss.
- **Restore to the timestamp:** Restores the database to the timestamp specified.
- **Restore to System Change Number (SCN):** Restores the database using the SCN specified. This SCN must be valid.



### Tip

You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

7. Confirm when prompted.

If the restore operation fails, the database will be in a "Restore Failed" state. You can try restoring again using a different restore option. However, Oracle recommends that you review the `RMAN` logs on the host and fix any issues before reattempting to restore the database.

### To restore a database using a specific backup from Object Storage

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose your **Compartment**.  
A list of DB systems is displayed.
3. Find the DB system where the database is located, and click the system name to display details about it.  
A list of databases is displayed.
4. Find the database you want to restore, and click its name to display details about it.
5. Under **Resources**, click **Backups**.

## CHAPTER 11 Database

---

A list of backups is displayed.

6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Restore**.
7. Confirm when prompted.

### *CREATING A NEW DATABASE FROM A BACKUP*

You can use a backup to create a database in an existing DB system or to launch a new DB system. See the following procedures for more information:

- [To create a database from a backup in an existing DB system](#)
- [To launch a new DB system from a backup](#)

### **USING THE API**

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to recover a database:

- [ListBackups](#)
- [GetBackup](#)
- [RestoreDatabase](#)
- [CreateDbHome](#) - For creating a DB system database from a standalone backup.

For the complete list of APIs for the Database service, see [Database Service API](#).

### **USING AN RMAN BACKUP**

This topic explains how to recover a Recovery Manager (RMAN) backup stored in Object Storage.

#### *PREREQUISITES*

You'll need the following:

## CHAPTER 11 Database

---

- A new DB system to restore the database to (see assumptions below). For more information, see [Creating Bare Metal and Virtual Machine DB Systems](#).
- The Oracle Database Cloud Backup Module must be installed on the DB system. For more information, see [Installing the Backup Module on the DB System](#).

### ASSUMPTIONS

The procedures below assume the following:

- A new DB system has been created to host the restored database and no other database exists on the new DB system. It is possible to restore to a DB system that has existing databases, but that is beyond the scope of this topic.
- The original database is lost and all that remains is the latest RMAN backup. For virtual machine DB systems, the procedure assumes the DB system (inclusive of the database) no longer exists.



### Warning

Any data not included in the most recent backup will be lost.

- The Oracle Wallet and/or encryption keys used by the original database at the time of the last backup is available.
- The RMAN backup contains a copy of the control file and spfile as of the most recent backup as well as all of the datafile and archivelog backups needed to perform a complete database recovery.
- An RMAN catalog will not be used during the restore.

### SETTING UP STORAGE ON THE DB SYSTEM

#### 1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
[opc@dbsys ~]$ sudo su -
```

3. You can use an existing empty database home or create a new one for the restore. Use the applicable commands to help you complete this step.

If you will be using an existing database home:

- Use the [dbcli list-dbhomes](#) command to list the database homes.

```
[root@dbsys ~]# dbcli list-dbhomes
ID Name DB Version Home Location

2e743050-b41d-4283-988f-f33d7b082bda OraDB12102_home1 12.1.0.2
/u01/app/oracle/product/12.1.0.2/dbhome_1
```

- Use the [dbcli list-databases](#) command to ensure the database home is not associated with any database.

If necessary, use the [dbcli create-dbhome](#) command to create a database home for the restore.

4. Use the [dbcli create-dbstorage](#) to set up directories for DATA, RECO, and REDO storage. The following example creates 10GB of ACFS storage for the rectest database.

```
[root@dbsys ~]# dbcli create-dbstorage --dbname rectest --dataSize 10 --dbstorage ACFS
```



**Note**

When restoring a version 11.2 database, ACFS storage must be specified.

## CHAPTER 11 Database

---

### PERFORMING THE DATABASE RESTORE AND RECOVERY

1. SSH to the DB system, log in as `opc`, and then become the oracle user.

```
sudo su - oracle
```

2. Create an entry in `/etc/oratab` for the database. Use the same SID as the original database.

```
db1:/u01/app/oracle/product/12.1.0.2/dbhome_1:N
```

3. Set the `ORACLE_HOME` and `ORACLE_SID` environment variables using the `oraenv` utility.

```
. oraenv
```

4. Obtain the DBID of the original database. This can be obtained from the file name of the `controlfile autobackup` on the backup media. The file name will include a string that contains the DBID. The typical format of the string is `c-XXXXXXXXXXXX-YYYYMMDD-NN` where `XXXXXXXXXXXX` is the DBID, `YYYYMMDD` is the date the backup was created, and `NN` is a sequence number to make the file name unique. The DBID in the following examples is 1508405000. Your DBID will be different.

Use the following `curl` syntax to perform a general query of Object Storage. The parameters in red are the same parameters you specified when installing the backup module as described in [Installing the Backup Module on the DB System](#).

```
curl -u '<user_ID>.com:<auth_token>' -v https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant_name>
```

See [Regions and Availability Domains](#) to look up the region name.

For example:

```
curl -u 'djones@mycompany.com:lcnk!d0++ptETd&C;THR' -v https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/mycompany
```

To get the DBID from the control file name, use the following syntax:

```
curl -u '<user_id>.com:<auth_token>' -v https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant_name>/<bucket_name>?prefix=sbt_catalog/c-
```

For example:

```
curl -u 'djones@mycompany.com:1cnk!d0++ptETd&C;THR' -v https://swiftobjectstorage.<region_
name>.oraclecloud.com/v1/mycompany/dbbackups/?prefix=sbt_catalog/c-
```

In the sample output below, 1508405000 is the DBID.

```
{
 "bytes": 1732,
 "content_type": "binary/octet-stream",
 "hash": "f1b61f08892734ed7af4f1ddaabae317",
 "last_modified": "2016-08-11T20:28:34.438000",
 "name": "sbt_catalog/c-1508405000-20160811-00/metadata.xml"
}
```

5. Run RMAN and connect to the target database. There is no need to create a pfile or spfile or use a backup controlfile. These will be restored in the following steps. Note that the target database is (not started). This is normal and expected at this point.

```
rman target /

Recovery Manager: Release 12.1.0.2.0 - Production on Wed Jun 22 18:36:40 2016

Copyright (c) 1982, 2014, Oracle and/or its affiliates. All rights reserved.

connected to target database (not started)
```

6. Set the DBID using the value obtained above.

```
RMAN> set dbid 1508405000;

executing command: SET DBID
```

7. Run the STARTUP NOMOUNT command. If the server parameter file is not available, RMAN attempts to start the instance with a dummy server parameter file. The ORA-01078 and LRM-00109 errors are normal and can be ignored.

```
RMAN> STARTUP NOMOUNT

startup failed: ORA-01078: failure in processing system parameters
LRM-00109: could not open parameter file '/u01/app/oracle/product/12.1.0.2/dbhome_
1/dbs/initdb1.ora'
```

```
starting Oracle instance without parameter file for retrieval of spfile
Oracle instance started

Total System Global Area 2147483648 bytes

Fixed Size 2944952 bytes
Variable Size 847249480 bytes
Database Buffers 1254096896 bytes
Redo Buffers 43192320 bytes
```

### 8. Restore the server parameter file from autobackup.

The `SBT_LIBRARY` is the same library specified with the `-libDir` parameter when the Backup Module was installed, for example `/home/oracle/lib/`.

The `OPC_PFILE` is the same file specified with the `-configfile` parameter when the Backup Module was installed, for example `/home/oracle/config`.

```
set controlfile autobackup format for device type sbt to '%F';
run {
 allocate channel c1 device type sbt PARMS 'SBT_LIBRARY=/home/oracle/lib/libopc.so, SBT_PARMS=
(OPC_PFILE=/home/oracle/config)';
 restore spfile from autobackup;
}
```

### 9. Create the directory for `audit_file_dest`. The default is

`/u01/app/oracle/admin/$ORACLE_SID/adump`. You can see the setting used by the original database by searching the spfile for the string, `audit_file_dest`.

```
strings ${ORACLE_HOME}/dbs/spfile${ORACLE_SID}.ora | grep audit_file_dest
*.audit_file_dest='/u01/app/oracle/admin/db1/adump'

mkdir -p /u01/app/oracle/admin/db1/adump
```

### 10. If block change tracking was enabled on the original database, create the directory for the block change tracking file. This will be a directory under `db_create_file_dest`. Search the `spfile` for the name of the directory.

```
strings ${ORACLE_HOME}/dbs/spfile${ORACLE_SID}.ora | grep db_create_file_dest
*.db_create_file_dest='/u02/app/oracle/oradata/db1'

mkdir -p /u02/app/oracle/oradata/db1/<$ORA_UNQNAME if available or database name>/changetracking
```

- Restart the instance with the restored server parameter file.

```
STARTUP FORCE NOMOUNT;
```

- Restore the controlfile from the RMAN autobackup and mount the database.

```
set controlfile autobackup format for device type sbt to '%F';
run {
 allocate channel c1 device type sbt PARMS 'SBT_LIBRARY=/home/oracle/lib/libopc.so, SBT_PARMS=
(OPC_PFILE=/home/oracle/config)';
 restore controlfile from autobackup;
 alter database mount;
}
```

- Restore and recover the database.

```
RESTORE DATABASE;
RECOVER DATABASE;
```

- RMAN will recover using archived redo logs until it can't find any more. It is normal for an error similar to the one below to occur when RMAN has applied the last archived redo log in the backup and can't find any more logs.

```
unable to find archived log
archived log thread=1 sequence=29
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of recover command at 06/28/2016 00:57:35
RMAN-06054: media recovery requesting unknown archived log for thread 1 with sequence 29 and
starting SCN of 2349563
```

- Open the database with resetlogs.

```
ALTER DATABASE OPEN RESETLOGS;
```

The recovery is complete. The database will have all of the committed transactions as of the last backed up archived redo log.

### Recovering a Database from a CLI Backup



#### Note

This topic is not applicable to Exadata DB Systems.

This topic explains how to perform a complete or point-in-time recovery of an existing database from a backup created with the [dbcli create-backup](#) command. The backup resides in the local Fast Recovery Area on the DB System.

To initiate the recovery, you'll use the [Oracle Database CLI Reference](#) command and specify the recovery type parameter (either `--recoverytype` or just `-t`). You can specify the following types of recovery:

- `-t Latest` for a complete recovery
- `-t SCN -s <scn>` for a recovery using a system change number (SCN) as the end point of the recovery
- `-t PITR <mm/dd/yyyy hh:mm:ss>` for a database point-in-time (incomplete) recovery based on a time stamp

The `dbcli create-recovery` attempts to perform a full recovery of the database. For information on performing a partial recovery (datafile, tablespace and PDB), see the *Oracle Database Backup Recovery Guide* for version [18.1](#), [12.2](#), [12.1](#), or [11.2](#).

#### PREREQUISITES

- The backup must have been created with the `dbcli create-backup` command.
- If the database is configured with Transparent Data Encryption (TDE), make sure the password-based and autologin TDE wallets are present in the following location:

```
/opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>
```

### RECOVERING THE DATABASE

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Find the ID of database you want to recover by using the [dbcli list-databases](#) command. You'll need the ID for the following step.

```
[root@dbsys ~]# dbcli list-databases
```

ID	DB Name	DB Version	CDB	Class	Shape
5a3e980b-e0fe-4909-9628-fcefe43b3326	prod	12.1.0.2	true	OLTP	odb1 ACFS

Configured

4. Initiate the recovery by using the [Oracle Database CLI Reference](#) command and specifying the database ID, recovery type parameter (`-t`), and any parameter required for the recover type, like the time stamp or system change number.

The following example initiates a complete recovery.

```
[root@dbsys ~]# dbcli create-recovery --dbid 5a3e980b-e0fe-4909-9628-fcefe43b3326 --recoverytype Latest
```

```
{
 "jobId" : "c9f81228-2ce9-43b4-88f6-b260d398cf06",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "August 08, 2016 18:20:47 PM UTC",
 "description" : "Create recovery for database id :5a3e980b-e0fe-4909-9628-fcefe43b3326",
 "updatedAt" : "August 08, 2016 18:20:47 PM UTC"
}
```

The following example initiates a point-in-time recovery of the specified database:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli create-recovery --dbid d4733796-dbea-4155-8606-24a85d64bd74 --recoverytype
PITR --recoveryTimeStamp 08/09/2016 5:12:15
```

Note the job ID in the command output.

5. Check the status of the recovery by using the [dbcli describe-job](#) command with the job ID from the previous step.

```
[root@dbsys ~]# dbcli describe-job -i c9f81228-2ce9-43b4-88f6-b260d398cf06

Job details

 ID: c9f81228-2ce9-43b4-88f6-b260d398cf06
Description: Create recovery for database id :5a3e980b-e0fe-4909-9628-fcefe43b3326
 Status: Success
 Created: August 8, 2016 6:20:47 PM UTC
 Message:

Task Name Status Start Time End Time

Database recovery validation
6:21:07 PM UTC Success August 8, 2016 6:20:47 PM UTC August 8, 2016
6:21:07 PM UTC
Database recovery
6:22:34 PM UTC Success August 8, 2016 6:21:07 PM UTC August 8, 2016
6:22:34 PM UTC
enable block change tracking
6:22:35 PM UTC Success August 8, 2016 6:22:34 PM UTC August 8, 2016
6:22:35 PM UTC
Open database
6:22:44 PM UTC Success August 8, 2016 6:22:35 PM UTC August 8, 2016
6:22:44 PM UTC
Restart database
6:23:41 PM UTC Success August 8, 2016 6:22:44 PM UTC August 8, 2016
6:23:41 PM UTC
Persist Recovery Metadata
6:23:41 PM UTC Success August 8, 2016 6:23:41 PM UTC August 8, 2016
6:23:41 PM UTC
```

You can also check the database restore report logs on the DB System at:

```
/opt/oracle/dcs/log/<nodename>/rman/bkup/<db_unique_name>
```

### Recovering a Database from the Oracle Cloud Infrastructure Classic Object Store



#### Note

This topic is not applicable to Exadata DB systems.

This topic explains how to recover a database using a backup created by the Oracle Database Backup Module and stored in Oracle Cloud Infrastructure Object Storage Classic.

The following terms are used throughout this topic:

- Source database: The database backup in Object Storage Classic.
- Target database: The new database on a DB system in Oracle Cloud Infrastructure.

#### PREREQUISITES

You'll need the following:

- The service name, identity name, container, user name, and password for Oracle Cloud Infrastructure Object Storage Classic.
- The backup password if password-based encryption was used when backing up to Object Storage Classic.
- The source database ID, database name, database unique name (required for setting up storage).
- If the source database is configured with Transparent Data Encryption (TDE), you'll need a backup of the wallet and the wallet password.
- Tnsnames to setup for any database links.
- The output of `Opatch lsinventory` for the source database `Oracle_home`, for reference.
- A copy of the `sqlpatch` directory from the source database home. This is required for rollback in case the target database does not include these patches.

### SETTING UP STORAGE ON THE DB SYSTEM

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Use the [dbcli create-dbstorage](#) to set up directories for DATA, RECO, and REDO storage. The following example creates 10GB of ACFS storage for the tdetest database.

```
[root@dbsys ~]# dbcli create-dbstorage --dbname tdetest --dataSize 10 --dbstorage ACFS
```



#### Note

When migrating a version 11.2 database, ACFS storage must be specified.

4. Use the [dbcli list-dbstorages](#) command to list the storage ID. You'll need the ID for the next step.

```
[root@dbsys ~]# dbcli list-dbstorages
```

ID	Type	DBUnique Name	Status
9dcdfb8e-e589-4d5f-861a-e5ba981616ed	Acfs	tdetest	Configured

5. Use the [dbcli describe-dbstorage](#) command with the storage ID from the previous step to list the DATA, RECO and REDO locations.

```
[root@dbsys ~]# dbcli describe-dbstorage --id 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
```

```
DBStorage details
```

```

ID: 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
DB Name: tdetest
```

## CHAPTER 11 Database

---

```
DBUnique Name: tdetest
DB Resource ID:
 Storage Type: Acfs
 DATA Location: /u02/app/oracle/oradata/tdetest
 RECO Location: /u03/app/oracle/fast_recovery_area/
 REDO Location: /u03/app/oracle/redo/
 State: ResourceState(status=Configured)
 Created: August 24, 2016 5:25:38 PM UTC
 UpdatedTime: August 24, 2016 5:25:53 PM UTC
```

6. Note the DATA, RECO and REDO locations. You'll need them later to set the `db_create_file_dest`, `db_create_online_log_dest`, and `db_recovery_file_dest` parameters for the database.

### CHOOSING AN ORACLE\_HOME

Decide which ORACLE\_HOME to use for the database restore and then switch to that home with the correct ORACLE\_BASE, ORACLE\_HOME, and PATH settings. The ORACLE\_HOME must not already be associated with a database.

To get a list of existing ORACLE\_HOMEs and to ensure that the ORACLE\_HOME is empty, use the [dbcli list-dbhomes](#) and the [dbcli list-databases](#) commands, respectively. To create a new ORACLE\_HOME, use the [dbcli create-dbhome](#) command.

### COPYING THE SOURCE DATABASE WALLETS

Skip this section if the source database is **not** configured with TDE.

1. On the DB system, become the oracle user:

```
sudo su - oracle
```

2. Create the following directory, if it does not already exist:

```
mkdir /opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>
```

3. Copy the ewallet.p12 file from the source database to the directory you created in the previous step.
4. On the target host, make sure that `$ORACLE_HOME/network/admin/sqlnet.ora` contains the following line:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=
(DIRECTORY=/opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

Add the line if it doesn't exist in the file. (The line might not be there if this is a new home and no database has been created yet on this host.)

5. Create the autologin wallet from the password-based wallet to allow auto-open of the wallet during restore and recovery operations.

For a version 12.1 or later database, use the `ADMINISTER KEY MANAGEMENT` command:

```
$cat create_autologin_12.sh

#!/bin/sh
if [$# -lt 2]; then
 echo "Usage: $0 <dbunique_name><remotewalletlocation>"
 exit 1;
fi

mkdir /opt/oracle/dcs/commonstore/wallets/tde/$1
cp $2/ewallet.p12* /opt/oracle/dcs/commonstore/wallets/tde/$1
rm -f autokey.ora
echo "db_name=$1" > autokey.ora
autokeystoreLog="autologinKeystore_`date +%Y%m%d_%H%M%S_%N`.log"
echo "Enter Keystore Password:"
read -s keystorePassword
echo "Creating AutoLoginKeystore -> "
sqlplus "/as sysdba" <<EOF
spool $autokeystoreLog
set echo on
startup nomount pfile=autokey.ora
ADMINISTER KEY MANAGEMENT CREATE AUTO_LOGIN KEYSTORE
FROM KEYSTORE '/opt/oracle/dcs/commonstore/wallets/tde/$1' -- Keystore location
IDENTIFIED BY "$keystorePassword";
shutdown immediate;
EOF
```

Adjust the `cwallet.sso` permissions from `oracle:asmadmin` to `oracle:oinstall`.

```
$ ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>

total 20
```

## CHAPTER 11 Database

---

```
-rw-r--r-- 1 oracle oinstall 5680 Jul 6 11:39 ewallet.p12
-rw-r--r-- 1 oracle asmadmin 5725 Jul 6 11:39 cwallet.sso
```

For a version 11.2 database, use the `orapki` command:

```
orapki wallet create -wallet wallet_location -auto_login [-pwd password]
```

### INSTALLING THE ORACLE DATABASE BACKUP MODULE

The backup module JAR file is included on the DB system but you need to install it.

1. SSH to the DB system, log in as `opc`, and then become the oracle user.

```
ssh -i <path to SSH key used when launching the DB System> opc@<DB System IP address or hostname>
sudo su - oracle
```

2. Change to the directory that contains the backup module `opc_install.jar` file.

```
cd /opt/oracle/oak/pkgrepos/orapks/oss/<version>/
```

3. Use the command syntax described in [Installing the Oracle Database Cloud Backup Module](#) to install the backup module.

### SETTING ENVIRONMENT VARIABLES

Set the following environment variables for the RMAN and SQL\*Plus sessions for the database:

```
ORACLE_HOME=<path of Oracle Home where the database is to be restored>
ORACLE_SID=<database instance name>
ORACLE_UNQNAME=<db_unique_name in lower case>
NLS_DATE_FORMAT="mm/dd/yyyy hh24:mi:ss"
```

### ALLOCATING AN RMAN SBT CHANNEL

For each restore operation, allocate an SBT channel and set the `SBT_LIBRARY` parameter to the location of the `libopc.so` file and the `OPC_FILE` parameter to the location of the `opc_sbt.ora` file, for example:

```
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/<ORACLE_HOME>/dbs/opc_sbt.ora)';
```

## CHAPTER 11 Database

---

(For more information about these files, see [Files Created When the Backup Module is Installed.](#))

### ENSURING DECRYPTION IS TURNED ON

Make sure that decryption is turned on for all the RMAN restore sessions.

```
set decryption wallet open identified by <keystore password>;
```

For more information, see [Providing Password Required to Decrypt Encrypted Backups.](#)

### RESTORING SPFILE

The following sample shell script restores the spfile. Set the `$dbID` variable to the dbid of the database being restored. By default, spfile is restored to `$ORACLE_HOME/dbs/spfile<sid>.ora`.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_dbid_${dbID}_restore_spfile.log"
startup nomount
set echo on
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
SET DBID=${dbID};
RESTORE SPFILE FROM AUTOBACKUP;
shutdown immediate;
EOF
```

### SETTING THE DATABASE PARAMETERS

1. Start the database in nomount mode.

```
startup nomount
```

2. Update spfile and modify the following parameters.
  - If the database storage type is ACFS, use the DATA, RECO, and REDO locations obtained from the `dbcli describe-dbstorage` command output, as described in [Setting Up Storage on the DB System](#):

## CHAPTER 11 Database

```
alter system set db_create_file_dest='/u02/app/oracle/oradata/' scope = spfile;
alter system set db_create_online_log_dest_1='/u03/app/oracle/redo' scope = spfile;
alter system set db_recovery_file_dest='/u03/app/oracle/fast_recovery_area' scope =
spfile;
```

- If the database storage type is ASM:

```
alter system set db_create_file_dest='+DATA' scope = spfile;
alter system set db_create_online_log_dest_1='+RECO' scope = spfile;
alter system set db_recovery_file_dest='+RECO' scope = spfile;
```

- Set `db_recovery_file_dest_size` is not set or is set incorrectly:

```
alter system set db_recovery_file_dest_size=<sizeG> scope=spfile;
```

- Set `audit_file_dest` to the correct value:

```
alter system set audit_file_dest=/u01/app/oracle/admin/<db_unique_name in lower
case>/adump
```

3. Remove the `control_files` parameter. The Oracle Managed Files (OMF) parameters will be used to create the control file.

```
alter system reset control_files scope=spfile;
```

4. Restart the database in nomount mode using the newly added parameters.

```
shutdown immediate
startup nomount
```

### RESTORING THE CONTROL FILE

Modify the following sample shell script for your environment to restore the control file. Set the `$dbID` variable to the dbid of the database being restored. Set `SBT_LIBRARY` to the location specified in the `-libDir` parameter when you installed the Backup Module. Set `OPC_PFILE` to the location specified in the `-configFile` parameter, which defaults to `ORACLE_HOME/dbs/opcSID.ora`.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S`_N`_dbid_${dbID}_restore_controlfile.log"
set echo on
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_LIBRARY=/Backup
Module libDir>/libopc.so ENV=(OPC_PFILE=/Backup Module configFile>/opcSID.ora)';
```

## CHAPTER 11 Database

```
SET DBID=$dbID;
RESTORE CONTROLFILE FROM AUTOBACKUP;
alter database mount;
}

exit;
EOF
```

### RESTORING THE DATABASE

1. Preview and validate the backup. The database is now mounted and RMAN should be able to locate the backup from the restored controlfile. This step helps ensure that the list of archivelogs is present and that the backup components can be restored . In the following examples, modify SBT\_LIBRARY and OPC\_PFILE as needed for your environment.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_restore_database_preview.log"
set echo on
run {
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
ALLOCATE CHANNEL c2 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
ALLOCATE CHANNEL c3 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
restore database validate header preview;
}
```

Review the output and if there are error messages, investigate the cause of the problem.

2. Redirect the restore using `set newname` to restore the data files in OMF format and use `switch datafile all` to allow the control file to update with the new data file copies.

```
rman target / <<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_restore_database_preview.log"
set echo on
run {
```

## CHAPTER 11 Database

```
ALLOCATE CHANNEL c1 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
ALLOCATE CHANNEL c2 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
ALLOCATE CHANNEL c3 DEVICE TYPE sbt MAXPIECESIZE 2 G FORMAT '%d_%I_%U' PARMS 'SBT_
LIBRARY=/tmp/oss/libopc.so ENV=(OPC_PFILE=/tmp/oss/opc_sbt.ora)';
set newname for database to new;
restore database;
switch datafile all;
switch tempfile all;
recover database;
}
```

This recovery will attempt to use the last available archive log backup and then fail with an error, for example:

```
RMAN-00571: =====
RMAN-00569: ===== ERROR MESSAGE STACK FOLLOWS =====
RMAN-00571: =====
RMAN-03002: failure of recover command at 07/20/2016 12:09:02
RMAN-06054: media recovery requesting unknown archived log for thread 1 with sequence 22 and
starting SCN of 878327
```

3. To complete the incomplete recovery, run a recovery using the sequence number and thread number shown in the RMAN-06054 message, for example:

```
Recover database until sequence 22 thread 1;
```

### RESETTING THE LOGS

Reset the logs.

```
alter database open resetlogs;
```

### PREPARING TO REGISTER THE DATABASE

Before you register the database:

1. Make sure the database COMPATIBLE parameter value is acceptable. If the value is less than the minimum, the database cannot be registered until you upgrade the database compatibility.

The minimum compatibility values are as follows:

## CHAPTER 11 Database

---

- For a version 18.1 database - 18.0.0.0
  - For a version 12.2 or 12.1 database - 12.1.0.2
  - For a version 11.2 database - 11.2.0.4
2. Verify that the database has registered with the listener and the service name.

```
lsnrctl services
```

3. Make sure the password file was restored or created for the new database.

```
ls -ltr $ORACLE_HOME/dbs/orapw<oracle sid>
```

If the file does not exist, create it using the orapwd utility.

```
orapwd file=<$ORACLE_HOME/dbs/orapw<$ORACLE_SID>> password=<sys password>
```

4. Make sure the restored database is open in read write mode.

```
select open_mode from v$database;
```

The command output should indicate read write mode. The `dbcli register-database` command will attempt to run `datapatch`, which requires read write mode. If there are PDBs, they should also be in read write mode to ensure that `datapatch` runs on them.

5. From oracle home on the restored database, use the following command verify the connection to SYS:

```
conn sys/<password>@//<hostname>:1521/<database service name>
```

This connection is required to register the database later. Fix any connection issues before continuing.

6. Make sure the database is running on spfile by using the SQL\*Plus command.

```
SHOW PARAMETERS SPFILE
```

7. (Optional) If you would like to manage the database backup with the `dbcli` command line interface, you can associate a new or existing backup configuration with the migrated database when you register it or after you register it. A backup configuration defines the backup destination and recovery window for the database. Use the following commands to create, list, and display backup configurations:

## CHAPTER 11 Database

---

- [dbcli update-backupconfig](#)
  - [dbcli list-backupconfigs](#)
  - [dbcli describe-backupconfig](#)
8. Copy the folder `$ORACLE_HOME/sqlpatch` from source database to the target database. This will enable the `dbcli register-database` command to roll back any conflicting patches.



### Note

If you are migrating a version 11.2 database, additional steps are required after you register the database. For more information, see [Rolling Back Patches on a Version 11.2 Database](#).

### REGISTERING THE DATABASE ON THE DB SYSTEM

The [dbcli register-database](#) command registers the restored database to the dcs-agent so it can be managed by the dcs-agent stack.



### Note

The `dbcli register-database` command is not available on 2-node RAC DB systems.

As the root user, use the `dbcli register-database` command to register the database on the DB system, for example:

```
[root@dbsys ~]# dbcli register-database --dbclass OLTP --dbshape odb1 --servicename tdetest --
syspassword
Password for SYS:
{
 "jobId" : "317b430f-ad5f-42ae-bb07-13f053d266e2",
 "status" : "Created",
 "message" : null,
```

## CHAPTER 11 Database

---

```
"reports" : [],
"createTimestamp" : "August 08, 2016 05:55:49 AM EDT",
"description" : "Database service registration with db service name: tdetest",
"updatedAt" : "August 08, 2016 05:55:49 AM EDT"
}
```

### UPDATING TNSNAMES.ORA

Check the `tnsnames.ora` in the backup location, check the database links used in the cloned database, and then add any relevant connection strings to the cloned database file at `$ORACLE_HOME/network/admin/tnsnames.ora`.

### ROLLING BACK PATCHES ON A VERSION 11.2 DATABASE

For version 11.2 databases, the `sqlpatch` application is not automated, so any interim patches (previously known as a "one-off" patches) applied to the source database that are not part of the installed PSU must be rolled back manually in the target database. After registering the database, execute the `catbundle.sql` script and then the `postinstall.sql` script with the corresponding PSU patch (or the overlay patch on top of the PSU patch), as described below.



#### Tip

Some interim patches may include files written to the `$ORACLE_HOME/rdbms/admin` directory as well as the `$ORACLE_HOME/sqlpatch` directory. Oracle recommends that you roll back these patches in the source database using the instructions in the patch read-me prior to migrating the database to OCI environment. Contact Oracle Support if you need assistance with rolling back these patches.

1. On the DB System, use the `dbcli list-dbhomes` command to find the PSU patch number for the version 11.2 database home. In the following sample command output, the PSU patch number is the second number in the DB Version column:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli list-dbhomes
ID Name DB Version
Home Location Status

59d9bc6f-3880-4d4f-b5a6-c140f16f8c64 OraDB11204_home1 11.2.0.4.160719 (23054319, 23054359)
/u01/app/oracle/product/11.2.0.4/dbhome_1 Configured
```

(The first patch number, 23054319 in the example above, is for the OCW component in the database home.)

2. Find the overlay patch, if any, by using the `lsinventory` command. In the following example, patch number **24460960** is the overlay patch on top of the 23054359 PSU patch.

```
$ $ORACLE_HOME/OPatch/opatch lsinventory
...
Installed Top-level Products (1):

Oracle Database 11g 11.2.0.4.0
There are 1 products installed in this Oracle Home.

Interim patches (5) :

Patch 24460960 : applied on Fri Sep 02 15:28:17 UTC 2016
Unique Patch ID: 20539912
 Created on 31 Aug 2016, 02:46:31 hrs PST8PDT
 Bugs fixed:
 23513711, 23065323, 21281607, 24006821, 23315889, 22551446, 21174504
 This patch overlays patches:
 23054359
 This patch needs patches:
 23054359
 as prerequisites
```

3. Start SQL\*Plus and execute the `catbundle.sql` script, for example:

```
SQL> startup
SQL> connect / as sysdba
```

```
SQL> @$ORACLE_HOME/rdbms/admin/catbundle.sql psu apply
exit
```

4. Apply the sqlpatch, using the overlay patch number from the previous step, for example:

```
SQL> connect / as sysdba
SQL> @$ORACLE_HOME/sqlpatch/24460960/postinstall.sql
exit
```



### Note

If the source database has one-off patches installed and those patches are **not** part of the installed PSU in the cloud environment, then the SQL changes that correspond to those one-off patches need to be rolled back. To rollback the SQL changes, copy the `ORACLE_HOME/sqlpatch/<patch#>/postdeinstall.sql` script from the source environment to the cloud environment and execute the `postdeinstall.sql` script.

### POST RESTORE CHECKLIST

After the database is restored and registered on the DB system, use the following checklist to verify the results and perform any post-restore customizations.

1. Make sure the database files were restored in OMF format.
2. Make sure the database is listed in the [dbcli list-databases](#) command output.
3. Check for the following external references in the database and update them if necessary:
  - External tables: If the source database uses external tables, back up that data and migrate it to the target host.

- Directories: Customize the default directories as needed for the restored database.
  - Database links: Make sure all the required TNS entries are updated in the `tnsnames.ora` file in `ORACLE_HOME`.
  - Email and URLs: Make sure any email addresses and URLs used in the database are still accessible from the DB system.
  - Scheduled jobs: Review the jobs scheduled in source database and schedule similar jobs as needed in the restored database.
4. If you associated a backup configuration when you registered the database, run a test back up using the [dbcli create-backup](#) command.
  5. If the restored database contains a CDB and PDBs, verify that patches have been applied to all PDBs.

### Using Oracle Data Guard



#### Note

This topic is not applicable to Exadata DB systems.

This topic explains how to use the Console to manage Data Guard associations in your DB system. To configure a Data Guard system across regions or between on-premises and Oracle Cloud Infrastructure DB systems, you must access the database host directly and use the DGMGRL utility.

For complete information on Oracle Data Guard, see the [Data Guard Concepts and Administration](#) documentation on the [Oracle Document Portal](#).

#### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Prerequisites

A Data Guard implementation requires two DB systems, one containing the primary database and one containing the standby database. When you enable Data Guard for a virtual machine DB system database, a new DB system with the standby database is created and associated with the primary database. For a bare metal DB system, the DB system with the database to be used as the standby must already exist before you enable Data Guard.



#### Tip

A Data Guard configuration on the Oracle Cloud Infrastructure is limited to one standby database per primary database.

Requirement details are as follows:

- Both DB systems must be in the same compartment, and they must be the same shape.
- The database versions and editions must be identical. Data Guard does not support Standard Edition. (Active Data Guard requires Enterprise Edition - Extreme Performance.)
- The database version determines whether Active Data Guard is enabled. If you are using the BYOL licensing model and if your license does not include Active Data Guard, you must either use Enterprise Edition - High Performance or set up Data Guard manually. See [Using Oracle Data Guard with the Database CLI](#).
- Both DB systems must use the same VCN, and port 1521 must be open.
- **Important!** Properly configure the security list ingress and egress rules for the subnets of both DB systems in the Data Guard association to allow TCP traffic to flow between

the applicable ports. Ensure that the rules you create are stateful (the default). For example, if the subnet of the primary DB System uses the source CIDR 10.0.0.0/24 and the subnet of the standby DB system uses the source CIDR 10.0.1.0/24, create rules as shown in the following example.



### Note

The egress rules in the example show how to enable TCP traffic only for port 1521, which is a minimum requirement for Data Guard to work. If TCP traffic is already enabled on all of your outgoing ports (0.0.0.0/0), then you need not explicitly add these specific egress rules.

### Security List for Primary DB System's Subnet

#### Ingress Rules:

```
Stateless: No
Source: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

#### Egress Rules:

```
Stateless: No
Destination: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

### Security List for Standby DB System's Subnet

#### Ingress Rules:

```
Stateless: No
Source: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

**Egress Rules:**

```
Stateless: No
Destination: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

For information about creating and editing rules, see [Security Lists](#).

### Availability Domain and Fault Domain Considerations for Data Guard

Oracle recommends that the DB system of the standby database be in a different availability domain from the DB system of the primary database to improve availability and disaster recovery. If you enable Data Guard for a database and your standby database is in the same availability domain as the primary (either by choice, or because you are working in a single availability domain region), Oracle recommends that you place the standby database in a different fault domain from that of the primary database. Note that if your primary and standby databases are 2-node RAC databases and both are in the same availability domain, only one of the two nodes of the standby database can be in a fault domain that does not include any other nodes from either the primary or standby database. This is because each availability domain has only three fault domains, and the primary and standby databases have a combined total of 4 nodes. For more information on availability domains and fault domains, see [Regions and Availability Domains](#).

### Working with Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data. The Oracle Cloud Infrastructure Database Data Guard implementation requires two databases, one in a primary role and one in a standby role. The two databases

compose a Data Guard association. Most of your applications access the primary database. The standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switch or fail over the standby database to the primary role.



### Tip

The standby databases in Oracle Cloud Infrastructure Database are physical standbys.

### SWITCHOVER

A switchover reverses the primary and standby database roles. Each database continues to participate in the Data Guard association in its new role. A switchover ensures no data loss. Performing planned maintenance on a DB system with a Data Guard association is typically done by switching the primary to the standby role, performing maintenance on the standby, and then switching it back to the primary role.

### FAILOVER

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable. A failover might result in some data loss when you use **Maximum Performance** protection mode.

### REINSTATE

Reinstates a database into the standby role in a Data Guard association. You can use the `reinstat` command to return a failed database into service after correcting the cause of failure.



### Note

You can't terminate a primary database that has a Data Guard association with a peer (standby) database. Delete the standby database first. Alternatively, you can switch over the primary database to the standby role, and then terminate it.

You can't terminate a DB system that includes Data Guard enabled databases. To remove the Data Guard association:

- For a bare metal DB system database - terminate the standby database.
- For a virtual machine DB system database - terminate the standby DB system.

### Using the Console

The Console allows you to enable a Data Guard association between databases, change the role of a database in a Data Guard association using either a *switchover* or a *failover* operation, and *reinststate* a failed database.

When you enable Data Guard, a separate Data Guard association is created for the primary and the standby database.

### To enable Data Guard on a bare metal DB system

If you don't already have bare metal DB systems with the databases that will assume the primary and standby roles, create them as described in [Creating Bare Metal and Virtual Machine DB Systems](#). A new DB system includes an initial database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system with the database for which you want to enable Data Guard.
3. Click the name of the DB system that contains the database you want to assume the primary role, and then click the name of that database.



### Tip

If Data Guard is already enabled, a shield icon appears next to the database name.

4. Under **Resources**, click **Data Guard Associations**.
5. Click **Enable Data Guard**.
6. In the **Enable Data Guard** dialog box, configure your Data Guard association.
  - **Protection Mode:** (Informational) The protection mode used for this Data Guard association. The Console supports only **Maximum Performance**.
  - **Availability Domain:** The availability domain of the peer DB system. If your database is in a single availability domain region, or if you choose to provision your peer (standby) database in the same availability domain as your primary database, the system provides a fault domain selector for your peer database. Oracle recommends that your peer DB system be in a different fault domain from your primary DB system in such cases. For more information on fault domains, see [Regions and Availability Domains](#).
  - **Peer DB System:** Select the DB system that contains the peer (standby) database.
  - **Transport Type:** (Informational) The redo transport type used for this Data Guard association. The Console supports only **Async**.
  - **Database Admin Password:** Enter the primary database admin password. The same password is used for the standby database.

7. Click **Enable**.

When the association is created, a shield icon appears next to the name of this database and its peer, and their respective roles (primary or standby) are displayed.

### To enable Data Guard on a virtual machine DB system

If you don't already have a virtual machine DB system with the database that will assume the primary role, create it as described in [Creating Bare Metal and Virtual Machine DB Systems](#). The new DB system will include the initial database. When you enable Data Guard on the primary database, a new virtual machine DB system will be launched for the standby database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system with the database for which you want to enable Data Guard.
3. Click the name of the DB system that contains the database you want to assume the primary role, and then click the name of that database.



#### Tip

If Data Guard is already enabled, a shield icon appears next to the database name.

4. Under **Resources**, click **Data Guard Associations**.
5. Click **Enable Data Guard**.
6. In the **Enable Data Guard** dialog box, configure your Data Guard association.
  - Display Name:** A friendly, display name for the DB system. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the DB system.
  - Availability Domain:** The availability domain in which the DB system resides.

- **Virtual Cloud Network:** (Informational) Shows the VCN in which the DB system will be launched. The VCN of the primary database and the standby database must be the same.
- **Client Subnet:** The subnet to which the DB system should attach.  
Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet will cause the private interconnect to malfunction.
- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for your standby database's DB system. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress [security rules](#) to your DB system. A maximum of five NSGs can be specified. For more information, see [Network Security Groups](#) and [Network Setup for DB Systems](#). Note that if you choose a subnet with a [security list](#), the security rules for the DB system will be a union of the rules in the security list and the NSG.

### To use network security groups

- a. Check the **Use Network Security Groups to Control Traffic** check box. Note that you must have a virtual cloud network selected to be able to assign NSGs to your DB system.
  - b. Specify the NSG to use with the DB system. You might need to specify more than one NSG. If you're not sure, contact your network administrator.
  - c. To use additional NSGs, click **+ Another Network Security Group**.
- **Hostname Prefix:** Your choice of host name for the DB system. The host name must begin with an alphabetic character, and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for a virtual machine DB system is 16.



### Important

The host name must be unique within the subnet. If it is not unique, the DB system will fail to provision.

- **Host Domain Name:** The domain name for the DB system. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of a domain name. Hyphens (-) are not permitted.  
**Host and Domain URL:** Combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 64 characters.
- **Protection Mode:** (Informational) The protection mode used for this Data Guard association. The Console supports only **Maximum Performance**.
- **Transport Type:** (Informational) The redo transport type used for this Data Guard association. The Console supports only **Async**.
- **Database Admin Password:** Enter the primary database admin password. The same password is used for the standby database.  
**Confirm Database Admin Password:** Re-enter the Database Admin Password you specified.

#### 7. Click **Enable**.

When the association is created, a shield icon appears next to the name of this database and its peer, and their respective roles (primary or standby) are displayed.

## To perform a database switchover

You initiate a switchover operation by using the Data Guard association of the primary

database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system with the primary database you want to switch over.
3. Click the DB system name, and then click the name of the primary database.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to perform a switchover, click the Actions icon (three dots), and then click **Switchover**.
6. In the **Switchover Database** dialog box, enter the database admin password, and then click **OK**.

This database should now assume the role of the standby, and the standby should assume the role of the primary in the Data Guard association.

### To perform a database failover

You initiate a failover operation by using the Data Guard association of the standby database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system with the primary database's peer standby you want to fail over to.
3. Click the DB system name, and then click the name of the standby database.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to perform a failover, click **Failover**.
6. In the **Failover Database** dialog box, enter the database admin password, and then click **OK**.

This database should now assume the role of the primary, and the old primary's role should display as **Disabled Standby**.

### To reinstate a database

After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary by using its Data Guard association.



#### Note

Before you can reinstate a version 12.2 database, you must perform some steps on the database host to stop the database or start it in `MOUNT` mode.

Set your `ORACLE_UNQNAME` environment variable to the value of the Database Unique Name (as seen in the Console), and then run these commands:

```
srvctl stop database -d db-unique-name -o abort
srvctl start database -d db-unique-name -o mount
```

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system with the failed database you want to reinstate.
3. Click the DB system name, and then click the database name.
4. Under **Resources**, click **Data Guard Associations**.
5. For the Data Guard association on which you want to reinstate this database, click the Actions icon (three dots), and then click **Reinstate**.
6. In the **Reinstate Database** dialog box, enter the database admin password, and then click **OK**.

This database should now be reinstated as the standby in the Data Guard association.

### To terminate a Data Guard association on a bare metal DB system

On a bare metal DB system, you remove a Data Guard association by terminating the standby database.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the DB system that includes the standby database you want to terminate.
3. Click the DB system name.
4. For the standby database you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
5. In the **Terminate Database** dialog box, enter the name of the database, and then click **OK**.

### To terminate a Data Guard association on a virtual machine DB system

On a virtual machine DB system, you remove a Data Guard association by terminating the standby DB system.

1. Open the navigation menu. Under **Database**, click **Bare Metal, VM, and Exadata**.
2. Choose the **Compartment** that contains the standby DB system that you want to terminate.
3. Click the DB system name, click the Actions icon (three dots), and then click **Terminate**.
4. Confirm when prompted.  
The DB system's icon indicates Terminating.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage Data Guard associations:

- [CreateDataGuardAssociation](#)
- [GetDataGuardAssociation](#)
- [ListDataGuardAssociations](#)
- [SwitchoverDataGuardAssociation](#)
- [FailoverDataGuardAssociation](#)
- [ReinstateDataGuardAssociation](#)
- [DeleteDbHome](#) - To terminate a bare metal DB system Data Guard association, delete the standby database.
- [TerminateDbSystem](#) - To terminate a virtual machine DB system Data Guard association, terminate the standby DB system.

For the complete list of APIs for the Database service, see [Database Service API](#).

### Using Oracle Data Guard with the Database CLI

Oracle recommends that you use the Console instead of the database CLI to set up and work with Data Guard in Oracle Cloud Infrastructure. See [Using Oracle Data Guard](#) for information and instructions.



#### Note

This topic is not applicable to Exadata DB systems. You can manually configure Data Guard on Exadata DB systems using native Oracle Database utilities and commands, however this topic explains how set up primary and standby databases using `dbcli`, which is not available on Exadata DB systems. For more information, see *Data Guard Concepts and Administration* for version [18.1](#), [12.2](#), [12.1](#), or [11.2](#).

This topic explains how to use the database CLI to set up Data Guard with Fast-Start Failover (FSFO) in Oracle Cloud Infrastructure. The following sections explain how to prepare the primary and standby databases, and then configure Data Guard to transmit redo data from the primary database and apply it to the standby database.



### Note

This topic assumes that you are familiar with Data Guard and FSFO. To learn more about them, see documentation at the [Oracle Document Portal](#).

### Prerequisites

To perform the procedures in this topic, you'll need the following information for the primary and standby databases.

- db\_name (or oracle\_sid)
- db\_unique\_name
- oracle home directory (or database home)

### To find the database information

After you've launched the primary and standby DB systems and created databases as described later in this topic, you can use the CLI on those systems to find the needed database information.

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

## CHAPTER 11 Database

```
login as: opc

[opc@dbsys ~]$ sudo su -
```

3. To find the `db_name` (or `oracle_sid`) and `db_uniqueName`, run the `dbcli list-databases -j` command.

```
[root@dbsys ~]# dbcli list-databases -j
[{
 "id" : "80ad855a-5145-4f8f-a08f-406c5e4684ff",
 "name" : "dbtst",
 "dbName" : "dbtst",
 "databaseUniqueName" : "dbtst_phx1cs",
 "dbVersion" : "12.1.0.2",
 "dbHomeId" : "2efe7af7-0b70-4e9b-ba8b-71f11c6fe287",
 "instanceOnly" : false,
 .
 .
 .
}
```

4. To find the oracle home directory (or database home), run the `dbcli list-dbhomes` command. If there are multiple database homes on the DB system, use the one that matches the `dbHomeId` in the `dbcli list-databases -j` command output shown above.

```
[root@dbtst ~]# dbcli list-dbhomes
```

ID	Name	DB Version
Home Location	Status	
2efe7af7-0b70-4e9b-ba8b-71f11c6fe287	OraDB12102_home1	12.1.0.2.160719 (23739960,
23144544)	<code>/u01/app/oracle/product/12.1.0.2/dbhome_1</code>	Configured
33ae99fe-5413-4392-88da-997f3cd24c0f	OraDB11204_home1	11.2.0.4.160719 (23054319,
23054359)	<code>/u01/app/oracle/product/11.2.0.4/dbhome_1</code>	Configured

### Creating a Primary DB System

If you don't already have a primary DB system, create one as described in [Creating Bare Metal and Virtual Machine DB Systems](#). The DB system will include an initial database. You

can create additional databases by using the [dbcli create-database](#) command available on the DB system.

### Creating a Standby DB System



#### Note

The standby database must have the same `db_name` as the primary database, but it must have a different `db_unique_name`. If you use the same database name for the standby and primary, you will have to delete the database from the standby DB system by using the `dbcli delete-database` command before you can run the `dbcli create-database` command described below. Deleting and creating the database will take several minutes to complete. The `dbcli` commands must be run as the root user.

1. Create a standby DB system as described in [Creating Bare Metal and Virtual Machine DB Systems](#) and wait for the DB system to finish provisioning and become available. You can create the standby DB system in a different availability domain from the primary DB system for availability and disaster recovery purposes (this is strongly recommended). You can create the standby DB system in the primary DB system's cloud network so that both systems are in a single, routable network.
2. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

3. Log in as `opc` and then `sudo` to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the `PATH` to the `dbcli` directory (`/opt/oracle/dcs/bin`).

```
login as: opc

[opc@dbsys ~]$ sudo su -
```

4. The DB system will include an initial database, but you'll need to create a standby database by using the `dbcli create-database` command with the `--instanceonly` parameter. This parameter creates only the database storage structure and starts the database in nomount mode (no other database files are created).

When using `--instanceonly`, both the `--dbname` and `--adminpassword` parameters are required and they should match the dbname and admin password of the primary database to avoid confusion.

The following sample command prompts for the admin password and then creates a storage structure for a database named dbname.

```
[root@dbsys ~]# dbcli create-database --dbname <same as primary dbname> --databaseUniqueName
<different from primary uniqueName> --instanceonly --adminpassword
```

If you are using pluggable databases, also specify the `--cdb` parameter.

For complete command syntax, see [dbcli create-database](#).

5. Wait a few minutes for the `dbcli create-database` command to create the standby database.

You can use the `dbcli list-jobs` command to verify that the creation job ran successfully, and then the `dbcli list-databases` command verify that the database is configured.

### Preparing the Primary DB System

To prepare the primary DB system, you'll need to configure static listeners, update `tnsnames.ora`, and configure some database settings and parameters.

#### CONFIGURING THE STATIC LISTENERS

Create static listeners to be used by RMAN and Data Guard Broker.

1. SSH to the primary DB system, log in as the `opc` or `root` user, and `sudo` to the grid OS user.

## CHAPTER 11 Database

```
sudo su - grid
```

2. Edit `/u01/app/<version>/grid/network/admin/listener.ora` and add the following content to it. The first static listener shown here is optional. The second DGMGRL static listener is optional for version 12.1 or later databases, but required for version 11.2 databases.

```
SID_LIST_LISTENER=
 (SID_LIST=
 (SID_DESC=
 (SDU=65535)
 (GLOBAL_DBNAME = <primary_db_unique_name>.<primary_db_domain>)
 (SID_NAME = <primary_oracle_sid>)
 (ORACLE_HOME=<oracle_home_directory>)
 (ENVS="TNS_ADMIN=<oracle_home_directory>/network/admin")
)
 (SID_DESC=
 (SDU=65535)
 (GLOBAL_DBNAME = <primary_db_unique_name>_DGMGRL.<primary_db_domain>)
 (SID_NAME = <primary_oracle_sid>)
 (ORACLE_HOME=<oracle_home_directory>)
 (ENVS="TNS_ADMIN=<oracle_home_directory>/network/admin")
)
)
)
```

3. Save your changes and then restart the listener.

```
$ srvctl stop listener
$ srvctl start listener
```

### ADDING NET SERVICE NAMES TO TNSNAMES.ORA

As the oracle user, edit `$ORACLE_HOME/network/admin/tnsnames.ora` and add the standby database net service name to it.

```
<standby_db_unique_name> =
 (DESCRIPTION =
 (SDU=65535)
 (ADDRESS = (PROTOCOL = TCP) (HOST = <standby_server>.<domain>) (PORT = 1521))
 (CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = <standby_db_unique_name>.<standby_db_domain>)
```

```
)
)
```

The sample above assumes that name resolution is working and that the `<standby_server>.<domain>` is resolvable at the primary database. You can also use the private IP address of the standby server if the IP addresses are routable within a single cloud network (VCN).

### CONFIGURING PRIMARY DATABASE PARAMETERS



#### Tip

If the primary and standby hosts have different directory structures, you might need to set additional parameters that are not discussed here, such as the `log_file_name_convert` parameter. See the RMAN documentation for more information about how to create standbys for hosts with different directory structures.

1. As the oracle user, enable automatic standby file management.

```
SQL> alter system set standby_file_management=AUTO;
```

2. Identify the Broker configuration file names and locations. The commands used for this depend on the type of database storage. If you're not sure of the database storage type, use the [dbcli list-databases](#) command on the DB system.

For ACFS database storage, use the following commands to set the Broker configuration files.

```
SQL> alter system set dg_broker_config_file1='/u02/app/oracle/oradata/<Primary db_unique_name>/dbs/dr1<Primary db_unique_name>.dat';
SQL> alter system set dg_broker_config_file2='/u02/app/oracle/oradata/<Primary db_unique_name>/dbs/dr2<Primary db_unique_name>.dat';
```

For ASM database storage, use the following commands to set the Broker configuration files.

## CHAPTER 11 Database

```
SQL> alter system set dg_broker_config_file1='+DATA/<Primary db_unique_name>/dr1<db_unique_name>.dat';
SQL> alter system set dg_broker_config_file2='+DATA/<Primary db_unique_name>/dr2<db_unique_name>.dat';
```

3. Enable Broker DMON process for the database.

```
SQL> alter system set dg_broker_start=true;
```

4. Force database logging for all database transactions.

```
SQL> alter database force logging ;
```

5. Add Standby Redo Logs (SRLs), based on the Online Redo Logs (ORLs). On a newly launched DB system, there will be three ORLs of size 1073741824, so create four SRLs of the same size.

You can use the query below to determine the number and size (in bytes) of the ORLs.

```
SQL> select group#, bytes from v$log;
```

GROUP#	BYTES
1	1073741824
2	1073741824
3	1073741824

All of the ORLs must be the same size.

The SRLs must be the same size as the ORLs, but there must be at least one more SRL than the ORLs. In the example above, there are three ORLs, so four SRLs are required. So specify the current redo logs plus one, and use the same size as the redo logs.

```
SQL> alter database add standby logfile thread 1 size <size>;
```

There should be only one member in the SRL group (by default, a DB system is created with only one member per SRL group). To ensure this, you can name the file with the following syntax.

```
alter database add standby logfile thread 1 group 4 (<logfile name with full path>) size 1073741824, group 5 (<logfile name with full path>) size 1073741824 ...
```

For ASM/OMF configurations, the above command uses the diskgroup instead of *<logfile name with full path>*.

```
alter database add standby logfile thread 1 group 4 (+RECO) size 1073741824, group 5(+RECO) size 1073741824 ...
```



### Tip

ORLs and SRLs should be sized so that log switches do not occur more frequently than every 10 minutes. This requires knowledge of the application and may need to be adjusted after deployment. For more information, see [Use Standby Redo Logs and Configure Size Appropriately](#).

6. Verify that you created the correct number of SRLs.

```
SQL> select group#, bytes from v$standby_log;
```

7. Make sure the database is in ARCHIVELOG mode.

```
SQL> archive log list
```

8. Enable database FLASHBACK. The minimum recommended value for `db_flashback_retention_target` is 120 minutes.

```
SQL> alter database flashback on ;
SQL> alter system set db_flashback_retention_target=120;
```

9. Perform a single switch redo log to activate archiving if database is newly created. (At least one log must be archived prior to running the RMAN duplicate.)

```
SQL> alter system switch logfile;
```

### Preparing the Standby Database

Before you prepare the standby database, make sure the database home on the standby is the same version as on the primary. (If the primary and standby databases are both newly created with the same database version, the database homes will be the same.) If it is not, create a database home that is the same version. You can use the [dbcli list-dbhomes](#) command to verify the versions and the [dbcli create-dbhome](#) command to create a new database home as needed.

To prepare the standby DB system, you'll need to configure static listeners, update `tnsnames.ora`, configure TDE Wallet, create a temporary password file, verify connectivity, run RMAN DUPLICATE, enable FLASHBACK, and then create the database service.

### CONFIGURING THE STATIC LISTENERS

Create static listeners to be used by RMAN and Data Guard Broker.

1. SSH to the standby DB system, log in as the `opc` or `root` user, and `sudo` to the `grid` OS user.

```
sudo su - grid
```

2. Append the following content to `/u01/app/<db_version>/grid/network/admin/listener.ora`.

The first static listener shown below is required for RMAN DUPLICATE. The second DGMGRL static listener is optional for database versions 12.2.0.1 and 12.1.0.2, but required for database version 11.2.0.4.

```
SID_LIST_LISTENER=
 (SID_LIST=
 (SID_DESC=
 (SDU=65535)
 (GLOBAL_DBNAME = <standby db_unique_name>.<standby db_domain>)
 (SID_NAME = <standby oracle_sid>)
 (ORACLE_HOME=<oracle home directory>)
 (ENVS="TNS_ADMIN=<oracle home directory>/network/admin")
)
 (SID_DESC=
 (SDU=65535)
 (GLOBAL_DBNAME = <standby db_unique_name>_DGMGRL.<standby db_domain>)
 (SID_NAME = <standby oracle_sid>)
 (ORACLE_HOME=<oracle home directory>)
 (ENVS="TNS_ADMIN=<oracle home directory>/network/admin")
)
)
)
```

3. Restart the listener.

```
$ srvctl stop listener
$ srvctl start listener
```

4. Verify that the static listeners are available. The sample output below is for database version 12.1.0.2. Note that the ...status UNKNOWN messages are expected at this point.

```

$ lsnrctl status

LSNRCTL for Linux: Version 12.1.0.2.0 - Production on 29-SEP-2016 21:09:25

Copyright (c) 1991, 2014, Oracle. All rights reserved.

Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=LISTENER)))
STATUS of the LISTENER

Alias LISTENER
Version TNSLSNR for Linux: Version 12.1.0.2.0 - Production
Start Date 29-SEP-2016 21:09:19
Uptime 0 days 0 hr. 0 min. 5 sec
Trace Level off
Security ON: Local OS Authentication
SNMP OFF
Listener Parameter File /u01/app/12.1.0.2/grid/network/admin/listener.ora
Listener Log File /u01/app/grid/diag/tnslsnr/dg2/listener/alert/log.xml
Listening Endpoints Summary...
 (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=LISTENER)))
 (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=10.0.1.24)(PORT=1521)))
Services Summary...
Service "dg2_phx2hx.oratst.org" has 1 instance(s).
 Instance "dg2", status UNKNOWN, has 1 handler(s) for this service...
Service "dg2_phx2hx_DGMGRL.oratst.org" has 1 instance(s).
 Instance "dg2", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully

```

#### ADDING NET SERVICE NAMES TO TNSNAMES.ORA

As the oracle user, add the standby database net service name to \$ORACLE\_HOME/network/admin/tnsnames.ora. \$ORACLE\_HOME is the database home where the standby database is running.

```

<Primary db_unique_name> =
 (DESCRIPTION =
 (SDU=65535)
 (ADDRESS = (PROTOCOL = TCP)(HOST = <primary_server>.<domain>) (PORT = 1521))

```

## CHAPTER 11 Database

---

```
(CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = <primary_db_unique_name>.<primary_db_domain>)
)
)

<Standby_db_unique_name> =
(DESCRIPTION =
(SDU=65535)
(ADDRESS = (PROTOCOL = TCP)(HOST = <standby_server>.<domain>) (PORT = 1521))
(CONNECT_DATA =
 (SERVER = DEDICATED)
 (SERVICE_NAME = <standby_db_unique_name>.<db_domain>)
)
)
```

### COPYING THE TDE WALLETS TO THE STANDBY SYSTEM

Copy the TDE wallet files from the primary DB system to standby DB system using SCP. The following sample command assumes the SCP command is being run by the oracle OS user and that the private key for oracle has been created and exists on the host where SCP is being run.

```
$ scp -i <private key> primary_server:/opt/oracle/dcs/commonstore/wallets/tde/<primary_db_unique_name>/*
standby_server:/opt/oracle/dcs/commonstore/wallets/tde/<standby_db_unique_name>
```

### SETTING UP THE STANDBY SYSTEM CONFIGURATION

As the oracle user, create the following directory for database version 11.2.0.4. This step is optional for version 12.2.0.1 and version 12.1.0.2.

```
[oracle@dbsys ~]$ mkdir -pv /u03/app/oracle/redo/<standby_db_unique_name uppercase>/controlfile
```

### CREATING THE AUDIT FILE DESTINATION

As the oracle user, create the following directory to use as the audit file destination.

```
[oracle@dbsys ~]$ mkdir -p /u01/app/oracle/admin/<db_name>/adump
```

Otherwise, the RMAN duplicate command used later will fail.

### CREATING A TEMPORARY PASSWORD FILE

As the oracle user, create a temporary password file.

## CHAPTER 11 Database

```
[oracle@dbsys ~]$ orapwd file=$ORACLE_HOME/dbs/orapw<standby oracle_sid> password=<admin password for primary> entries=5
```

The password **must** be the same as the admin password of the primary database. Otherwise, the RMAN duplicate step below will fail with: RMAN-05614: Passwords for target and auxiliary connections must be the same when using active duplicate.

### VERIFYING THE STANDBY DATABASE IS AVAILABLE

1. As the oracle user, set the environment variables.

```
[oracle@dbsys ~]$. oraenv
<enter the db_name>
```

2. Replace \$ORACLE\_HOME/dbs/init<standby sid\_name>.ora with the following content:

```
db_name=<Primary db_name>
db_unique_name=<standby db_unique_name>
db_domain=<standby db_domain>
```

3. Remove the spfile from the standby.

```
/u02/app/oracle/oradata/<standby db_unique_name>/dbs/spfile$ORACLE_SID.ora
```

The database needs to be started in nomount mode with no spfile specified, but the original init file contains an spfile parameter which will prevent the RMAN duplicate step from working.

4. Set the ORACLE\_UNQNAME environment variable to point to your DB\_UNIQUE\_NAME.

```
$ export ORACLE_UNQNAME =db_unique_name
```



#### Important

If you do not perform this step, the wallet will not be opened, and running the RMAN DUPLICATE command in the subsequent step will fail.

5. The dbcli create-database --instanceonly command used earlier opens the

## CHAPTER 11 Database

---

standby database as a primary in read/write mode, so the database needs to be brought down before proceeding to the nomount step below.

```
$ sqlplus / as sysdba
SQL> shutdown immediate
```

6. Start the database in nomount mode.

```
SQL> startup nomount
```

### VERIFYING THE DATABASE CONNECTIONS

Verify the connection between the primary and standby databases.

1. Make sure that the listener port 1521 is open in the security list(s) used for the primary and standby DB systems. For more information, see [Updating the Security List for the DB System](#).
2. From the primary database, connect to standby database.

```
$ sqlplus sys/<password>@<standby net service name> as sysdba
```

3. From standby database, connect to primary database.

```
$ sqlplus sys/<password>@<primary net service name> as sysdba
```

### RUNNING THE RMAN DUPLICATE COMMAND

Run the RMAN DUPLICATE command on the standby DB system, as the oracle user.

If the primary database is large, you can allocate additional channels to improve performance. For a newly installed database, one channel typically runs the database duplication in a couple of minutes.

Make sure that there are no errors generated by the RMAN DUPLICATE command. If errors occur, restart the database using the `init.ora` file (not `spfile`) in case it is generated under `$ORACLE_HOME/dbs` as part of RMAN DUPLICATE.

In the following examples, use lowercase for the `<Standby db_unique_name>` unless otherwise specified.

For ACFS storage layout, run the following commands.

## CHAPTER 11 Database

```
$ rman target sys/<password>@<primary alias> auxiliary sys/<password>@<standby alias> log=rman.out

RMAN> run { allocate channel prim1 type disk;
 allocate auxiliary channel sby type disk;
 duplicate target database for standby from active database
 dorecover
 spfile
 parameter_value_convert '/<Primary db_unique_name>/','/<Standby db_unique_name>/','/<Primary
db_unique_name uppercase>/','/<Standby db_unique_name uppercase >/'
 set db_unique_name='<Standby db_unique_name>'
 set db_create_file_dest='/u02/app/oracle/oradata/<Standby db_unique_name>'
 set dg_broker_config_file1='/u02/app/oracle/oradata/<Standby db_unique_name>/dbs/dr1<Standby
db_unique_name>.dat'
 set dg_broker_config_file2='/u02/app/oracle/oradata/<Standby db_unique_name>/dbs/dr2<Standby
db_unique_name>.dat'
 set dispatchers ='(PROTOCOL=TCP) (SERVICE=<Standby db_unique_name>XDB) '
 set instance_name='<Standby db_unique_name>'
 ;
 }
```

For ASM storage layout, run the following commands.

```
$ rman target sys/<password>@<primary alias> auxiliary sys/<password>@<standby alias> log=rman.out

RMAN> run { allocate channel prim1 type disk;
 allocate auxiliary channel sby type disk;
 duplicate target database for standby from active database
 dorecover
 spfile
 parameter_value_convert '/<Primary db_unique_name>/','/<Standby db_unique_name>/','/<Primary
db_unique_name uppercase>/','/<Standby db_unique_name uppercase>/'
 set db_unique_name='<Standby db_unique_name>'
 set dg_broker_config_file1='+DATA/<Standby db_unique_name>/dr1<Standby db_unique_name>.dat'
 set dg_broker_config_file2='+DATA/<Standby db_unique_name>/dr2<Standby db_unique_name>.dat'
 set dispatchers ='(PROTOCOL=TCP) (SERVICE=<Standby db_unique_name>XDB) '
 set instance_name='<Standby db_unique_name>'
 ;
 }
```

## CHAPTER 11 Database

---

### ENABLING DATABASE FLASHBACK

1. As a Data Guard best practice, enable flashback and set `db_flashback_retention_target` to at least 120 minutes on both the primary and standby databases.

```
SQL> alter database flashback on;
SQL> alter system set db_flashback_retention_target=120;
```

2. Verify that the standby database is created properly.

```
SQL> select FORCE_LOGGING, FLASHBACK_ON, OPEN_MODE, DATABASE_ROLE, SWITCHOVER_STATUS, DATAGUARD_
BROKER, PROTECTION_MODE from v$database ;
```

### CREATING A DATABASE SERVICE

Oracle recommends creating a database service for the standby database by using `srvctl`. For ACFS storage layout.

1. Create a shared directory and copy the spfile file to it.

```
$ mkdir -pv /u02/app/oracle/oradata/<Standby db_unique_name>/dbs
$ cp $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora /u02/app/oracle/oradata/<Standby db_unique_
name>/dbs
```

2. Stop and remove the existing database service.

```
$ srvctl stop database -d <standby db_unique_name>
$ srvctl remove database -d <standby db_unique_name>
```

3. Create the database service.

```
$ srvctl add database -d <standby db_unique_name> -n <standby db_name> -o $ORACLE_HOME -c SINGLE
-p '/u02/app/oracle/oradata/<standby db_unique_name>/dbs/spfile<standby db_name>.ora'
-x <standby hostname> -s "READ ONLY" -r PHYSICAL_STANDBY -i <db_name>
$ srvctl setenv database -d <standby db_unique_name> -t "ORACLE_
UNQNAME=<standby db_unique_name>"
$ srvctl config database -d <standby db_unique_name>
```

4. Start the database service.

```
$ srvctl start database -d <standby db_unique_name>
```

5. Clean up the files from `$ORACLE_HOME/dbs`.

## CHAPTER 11 Database

```
$ rm $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora
$ rm $ORACLE_HOME/dbs/init<standby oracle_sid>.ora
```

6. Create the `$ORACLE_HOME/dbs/init<standby oracle_sid>.ora` file to reference the new location of the spfile file.

```
SPFILE='/u02/app/oracle/oradata/<standby db_unique_name>/dbs/spfile<standby db_name>.ora'
```

7. Stop the standby database and then start it by using `srvctl`.

```
srvctl stop database -d <standby db_unique_name>
srvctl start database -d <standby db_unique_name>
```

For ASM storage layout.

1. Consider generating the spfile file under `+DATA`.

```
SQL> create pfile='init<standby oracle_sid>.ora' from spfile ;
SQL> create spfile='+DATA' from pfile='init<standby oracle_sid>.ora' ;
```

2. Stop and remove the existing database service.

```
$ srvctl stop database -d <standby db_unique_name>
$ srvctl remove database -d <standby db_unique_name>
```

3. Create the database service.

```
$ srvctl add database -d <standby db_unique_name> -n <standby db_name> -o $ORACLE_HOME -c
SINGLE -p '+DATA/<standby db_unique_name>/PARAMETERFILE/spfile.xxx.xxxxxx'
-x <standby hostname> -s "READ ONLY" -r PHYSICAL_STANDBY -i <db_name>
$ srvctl setenv database -d <standby db_unique_name> -t "ORACLE_UNQNAME=<standby db_unique_name>"
$ srvctl config database -d <standby db_unique_name>
```

4. Start the database service.

```
$ srvctl start database -d <standby db_unique_name>
```

5. Clean up the files from `$ORACLE_HOME/dbs`.

```
$ rm $ORACLE_HOME/dbs/init<standby oracle_sid>.ora
$ rm $ORACLE_HOME/dbs/spfile<standby oracle_sid>.ora
```

6. Create `$ORACLE_HOME/dbs/init<standby oracle_sid>.ora` file to reference the new location of the spfile file.

```
SPFILE='+DATA/<standby db_unique_name>/PARAMETERFILE/spfile.xxx.xxxxxx'
```

7. Stop the database and start the standby database by using srvctl.

```
$ srvctl start database -d <standby db_unique_name>
```

### Configuring Data Guard

Perform the following steps to complete the configuration of Data Guard and enable redo transport from the primary database and redo apply in the standby database.

1. Run the dgmgrl command line utility from either the primary or standby DB system and connect to the primary database using sys credentials.

```
DGMGRL> connect sys/<sys_password>@<primary tns alias>
```

2. Create the Data Guard configuration and identify for the primary and standby databases.

```
DGMGRL> create configuration mystby as primary database is <primary db_unique_name> connect
identifier is <primary tns alias>;
add database <standby db_unique_name> as connect identifier is <standby tns alias> maintained
as physical;
```

3. Enable Data Guard configuration.

```
DGMGRL> enable configuration;
```

4. Verify that Data Guard setup was done properly. Run the following SQL in **both** the primary and standby databases.

```
SQL> select FORCE_LOGGING, FLASHBACK_ON, OPEN_MODE, DATABASE_ROLE, SWITCHOVER_STATUS, DATAGUARD_
BROKER, PROTECTION_MODE from v$database;
```

5. Verify that Data Guard processes are initiated in the standby database.

```
SQL> select PROCESS,PID,DELAY_MINS from V$MANAGED_STANDBY;
```

6. Verify parameter configuration on primary and standby.

```
SQL> show parameter log_archive_dest_
SQL> show parameter log_archive_config
SQL> show parameter fal_server
SQL> show parameter log_archive_format
```

7. Verify that the Data Guard configuration is working. Specifically, make sure redo

shipping and redo apply are working and that the standby is not unreasonably lagging behind the primary.

```
DGMGRL> show configuration verbose
DGMGRL> show database verbose <standby db_unique_name>
DGMGRL> show database verbose <primary db_unique_name>
```

Any discrepancies, errors, or warnings should be resolved. You can also run a transaction on the primary and verify that it's visible in the standby.

8. Verify that the Data Guard configuration is functioning as expected by performing switchover and failover in both directions. Run `show configuration` after each operation and make sure there are no errors or warnings.



### Warning

This step is optional, based on your discretion. If for any reason the configuration is not valid, the switchover and/or failover will fail and it might be difficult or impossible to start the primary database. A recovery of the primary might be required, which will affect availability.

```
DGMGRL> switchover to <standby db_unique_name>
DGMGRL> switchover to <primary db_unique_name>

#connect to standby before failover:

DGMGRL> connect sys/<sys password>@<standby db_unique_name>
DGMGRL> failover to <standby db_unique_name>
DGMGRL> reinstate database <primary db_unique_name>

#connect to primary before failover:

DGMGRL> connect sys/<sys password>@<primary db_unique_name>
DGMGRL> failover to <primary db_unique_name>
DGMGRL> reinstate database <standby db_unique_name>
```

## Configuring Observer (Optional)

The best practice for high availability and durability is to run the primary, standby, and observer in separate availability domains. The observer determines whether or not to failover to a specific target standby database. The server used for observer requires the Oracle Client Administrator software, which includes the Oracle SQL NET and Broker.

1. Configure TNS alias names for both the primary and standby databases as described previously, and verify the connection to both databases.
2. Change protection mode to either maxavailability or maxperformance (maxprotection is not supported for FSFO).

To enable maxavailability:

```
DGMGRL> edit database <standby db_unique_name> set property 'logXptMode'='SYNC';
DGMGRL> edit database <primary db_unique_name> set property 'logXptMode'='SYNC';
DGMGRL> edit configuration set protection mode as maxavailability;
```

To enable maxperformance:

```
DGMGRL> edit configuration set protection mode as maxperformance;
DGMGRL> edit database <standby db_unique_name> set property 'logXptMode'='ASYNC';
DGMGRL> edit database <primary db_unique_name> set property 'logXptMode'='ASYNC';
```

For maxperformance, the FastStartFailoverLaglimit property limits the maximum amount of permitted data loss to 30 seconds by default.

3. The following properties should also be considered. Run `show configuration verbose` to see their current values.
  - FastStartFailoverPmyShutdown
  - FastStartFailoverThreshold
  - FastStartFailoverTarget
  - FastStartFailoverAutoReinstate

(Running `show configuration` will result in the following error until the observer is started: Warning : ORA-16819: fast-start failover observer not started.)

4. Enable fast-start failover from Broker:

```
DGMGRL> Enable fast_start failover
```

5. Verify the fast-start failover and associated settings.

```
DGMGRL> show fast_start failover
```

6. Start the observer from Broker (it will run in the foreground, but can also be run in the background).

```
DGMGRL> start observer
```

7. Verify fast-start failover is enabled and without errors or warnings.

```
DGMGRL> show configuration verbose
```

8. Always test failover in both directions to ensure that everything is working as expected. Verify that FSFO is running properly by performing a shutdown abort of the primary database.

The observer should start the failover to the standby database. If protection mode is set to maxprotection, some loss of data can occur, based on the FastStartFailoverLaglimit value.

## Oracle Database CLI Reference

The database CLI (dbcli) is a command line interface available on bare metal and virtual machine DB systems. After you connect to the DB system, you can use the database CLI to perform tasks such as creating Oracle database homes and databases.



### Note

The database CLI is **not** for use on Exadata DB systems.

### Operational Notes

- The database CLI commands must be run as the root user.
- dbcli is in the `/opt/oracle/dcs/bin/` directory.  
This directory is included in the path for the root user's environment.

- Oracle Database maintains logs of the `dbcli` command output in the `dcsccli.log` and `dcsc-agent.log` files in the `/opt/oracle/dcs/log/` directory.
- The database CLI commands and most parameters are case sensitive and should be typed as shown. A few parameters are not case sensitive, as indicated in the parameter descriptions, and can be typed in uppercase or lowercase.



### Warning

Oracle recommends that you avoid specifying parameter values that include confidential information when you use the database CLI commands.

## Syntax

The database CLI commands use the following syntax:

```
dbcli command [parameters]
```

where:

- `command` is a verb-object combination such as `create-database`.
- `parameters` include additional options for the command. Most parameter names are preceded with two dashes, for example, `--help`. Abbreviated parameter names are preceded with one dash, for example, `-h`.
- User-specified parameter values are shown in red text within angle brackets, for example, `<db_home_id>`. Omit the angle brackets when specifying these values.
- The help parameter is available with every command.

The remainder of this topic contains syntax and other details about the commands.

## CLI Update Command

Occasionally, new commands are added to the database CLI and other commands are updated to support new features. You can use the following command to update the database CLI:

## CHAPTER 11 Database

### CLIADM UPDATE-DBCLI

Use the `cliadm update-dbcli` command to update the database CLI with the latest new and updated commands.



#### Note

The `cliadm update-dbcli` command is not available on 2-node RAC DB systems.

### SYNTAX

```
cliadm update-dbcli [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command updates the `dbcli`:

```
[root@dbsys ~]# cliadm update-dbcli
{
 "jobId" : "dc9ce73d-ed71-4473-99cd-9663b9d79bfd",
 "status" : "Created",
 "message" : "Dcs cli will be updated",
 "reports" : [],
 "createTimestamp" : "January 18, 2017 10:19:34 AM PST",
 "resourceList" : [],
 "description" : "dbcli patching",
 "updatedAt" : "January 18, 2017 10:19:34 AM PST"
}
```

### Agent Commands

The following commands are available to manage agents:

- [dbcli ping-agent](#)
- [dbcli list-agentConfigParameters](#)
- [dbcli update-agentConfigParameters](#)

#### DBCLI PING-AGENT

Use the `dbcli ping-agent` command to test the reachability of an agent.

##### SYNTAX

```
dbcli ping-agent [-h] [-j]
```

##### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### DBCLI LIST-AGENTCONFIGPARAMETERS

Use the `dbcli list-agentConfigParameters` command to list agent configuration parameters.

##### SYNTAX

```
dbcli list-agentConfigParameters [-n] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	-name	(Optional) Parameter name.

### DBCLI UPDATE-AGENTCONFIGPARAMETERS

Use the `dbcli update-agentConfigParameters` command to update agent configuration parameters.

### SYNTAX

```
dbcli update-agentConfigParameters -n <parameter> [-v <value>] [-a] [-c] [-d] [-u] [-r] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-a	--append	(Optional) Appends the specified values to the specified parameters. Example with multiple parameter names and values: <code>-n p1 -v v1 -n p2 -v v2 -a</code>
-c	--comment	(Optional) Adds a comment for the parameter. Default: [ ]
-d	--description	(Optional) Adds a description for the parameter. Default: [ ]
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	Parameter name. Example with multiple parameter names and values: <code>-n p1 -v v1 -n p2 -v v2</code> Default: [ ]

Parameter	Full Name	Description
-r	--reset	(Optional) Resets the parameter to the default value. Example resetting multiple parameters: -n p1 -n p2 -r Default: false
-u	--update	(Optional) Replaces the specified parameter values as directed. Example with multiple parameter names and values: -n p1 -v v1 -n p2 -v v2 -u Default: false
-v	--value	(Optional) Parameter value. Example with multiple parameter names and values: -n p1 -v v1 -n p2 -v v2 Default: [ ]

### Autologcleanpolicy Commands

The following commands are available to manage policies for automatic cleaning (purging) of logs.

- [dbcli create-autoLogCleanPolicy](#)
- [dbcli list-autoLogCleanPolicy](#)

#### DBCLI CREATE-AUTOLOGCLEANPOLICY

Use the `dbcli create-autoLogCleanPolicy` command to create policies for automatic cleaning (purging) of logs.

#### SYNTAX

```
dbcli create-autoLogCleanPolicy [-c {gi|database|dcs}] [-f <number>] [-o <number>] [-u {Day|Hour|Minute}] [-uMB <number>] [-uPer <number>] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-c	--components	(Optional) Components to purge. Possible values are <code>gi</code> , <code>database</code> , and <code>dcs</code> . Separate multiple values with commas. Example: <code>gi,dcs</code>
-f	--freeSpaceBelowPercentage	(Optional) Purges logs when the free disk space is below the specified percentage of the total partition size. Valid range: 20-50. Default: 20.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-o	--olderthan	(Optional) Quantity portion of time interval. Default: 30. Cleans logs older than the specified time interval ( <code>-o</code> and <code>-u</code> ).
-u	--olderThanUnit	(Optional) Unit portion of time interval. Possible values: <code>Day</code> , <code>Hour</code> , or <code>Minute</code> . Default: <code>Day</code> . Cleans logs older than the specified time interval ( <code>-o</code> and <code>-u</code> ).
-uMB	--usageOverMB	(Optional) Purges logs when log usage exceeds the specified number of MegaBytes (MB). Valid range: 10 to 50% of total partition size.
-uPer	--usageOverPercentage	(Optional) Purges logs when log usage exceeds the specified percentage of the total partition size. Valid range: 10-50.

## CHAPTER 11 Database

---

### DBCLI LIST-AUTOLOGCLEANPOLICY

Use the `dbcli list-autoLogCleanPolicy` command to list policies for automatic cleaning of logs.

#### SYNTAX

```
dbcli list-autoLogCleanPolicy [-c {gi|database|dcs}] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-c	--components	(Optional) Components. Possible values are <code>gi</code> , <code>database</code> , and <code>dcs</code> . Separate multiple values with commas. Example: <code>gi,dcs</code>
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Backup Commands

The following commands are available to back up databases:

- [dbcli create-backup](#)
- [dbcli getstatus-backup](#)
- [dbcli schedule-backup](#)



### Note

Instead of using `dbcli`, you can use the Console or the API to manage backing up your bare metal or virtual machine DB system databases to Object Storage. However, if you switch from using `dbcli` to using managed backups, a new backup configuration is created and associated with your database, and backups you created by using `dbcli` will not be accessible from the managed backup interfaces. For information about managed backups, see [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#).

Before you can back up a database by using the [dbcli create-backup](#) command, you'll need to:

1. Create a backup configuration by using the [dbcli create-backupconfig](#) command.
2. Associate the backup configuration with the database by using the [dbcli update-database](#) command.

After a database is associated with a backup configuration, you can use the `dbcli create-backup` command in a `cron` job to run backups automatically. You can use a cron utility such as CronMaker to help build expressions. For more information, see <http://www.cronmaker.com>.

### DBCLI CREATE-BACKUP

Use the `dbcli create-backup` command to create a backup of a database.

#### SYNTAX

```
dbcli create-backup -in <db_name> -i <db_id> [-bt {Regular-L0|Regular-L1|Longterm|ArchiveLog}] [-c {Database|TdeWallet}] [-k <n>] [-t <tag>] [-h] [-j]
```

## CHAPTER 11 Database

---

### *PARAMETERS*

<b>Parameter</b>	<b>Full Name</b>	<b>Description</b>
-bt	-- backupType	(Optional) Backup type. Possible values are Regular-L0, Regular-L1, Longterm, and ArchiveLog. Regular-L0 and Regular L1 correspond to incremental L0 and L1 backups. Longterm corresponds to Full backup. ArchiveLog corresponds to archived redo logs backup. The default value is Regular-L1. Values are not case-sensitive. If omitted, the default value is used.

Parameter	Full Name	Description
-c	-- component	<p>(Optional) Component. Possible values are Database and TdeWallet. The default value is Database. The value TdeWallet backs up TDE wallets. Values are not case-sensitive. If omitted, the default value is used.</p> <div data-bbox="704 617 1291 1556" style="background-color: #e1f5fe; padding: 10px; border: 1px solid #cfcfcf;">  <p><b>Note</b></p> <p>TDE wallets are automatically backed up in the following situations:</p> <ul style="list-style-type: none"> <li>• A database is created with an Object Storage backup configuration.</li> <li>• A database that has an Object Storage backup configuration is updated.</li> <li>• An Object Storage backup configuration is updated.</li> <li>• A backup of the type Longterm is created.</li> <li>• The TDE key for a database is rotated.</li> <li>• A database is backed up and no TDE wallet backups exist yet.</li> </ul> </div>
-h	--help	(Optional) Displays help for using the command.

## CHAPTER 11 Database

Parameter	Full Name	Description
-i	--dbid	The ID of the database to back up. Use the <code>dbcli list-databases</code> command to get the database's ID.
-in	--dbName	The name of the database to back up. Use the <code>dbcli list-databases</code> command to get the database's name.
-j	--json	(Optional) Displays JSON output.
-k	--keepDays	(Optional) Specifies the time until which the backup or copy must be kept. After this time the backup is obsolete, regardless of the backup retention policy settings. For Longterm backup type only.
-t	--tag	(Required for Longterm backup type) Specifies a user-specified tag name for a backup set and applies this tag to the output files generated by the command. This value is not case sensitive. Valid number of characters: 1 to 30. The characters are limited to the characters that are valid in file names on the target file system. For example, ASM does not support the use of the hyphen (-) character in the file names it uses internally, so <code>weekly-incremental</code> is not a valid tag name for backups in ASM disk groups. Environment variables are not valid in the TAG parameter.

### EXAMPLES

The following command creates a backup of the specified database using the database ID.

```
[root@dbsys ~]# dbcli create-backup -i 573cadb2-0cc2-4c1c-9c31-595ab8963d5b
```

The following command creates a backup of the specified database using the database name ("mydb").

```
[root@dbsys ~]# dbcli create-backup -in mydb
```

## CHAPTER 11 Database

---

### DBCLI GETSTATUS-BACKUP

Use the `dbcli getstatus-backup` command to display the status of a backup.

#### SYNTAX

```
dbcli getstatus-backup -t <backup_type> [i <id>] [-in <name>] [-l] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbId	(Optional) Database Resource ID.
-in	--dbName	(Optional) Database Resource Name.
-j	--json	(Optional) Displays JSON output.
-l	--isLatestBackupReport	(Optional) Latest backup report. Default: true.
-t	--backupType	Backup type.

### DBCLI SCHEDULE-BACKUP

Use the `dbcli schedule-backup` command to schedule a backup of a database.

#### SYNTAX

```
dbcli schedule-backup -t <backup_type> -f <number> [i <id>] [-in <name>] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-f	--frequency	Frequency in minutes.
-h	--help	(Optional) Displays help for using the command.

Parameter	Full Name	Description
-i	--dbId	(Optional) Database Resource ID.
-in	--dbName	(Optional) Database Resource Name.
-j	--json	(Optional) Displays JSON output.
-t	--backupType	Backup type.

### Backupconfig Commands

A backup configuration determines the backup destination and recovery window for database backups. You create the backup configuration and then associate it with a database by using the `dbcli update-database` command.



#### Warning

Backups that were configured using the Console may become unusable if you make changes using these commands. For backups configured using the Console, use these commands with support guidance only.



#### Note

Instead of using `dbcli`, you can use the Console or the API to manage backing up your bare metal or virtual machine DB system databases to Object Storage. For information about managed backups, see [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#).

After a database is associated with a backup configuration, you can use the `dbcli create-backup` command in a `cron` job to run backups automatically. You can use a cron utility such

as CronMaker to help build expressions. For more information, see <http://www.cronmaker.com>.

The following commands are available to manage backup configurations:

- [dbcli create-backupconfig](#)
- [dbcli list-backupconfigs](#)
- [dbcli describe-backupconfig](#)
- [dbcli update-backupconfig](#)
- [dbcli delete-backupconfig](#)

### DBCLI CREATE-BACKUPCONFIG

Use the `dbcli create-backupconfig` command to create a backup configuration that defines the backup destination and recovery windows.

#### SYNTAX

```
dbcli create-backupconfig -d {DISK|OBJECTSTORE|NONE} -c <bucket> -o <object_store_swift_id> -on
<object_store_swift_name> -w <n> -n <name> [-cr|-no-cr] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-c	--container	The name of an existing bucket in the Oracle Cloud Infrastructure Object Storage service. You can use the Console or the Object Storage API to create the bucket. For more information, see <a href="#">Managing Buckets</a> .  You must also specify <code>--backupdestination objectstore</code> and the <code>--objectstoreswiftId</code> parameter.
-cr -no-cr	--crosscheck --no-crosscheck	(Optional) Indicates whether to enable the crosscheck operation. This operation determines if the files on the disk or in the media management catalog correspond to data in the RMAN repository. If omitted, the default setting is used (crosscheck is enabled by default).
-d	--backupdestination	The backup destination as one of the following (these values are <b>not</b> case sensitive):  DISK - The local Fast Recovery Area.  OBJECTSTORE - The Oracle Cloud Infrastructure Object Storage service. You must also specify the <code>--container</code> and <code>--objectstoreswiftId</code> parameters.  NONE - Disables the backup.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	The name of the backup configuration.

Parameter	Full Name	Description
-o	--objectstoreswiftId	<p>The ID of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the <a href="#">dbcli list-objectstoreswifts</a> command to get the object store ID. Use the <a href="#">dbcli create-objectstoreswift</a> command to create an object store.</p> <p>You must also specify <code>--backupdestination objectstore</code> and the <code>--container</code> parameter.</p>
-on	--objectstoreswiftName	<p>The name of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the <a href="#">dbcli list-objectstoreswifts</a> command to get the object store name. Use the <a href="#">dbcli create-objectstoreswift</a> command to create an object store.</p> <p>You must also specify <code>--backupdestination objectstore</code> and the <code>--container</code> parameter.</p>
-w	--recoverywindow	<p>The number of days for which backups and archived redo logs are maintained. The interval always ends with the current time and extends back in time for the number of days specified.</p> <p>For a DISK backup destination, specify 1 to 14 days.</p> <p>For an OBJECTSTORE backup destination, specify 1 to 30 days.</p>

*EXAMPLE*

The following command creates a backup configuration named `dbbkcfg1`:

```
[root@dbsys ~]# dbcli create-backupconfig -d Disk -w 7 -n dbbkcfg1
{
```

## CHAPTER 11 Database

```
"jobId" : "4e0e6011-db53-4142-82ef-eb561658a0a9",
"status" : "Success",
"message" : null,
"reports" : [{
 "taskId" : "TaskParallel_919",
 "taskName" : "persisting backup config metadata",
 "taskResult" : "Success",
 "startTime" : "November 18, 2016 20:21:25 PM UTC",
 "endTime" : "November 18, 2016 20:21:25 PM UTC",
 "status" : "Success",
 "taskDescription" : null,
 "parentTaskId" : "TaskSequential_915",
 "jobId" : "4e0e6011-db53-4142-82ef-eb561658a0a9",
 "tags" : [],
 "reportLevel" : "Info",
 "updatedAt" : "November 18, 2016 20:21:25 PM UTC"
}],
"createTimestamp" : "November 18, 2016 20:21:25 PM UTC",
"description" : "create backup config:dbbkcfg1",
"updatedAt" : "November 18, 2016 20:21:25 PM UTC"
}
```

### DBCLI LIST-BACKUPCONFIGS

Use the `dbcli list-backupconfigs` command to list all the backup configurations in the DB system.

#### SYNTAX

```
dbcli list-backupconfigs [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command lists a backup configuration:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli list-backupconfigs
```

ID	Name	RecoveryWindow	BackupDestination
ccdd56fe-a40b-4e82-b38d-5f76c265282d	dbbkcfg1	7	Disk

-----  
-----  
CreateTime  
-----  
-----  
10, 2016 12:24:08 PM UTC

### DBCLI DESCRIBE-BACKUPCONFIG

Use the `dbcli describe-backupconfig` command to show details about a specific backup configuration.

#### SYNTAX

```
dbcli describe-backupconfig -i <id> -in <name> [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--backupconfigid	The backup configuration ID. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-in	--backupconfigname	The backup configuration name. Use the <code>dbcli list-backupconfigs</code> command to get the name.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays details about a backup configuration:

```
[root@dbsys ~]# dbcli describe-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d
```

```
Backup Config details

```

## CHAPTER 11 Database

```
ID: ccdd56fe-a40b-4e82-b38d-5f76c265282d
Name: dbbkcfg1
RecoveryWindow: 7
BackupDestination: Disk
CreatedTime: July 10, 2016 12:24:08 PM UTC
UpdatedTime: July 10, 2016 12:24:08 PM UTC
```

### DBCLI UPDATE-BACKUPCONFIG

Use the `dbcli update-backupconfig` command to update an existing backup configuration.

#### SYNTAX

```
dbcli update-backupconfig -i <id> -in <name> -w <n> -c <bucket> -o <object_store_swift_id> -on <object_
store_swift_name> [-cr|-no-cr] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-c	--container	The name of an existing bucket in the Oracle Cloud Infrastructure Object Storage service. You can use the Console or the Object Storage API to create the bucket. For more information, see <a href="#">Managing Buckets</a> .  You must also specify <code>--backupdestination objectstore</code> and the <code>--objectstoreswiftId</code> parameter.
-cr -no-cr	--crosscheck --no-crosscheck	(Optional) Indicates whether to enable the crosscheck operation. This operation determines if the files on the disk on in the media management catalog correspond to data in the RMAN repository. If omitted, the default setting is used (crosscheck is enabled by default).
-h	--help	(Optional) Displays help for using the command.

Parameter	Full Name	Description
-i	--backupconfigid	The ID of the backup configuration to update. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-in	--backupconfigname	The name of the backup configuration to update. Use the <code>dbcli list-backupconfigs</code> command to get the name.
-j	--json	(Optional) Displays JSON output.
-o	--objectstoreswiftId	The ID of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the <a href="#">dbcli list-objectstoreswifts</a> command to get the object store ID. Use the <a href="#">dbcli create-objectstoreswift</a> command to create an object store.  You must also specify <code>--backupdestination objectstore</code> and the <code>--container</code> parameter.
-on	--objectstoreswiftname	The name of the object store that contains the endpoint and credentials for the Oracle Cloud Infrastructure Object Storage service. Use the <a href="#">dbcli list-objectstoreswifts</a> command to get the object store name. Use the <a href="#">dbcli create-objectstoreswift</a> command to create an object store.  You must also specify <code>--backupdestination objectstore</code> and the <code>--container</code> parameter.
-w	--recoverywindow	The new disk recovery window.  For a DISK backup destination, specify 1 to 14 days.  For an OBJECTSTORE backup destination, specify 1 to 30 days.

## CHAPTER 11 Database

### EXAMPLE

The following command updates the recovery window for a backup configuration:

```
[root@dbsys ~]# dbcli update-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d -w 5
{
 "jobId" : "0e849291-e1e1-4c7a-8dd2-62b522b9b807",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1468153731699,
 "description" : "update backup config: dbbkcfg1",
 "updatedAt" : 1468153731700
}
```

### DBCLI DELETE-BACKUPCONFIG

Use the `dbcli delete-backupconfig` command to delete a backup configuration.

### SYNTAX

```
dbcli delete-backupconfig -i <id> -in <name> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	The backup configuration ID to delete. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-in	--backupconfigname	The name of the backup configuration to delete. Use the <code>dbcli list-backupconfigs</code> command to get the name.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command deletes the specified backup configuration:

```
[root@dbsys ~]# dbcli delete-backupconfig -i ccdd56fe-a40b-4e82-b38d-5f76c265282d
```

### Bmccredential Commands

The following commands are available to manage credentials configurations, which are required for downloading DB system patches from the Oracle Cloud Infrastructure Object Storage service. For more information, see [Patching a DB System](#).

- [dbcli create-bmccredential](#)
- [dbcli list-bmccredentials](#)
- [dbcli describe-bmccredential](#)
- [dbcli delete-bmccredential](#)
- [dbcli update-bmccredential](#)



#### Note

The `bmccredential` commands are not available on 2-node RAC DB systems.

#### DBCLI CREATE-BMCCREDENTIAL

Use the `dbcli create-bmccredential` command to create a credentials configuration.

#### PREREQUISITES

Before you can create a credentials configuration, you'll need these items:

- An RSA key pair **in PEM format** (minimum 2048 bits). See [How to Generate an API Signing Key](#).
- The fingerprint of the public key. See [How to Get the Key's Fingerprint](#).
- Your tenancy's OCID and user name's OCID. See [Where to Get the Tenancy's OCID and User's OCID](#).

Then you'll need to upload the public key in the Console. See [How to Upload the Public Key](#).

## CHAPTER 11 Database

### SYNTAX

```
dbcli create-bmccredential -c [backup|patching|other] -t <tenant_ocid> -u <user_ocid> -f <fingerprint>
-k <private_key_path> -p|-hp <passphrase> [-n <credentials_name>] [-e <object_store_url>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-c	--credentialsType	The type of Object Storage credentials configuration to create (these values are <b>not</b> case sensitive):  BACKUP - Reserved for the future use.  PATCHING - For downloading patches from the service.  OTHER - Reserved for the future use.
-e	--objectStoreUrl	(Optional) The Object Storage endpoint URL.  Omit this parameter when --credentialsType PATCHING is specified. The following URL is assumed:  https://objectstorage.<region_name>.oraclecloud.com  See <a href="#">Regions and Availability Domains</a> for region name strings.
-f	--fingerPrint	The public key fingerprint. You can find the fingerprint in the Console by clicking your user name in the upper right corner and then clicking <b>User Settings</b> . The fingerprint looks something like this:  <pre>-f 61:9e:52:26:4b:dd:46:dc:8c:a8:05:6b:9f:0a:30:d2</pre>
-k	--privateKey	The path to the private key file in PEM format, for example:  <pre>-k /root/.ssh/privkey</pre>
-h	--help	(Optional) Displays help for using the command.

Parameter	Full Name	Description
-j	--json	(Optional) Displays JSON output.
-n	--name	(Optional) The name for the new credentials configuration. The name is useful for tracking the configuration.
-p -hp	--passPhrase	The passphrase for the public/private key pair, if you specified one when creating the key pair.  Specify <code>-p</code> (with no passphrase) to be prompted.  Specify <code>-hp &lt;passphrase&gt;</code> to provide the passphrase in the command.
-t	--tenantOcid	Your tenancy OCID. See <a href="#">Where to Find Your Tenancy's OCID</a> . The tenancy OCID looks something like this:  <code>ocidl.tenancy.oc1..&lt;unique_ID&gt;</code>
-u	--userOcid	The user name OCID for your Oracle Cloud Infrastructure user account. You can find the OCID in the Console: Open the <b>Profile</b> menu (  ) and click <b>User Settings</b> . The user name OCID looks something like this:  <code>ocidl.user.oc1..&lt;unique_ID&gt;</code>

**EXAMPLE**

The following command creates a credentials configuration:

```
[root@dbsys ~]# dbcli create-bmccredential -c patching -hp mypass -t
ocidl.tenancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdsq -u
ocidl.user.oc1..aaaaaaaaalhdvixuqi7xevqsksccl6edokgldvuf6raskcioq4x2z7watsfa -f
60:9e:56:26:4b:dd:46:dc:8c:a8:05:6d:9f:0a:30:d2 -k /root/.ssh/privkey

{
 "jobId" : "f8c80510-b717-4ee2-a47e-cd380480b28b",
 "status" : "Created",
 "message" : null,
 "reports" : [],
```

## CHAPTER 11 Database

```
"createTimestamp" : "December 26, 2016 22:46:38 PM PST",
"resourceList" : [],
"description" : "BMC Credentials Creation",
"updatedAt" : "December 26, 2016 22:46:38 PM PST"
}
```

### DBCLI LIST-BMCCREDENTIALS

Use the `dbcli list-bmccredentials` command to list the credentials configurations on the DB system.

#### SYNTAX

```
dbcli list-bmccredentials [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command lists the credentials configurations on the DB system:

```
[root@dbsys ~]# dbcli list-bmccredentials
 ID Name Type End Point
 Status

f19d7c8b-d0d5-4jhf-852b-eb2a81cb7ce5 patch1 Patching https://objectstorage.us-
phoenix-1.oraclecloud.com Configured
f1a8741c-b0c4-4jhf-239b-ab2a81jhfde4 patch2 Patching https://objectstorage.us-
phoenix-1.oraclecloud.com Configured
```

### DBCLI DESCRIBE-BMCCREDENTIAL

Use the `dbcli describe-bmccredential` command to display details about a credentials configuration.

## CHAPTER 11 Database

### SYNTAX

```
dbcli describe-bmccredential -i <credentials_id> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	The ID for the credentials configuration. Use the <a href="#">dbcli list-bmccredentials</a> command to get the ID.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays details about the specified credentials configuration:

```
[root@dbsys ~]# dbcli describe-bmccredential -i 09f9988e-eed5-4dde-8814-890828d1c763

BMC Credentials details

 ID: 09f9988e-eed5-4dde-8814-890678d1c763
 Name: patch23
 Tenant OCID: ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdsq
 User OCID: ocidl.user.oc1..aaaaaaaalhjhfiuxqi7xevqsksccl6edokgldvuf6raskcioq4x2z7watjhf
 Credentials Type: Patching
 objectStore URL: https://objectstorage.us-phoenix-1.oraclecloud.com
 Status: Configured
 Created: January 9, 2017 1:19:11 AM PST
 UpdatedTime: January 9, 2017 1:41:46 AM PST
```

### DBCLI DELETE-BMCCREDENTIAL

Use the `dbcli delete-bmccredential` command to delete a credentials configuration.

### SYNTAX

```
dbcli delete-bmccredential -i <credentials_id> [-h] [-j]
```

## CHAPTER 11 Database

---

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	The ID for the credentials configuration. Use the <a href="#">dbcli list-bmccredentials</a> command to get the ID.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command deletes the specified credentials configuration:

```
[root@dbsys ~]# dbcli delete-bmccredential -i f19d7c8b-d0d5-4jhf-852b-eb2a81cb7ce5
```

### DBCLI UPDATE-BMCCREDENTIAL

Use the `dbcli update-bmccredential` command to update a credentials configuration.

### SYNTAX

```
dbcli update-bmccredential -i <credentials_id> -n <credentials_name> -c [backup|patching|other] -p|-hp <passphrase> -f <fingerprint> -k <private_key_path> [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-c	--credentialsType	The type of Object Storage credentials configuration (these values are <b>not</b> case sensitive):  BACKUP - Reserved for the future use.  PATCHING - For downloading patches from the service.  OTHER - Reserved for the future use.
-i	--id	The ID for the credentials configuration. Use the <a href="#">dbcli list-bmccredentials</a> command to get the ID.
-f	--fingerPrint	The public key fingerprint, for example:  <code>-f 61:9e:52:26:4b:dd:46:dc:8c:a8:05:6b:9f:0a:30:d2</code>
-k	--privateKey	The path to the private key file in PEM format, for example:  <code>-k /root/.ssh/privkey</code>
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	(Optional) The name for the credentials configuration. Use the <a href="#">dbcli list-bmccredentials</a> command to get the name.
-p -hp	--passPhrase	The passphrase for the public/private key pair, if you specified one when creating the key pair.  Specify <code>-p</code> (with no passphrase) to be prompted.  Specify <code>-hp &lt;passphrase&gt;</code> to provide the passphrase in the command.

### EXAMPLE

The following command updates a credentials configuration:

## CHAPTER 11 Database

---

```
[root@dbsys ~]# dbcli update-bmccredential -c OTHER -i 6f921b29-61b6-56f4-889a-ce9270621956
{
 "jobId" : "6e95a69e-cf73-4e51-a444-c7e4b9631c27",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "January 19, 2017 12:01:10 PM PST",
 "resourceList" : [],
 "description" : "Update BMC Credentials of object 6f921b29-61b6-48f4-889a-ce9270621945",
 "updatedAt" : "January 19, 2017 12:01:10 PM PST"
```

### Component Command

#### DBCLI DESCRIBE-COMPONENT



#### Tip

Your DB system might not include this newer command. If you have trouble running the command, use the [cliadm update-dbcli](#) command to update the database CLI and then retry the command.



#### Note

The `dbcli describe-component` command is not available on 2-node RAC DB systems. Patching 2-node systems from Object Storage is not supported.

Use the `dbcli describe-component` command to show the installed and available patch versions for the server, storage, and/or database home components in the DB system.

This command requires a valid Object Storage credentials configuration. Use the [dbcli create-bmccredential](#) command to create the configuration if you haven't already done so. If the

## CHAPTER 11 Database

configuration is missing or invalid, the command fails with the error: Failed to connect to the object store. Please provide valid details.

For more information about updating the CLI, creating the credentials configuration, and applying patches, see [Patching a DB System](#).

### SYNTAX

```
dbcli describe-component [-s <server_group>] [-d <db_group>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-d	--dbhomes	(Optional) Lists the installed and available patch versions for only the database home components.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-s	--server	(Optional) Lists the installed and available patch versions for only the server components.

### EXAMPLE

The following command to show the current component versions and the available patch versions in the object store:

```
[root@dbsys ~]# dbcli describe-component
System Version

12.1.2.10.0

Component Installed Version Available Version

OAK 12.1.2.10.0 up-to-date
GI 12.1.0.2.161018 up-to-date
ORADB12102_HOME1 12.1.0.2.161018 up-to-date
ORADB12102_HOME2, ORADB12102_HOME3 12.1.0.2.160719 12.1.0.2.161018
```

## Database Commands

The following commands are available to manage databases:

- [dbcli clone-database](#)
- [dbcli create-database](#)
- [dbcli delete-database](#)
- [dbcli describe-database](#)
- [dbcli list-databases](#)
- [dbcli modify-database](#)
- [dbcli recover-database](#)
- [dbcli register-database](#)
- [dbcli update-database](#)

### DBCLI CLONE-DATABASE

Use the `dbcli clone-database` command to clone a database.

#### SYNTAX

```
dbcli clone-database -f <name> -u <name> -n <name> [-s <shape>] [-t <type>] [m <sys_password>] [-p <tde_password>] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-f	--sourcedbname	Source database name.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-m	--syspassword	(Optional) Password for SYS.
-n	--dbname	Database name.

Parameter	Full Name	Description
-p	--tdepassword	(Optional) Password for source TDE wallet.
-s	--dbshape	(Optional) Database shape. Examples: odb1, odb2.
-t	--dbtype	(Optional) Database Type: SI
-u	--databaseUniqueName	Database unique name.

**DBCLI CREATE-DATABASE**

Use the `dbcli create-database` command to create a new database. You can create a database with a new or existing Oracle Database home, however each database home can have only one database.

It takes a few minutes to create the database. After you run the `dbcli create-database` command, you can use the `dbcli list-jobs` command to check the status of the database creation job.

**Tip**

Wait for the database creation job to complete before you attempt to create another database. Running multiple `dbcli create-database` commands at the same time can result in some of the creation jobs not completing.

Once the database is created, you can use the `dbcli list-databases -j` command to see additional information about the database.



### Note

The `dbcli create-database` command is available on bare metal DB systems only.

You must create and activate a master encryption key for any PDBs that you create. After creating or plugging in a new PDB on a 1- or 2-node RAC DB System, use the `dbcli update-tdekey` command to create and activate a master encryption key for the PDB. Otherwise, you might encounter the error `ORA-28374: typed master key not found in wallet` when attempting to create tablespaces in the PDB. In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers. For more information, see "Overview of Managing a Multitenant Environment" in the [Oracle Database Administrator's Guide](#).

### SYNTAX

```
dbcli create-database -dh <db_home_id> -cl {OLTP|DSS|IMDB} -n <db_name> -u <unique_name> -bi <bkup_
config_id> -bn <bkup_config_name> -m -s <db_shape> -r {ACFS|ASM} -y {SI|RAC|RACOne} [-dn <name>] -io -d
<pdb_admin_user> [-p <pdb>] [-ns <nlscharset>] [-cs <charset>] [-l <language>] [-dt <territory>] -v
<version> [-c|-no-c] [-co|-no-co] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-bi	--backupconfigid	Defines the backup configuration identifier for future use. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-bn	--backupconfigname	Defines the backup configuration name for future use. Use the <code>dbcli list-backupconfigs</code> command to get the name.
-c -no-c	--cdb --no-cdb	(Optional) Indicates whether to create a Container Database. If omitted, a Container Database is not created.
-cs	--characterset	(Optional) Defines the character set for the database. The default is AL32UTF8.
-cl	--dbclass	Defines the database class. The options are OLTP, DSS, or IMDB. The default is OLTP. For Enterprise Editions, all three classes are supported. For Standard Edition, only OLTP is supported.
-co -no-co	--dbconsole --no-dbconsole	(Optional) Indicates whether the Database Console is enabled. If omitted, the console is not enabled.  This parameter is not available for a version 11.2.0.4 database on a 2-node RAC DB system. For more information, see <a href="#">To enable the console for a version 11.2.0.4 database on a multi-node DB system</a> .
-d	--pdbadmin	Defines the name of the Pluggable Database (PDB) Admin User. The default value is <code>pdbadmin</code> .

Parameter	Full Name	Description
-dn	--dbdomainname	(Optional) Database domain name (indicates the logical location of the database within the network structure).
-dt	--dbterritory	(Optional) Defines the territory for the database. The default is AMERICA.
-dh	--dbhomeid	Identifies the database home in which to create the database. The database home must be empty because each database home can have only one database. You can use the <code>dbcli list dbhomes</code> command to get the DB home ID.  If this parameter is omitted, the database is created with a new Oracle home.
-h	--help	(Optional) Displays help for the command.
-j	--json	(Optional) Displays JSON output.
-l	--dblanguage	(Optional) Defines the language for the database. The default is AMERICAN.
-m	--adminpassword	A strong password for SYS, SYSTEM, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least 2 uppercase, 2 lowercase, 2 numeric, and 2 special characters. The special characters must be <code>_</code> , <code>#</code> , or <code>-</code> . The password must not contain the username (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reversed order and regardless of casing.  Specify <code>-m</code> (with no password) to be prompted for the password.

## CHAPTER 11 Database

Parameter	Full Name	Description
-n	--dbname	Defines the name given to the new database. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
-ns	--nationalcharacterset	(Optional) Defines the national character set for the database. The default is AL16UTF16.
-p	--pdbname	(Optional) Defines a unique name for the PDB. The PDB name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters. The only special character permitted is the underscore ( _). The default value is pdb1.  PDB names must be unique within a CDB and within the listener to which they are registered. Make sure the PDB name is unique on the system. To ensure uniqueness, do not use the default name value (pdb1).
-r	--dbstorage	Defines the database storage, either ACFS or ASM. The default value is ASM.  See <a href="#">Usage Notes</a> for more information.
-s	--dbshape	Identifies the database sizing template to use for the database. For example, odb1, odb2, or odb3. The default is odb1. For more information, see <a href="#">Database Sizing Templates</a> .

## CHAPTER 11 Database

Parameter	Full Name	Description
-u	-- databaseUniqueName	Defines a unique name for the database to ensure uniqueness within an Oracle Data Guard group (a primary database and its standby databases). The unique name can contain only alphanumeric and underscore ( <code>_</code> ) characters. The unique name cannot be changed. The unique name defaults to the name specified in the <code>--dbname</code> parameter.
-v	--version	Defines the database version as one of the following: <ul style="list-style-type: none"><li>• 18.1.0.0</li><li>• 12.2.0.1</li><li>• 12.1.0.2 (the default)</li><li>• 11.2.0.4</li></ul>
-y	--dbtype	Defines the database type. Specify SI for a 1-node instance, RAC for a 2-node cluster, or RACOne for 1-node instance with a second node in cold standby mode. The default value is RAC. These values are not case sensitive.

### USAGE NOTES

- You cannot mix Oracle Database Standard Edition and Enterprise Edition databases on the same DB system. (You can mix supported database versions on the DB system, but not editions.)
- When `--dbhomeid` is not provided, the `dbcli create-database` command will create a new Oracle Database home.



### Note

Bare metal DB systems allow only one database per database home.

- When `--dbhomeid` is provided, the `dbcli create-database` command creates the database using the Oracle home specified. Use the `dbcli list-dbhomes` command to get the `dbhomeid`. The database home you specify must be empty.
- Oracle Database 12.1 or later databases are supported on both Oracle Automatic Storage Management (ASM) and Oracle ASM Cluster file system (ACFS). The default is Oracle ACFS.
- Oracle Database 11.2 is supported on Oracle ACFS.
- Each database is configured with its own Oracle ACFS file system for the datafiles and uses the following naming convention: `/u02/app/db user/oradata/db name`. The default size of this mount point is 100G.
- Online logs are stored in the `/u03/app/db user/redo/` directory.
- The Oracle Fast Recovery Area (FRA) is located in the `/u03/app/db user/fast_recovery_area` directory.

### EXAMPLES

To create a database and be prompted for the password interactively:

```
[root@dbsys ~]# dbcli create-database -n hrdb -c -m -cl OLTP -s odb2 -p pdb1
```

```
Password for SYS,SYSTEM and PDB Admin:
```

```
{
 "jobId" : "f12485f2-dcbe-4ddf-aeel-de24d37037b6",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "August 08, 2016 03:54:03 AM EDT",
 "description" : "Database service creation with db name: hrdb",
 "updatedAt" : "August 08, 2016 03:54:03 AM EDT"
}
```

## CHAPTER 11 Database

To create a database non-interactively, providing the password on the command line:

```
[root@dbsys ~]# dbcli create-database -n crmdb -hm <password> -cl OLTP -s odb2
{
 "jobId" : "30b5e2a6-493b-4461-98b8-78e9a15f8cdd",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "August 08, 2016 03:59:22 AM EDT",
 "description" : "Database service creation with db name: crmdb",
 "updatedAt" : "August 08, 2016 03:59:22 AM EDT"
}
```

### DBCLI DELETE-DATABASE

Use the `dbcli delete-database` command to delete a database.



#### Note

The `dbcli create-database` command is available on bare metal DB systems only.

### SYNTAX

```
dbcli delete-database -i <db_id> -in <db_name> [-fd] [-j] [-h]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-fd	--force	(Optional) Forces the delete operation.
-i	--dbid	The ID of the database to delete. Use the <code>dbcli list-databases</code> command to get the database ID.

## CHAPTER 11 Database

Parameter	Full Name	Description
-in	--dbName	The name of the database to delete. Use the <code>dbcli list-databases</code> command to get the database name.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command deletes the database named 625d9b8a-baea-4994-94e7-4c4a857a17f9:

```
[root@dbsys ~]# dbcli delete-database -i 625d9b8a-baea-4994-94e7-4c4a857a17f9
```

### DBCLI DESCRIBE-DATABASE

Use the `dbcli describe-database` command to display database details.

### SYNTAX

```
dbcli describe-database -i <db_id> -in <db_name> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	The ID of the database to display. Use the <code>dbcli list-databases</code> command to get the database ID.
-in	--dbName	The name of the database to display. Use the <code>dbcli list-databases</code> command to get the database name.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays information for a database named b727bf80-c99e-4846-ac1f-28a81a725df6:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli describe-dbhome -i b727bf80-c99e-4846-ac1f-28a81a725df6
```

DB Home details

```

ID: b727bf80-c99e-4846-ac1f-28a81a725df6
Name: OraDB12102_home1
Version: 12.1.0.2
Home Location: /u01/app/orauser/product/12.1.0.2/dbhome_1
Created: Jun 2, 2016 10:19:23 AM
```

### DBCLI LIST-DATABASES

Use the `dbcli list-databases` command to list all databases on the DB system.

#### SYNTAX

```
dbcli list-databases [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays a list of databases:

```
[root@dbsys ~]# dbcli list-databases
```

```
ID DB Name DB Version CDB Class Shape
Storage Status

80ad855a-5145-4f8f-a08f-406c5e4684ff dbst 12.1.0.2 true OLTP odb2
ACFS Configured
6f4e36ae-120b-4436-b0bf-d0c4aef9f7c9 db11tsta 11.2.0.4 false OLTP odb1
ACFS Configured
d8e31790-84e6-479c-beb0-ef97207091a2 db11tstb 11.2.0.4 false OLTP odb1
ACFS Configured
```

## CHAPTER 11 Database

```
cce096c7-737b-447a-baa1-f4c2a330c030 pdbstst 12.1.0.2 true OLTP odb1
ACFS Configured
```

The following command displays the JSON output for a database:

```
[root@dbsys ~]# dbcli list-databases -j
[{
 "id" : "80ad855a-5145-4f8f-a08f-406c5e4684ff",
 "name" : "dbtst",
 "dbName" : "dbtst",
 "databaseUniqueName" : "dbtst_phx1cs",
 "dbVersion" : "12.1.0.2",
 "dbHomeId" : "2efe7af7-0b70-4e9b-ba8b-71f11c6fe287",
 "instanceOnly" : false,
 "registerOnly" : false,
 "dbId" : "167525515",
 "isCdb" : true,
 "pdbName" : "pdb1",
 "pdbAdminUserName" : "pdbuser",
 "enableTDE" : true,
 "dbType" : "SI",
 "dbTargetNodeNumber" : "0",
 "dbClass" : "OLTP",
 "dbShape" : "odb2",
 "dbStorage" : "ACFS",
 "dbCharacterSet" : {
 "characterSet" : "US7ASCII",
 "nlsCharacterSet" : "AL16UTF16",
 "dbTerritory" : "AMERICA",
 "dbLanguage" : "AMERICAN"
 },
 "dbConsoleEnable" : false,
 "backupConfigId" : null,
 "backupDestination" : "NONE",
 "cloudStorageContainer" : null,
 "state" : {
 "status" : "CONFIGURED"
 },
 "createTime" : "November 09, 2016 17:23:05 PM UTC",
 "updateTime" : "November 09, 2016 18:00:47 PM UTC"
}
```

## CHAPTER 11 Database

### DBCLI MODIFY-DATABASE

Use the `dbcli modify-database` command to modify a database.

#### SYNTAX

```
dbcli modify-database -i <db_id> -dh <destination_db_home_id> [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-dh	--destdbhomeid	Destination database home ID.
-h	--help	(Optional) Displays help for using the command.
-i	--databaseid	Database ID.
-j	--json	(Optional) Displays JSON output.

### DBCLI RECOVER-DATABASE

Use the `dbcli recover-database` command to recover a database.

#### SYNTAX

```
dbcli recover-database [-br <json>] [-in <db_name>] [-i <db_id>] [-r <time>] [-t {Latest|PITR|SCN}] [-s] [-l <location>] [-tp <tde_password>] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-br	--backupReport	(Optional) JSON input for backup report.
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbName	(Optional) Database name.

## CHAPTER 11 Database

Parameter	Full Name	Description
-j	--json	(Optional) Displays JSON output.
-l	--tdeWalletLocation	(Optional) TDE wallet backup location. TDE wallet should be backed up in tar.gz format.
-r	--recoveryTimeStamp	(Required when recovery type is PITR) Recovery timestamp in the format mm/dd/yyyy hh:mi:ss. Default: [ ]
-s	--scn	(Required when recovery type is SCN) SCN.
-t	--recoverytype	(Required when backup report is provided) Recovery type. Possible values are Latest, PITR, and SCN.
-tp	--tdeWalletPassword	(Optional) TDE wallet password.

### DBCLI REGISTER-DATABASE

Use the `dbcli register-database` command to register a database that has been migrated to Oracle Cloud Infrastructure. The command registers the database to the `dc-agent` so it can be managed by the `dc-agent` stack.



#### Note

The `dbcli register-database` command is not available on 2-node RAC DB systems.

### SYNTAX

```
dbcli register-database -bi <bkup_config_id> -c {OLTP|DSS|IMDB} [-co|-no-co] -s {odbl|odb2|...} -t SI [-o <db_host_name>] [-tp <password>] -sn <service_name> -p [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-bi	--backupconfigid	Defines the backup configuration ID. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-c	--dbclass	Defines the database class. The options are OLTP, DSS, or IMDB. The default is OLTP. For Enterprise Editions, all three classes are supported. For Standard Edition, only OLTP is supported.
-co -no-co	--dbconsole --no-dbconsole	(Optional) Indicates whether the Database Console is enabled or not. If omitted, the console is not enabled.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-o	--hostname	(Optional) Defines the database host name. The default is <code>Local host name</code> .
-p	--syspassword	Defines a strong password for SYS. Specify <code>-p</code> with no password. You will be prompted for the password.  If you must provide the password in the command, for example in a script, use <code>-hp &lt;password&gt;</code> instead of <code>-p</code> .
-s	--dbshape	Defines the database sizing template to use for the database. For example, <code>odb1</code> , <code>odb2</code> , and <code>odb3</code> . For more information, see <a href="#">Database Sizing Templates</a> .
-sn	--servicename	Defines the database service name used to build the EZCONNECT string for connecting to the database. The connect string format is <code>hostname:port/servicename</code> .

## CHAPTER 11 Database

Parameter	Full Name	Description
-t	--dbtype	(Optional) Defines the Database Type as single node (SI). The default value is SI.
-tp	-- tdeWalletPassword	(Optional) Password for TDE wallet. Required if TDE is enabled on the migrated database.

### EXAMPLE

The following command registers the database with the specified database class, service name, and database sizing template.

```
[root@dbsys ~]# dbcli register-database -c OLTP -s odb1 -sn crmdb.example.com -p
Password for SYS:
{
 "jobId" : "317b430f-ad5f-42ae-bb07-13f053d266e2",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "August 08, 2016 05:55:49 AM EDT",
 "description" : "Database service registration with db service name: crmdb.example.com",
 "updatedAt" : "August 08, 2016 05:55:49 AM EDT"
}
```

### DBCLI UPDATE-DATABASE

Use the `dbcli update-database` command to associate a backup configuration with a database.

### SYNTAX

```
dbcli update-database -i <db_id> -bi <bkup_config_id> -bin <bkup_config_name> [-id <id>] -in <name> [-no-ab] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-bi	--backupconfigid	Defines the backup configuration ID. Use the <code>dbcli list-backupconfigs</code> command to get the ID.
-bin	-- backupconfigname	Defines the backup configuration name for future use. Use the <code>dbcli list-backupconfigs</code> command to get the name.
-id	--databaseid	(Optional.) Specifies the DBID, which is a unique 32-bit identification number computed when the database is created. RMAN displays the DBID upon connection to the target database. You can obtain the DBID by querying the <code>V\$DATABASE</code> view or the <code>RC_DATABASE</code> and <code>RC_DATABASE_INCARNATION</code> recovery catalog views.
-in	--dbName	Defines the database name to be updated. Use the <code>dbcli list-databases</code> command to get the database name.
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	Defines the database ID to be updated. Use the <code>dbcli list-databases</code> command to get the database ID.

Parameter	Full Name	Description
-j	--json	(Optional) Displays JSON output.
-no-ab	--noautobackup	(Optional) Disables automatic backups for the specified database.  <div style="border: 1px solid #0070c0; background-color: #e1f5fe; padding: 10px; margin-top: 10px;">  <p><b>Note</b></p> <p>Once disabled, automatic backup cannot be re-enabled using the CLI. To re-enable automatic backup, use the Console.</p> </div>

**EXAMPLE**

The following command associates a backup configuration file with a database:

```
[root@dbsys ~]# dbcli update-database -bi 78a2a5f0-72b1-448f-bd86-cf41b30b64ee -i 71ec8335-113a-46e3-b81f-235f4d1b6fde
{
 "jobId" : "2b104028-a0a4-4855-b32a-b97a37f5f9c5",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1467775842977,
 "description" : "update database id:71ec8335-113a-46e3-b81f-235f4d1b6fde",
 "updatedAt" : 1467775842978
}
```

**Dbhome Commands**

The following commands are available to manage database homes:

- [dbcli create-dbhome](#)
- [dbcli describe-dbhome](#)

## CHAPTER 11 Database

---

- [dbcli delete-dbhome](#)
- [dbcli list-dbhomes](#)
- [dbcli update-dbhome](#)

### DBCLI CREATE-DBHOME

Use the `dbcli create-dbhome` command to create an Oracle Database Home.

#### SYNTAX

```
dbcli create-dbhome -v <version> [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-v	--version	Defines the Database Home version. Specify one of the supported versions: <ul style="list-style-type: none"><li>• 18.1.0.0</li><li>• 12.2.0.1</li><li>• 12.1.0.2</li><li>• 11.2.0.4</li></ul>

#### EXAMPLE

The following command creates an Oracle Database Home version 12.1.0.2:

```
[root@dbsys ~]# dbcli create-dbhome -v 12.1.0.2
```

### DBCLI DESCRIBE-DBHOME

Use the `dbcli describe-dbhome` command to display Oracle Database Home details.

## CHAPTER 11 Database

### SYNTAX

```
dbcli describe-dbhome -i <db_home_id> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	Identifies the database home ID. Use the <code>dbcli list-dbhomes</code> command to get the ID.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following output is an example of using the display Oracle Database Home details command.

```
[root@dbsys ~]# dbcli describe-dbhome -i 52850389-228d-4397-bbe6-102fda65922b

DB Home details

 ID: 52850389-228d-4397-bbe6-102fda65922b
 Name: OraDB12102_home1
 Version: 12.1.0.2
 Home Location: /u01/app/oracle/product/12.1.0.2/dbhome_1
 Created: June 29, 2016 4:36:31 AM UTC
```

### DBCLI DELETE-DBHOME

Use the `dbcli delete-dbhome` command to delete a database home from the DB system.

### SYNTAX

```
dbcli delete-dbhome -i <db_home_id> [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	Identifies the database home ID to be deleted. Use the <code>dbcli list-dbhomes</code> command to get the ID.
-j	--json	(Optional) Displays JSON output.

### DBCLI LIST-DBHOMES

Use the `dbcli list-dbhomes` command to display a list of Oracle Home directories.

### SYNTAX

```
dbcli list-dbhomes [-h] [-j]
```

### PARAMETER

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays a list of Oracle Home directories.

```
[root@dbsys ~]# dbcli list-dbhomes
ID Name DB Version Home Location

b727bf80-c99e-4846-ac1f-28a81a725df6 OraDB12102_home1 12.1.0.2
/u01/app/orauser/product/12.1.0.2/dbhome_1
```

### DBCLI UPDATE-DBHOME



#### Tip

Your DB system might not include this newer command. If you have trouble running the command, use the [cliadm update-dbcli](#) command to update the database CLI and then retry the command.

Use the `dbcli update-dbhome` command to apply the DBBP bundle patch to a database home. For more information about applying patches, see [Patching a DB System](#).

#### SYNTAX

```
dbcli update-dbhome -i <db_home_id> -n <node> [--local] [--precheck] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbhomeid	The ID of the database home. Use the <code>dbcli list-dbhomes</code> command to get the ID.
-j	--json	(Optional) Displays JSON output.
-n	--node	(Optional) Node number to be updated. Use the <code>dbcli list-nodes</code> command to get the node number.
	--local	(Optional) Performs the operation on the local node of a multi-node high availability (HA) system. This parameter is not needed to perform the operation on a single-node system.
	--precheck	(Optional) Runs precheck operations to check prerequisites.

## CHAPTER 11 Database

### EXAMPLE

The following commands update the database home and show the output from the update job:

```
[root@dbsys ~]# dbcli update-dbhome -i e1877dac-a69a-40a1-b65a-d5e190e671e6
{
 "jobId" : "493e703b-46ef-4a3f-909d-bbd123469bea",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "January 19, 2017 10:03:21 AM PST",
 "resourceList" : [],
 "description" : "DB Home Patching: Home Id is e1877dac-a69a-40a1-b65a-d5e190e671e6",
 "updatedAt" : "January 19, 2017 10:03:21 AM PST"
}

dbcli describe-job -i 493e703b-46ef-4a3f-909d-bbd123469bea

Job details

 ID: 493e703b-46ef-4a3f-909d-bbd123469bea
Description: DB Home Patching: Home Id is e1877dac-a69a-40a1-b65a-d5e190e671e6
 Status: Running
 Created: January 19, 2017 10:03:21 AM PST
 Message:

Task Name Start Time End Time

 Status

Create Patching Repository Directories January 19, 2017 10:03:21 AM PST January 19, 2017 10:03:21
AM PST Success
Download latest patch metadata January 19, 2017 10:03:21 AM PST January 19, 2017 10:03:21
AM PST Success
Update System version January 19, 2017 10:03:21 AM PST January 19, 2017 10:03:21
AM PST Success
Update Patching Repository January 19, 2017 10:03:21 AM PST January 19, 2017 10:03:31
AM PST Success
Opatch updation January 19, 2017 10:03:31 AM PST January 19, 2017 10:03:31
AM PST Success
Patch conflict check January 19, 2017 10:03:31 AM PST January 19, 2017 10:03:31
AM PST Running
```

## Dbstorage Commands

The following commands are available to manage database storage:

- [dbcli list-dbstorages](#)
- [dbcli describe-dbstorage](#)
- [dbcli create-dbstorage](#)
- [dbcli delete-dbstorage](#)

### DBCLI LIST-DBSTORAGES

Use the `dbcli list-dbstorages` command to list the database storage in the DB system.

#### SYNTAX

```
dbcli list-dbstorages [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays details about database storage:

```
[root@dbsys ~]# dbcli list-dbstorages
```

ID	Type	DBUnique Name	Status
afb4alce-d54d-4993-a149-0f28c9fb33a4	Acfs	db1_2e56b3a9b815	Configured
d81e8013-4551-4d10-880b-d1a796bca1bc	Acfs	db11xp	Configured

### DBCLI DESCRIBE-DBSTORAGE

Use the `dbcli describe-dbstorage` command to show detailed information about a specific database storage resource.

## CHAPTER 11 Database

### SYNTAX

```
dbcli describe-dbstorage -i <db_storage_id> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	Defines the database storage ID. Use the <code>dbcli list-dbstorages</code> command to get the database storage ID.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays the database storage details for 105a2db2-625a-45ba-8bdd-ee46da0fd83a:

```
[root@dbsys ~]# dbcli describe-dbstorage -i 105a2db2-625a-45ba-8bdd-ee46da0fd83a
```

```
DBStorage details
```

```

ID: 105a2db2-625a-45ba-8bdd-ee46da0fd83a
DB Name: db1
DBUnique Name: db1
DB Resource ID: 439e7bd7-f717-447a-8046-08b5f6493df0
Storage Type:
DATA Location: /u02/app/oracle/oradata/db1
RECO Location: /u03/app/oracle/fast_recovery_area/
REDO Location: /u03/app/oracle/redo/
State: ResourceState(status=Configured)
Created: July 3, 2016 4:19:21 AM UTC
UpdateTime: July 3, 2016 4:41:29 AM UTC
```

### DBCLI CREATE-DBSTORAGE

Use the `dbcli create-dbstorage` command to create the database storage layout without creating the complete database. This is useful for database migration and standby database

## CHAPTER 11 Database

creation.

### SYNTAX

```
dbcli create-dbstorage -n <db_name> [-u <db_unique_name>] [-r {ACFS|ASM}] [-s <datasize>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--dbname	Defines the database name. The database name must begin with an alphabetic character and can contain a maximum of eight alphanumeric characters. Special characters are not permitted.
-r	--dbstorage	(Optional) Defines the type of database storage as ACFS or ASM. The default value is ASM.
-s	--dataSize	(Optional) Defines the data size in GBs. The minimum size is 10GB. The default size is 100GB.
-u	--databaseUniqueName	(Optional) Defines the unique name for the database. The default is the database name specified in --dbname.

### EXAMPLE

The following command creates database storage with a storage type of ACFS:

```
[root@dbsys ~]# dbcli create-dbstorage -r ACFS -n testdb -u testdbname
```

```
{
 "jobId" : "5884a77a-0577-414f-8c36-1e9d8a1e9cee",
 "status" : "Created",
 "message" : null,
 "reports" : [],
}
```

## CHAPTER 11 Database

```
"createTimestamp" : 1467952215102,
"description" : "Database storage service creation with db name: testdb",
"updatedAt" : 1467952215103
}
```

### DBCLI DELETE-DBSTORAGE

Use the `dbcli delete-dbstorage` command to delete database storage that is not being used by the database. A error occurs if the resource is in use.

#### SYNTAX

```
dbcli delete-dbstorage -i <dbstorageID> [-h] [-j]
```

#### PARAMETERS

Parameter	Parameter	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	The database storage ID to delete. Use the <code>dbcli list-dbstorages</code> command to get the database storage ID.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command deletes the specified database storage:

```
[root@dbsys ~]# dbcli delete-dbstorage -i f444dd87-86c9-4969-a72c-fb2026e7384b

{
 "jobId" : "467c9388-18c6-4e1a-8655-2fd3603856ef",
 "status" : "Running",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1467952336843,
 "description" : "Database storage service deletion with id: f444dd87-86c9-4969-a72c-fb2026e7384b",
 "updatedAt" : 1467952336856
}
```

### Dgconfig Commands

#### DBCLI LIST-DGCONFIGS

Use the `dbcli list-dgconfigs` command to list DG configurations.

#### SYNTAX

```
dbcli list-dgconfigs [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Featuretracking Commands

#### DBCLI LIST-FEATURETRACKING

Use the `dbcli list-featuretracking` command to list tracked features.

#### SYNTAX

```
dbcli list-featuretracking[-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Job Commands

The following commands are available to manage jobs:

## CHAPTER 11 Database

- [dbcli describe-job](#)
- [dbcli list-jobs](#)

### DBCLI DESCRIBE-JOB

Use the `dbcli describe-job` command to display details about a specific job.

#### SYNTAX

```
dbcli describe-job -i <job_id> [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--jobid	Identifies the job. Use the <code>dbcli list-jobs</code> command to get the jobid.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays details about the specified job ID:

```
[root@dbsys ~]# dbcli describe-job -i 74731897-fb6b-4379-9a37-246912025c17

Job details

 ID: 74731897-fb6b-4379-9a37-246912025c17
Description: Backup service creation with db name: dbtst
 Status: Success
 Created: November 18, 2016 8:33:04 PM UTC
 Message:

Task Name Start Time End Time

 Status

Backup Validations November 18, 2016 8:33:04 PM UTC November 18, 2016 8:33:13
```

## CHAPTER 11 Database

```
PM UTC Success
validate recovery window November 18, 2016 8:33:13 PM UTC November 18, 2016 8:33:17
PM UTC Success
Db cross check November 18, 2016 8:33:17 PM UTC November 18, 2016 8:33:23
PM UTC Success
Database Backup November 18, 2016 8:33:23 PM UTC November 18, 2016 8:34:22
PM UTC Success
Backup metadata November 18, 2016 8:34:22 PM UTC November 18, 2016 8:34:22
PM UTC Success
```

### DBCLI LIST-JOBS

Use the `dbcli list-jobs` command to display a list of jobs, including the job IDs, status, and the job

created date and time stamp.

#### SYNTAX

```
dbcli list-jobs [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays a list of jobs:

```
[root@dbsys ~]# dbcli list-jobs
ID Description
Created Status

0a362dac-0339-41b5-9c9c-4d229e363eaa Database service creation with db name: db11
November 10, 2016 11:37:54 AM UTC Success
9157cc78-b487-4ee9-9f46-0159f10236e4 Database service creation with db name: jhfpdb
```

## CHAPTER 11 Database

```
November 17, 2016 7:19:59 PM UTC Success
013c408d-37ca-4f58-a053-02d4efdc42d0 create backup config:myBackupConfig
November 18, 2016 8:28:14 PM UTC Success
921a54e3-c359-4aea-9efc-6ae7346cb0c2 update database id:80ad855a-5145-4f8f-a08f-406c5e4684ff
November 18, 2016 8:32:16 PM UTC Success
74731897-fb6b-4379-9a37-246912025c17 Backup service creation with db name: dbtst
November 18, 2016 8:33:04 PM UTC Success
40a227b1-8c47-46b9-a116-48cc1476fc12 Creating a report for database 80ad855a-5145-4f8f-a08f-
406c5e4684ff November 18, 2016 8:41:39 PM UTC Success
```

### Latestpatch Command

#### DBCLI DESCRIBE-LATESTPATCH



#### Tip

Your DB system might not include this newer command. If you have trouble running the command, use the [cliadm update-dbcli](#) command to update the database CLI and then retry the command.



#### Note

The `dbcli describe-latestpatch` command is not available on 2-node RAC DB systems. Patching 2-node systems from Object Storage is not supported.

Use the `dbcli describe-latestpatch` command show the latest patches applicable to the DB system and available in Oracle Cloud Infrastructure Object Storage.

This command requires a valid Object Storage credentials configuration. Use the [dbcli create-bmcredential](#) command to create the configuration if you haven't already done so. If the configuration is missing or invalid, the command fails with the error: Failed to connect to the object store. Please provide valid details.

## CHAPTER 11 Database

For more information about updating the CLI, creating the credentials configuration, and applying patches, see [Patching a DB System](#).

### SYNTAX

```
dbcli describe-latestpatch [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays patches available in the object store:

```
[root@dbsys ~]# dbcli describe-latestpatch

componentType availableVersion

gi 12.1.1.0.2.161018
db 11.2.0.4.161018
db 12.1.1.0.2.161018
oak 12.1.2.10.0
```

## Logcleanjob Commands

The following commands are available to manage log cleaning jobs:

- [dbcli create-logCleanJob](#)
- [dbcli describe-logCleanJob](#)
- [dbcli list-logCleanJobs](#)

### DBCLI CREATE-LOGCLEANJOB

Use the `dbcli create-logCleanJob` command to create a log cleaning job.

## CHAPTER 11 Database

### SYNTAX

```
dbcli create-logCleanJob [-c {gi|database|dcs}] [-o <number>] [u {Day|Hour|Minute}] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-c	--components	(Optional) Components. Possible values are gi, database, and dcs. Separate multiple values by commas.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-o	--olderThan	(Optional) Quantity portion of time interval. Default: 30. Cleans logs older than the specified time interval (-o and -u).
-u	--unit	(Optional) Unit portion of time interval. Possible values: Day, Hour, or Minute. Default: Day. Cleans logs older than the specified time interval (-o and -u).

### DBCLI DESCRIBE-LOGCLEANJOB

Use the `dbcli describe-logCleanJob` command to display the summary for a log cleaning job.

### SYNTAX

```
dbcli describe-logCleanJob -i <job_id> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--jobid	ID of log cleaning job for which to display the summary.
-j	--json	(Optional) Displays JSON output.

## CHAPTER 11 Database

---

### DBCLI LIST-LOGCLEANJOBS

Use the `dbcli list-logCleanJobs` command to list log cleaning jobs.

#### SYNTAX

```
dbcli list-logCleanJobs [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Logspaceusage Command

#### DBCLI LIST-LOGSPACEUSAGE

Use the `dbcli list-logSpaceUsage` command to list log space usage.

#### SYNTAX

```
dbcli list-logSpaceUsage [-c {gi|database|dcs}] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-c	-- components	(Optional) Components. Possible values: gi, database, and dcs. Separate multiple values by commas.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Netsecurity Commands

The following commands are available to manage network encryption on the DB system:

## CHAPTER 11 Database

- [dbcli describe-netsecurity](#)
- [dbcli update-netsecurity](#)

### DBCLI DESCRIBE-NETSECURITY

Use the `dbcli describe-netsecurity` command to display the current network encryption setting for a database home.

#### SYNTAX

```
dbcli describe-netsecurity -H <db_home_id> [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-H	--dbHomeId	Defines the database home ID. Use the <code>dbcli list-dbhomes</code> command to get the <code>dbhomeid</code> .
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

#### EXAMPLE

The following command displays the encryption setting for specified database home:

```
[root@dbsys ~]# dbcli describe-netsecurity -H 16c96a9c-f579-4a4c-a645-8d4d22d6889d
```

```
NetSecurity Rules
```

```

DatabaseHomeID: 16c96a9c-f579-4a4c-a645-8d4d22d6889d

Role: Server
EncryptionAlgorithms: AES256 AES192 AES128
IntegrityAlgorithms: SHA1
ConnectionType: Required

Role: Client
EncryptionAlgorithms: AES256 AES192 AES128
```

## CHAPTER 11 Database

```
IntegrityAlgorithms: SHA1
ConnectionType: Required
```

### DBCLI UPDATE-NETSECURITY

Use the `dbcli update-netsecurity` command to update the Oracle Net security configuration on the DB system.

#### SYNTAX

```
dbcli update-netsecurity {-c|-s} -t {REJECTED|ACCEPTED|REQUESTED|REQUIRED} -H db_home_id -e
{AES256|AES192|AES128} -i {SHA1|SHA512|SHA384|SHA256} [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-c	--client	Indicates that the specified data encryption or data integrity configuration is for the client. ( <code>--client</code> and <code>--server</code> are mutually exclusive.)
-e	--encryptionAlgorithms	Defines the algorithm to be used for encryption. Specify either AES256, AES192, or AES128.
-H	--dbHomeId	Defines the database home ID. Use the <code>dbcli list-dbhomes</code> command to get the <code>dbHomeId</code> .
-h	--help	(Optional) Displays help for using the command.
-i	--integrityAlgorithms	Defines the algorithm to be used for integrity. Specify either SHA1, SHA512, SHA384, or SHA256. For Oracle Database 11g, the only accepted value is SHA1.
-j	--json	(Optional) Displays JSON output.

Parameter	Full Name	Description
-s	--server	Indicates that the specified data encryption or data integrity configuration is for the server. ( <code>--client</code> and <code>--server</code> are mutually exclusive.)
-t	--connectionType	<p>Specifies how Oracle Net Services data encryption or data integrity is negotiated with clients. The following values are listed in the order of increasing security:</p> <p>REJECTED - Do not enable data encryption or data integrity, even if required by the client.</p> <p>ACCEPTED - Enable data encryption or data integrity if required or requested by the client.</p> <p>REQUESTED - Enable data encryption or data integrity if the client permits it.</p> <p>REQUIRED - Enable data encryption or data integrity or preclude the connection.</p> <p>For detailed information about network data encryption and integrity, see <a href="https://docs.oracle.com/database/121/DBSEG/asoconfig.htm#DBSEG1047">https://docs.oracle.com/database/121/DBSEG/asoconfig.htm#DBSEG1047</a>.</p>

*EXAMPLE*

The following command updates the connection type to ACCEPTED:

```
[root@dbsys ~]# dbcli update-netsecurity -H a2ffbb07-c9c0-4467-a458-bce4d3b76cd5 -t ACCEPTED
```

**Node Command****DBCLI LIST-NODES**

Use the `dbcli list-nodes` command to display a list of nodes, including the node numbers.

## CHAPTER 11 Database

### SYNTAX

```
dbcli list-nodes [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays a list of nodes:

```
[root@dbsys ~]# dbcli list-nodes
node Number node Name ilom Name IP Address Subnet Mask Gate

0 rac21 N/A N/A N/A N/A
1 rac22 N/A N/A N/A N/A
```

## Objectstoreswift Commands

You can back up a database to an existing bucket in the Oracle Cloud Infrastructure Object Storage service by using the [dbcli create-backup](#) command, but first you'll need to:

1. Create an object store on the DB system, which contains the endpoint and credentials to access Object Storage, by using the [dbcli create-objectstoreswift](#) command.
2. Create a backup configuration that refers to the object store ID and the bucket name by using the [dbcli create-backupconfig](#) command.
3. Associate the backup configuration with the database by using the [dbcli update-database](#) command.

The following commands are available to manage object stores.

## CHAPTER 11 Database

---

- [dbcli create-objectstoreswift](#)
- [dbcli describe-objectstoreswift](#)
- [dbcli list-objectstoreswifts](#)

### DBCLI CREATE-OBJECTSTORESWIFT

Use the `dbcli create-objectstoreswift` command to create an object store.

#### SYNTAX

```
dbcli create-objectstoreswift -n <object_store_name> -t <tenant_name> -u <user_name> -e
https://swiftobjectstorage.<region_name>.oraclecloud.com/v1 -p [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-e	--endpointurl	The following endpoint URL.  https://swiftobjectstorage.<region_name>.oraclecloud.com/v1  See <a href="#">Regions and Availability Domains</a> for region name strings.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.
-n	--name	The name for the object store to be created.

Parameter	Full Name	Description
-p	-- swiftpassword	<p>The auth token that you generated by using the Console or IAM API. For information about generating an auth token for use with Swift, see <a href="#">Managing User Credentials</a>.</p> <p>This is <b>not</b> the password for the Oracle Cloud Infrastructure user.</p> <p>Specify <code>-p</code> (with no password) to be prompted.</p> <p>Specify <code>-hp "&lt;password&gt;"</code> in quotes to provide the password (auth token) in the command.</p>
-t	--tenantname	<p>The case-sensitive tenant name that you specify when signing in to the Console.</p>
-u	--username	<p>The user name for the Oracle Cloud Infrastructure user account, for example:</p> <pre>-u djones@example.com</pre> <p>This is the user name you use to sign in to the Console.</p> <p>The user name must have tenancy-level access to the Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to <i>all</i> of the cloud services. Instead, an administrator can create a policy that allows tenancy-level access to just Object Storage. The following is an example of such a policy.</p> <pre>Allow group DBAdmins to manage buckets in tenancy</pre> <pre>Allow group DBAdmins to manage objects in tenancy</pre> <p>For more information about adding a user to a group, see <a href="#">Managing Groups</a>. For more information about policies, see <a href="#">Getting Started with Policies</a>.</p>

## CHAPTER 11 Database

### EXAMPLE

The following command creates an object store and prompts for the Swift password:

```
[root@dbsys ~]# dbcli create-objectstoreswift -n r2swift -t CompanyABC -u djones@example.com -e
https://swiftobjectstorage.<region_name>.oraclecloud.com/v1 -p
Password for Swift:
{
 "jobId" : "c565bb71-f67b-4fab-9d6f-a34eae36feb7",
 "status" : "Created",
 "message" : "Create object store swift",
 "reports" : [],
 "createTimestamp" : "January 19, 2017 11:11:33 AM PST",
 "resourceList" : [{
 "resourceId" : "8a0fe039-f5d4-426a-8707-256c612b3a30",
 "resourceType" : "ObjectStoreSwift",
 "jobId" : "c565bb71-f67b-4fab-9d6f-a34eae36feb7",
 "updatedAt" : "January 19, 2017 11:11:33 AM PST"
 }],
 "description" : "create object store:biyanr2swift",
 "updatedAt" : "January 19, 2017 11:11:33 AM PST"
}
```

### DBCLI DESCRIBE-OBJECTSTORESWIFT

Use the `dbcli describe-objectstoreswift` command to display details about an object store.

### SYNTAX

```
dbcli describe-objectstoreswift -i <object_store_swift_id> -in <object_store_swift_name> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--objectstoreswiftid	The object store ID. Use the <code>dbcli list-objectstoreswifts</code> command to get the ID.

## CHAPTER 11 Database

Parameter	Full Name	Description
-in	--objectstoreswiftName	The object store name. Use the <code>dbcli list-objectstoreswifts</code> command to get the name.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command displays details about an object store:

```
[root@dbsys ~]# dbcli describe-objectstoreswift -i 910e9e2d-25b4-49b4-b88e-ff0332f7df87
Object Store details

 ID: 910e9e2d-25b4-49b4-b88e-ff0332f7df87
 Name: objstrswift15
 UserName: djones@example.com
 TenantName: CompanyABC
endpoint URL: https://swiftobjectstorage.<region_name>.oraclecloud.com/v1
 CreatedTime: November 16, 2016 11:25:34 PM UTC
 UpdatedTime: November 16, 2016 11:25:34 PM UTC
```

### DBCLI LIST-OBJECTSTORESWIFTS

Use the `dbcli list-objectstoreswifts` command to list the object stores on a DB system.

### SYNTAX

```
dbcli list-objectstoreswifts [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command lists the object stores on the DB system:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli list-objectstoreswifts
```

ID	Name	UserName	TenantName	Url
2915bc6a-6866-436a-a38c-32302c7c4d8b	swiftobjstr1	djones@example.com	LargeComputers	https://swiftobjectstorage.<region_name>.oraclecloud.com/v1
910e9e2d-25b4-49b4-b88e-ff0332f7df87	objstrswift15	djones@example.com	LargeComputers	https://swiftobjectstorage.<region_name>.oraclecloud.com/v1

### Pendingjob Command

#### DBCLI LIST-PENDINGJOBS

Use the `dbcli list-pendingjobs` command to display a list of pending jobs.

#### SYNTAX

```
dbcli list-pendingjobs [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### Rmanbackupreport Commands

The following commands are available to manage RMAN backup reports:

- [dbcli create-rmanbackupreport](#)
- [dbcli delete-rmanbackupreport](#)
- [dbcli describe-rmanbackupreport](#)
- [dbcli list-rmanbackupreports](#)

## CHAPTER 11 Database

---

### DBCLI CREATE-RMANBACKUPREPORT

Use the `dbcli create-rmanbackupreport` command to create an RMAN backup report.

#### SYNTAX

```
dbcli create-rmanbackupreport -w {summary|detailed} -rn <name> [-i <db_id>] [-in <db_name>] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbname	(Optional) Database resource name.
-j	--json	(Optional) Displays JSON output.
-rn	--rptname	RMAN backup report name. Maximum number of characters: 30. Wrap name in single quotes when special characters are used.
-w	--reporttype	RMAN backup report type. Possible values: summary or detailed.

### DBCLI DELETE-RMANBACKUPREPORT

Use the `dbcli delete-rmanbackupreport` command to delete an RMAN backup report.

#### SYNTAX

```
dbcli delete-rmanbackupreport [-d <db_id>] [-dn <db_name>] [-n <number>] [-i <rpt_id>] [-in <rpt_name>] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-d	--dbid	(Optional) Database resource ID.
-dn	--dbname	(Optional) Database resource name.
-h	--help	(Optional) Displays help for using the command.
-i	--reportid	(Optional) RMAN backup report ID
-in	--rptname	(Optional) RMAN backup report name
-j	--json	(Optional) Displays JSON output.
-n	--numofday	(Optional) Number of days since created (provided with Database ID/Database Name)

### DBCLI DESCRIBE-RMANBACKUPREPORT

Use the `dbcli describe-rmanbackupreport` command to

### SYNTAX

```
dbcli describe-rmanbackupreport [-i <rpt_id>] [-in <rpt_name>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--id	(Optional) RMAN backup report ID
-in	--name	(Optional) RMAN backup report name
-j	--json	(Optional) Displays JSON output.

## CHAPTER 11 Database

---

### DBCLI LIST-RMANBACKUPREPORTS

Use the `dbcli list-rmanbackupreports` command to

#### SYNTAX

```
dbcli list-rmanbackupreports [-i <db_id>] [-in <db_name>] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbid	(Optional) Database resource ID.
-in	--dbName	(Optional) Database resource name.
-j	--json	(Optional) Displays JSON output.

### Schedule Commands

The following commands are available to manage schedules:

- [dbcli describe-schedule](#)
- [dbcli list-schedules](#)
- [dbcli update-schedule](#)

### DBCLI DESCRIBE-SCHEDULE

Use the `dbcli describe-schedule` command to describe a schedule.

#### SYNTAX

```
dbcli describe-schedule -i <id> [-h] [-j]
```

## CHAPTER 11 Database

---

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--scheduleid	Schedule ID.
-j	--json	(Optional) Displays JSON output.

### DBCLI LIST-SCHEDULES

Use the `dbcli list-schedules` command to list schedules.

### SYNTAX

```
dbcli list-schedules [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### DBCLI UPDATE-SCHEDULE

Use the `dbcli update-schedule` command to update a schedule.

### SYNTAX

```
dbcli update-schedule -i <id> [-x <expression>] [-t <description>] [-d] [-e] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-d	--disable	(Optional) Disables the schedule.
-e	--enable	(Optional) Enables the schedule.
-h	--help	(Optional) Displays help for using the command.
-i	--scheduleid	Schedule ID.
-j	--json	(Optional) Displays JSON output.
-t	--description	(Optional) Description
-x	-- cronExpression	(Optional) Cron expression. Use <a href="https://cronmaker.com">cronmaker.com</a> to generate a valid cron expression.

### Scheduledexecution Command

#### DBCLI LIST-SCHEDULEDEXECUTIONS

Use the `dbcli list-scheduledExecutions` command to list scheduled executions.

#### SYNTAX

```
dbcli list-scheduledExecutions [-e <execution_id>] [-i <schedule_id>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-e	--executionid	(Optional) Execution ID.
-h	--help	(Optional) Displays help for using the command.

Parameter	Full Name	Description
-i	--scheduleid	(Optional) Schedule ID.
-j	--json	(Optional) Displays JSON output.

## Server Command

### DBCLI UPDATE-SERVER



#### Tip

Your DB system might not include this newer command. If you have trouble running the command, use the [cliadm update-dbcli](#) command to update the database CLI and then retry the command.

Use the `dbcli update-server` command to apply patches to the server components in the DB system. For more information about applying patches, see [Patching a DB System](#).

#### SYNTAX

```
dbcli update-server [-n <number>] [--local] [--precheck] [-h] [-j]
```

#### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

## CHAPTER 11 Database

Parameter	Full Name	Description
-l	--local	(Optional) Performs the operation on the local node of a multi-node high availability (HA) system. This parameter is not needed to perform the operation on a single-node system.
-n	--node	(Optional) Node number to be updated. Use the <code>dbcli list-nodes</code> command to get the node number.
-p	--precheck	(Optional) Runs precheck operations to check prerequisites.

### EXAMPLES

The following commands update the server and show the output from the update job:

```
[root@dbsys ~]# dbcli update-server
{
 "jobId" : "9a02d111-e902-4e94-bc6b-9b820ddf6ed8",
 "status" : "Created",
 "reports" : [],
 "createTimestamp" : "January 19, 2017 09:37:11 AM PST",
 "resourceList" : [],
 "description" : "Server Patching",
 "updatedAt" : "January 19, 2017 09:37:11 AM PST"
}

dbcli describe-job -i 9a02d111-e902-4e94-bc6b-9b820ddf6ed8

Job details

 ID: 9a02d111-e902-4e94-bc6b-9b820ddf6ed8
Description: Server Patching
 Status: Running
Created: January 19, 2017 9:37:11 AM PST
Message:

Task Name Start Time End Time

 Status

Create Patching Repository Directories January 19, 2017 9:37:11 AM PST January 19, 2017 9:37:11 AM
```

## CHAPTER 11 Database

---

PST	Success		
Download latest patch metadata		January 19, 2017 9:37:11 AM PST	January 19, 2017 9:37:11 AM
PST	Success		
Update System version		January 19, 2017 9:37:11 AM PST	January 19, 2017 9:37:11 AM
PST	Success		
Update Patching Repository		January 19, 2017 9:37:11 AM PST	January 19, 2017 9:38:35 AM
PST	Success		
oda-hw-mgmt upgrade		January 19, 2017 9:38:35 AM PST	January 19, 2017 9:38:58 AM
PST	Success		
Opatch updation		January 19, 2017 9:38:58 AM PST	January 19, 2017 9:38:58 AM
PST	Success		
Patch conflict check		January 19, 2017 9:38:58 AM PST	January 19, 2017 9:42:06 AM
PST	Success		
apply clusterware patch		January 19, 2017 9:42:06 AM PST	January 19, 2017 10:02:32
AM PST	Success		
Updating GiHome version		January 19, 2017 10:02:32 AM PST	January 19, 2017 10:02:38
AM PST	Success		

The following command updates node 0 of the server only, with precheck:

```
dbcli update-server -n 0 -p
{
 "jobId" : "3e2ale3c-83d3-4101-86b8-4d525f3f8c18",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "April 26, 2019 06:07:27 AM UTC",
 "resourceList" : [],
 "description" : "Server Patching Prechecks",
 "updatedAt" : "April 26, 2019 06:07:27 AM UTC"
}
```

### System Command

#### DBCLI DESCRIBE-SYSTEM

Use the `dbcli describe-system` command to display details about the system. On a 2-node RAC DB system, the command provides information about the local node.

#### SYNTAX

```
dbcli describe-system [-b] [-d] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-b	--bom	(Optional) Displays BOM information.
-d	--details	(Optional) Displays additional information about the DB system, including dcs CLI and agent version information.
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### TDE Commands

The following commands are available to manage TDE-related items (backup reports, keys, and wallets):

- [dbcli list-tdebackupreports](#)
- [dbcli update-tdekey](#)
- [dbcli recover-tdewallet](#)

#### DBCLI LIST-TDEBACKUPREPORTS

Use the `dbcli list-tdebackupreports` command to list backup reports for TDE wallets.

#### SYNTAX

```
dbcli list-tdebackupreports [-i <db_id>] [-in <db_name>] [-h] [-j]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-i	--dbResid	(Optional) Displays the TDE Wallet backup reports for the specified database resource ID. Use the <code>dbcli list-databases</code> command to get the database resource ID.
-in	--dbResname	(Optional) Displays the TDE Wallet backup reports for the specified database resource name. Use the <code>dbcli list-databases</code> command to get the database resource name.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command lists the backup reports for TDE wallets:

```
[root@dbsys ~]# dbcli list-tdebackupreports
DbResID OraDbId BackupLocation

538ca5b1-654d-4418-8ce1-f49b6c987a60 1257156075 https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaasimage/backupptest/host724007/tdewallet/Testdb5/1257156075/2017-08-17/TDEWALLET_
BMC60_2017-08-17_10-58-17.0990.tar.gz
538ca5b1-9fb2-4245-b157-6e25d7c988c5 704287483 https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaasimage/backupptest/host724007/tdewallet/Testdb1/704287483/2017-08-17/TDEWALLET_
AUTO_2017-08-17_11-03-25.0953.tar.gz
538ca5b1-9fb2-4245-b157-6e25d7c988c5 704287483 https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaasimage/backupptest/host724007/tdewallet/Testdb1/704287483/2017-08-17/TDEWALLET_
BMC62_2017-08-17_11-04-41.0264.tar.gz
19714ffa-de1b-4433-9188-c0592887e609 1157116855 https://swiftobjectstorage.us-phoenix-
1.oraclecloud.com/v1/dbaasimage/backupptest/host724007/tdewallet/Testdb7/1157116855/2017-08-17/TDEWALLET_
AUTO_2017-08-17_11-57-47.0605.tar.gz
```

### DBCLI UPDATE-TDEKEY

Use the `dbcli update-tdekey` command to update the TDE encryption key inside the TDE wallet. You can update the encryption key for Pluggable Databases (if `-pdbNames` are

## CHAPTER 11 Database

specified), and/or the Container Database (if `-rootDatabase` is specified).

### SYNTAX

```
dbcli update-tdekey -i <db_id> -p [-all] -n <pdbname1,pdbname2> [-r|-no-r] -t <tag_name> [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-all	--allPdbNames	(Optional) Flag to rotate (update) all PDB names. To update all instead of specified PDB names, use this parameter instead of <code>-n</code> . Default: false.
-i	--databaseId	Defines the database ID for which to update the key.
-p	--password	Defines the TDE Admin wallet password. Specify <code>-p</code> with no password. You will be prompted for the password.  If you must provide the password in the command, for example in a script, use <code>-hp &lt;password&gt;</code> instead of <code>-p</code> .
-n	--pdbNames	Defines the PDB names to be rotated (updated).
-r -no-r	--rootDatabase  --no-rootDatabase	Indicates whether to rotate the key for the root database if it is a container database.
-t	-tagName	Defines the TagName used to backup the wallet. The default is <code>OdaRotateKey</code> .
-h	--help	(Optional) Displays help for using the command.
-j	--json	(Optional) Displays JSON output.

### EXAMPLE

The following command updates the key for `pdb1` and `pdb2` only:

## CHAPTER 11 Database

---

```
[root@dbsys ~]# dbcli update-tdekey -dbid ee3eaab6-a45b-4e61-a218-c4ba665503d9 -p -n pdb1,pdb2

TDE Admin wallet password:
{
 "jobId" : "08e5edb1-42e1-4d16-a47f-783c0afa4778",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1467876407035,
 "description" : "TDE update",
 "updatedAtTime" : 1467876407035
}
```

The following command updates pdb1, pdb2, and the container database:

```
[root@dbsys ~]# dbcli update-tdekey -dbid ee3eaab6-a45b-4e61-a218-c4ba665503d9 -p -n pdb1,pdb2 -r

TDE Admin wallet password:
{
 "jobId" : "c72385f0-cd81-42df-a8e8-3a1e7cab1278",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : 1467876433783,
 "description" : "TDE update",
 "updatedAtTime" : 1467876433783
}
```

### DBCLI RECOVER-TDEWALLET

Use the `dbcli recover-tdewallet` command to recover a TDE wallet.

#### SYNTAX

```
dbcli recover-tdewallet -in <db_name> -tp <password> [-l <location>] [-h] [-j]
```

### PARAMETERS

Parameter	Full Name	Description
-h	--help	(Optional) Displays help for using the command.
-in	--dbName	Database name.
-j	--json	(Optional) Displays JSON output.
-l	-- tdeWalletBackuplocation	(Optional) TDE wallet backup location. TDE wallet should be backed up in tar.gz format.
-tp	--tdeWalletPassword	Defines the TDE Admin wallet password.

### Admin Commands

The following commands are to perform administrative actions on the DB system:

- [dbadmcli manage diagcollect](#)
- [dbadmcli power](#)
- [dbadmcli power disk status](#)
- [dbadmcli show controller](#)
- [dbadmcli show disk](#)
- [dbadmcli show diskgroup](#)
- [dbadmcli show env\\_hw](#) (environment type and hardware version)
- [dbadmcli show fs](#) (file system details)
- [dbadmcli show storage](#)
- [dbadmcli storddiag](#)

#### DBADMCLI MANAGE DIAGCOLLECT

Use the `dbadmcli manage diagcollect` command to collect diagnostic information about a DB system for troubleshooting purposes, and for working with Oracle Support Services.

## CHAPTER 11 Database

---

### SYNTAX

```
dbadmcli manage diagcollect --storage [-h]
```

### PARAMETERS

Parameter	Description
-h	(Optional) Displays help for using the command.
--storage	Collects all of the logs for any storage issues.

### EXAMPLE

```
[root@dbsys ~]# dbadmcli manage diagcollect --storage
Collecting storage log data. It will take a while, please wait...
Collecting oak data. It will take a while, please wait...
tar: Removing leading `/' from member names
tar: /opt/oracle/oak/onecmd/tmp/OakCli-Command-Output.log: file changed as we read it

Logs are collected to : /opt/oracle/oak/log/dbsys/oakdiag/oakStorage-dbsys-20161118_2101.tar.gz
```

### DBADMCLI POWER

Use the `dbadmcli power` command to power a disk on or off.



#### Note

The `dbadmcli power` command is not available on 2-node RAC DB systems.

### SYNTAX

```
dbadmcli power {-on|-off} <name> [-h]
```

## CHAPTER 11 Database

### PARAMETERS

Parameter	Description
-h	(Optional) Displays help for using the command.
<i>name</i>	Defines the disk resource name. The resource name format is pd_[0..3]. Use the <code>dbadmcli show disk</code> command to get the disk resource name.
-off	Powers off the disk.
-on	Powers on the disk.

### DBADMCLI POWER DISK STATUS

Use the `dbadmcli power disk status` command to display the current power status of a disk.

### SYNTAX

```
dbadmcli power disk status <name> [-h]
```

### PARAMETERS

Parameter	Description
-h	(Optional) Displays help for using the command.
<i>name</i>	Identifies a specific disk resource name. The resource name format is pd_[0..3]. For example, <code>pd_01</code> .

### EXAMPLE

```
[root@dbsys ~]# dbadmcli power disk status pd_00
```

```
The disk is powered ON
```

### DBADMCLI SHOW CONTROLLER

Use the `dbadmcli show controller` command to display details of the controller.

## CHAPTER 11 Database

### SYNTAX

```
dbadmcli show controller <controller_id> [-h]
```

### PARAMETER

Parameter	Description
<i>controller_id</i>	The ID number of the controller. Use the <code>dbadmcli show storage</code> command to get the ID.
-h	(Optional) Displays help for using the command.

### DBADMCLI SHOW DISK

Use the `dbadmcli show disk` command to display the status of a single disk or all disks on the DB system.

### SYNTAX

```
dbadmcli show disk [<name>] [-shared] [-all] [-getlog] [-h]
```

### PARAMETERS

Parameter	Description
-all	(Optional) Displays detailed information for the named disk.
-h	(Optional) Displays help for using the command.
-getlog	(Optional) Displays all the SMART log entries for an NVMe disk.
<i>name</i>	(Optional) Identifies a specific disk resource name. The resource name format is <code>pd_[0..3]</code> . If omitted, the command displays information about all disks on the system.
-shared	(Optional) Displays all the shared disks.

### EXAMPLES

To display the status of all the disks on the system:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbadmcli show disk
```

NAME	PATH	TYPE	STATE	STATE_DETAILS
pd_00	/dev/nvme2n1	NVD	ONLINE	Good
pd_01	/dev/nvme3n1	NVD	ONLINE	Good
pd_02	/dev/nvme1n1	NVD	ONLINE	Good
pd_03	/dev/nvme0n1	NVD	ONLINE	Good

To display the status of a disk named pd\_00:

```
[root@dbsys ~]# dbadmcli show disk pd_00
```

The Resource is : pd\_00

```
ActionTimeout : 1500
ActivePath : /dev/nvme2n1
AsmDiskList : |data_00||reco_00|
AutoDiscovery : 1
AutoDiscoveryHi : |data:70:NVD||reco:30:NVD|
CheckInterval : 300
ColNum : 0
CriticalWarning : 0
DependListOpr : add
Dependency : |0|
DiskId : 360025380144d5332
DiskType : NVD
Enabled : 1
ExpNum : 29
HbaPortNum : 10
IState : 0
Initialized : 0
IsConfigDepende : false
ModelNum : MS1PC2DD30RA3.2T
MonitorFlag : 1
MultiPathList : |/dev/nvme2n1|
Name : pd_00
NewPartAddr : 0
OSUserType : |userType:Multiuser|
PlatformName : X5_2_LITE_IAAS
PrevState : Invalid
PrevUsrDevName :
SectorSize : 512
SerialNum : S2LHNAAH502855
Size : 3200631791616
SlotNum : 0
```

## CHAPTER 11 Database

---

```
SmartDiskWarnin : 0
SmartTemperatur : 32
State : Online
StateChangeTs : 1467176081
StateDetails : Good
TotalSectors : 6251233968
TypeName : 0
UsrDevName : NVD_S00_S2LHNAAH502855
VendorName : Samsung
gid : 0
mode : 660
uid : 0
```

To display the SMART logs for an NVMe disk:

```
[root@dbsys ~]# dbadmcli show disk pd_00 -getlog
SMART / Health Information :

Critical Warning : Available Spare below Threshold : FALSE
Critical Warning : Temperature above Threshold : FALSE
Critical Warning : Reliability Degraded : FALSE
Critical Warning : Read-Only Mode : FALSE
Critical Warning : Volatile Memory Backup Device Failure : FALSE
Temperature : 32 degree
Celsius
Available Spare : 100%
Available Spare Threshold : 10%
Device Life Used : 0%
Data Units Read (in 512k byte data unit) : 89493
Data Units Written (in 512k byte data unit) : 270387
Number of Host Read Commands : 4588381
Number of Host Write Commands : 6237344
Controller Busy Time : 3 minutes
Number of Power Cycles : 227
Number of Power On Hours : 1115
Number of Unsafe Shutdowns : 218
Number of Media Errors : 0
Number of Error Info Log Entries : 0
```

### DBADMCLI SHOW DISKGROUP

Use the `dbadmcli show diskgroup` command to list configured diskgroups or display a specific diskgroup configuration.

## CHAPTER 11 Database

---

### SYNTAX

To list configured diskgroups:

```
dbadmcli show diskgroup [-h]
```

To display DATA configurations:

```
dbadmcli show diskgroup [DATA] [-h]
```

To display RECO configurations:

```
dbadmcli show diskgroup [RECO] [-h]
```

### PARAMETERS

Parameter	Description
DATA	(Optional) Displays the DATA diskgroup configurations.
-h	(Optional) Displays help for using the command.
RECO	(Optional) Displays the RECO diskgroup configurations.

### EXAMPLES

To list all diskgroups:

```
[root@dbsys ~]# dbadmcli show diskgroup
```

```
DiskGroups

DATA
RECO
```

To display DATA configurations:

```
[root@dbsys ~]# dbadmcli show diskgroup DATA
```

```
ASM_DISK PATH DISK STATE STATE_DETAILS
data_00 /dev/NVD_S00_S2LHNAAH101026p1 pd_00 ONLINE Good
data_01 /dev/NVD_S01_S2LHNAAH101008p1 pd_01 ONLINE Good
```

## CHAPTER 11 Database

---

### DBADMCLI SHOW ENV\_HW

Use the `dbadmcli show env_hw` command to display the environment type and hardware version of the current DB system.

#### SYNTAX

```
dbadmcli show env_hw [-h]
```

#### PARAMETER

Parameter	Description
-h	(Optional) Displays help for using the command.

### DBADMCLI SHOW FS

Use the `dbadmcli show fs` command to display file system details.

#### SYNTAX

```
dbadmcli show fs [-h]
```

#### PARAMETER

Parameter	Description
-h	(Optional) Displays help for using the command.

### DBADMCLI SHOW STORAGE

Use the `dbadmcli show storage` command to show the storage controllers, expanders, and disks.

#### SYNTAX

```
dbadmcli show storage [-h]
```

To show storage errors:

```
dbadmcli show storage -errors [-h]
```

## CHAPTER 11 Database

---

### PARAMETERS

Parameter	Description
-errors	(Optional) Shows storage errors.
-h	(Optional) Displays help for using the command.

### EXAMPLE

To display storage devices:

```
[root@dbsys ~]# dbadmcli show storage
==== BEGIN STORAGE DUMP =====
Host Description: Oracle Corporation:ORACLE SERVER X5-2
Total number of controllers: 5

 Id = 4
 Pci Slot = -1
 Serial Num =
 Vendor =
 Model =
 FwVers =
 strId = iscsi_tcp:00:00.0
 Pci Address = 00:00.0

 Id = 0
 Pci Slot = 13
 Serial Num = S2LHNAAH504431
 Vendor = Samsung
 Model = MS1PC2DD3ORA3.2T
 FwVers = KPYA8R3Q
 strId = nvme:25:00.00
 Pci Address = 25:00.0

 Id = 1
 Pci Slot = 12
 Serial Num = S2LHNAAH505449
 Vendor = Samsung
 Model = MS1PC2DD3ORA3.2T
 FwVers = KPYA8R3Q
 strId = nvme:27:00.00
 Pci Address = 27:00.0
```

## CHAPTER 11 Database

```
Id = 2
Pci Slot = 10
Serial Num = S2LHNAAH503573
Vendor = Samsung
Model = MS1PC2DD3ORA3.2T
FwVers = KPYA8R3Q
strId = nvme:29:00.00
Pci Address = 29:00.0

Id = 3
Pci Slot = 11
Serial Num = S2LHNAAH503538
Vendor = Samsung
Model = MS1PC2DD3ORA3.2T
FwVers = KPYA8R3Q
strId = nvme:2b:00.00
Pci Address = 2b:00.0

Total number of expanders: 0
Total number of PDs: 4
 /dev/nvme2n1 Samsung NVD 3200gb slot: 0 pci : 29
 /dev/nvme3n1 Samsung NVD 3200gb slot: 1 pci : 2
 /dev/nvme1n1 Samsung NVD 3200gb slot: 2 pci : 27
 /dev/nvme0n1 Samsung NVD 3200gb slot: 3 pci : 25
==== END STORAGE DUMP =====
```

### DBADMCLI STORDIAG

Use the `dbadmcli stordiag` command to collect detailed information for each disk or NVM Express (NVMe).

#### SYNTAX

```
dbadmcli stordiag <name> [-h]
```

### PARAMETERS

Parameter	Description
<i>name</i>	Defines the disk resource name. The resource name format is pd_[0..3].
-h	(Optional) Displays help for using the command.

### EXAMPLE

To display detailed information for NVMe pd\_00:

```
[root@dbsys ~]# dbadmcli stordiag pd_0
```

## Database Sizing Templates

When you create a database using the `dbcli create-database` command, you can specify a database sizing template with the `--dbshape` parameter. The sizing templates are configured for different types of database workloads. Choose the template that best matches the most common workload your database performs:

- Use the OLTP templates if your database workload is primarily online transaction processing (OLTP).
- Use the DSS templates if your database workload is primarily decision support (DSS) or data warehousing.
- Use the in-memory (IMDB) templates if your database workload can fit in memory, and can benefit from in-memory performance capabilities.

The following tables describe the templates for each type of workload.

## CHAPTER 11 Database

### OLTP DATABASE SIZING TEMPLATES

Template	CPU Cores	SGA (GB)	PGA (GB)	Flash (GB)	Processes	Redo Log File Size (GB)	Log Buffer (MB)
odb1s	1	2	1	6	200	1	16
odb1	1	4	2	12	200	1	16
odb2	2	8	4	24	400	1	16
odb4	4	16	8	48	800	1	32
odb6	6	24	12	72	1200	2	64
odb8	8	32	16	n/a	1600	2	64
odb10	10	40	20	n/a	2000	2	64
odb12	12	48	24	144	2400	4	64
odb16	16	64	32	192	3200	4	64
odb20	20	80	40	n/a	4000	4	64
odb24	24	96	48	192	4800	4	64
odb32	32	128	64	256	6400	4	64
odb36	36	128	64	256	7200	4	64

## CHAPTER 11 Database

### DSS DATABASE SIZING TEMPLATES

Template	CPU Cores	SGA (GB)	PGA (GB)	Processes	Redo Log File Size (GB)	Log Buffer (MB)
odb1s	1	1	2	200	1	16
odb1	1	2	4	200	1	16
odb2	2	4	8	400	1	16
odb4	4	8	16	800	1	32
odb6	6	12	24	1200	2	64
odb8	8	16	32	1600	2	64
odb10	10	20	40	2000	2	64
odb12	12	24	48	2400	4	64
odb16	16	32	64	3200	4	64
odb20	20	40	80	4000	4	64
odb24	24	48	96	4800	4	64
odb32	32	64	128	6400	4	64
odb36	36	64	128	7200	4	64

## IN-MEMORY DATABASE SIZING TEMPLATES

Template	CPU Cores	SGA (GB)	PGA (GB)	In-Memory (GB)	Processes	Redo Log File Size (GB)	Log Buffer (MB)
odb1s	1	2	1	1	200	1	16
odb1	1	4	2	2	200	1	16
odb2	2	8	4	4	400	1	16
odb4	4	16	8	8	800	1	32
odb6	6	24	12	12	1200	2	64
odb8	8	32	16	16	1600	2	64
odb10	10	40	20	20	2000	2	64
odb12	12	48	24	24	2400	4	64
odb16	16	64	32	32	3200	4	64
odb20	20	80	40	40	4000	4	64
odb24	24	96	48	48	4800	4	64
odb32	32	128	64	64	6400	4	64
odb36	36	128	64	64	7200	4	64

## Storage Scaling Considerations for Virtual Machine Databases Using Fast Provisioning



### Note

This topic applies only to 1-node virtual machine DB systems.

When you provision a virtual machine DB system using the [fast provisioning option](#), the **Available storage (GB)** value you specify during provisioning determines the maximum total storage available through scaling. The following table details the maximum storage value available through scaling for each setting offered in the provisioning workflow:

Initial storage specified during provisioning (GB)	Maximum storage available through scaling (GB)
256	2560
512	2560
1024	5120
2048	10240
4096	20480
8192	40960

For more information on creating a virtual machine DB system, see [Creating Bare Metal and Virtual Machine DB Systems](#).

## Database Metrics

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure Database service resources by using metrics, alarms, and [notifications](#). For more

information, see [Monitoring Overview](#) and [Notifications Overview](#).



### Note

Database metrics are currently available only for Autonomous Databases in serverless deployments.

This topic describes the metrics emitted by the Database service in the `oci_autonomous_database` namespace.

Resources: Autonomous Databases.

## Overview of the Autonomous Database Metrics

The Database service metrics help you measure useful quantitative data about your Autonomous Databases, such as CPU and storage utilization, the number of successful and failed database logon and connection attempts, database operations, SQL queries, and transactions, and so on. You can use metrics data to diagnose and troubleshoot problems with Autonomous Databases. For a complete list of available metrics for Autonomous Databases, see [Available Metrics: oci\\_autonomous\\_database](#).

To view a default set of metrics charts in the Console, navigate to the Autonomous Database that you're interested in, and then click **Metrics**. You also can use the Monitoring service to create [custom queries](#).

## Prerequisites

**IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on

user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_autonomous\_database

The metrics listed in the following table are automatically available for any Autonomous Database that you create. You do not need to enable monitoring on the resource to get these metrics.

Database service metrics for Autonomous Databases include the following dimensions:

#### **AUTONOMOUSDBTYPE**

The type of Autonomous Database, Autonomous Data Warehouse (ADW) or Autonomous Transaction Processing (ATP).

#### **DISPLAYNAME**

The friendly name of the Autonomous Database.

#### **REGION**

The region in which the Autonomous Database resides.

#### **RESOURCEID**

The OCID of the Autonomous Database.

#### **RESOURCENAME**

The name of the Autonomous Database.

In the following table, metrics that are marked with an asterisk (\*) can be viewed only on the **Service Metrics** page of the Oracle Cloud Infrastructure Console.

Metric	Metric Display Name	Unit	Description	Dimensions
CurrentLogons*	<b>Current Logons</b>	count	The number of successful logons during the selected interval.	AutonomousDBType displayName region
CpuUtilization	<b>CPU Utilization</b>	percent	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.	resourceId resourceName
ExecuteCount	<b>Execute Count</b>	count	The number of user and recursive calls that executed SQL statements during the selected interval.	
FailedConnections*	<b>Failed Connections</b>	count	The number of failed database connections.	

Metric	Metric Display Name	Unit	Description	Dimensions
FailedLogons	<b>Failed Logons</b>	count	The number of logons that failed because of an invalid username and/or password, during the selected interval.	
ParseCount*	<b>Parse Count (Total)</b>	count	The number of hard and soft parses during the selected interval.	
QueuedStatements	<b>Queued Statements</b>	count	The number of queued SQL statements, aggregated across all consumer groups, during the selected interval.	
RunningStatements	<b>Running Statements</b>	count	The number of running SQL statements, aggregated across all consumer groups, during the selected interval.	

Metric	Metric Display Name	Unit	Description	Dimensions
Sessions	<b>Sessions</b>	count	The number of sessions in the database.	
StorageUtilization	<b>Storage Utilization</b>	percent	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	
TransactionCount*	<b>Transaction Count</b>	count	The combined number of user commits and user rollbacks during the selected interval.	
UserCalls*	<b>User Calls</b>	count	The combined number of logons, parses, and execute calls during the selected interval.	

## Using the Console

### To view default metric charts for a single Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction**

### **Processing or Autonomous Data Warehouse.**

2. Choose the **Compartment** that contains the Autonomous Database you want to view, and then click display name of the database to view its details.
3. Under **Resources**, click **Metrics**.  
The **Metrics** page displays a default set of charts for the current Autonomous Database. See [Available Metrics: oci\\_autonomous\\_database](#) for information about the default charts.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for multiple Autonomous Databases

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that contains the Autonomous Databases that you're interested in.
3. For **Metric Namespace**, select **oci\_autonomous\_database**.  
The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.



#### **Tip**

If there are multiple Autonomous Databases in the compartment, the charts default to show a separate line for each master encryption key. You can instead show a single line aggregated across all Autonomous Databases in the compartment by selecting the **Aggregate Metric Streams** check box.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

## Using Performance Hub to Analyze Database Performance in Oracle Cloud Infrastructure



### Note

Performance Hub is currently available for Autonomous Databases.

This topic describes how to use Oracle Cloud Infrastructure's Performance Hub tool for Oracle Database performance analysis and tuning. Performance Hub offers two performance monitoring tools in the Console, Active Session History (ASH) Analytics and SQL Monitoring.

You can view real-time and historical performance data in Performance Hub. When you view historical data in the Performance Hub, you are viewing statistics collected as part of the hourly snapshots of your database.

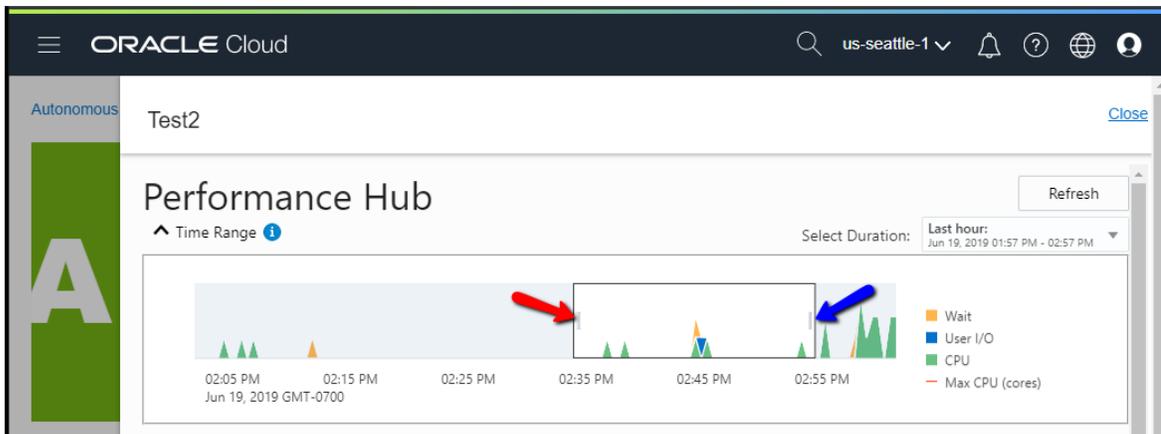
The Performance Hub page consists of the following sections:

- **Time Range** field and time slider:

The **Time Range** selector is displayed on the top of the Performance Hub page. Use the **Select Duration** field to set the time duration. By default, **Last 60 mins** is selected. You can choose to view **Last 8 hours**, **Last 24 hours**, **Last week**, or specify a custom time range using the **Custom** option.

The Time Range field shows active sessions in chart form for the time period selected. The active sessions chart displays the average number of active sessions broken down by **CPU**, **User I/O**, and **Wait**.

The sliding box on the time range chart is known as the **time slider**. Use the time slider to select the exact period of time for which data is displayed in the Performance Hub tables and graphs. This is a subsection of the period of time shown in the **Time Range** field. In the image that follows, red and blue arrows point to the vertical 'handlebar' elements on the left and right boundaries of the slider box.



You can slide the box to the left or the right to shift the time period under analysis. To slide the entire box, left-click anywhere inside the box and drag the box to the left or the right. You can widen or narrow the box to increase or decrease the length of time under analysis. To widen or narrow the box, left-click and hold the handlebar on either side of the box, then drag to the left or the right to increase or decrease the size of the box.

To refresh the data in Performance Hub according to the time range chosen, click the **Refresh** button.

- **Active Session History (ASH) Analytics** tab:

Displayed by default. The ASH Analytics tab shows Active Session History (ASH) analytics charts to explore ASH data. It allows you to drill down into database performance across multiple dimensions such as **Consumer Group**, **Wait Class**, **SQL ID**, and **User Name**. On the ASH Analytics tab, you can select an Average Active Sessions dimension and view the top activity for that dimension for the selected time period. For information on ASH, see [Active Session History \(ASH\)](#) in *Oracle Database Concepts*.

- **SQL Monitoring** tab:

The SQL Monitoring tab is not displayed by default. To view it, click **SQL Monitoring** on the Performance Hub page.

SQL statements are only monitored if they've been running for at least five seconds or if they're run in parallel. The table in this section displays monitored SQL statement executions by dimensions including **Last Active Time**, **CPU Time**, and **Database Time**. The table displays currently running SQL statements and SQL statements that completed, failed, or were terminated. The columns in the table provide information for monitored SQL statements including **Status**, **Duration**, and **SQL ID**.

The **Status** column has the following icons:

- A spinning icon indicates that the SQL statement is executing.
- A green check mark icon indicates that the SQL statement completed its execution during the specified time period.
- A red cross icon indicates that the SQL statement did not complete, either due to an error, or due to the session being terminated.
- A clock icon indicates that the SQL statement is queued.

To terminate a running or queued SQL statement, click **Kill Session**.

You can also click an SQL ID to go to the corresponding **Real-time SQL Monitoring** page. This page provides additional details to help you tune the selected SQL statement.

### Using the Oracle Cloud Infrastructure Console

#### To navigate to Performance Hub in the Oracle Cloud Infrastructure Console interface of an Autonomous Database

1. Open the navigation menu. Under **Database**, click **Autonomous Transaction Processing** or **Autonomous Data Warehouse**.
2. Choose your **Compartment**.
3. In the list of Autonomous Databases, click on the display name of the database you wish to analyze using Performance Hub reports.
4. Click **Performance Hub**.

#### To view the average active session data by a selected dimension

1. Go to the **Performance Hub** page of the Oracle Cloud Infrastructure Console for the database which you wish to manage. See [To navigate to Performance Hub in the Oracle Cloud Infrastructure Console interface of an Autonomous Database](#) for more information.
  - The database name is displayed at the top of the Performance Hub page.
  - The time period for which information is available on the Performance Hub is displayed in the **Time Range** field. The selected time period is indicated on the time slider graph by the adjustable blue-colored block.  
The **ASH Analytics** tab is displayed with the top activity for a selected dimension in the selected time period.
2. Use the **Select Duration** selector to set the exact period of time for which data is displayed in the ASH Analytics tables and graphs. By default, the last hour is selected. The time duration is the total amount of time available for analysis.
3. Use the blue box on the **time slider** to further narrow down the time period for which performance data is displayed on the **ASH Analytics** tab.

4. Select a dimension in the **Average Active Sessions** drop-down list to display ASH analytics by that particular dimension. By default, the **Consumer Group** dimension is selected and the data is categorized by the **High, Medium, or Low** service name associated with the Autonomous Database.

Optionally, you can:

- Click the **Max Threads** check box to view the number of Max CPU Threads. These are denoted by a red line on the chart.
  - Click the **Total Activity** check box to view a black border that denotes total activity of all the components of the selected dimension on the chart. This option is selected by default when you use the filtering capabilities to only view the data for a particular component within a dimension. For information on filtering Average Active Sessions data, see [Filter Average Active Sessions Data](#).
5. For the dimension selected in the **Average Active Sessions** drop-down list, you can further drill down into session details by selecting dimensions in the two sections at the bottom of the **ASH Analytics** tab. By default, the following dimensions are selected:
    - **SQL ID by Consumer Group**, which displays the SQL statements with the top average active sessions activity for consumer groups for the selected time period. You can right-click the bar charts to sort the SQL statements in ascending or descending order or click the SQL ID to go the SQL Details page.
    - **User Session by Consumer Group**, which displays the user sessions with the top average active sessions activity for consumer groups for the selected time period. You can right-click the bar charts to sort the user sessions in ascending or descending order or click the user session to go to the User Session page.

### To filter average active sessions data

1. Go to the **Performance Hub** page of the Oracle Cloud Infrastructure Console for the database which you wish to manage. See [To navigate to Performance Hub in the Oracle Cloud Infrastructure Console interface of an Autonomous Database](#) for more information.

- The database name is displayed at the top of the Performance Hub page.
  - The time period for which information is available on the Performance Hub is displayed in the **Time Range** field. The selected time period is indicated on the time slider graph by the adjustable blue-colored block.  
The **ASH Analytics** tab is displayed with the top activity for a selected dimension in the selected time period.
2. Use the **Select Duration** selector to set the exact period of time for which data is displayed in the ASH Analytics tables and graphs. By default, the last hour is selected. The time duration is the total amount of time available for analysis.
  3. Use the blue box on the **time slider** to further narrow down the time period for which performance data is displayed on the **ASH Analytics** tab.
  4. In the **ASH Analytics** tab, select a dimension in the Average Active Sessions by drop-down list. By default, **Consumer Group** is selected.  
The chart is displayed. Each color in the chart denotes a component of the selected dimension. For example, the Consumer Group dimension has **High, Medium, and Low**, which are predefined service names assigned to your Autonomous Database to provide different levels of concurrency and performance.
  5. Click a component in the legend. The selected component is displayed in the **Applied Filters** field and the chart is updated to only display data pertaining to that component. The total activity, which includes all the components of the dimension, is denoted by a black outline and is displayed by default when you filter data.

### To view the SQL Monitoring report

1. Go to the **Performance Hub** page of the Oracle Cloud Infrastructure Console for the database which you wish to manage. See [To navigate to Performance Hub in the Oracle Cloud Infrastructure Console interface of an Autonomous Database](#) for more information.

- The database name is displayed at the top of the Performance Hub page.
  - The time period for which information is available on the Performance Hub is displayed in the **Time Range** field. The selected time period is indicated on the time slider graph by the adjustable blue-colored block.
2. Click **SQL Monitoring** to view the SQL monitoring tab.
  3. Optionally, you can get detailed information on a specific SQL statements by clicking an ID number in the **SQL ID column**. When you click an ID number, the Real-time SQL Monitoring page is displayed.
  4. Click **Download Report** to download the report data for your selected SQL statement.

## Migrating Databases to the Cloud

You can migrate your on-premises Oracle Database to an Oracle Cloud Infrastructure Database service database using a number of different methods that use several different tools. The method that applies to a given migration scenario depends on several factors, including the version, character set, and platform endian format of the source and target databases.

### Choosing a Migration Method

Not all migration methods apply to all migration scenarios. Many of the migration methods apply only if specific characteristics of the source and destination databases match or are compatible. Moreover, additional factors can affect which method you choose for your migration from among the methods that are technically applicable to your migration scenario.

Some of the characteristics and factors to consider when choosing a migration method are:

- On-premises database version
- Database service database version
- On-premises host operating system and version
- On-premises database character set
- Quantity of data, including indexes

- Data types used in the on-premises database
- Storage for data staging
- Acceptable length of system outage
- Network bandwidth

To determine which migration methods are applicable to your migration scenario, gather the following information.

1. Database version of your on-premises database:
  - Oracle Database 12c Release 2 version 12.2.0.1
  - Oracle Database 12c Release 1 version 12.1.0.2 or higher
  - Oracle Database 12c Release 1 version lower than 12.1.0.2
  - Oracle Database 11g Release 2 version 11.2.0.3 or higher
  - Oracle Database 11g Release 2 version lower than 11.2.0.3
2. For on-premises Oracle Database 12c Release 2 and Oracle Database 12c Release 1 databases, the architecture of the database:
  - Multitenant container database (CDB)
  - Non-CDB
3. Endian format (byte ordering) of your on-premises database's host platform  
Some platforms are little endian and others are big endian. Query `V$TRANSPORTABLE_PLATFORM` to identify the endian format, and to determine whether cross-platform tablespace transport is supported.  
The Oracle Cloud Infrastructure Database uses the Linux platform, which is little endian.
4. Database character set of your on-premises database and the Oracle Cloud Infrastructure Database database.  
Some migration methods require that the source and target databases use compatible database character sets.
5. Database version of the Oracle Cloud Infrastructure Database database you are migrating to:

- Oracle Database 12c Release 2
- Oracle Database 12c Release 1
- Oracle Database 11g Release 2

Oracle Database 12c Release 2 and Oracle Database 12c Release 1 databases created on the Database service use CDB architecture. Databases created using the Enterprise Edition software edition are single-tenant, and databases created using the High Performance or Extreme Performance software editions are multitenant.

After gathering this information, use the “source” and “destination” database versions as your guide to see which migration methods apply to your migration scenario:

- [Migrating from Oracle Database 11g to Oracle Database 11g in the Cloud](#)
- [Migrating from Oracle Database 11g to Oracle Database 12c in the Cloud](#)
- [Migrating from Oracle Database 12c CDB to Oracle Database 12c in the Cloud](#)
- [Migrating from Oracle Database 12c Non-CDB to Oracle Database 12c in the Cloud](#)

### Migration Connectivity Options

You have several connectivity options when migrating your on-premises databases to the Oracle Cloud Infrastructure. The options are listed below in order of preference.

1. **FastConnect:** Provides a secure connection between your existing network and your virtual cloud network (VCN) over a private physical network instead of the internet. For more information, see [FastConnect](#).
2. **IPSec VPN:** Provides a secure connection between a dynamic routing gateway (DRG) and customer-premise equipment (CPE), consisting of multiple IPSec tunnels. The IPSec connection is one of the components forming a site-to-site VPN between a VCN and your on-premises network. For more information, see [VPN Connect](#).
3. **Internet gateway:** Provides a path for network traffic between your VCN and the internet. For more information, see [Internet Gateway](#).

### Migration Methods

Many methods exist to migrate Oracle databases to the Oracle Cloud Infrastructure Database service. Which of these methods apply to a given migration scenario depends on several factors, including the version, character set, and platform endian format of the source and target databases.

- [Data Pump Conventional Export/Import](#)
- [Data Pump Full Transportable](#)
- [Data Pump Transportable Tablespace](#)
- [Remote Cloning a PDB](#)
- [Remote Cloning Non-CDB](#)
- [RMAN Cross-Platform Transportable PDB](#)
- [RMAN Cross-Platform Transportable Tablespace Backup Sets](#)
- [RMAN Transportable Tablespace with Data Pump](#)
- [RMAN DUPLICATE from an Active Database](#)
- [RMAN CONVERT Transportable Tablespace with Data Pump](#)
- [SQL Developer and INSERT Statements to Migrate Selected Objects](#)
- [SQL Developer and SQL\\*Loader to Migrate Selected Objects](#)
- [Unplugging/Plugging a PDB](#)
- [Unplugging/Plugging Non-CDB](#)

### Migrating an On-Premises Database to Oracle Cloud Infrastructure by Creating a Backup in the Cloud



#### Note

This topic is not applicable to Exadata DB systems.

## CHAPTER 11 Database

---

You can migrate an on-premises database to Oracle Cloud Infrastructure by creating a backup of your on-premises database in Oracle Cloud Infrastructure's Database service.

Oracle provides a Python script to create a backup of your database. The script invokes an API call to create the backup and then places the backup in Oracle Cloud Infrastructure. You can then use the Console or the API to [create a new database or DB system](#) from that backup. Backups created using the instructions in this topic appear under Standalone Backups in the console.

The Python script is bundled as a part of the Oracle Cloud Infrastructure CLI installation. Oracle provides the migration script and associated files at no cost. Normal Object Storage charges apply for the storage of your backup in Oracle Cloud Infrastructure.

### **Compatibility**

The scripted migration process is compatible with the following bare metal and virtual machine DB system configurations:

Configuration	Version or Type	Notes
Database Version	18.x 12.2.0.1 12.1.0.2 11.2.0.4	<ul style="list-style-type: none"> <li>• For versions 18c, 12.2.0.1, and 12.1.0.2:               <ul style="list-style-type: none"> <li>◦ Only Container Databases (CDBs) are supported. The scripted migration process may work with non-CDB databases for these database versions, but Oracle does not provide support for the migration of non-CDB databases using the script described in this topic.</li> </ul> <p>For information on creating an on-premises pluggable database (PDB) by cloning a non-CDB in Oracle Database 18c, see <a href="#">About Cloning a Non-CDB</a>. For an overview of multitenant architecture in Oracle Database 18c, see <a href="#">Introduction to the Multitenant Architecture</a>.</p> <p>For information on creating an on-premises pluggable database (PDB) from a non-CDB database in Oracle Database 12c Release 2 (12.2), see <a href="#">Upgrading a Non-CDB Oracle Database To a PDB on a CDB</a>. For an overview of multitenant architecture in 12c Release 2, see <a href="#">Overview of Managing a Multitenant Environment</a>.</p> <li>◦ The Oracle Cloud Infrastructure Database service will attempt to run datapatch, which requires read/write mode. If there are pluggable databases (PDBs), they should also be in read/write mode to ensure that datapatch runs on them.</li> </li></ul> <li>• For version 11.2.0.4, depending on the source</li>

Configuration	Version or Type	Notes
		<p>database patch level, you may need to roll back patches prior to migrating. See <a href="#">Rolling Back Patches on a Version 11.2 Database</a> for more information.</p> <ul style="list-style-type: none"> <li>If your on-premises database has an interim patch (previous known as a one-off patch), see <a href="#">Applying Interim Patches</a> for details on applying the patch in Oracle Cloud Infrastructure.</li> </ul>
Source Database Platform	<p>Oracle Enterprise Linux / Red Hat Enterprise Linux 5.x</p> <p>Oracle Linux / Red Hat Enterprise Linux 6.x</p> <p>Oracle Linux / Red Hat Enterprise 7.x</p>	<ul style="list-style-type: none"> <li>The scripted migration described in this topic may work in Microsoft Windows environments, but Oracle currently does not provide support for this script in Windows.</li> <li>For Oracle Linux 6.x users, see <a href="#">Configuring Oracle Linux 6 to install Python</a> for details on configuring the operating system to install a compatible version of Python. See <a href="#">Installing the CLI</a> for more information regarding Oracle Linux 6.</li> </ul>
Encryption	TDE Non-TDE	<ul style="list-style-type: none"> <li>In a non-TDE configuration, the RMAN encryption password is required.</li> <li>Unencrypted on-premises databases remain unencrypted when restored to Oracle Cloud Infrastructure. The stored RMAN standalone backups are always encrypted.</li> </ul>

## CHAPTER 11 Database

---

Configuration	Version or Type	Notes
Target Database Edition	Standard Edition Enterprise Edition Enterprise Edition - High Performance Enterprise Edition - Extreme Performance	
Cluster	Single RAC	

### Prerequisites

On the source database host:

- Outbound internet connectivity for installing Python packages, running yum install, and access to the Oracle Cloud Infrastructure API and Object Storage.
- RMAN configuration to autobackup `controlfile` and `spfile`:

```
RMAN> CONFIGURE CONTROLFILE AUTOBACKUP ON;
```



### Note

RMAN configuration changes must be completed prior to running the script. The script may modify RMAN parameters as required to complete the backup and migration tasks.

### To Migrate an On-Premises Database Using a Standalone Backup

Perform the following tasks on the source database host:

1. Create a directory named `/home/oracle/migrate`.



### Tip

You can name the `migrate` portion of the directory path anything you want. If you use a different name, you must adjust all of the paths that appear in this task accordingly. The following examples assume the name `migrate` for simplicity and clarity.

2. As *root*, run the CLI installer in the directory you created in step 1. (For example, `/home/oracle/migrate`.) See [Installing the CLI: Windows](#) for instructions on running the installer script in Windows. See [Installing the CLI: MacOS, Linux, and Unix](#) for instructions on running the installer script in the Bash environment.  
The installer installs Python 3.6.0 if either Python 2.7 or Python 3.6 does not exist on the machine. The installer also installs the Python script required to create and migrate a standalone backup from an on-premises database.

On Oracle Linux 6, a newer version of Python (such as Python 3.6.0) is usually required. **Use the following instructions to configure Oracle Linux 6 before running the backup script.**

### Configuring Oracle Linux 6 to install Python

In Oracle Linux 6 use the following `/etc/yum.repos.d/ol6.repo` file to ensure that a compatible version of Python is installed by the script if a compatible version is not already installed. Include this file before attempting to run the script with the `./install.sh` command.

```
[ol6_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1
```

3. Copy the following files into the new directory:
  - [Oracle Database Backup Module](#) (`opc_install.jar`)
  - Your API `*.pem` key file.
4. Respond to the prompts as follows:

```
(yum install)
Is this ok [y/N]: y

===> Missing native dependencies. Continue and install the following dependencies: gcc, libffi-devel, python36u-devel, openssl-devel? (Y/n): Y

===> In what directory would you like to place the install? (leave blank to use '/root/lib/oracle-cli'): /home/oracle/migrate/lib/oracle-cli

===> In what directory would you like to place the 'oci' executable? (leave blank to use '/root/bin'): /home/oracle/migrate/bin

===> In what directory would you like to place the OCI scripts? (leave blank to use '/root/bin/oci-cli-scripts'): /home/oracle/migrate/bin/oci-cli-scripts
```

```
===> Currently supported optional packages are: ['db (will install cx_Oracle)'] What optional CLI
packages would you like to be installed (comma separated names; press enter if you don't need any
optional packages)?: db

===> Modify profile to update your $PATH and enable shell/tab completion now? (Y/n): Y

===> Enter a path to an rc file to update (leave blank to use '/root/.bashrc'):
/home/oracle/.bashrc
```

### 5. Perform the following file operations:

```
chown -R oracle:oinstall /home/oracle/migrate
```

### 6. Edit the /home/oracle/migrate/config.txt file

```
[DEFAULT]
tenancy=<your_tenancy_OCID>
user=<your_user_OCID>
fingerprint=<fingerprint>
key_file=/home/oracle/migrate/<your_api_key>.pem
region=<region>
```

If you do not know your API signing key's fingerprint, see [How to Get the Key's Fingerprint](#).

### 7. As *oracle user* (not *root*), run one of the following sets of commands, depending on the type of database you are migrating.

*For a non-TDE database:*

```
export AD=<destination_availability_domain>
export C=<destination_compartment_OCID>
export ORACLE_SID=<ORACLE_SID>
export ORACLE_HOME=<ORACLE_HOME>
export PATH=$PATH:$ORACLE_HOME/bin
export LC_ALL=en_US.UTF-8
export ORACLE_UNQNAME=<source_DB_unique_name>
rm -rf /home/oracle/migrate/onprem_upload
cd /home/oracle/migrate/bin/oci-cli-scripts/
./create_backup_from_onprem --config-file /home/oracle/migrate/config.txt --display-name
<example_display_name> --availability-domain $AD --edition ENTERPRISE_EDITION_EXTREME_PERFORMANCE
```

## CHAPTER 11 Database

```
--opc-installer-dir /home/oracle/migrate --tmp-dir /home/oracle/migrate/onprem_upload --
compartment-id $C --rman-password <password>
```

*For a TDE-enabled database:*

```
export AD=<destination_availability_domain>
export C=<destination_compartment_OCID>
export ORACLE_SID=<ORACLE_SID>
export ORACLE_HOME=<ORACLE_HOME>
export PATH=$PATH:$ORACLE_HOME/bin
rm -rf /home/oracle/migrate/onprem_upload
cd /home/oracle/migrate/bin/oci-cli-scripts/
./create_backup_from_onprem --config-file /home/oracle/migrate/config.txt --display-name
<example_display_name> --availability-domain $AD --edition ENTERPRISE_EDITION_EXTREME_PERFORMANCE
--opc-installer-dir /home/oracle/migrate --tmp-dir /home/oracle/migrate/onprem_upload --
compartment-id $C
```

See the following list of parameters used by the script for more details.

### Parameters used by the script

Parameter	Description	Required
--config-file	The path to the oci-cli config file. The default path is as follows: ~/.oci/config	No
--profile	The profile in the config file to load. This profile will also be used to locate any default parameter values which have been specified in the OCI CLI-specific configuration file. The default value is DEFAULT.	No
-- compartment- id	The compartment OCID of the Oracle Cloud Infrastructure compartment that will contain your standalone backup.	Yes

Parameter	Description	Required
<code>--display-name</code>	The name of the backup, as you wish it to be displayed in the OCI Console under Standalone Backups.	Yes
<code>--availability-domain</code>	The availability domain where the backup is to be stored.	Yes
<code>--edition</code>	The edition of the Oracle Cloud Infrastructure DB system that will contain the database created from the standalone backup. You can choose the same edition as the on-premises database, or any addition above the on-premises database. The choices, listed from lowest to highest, are the following: <ul style="list-style-type: none"> <li>• STANDARD_EDITION</li> <li>• ENTERPRISE_EDITION</li> <li>• ENTERPRISE_EDITION_HIGH_PERFORMANCE</li> <li>• ENTERPRISE_EDITION_EXTREME_PERFORMANCE</li> </ul>	Yes
<code>--opc-installer-dir</code>	The directory containing the <code>opc_installer.jar</code> file. This is the directory you created in step 1 of this procedure.	Yes
<code>--additional-opc-args</code>	Optional additional arguments for the <code>opc</code> installer.	No
<code>--tmp-dir</code>	Optional temporary directory for intermediate files.	No

Parameter	Description	Required
<code>--rman-password</code>	The RMAN password to use for the standalone backup.	Required if TDE is not enabled
<code>--rman-channels</code>	RMAN channels. The default value is 5.	No
<code>--help</code>	Displays in-line help for the script in the OCI-CLI environment.	No

The script will produce a standalone backup of your on-premises database in your Oracle Cloud Infrastructure tenancy. You can check the Console for your backup by viewing the **Standalone Backups** page in the Database service, under **Bare Metal, VM, and Exadata**.



#### Tip

To access command line help for the backup script, run the following command in the `/home/oracle/migrate/bin/oci-cli-scripts/` directory:

```
create_backup_from_onprem --help
```

8. Create a new database or launch a new DB system using the backup you created in the preceding step. See the following instructions to perform these tasks:
  - [Recovering a Database from Object Storage](#)
  - [Recovering a Database from Object Storage](#)

### Migrating from Oracle Database 11g to Oracle Database 11g in the Cloud

You can migrate Oracle Database 11g databases from on-premises to Oracle Database 11g databases in the Database service using several different methods.

The applicability of some of the migration methods depends on the on-premises database's character set and platform endian format.

If you have not already done so, determine the database character set of your on-premises database, and determine the endian format of the platform your on-premises database resides on. Use this information to help you choose an appropriate method.

- **Data Pump Conventional Export/Import**  
This method can be used regardless of the endian format and database character set of the on-premises database.  
For the steps this method entails, see [Data Pump Conventional Export/Import](#).
- **Data Pump Transportable Tablespace**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database database are compatible.  
For the steps this method entails, see [Data Pump Transportable Tablespace](#).
- **RMAN Transportable Tablespace with Data Pump**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database database are compatible.  
For the steps this method entails, see [RMAN Transportable Tablespace with Data Pump](#).
- **RMAN `CONVERT` Transportable Tablespace with Data Pump**  
This method can be used only if the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database database are compatible.  
This method is similar to the Data Pump Transportable Tablespace method, with the addition of the RMAN `CONVERT` command to enable transport between platforms with different endianness. Query `V$TRANSPORTABLE_PLATFORM` to determine if the on-premises database platform supports cross-platform tablespace transport and to

determine the endian format of the platform. The Database service platform is little-endian format.

For the steps this method entails, see [RMAN CONVERT Transportable Tablespace with Data Pump](#).

### Migrating from Oracle Database 11g to Oracle Database 12c in the Cloud

You can migrate Oracle Database 11g databases from on-premises to Oracle Database 12c databases in the Database service using several different methods.

The applicability of some of the migration methods depends on the on-premises database's version, database character set and platform endian format.

If you have not already done so, determine the database version and database character set of your on-premises database, and determine the endian format of the platform your on-premises database resides on. Use this information to help you choose an appropriate method.

- **Data Pump Conventional Export/Import**  
This method can be used regardless of the endian format and database character set of the on-premises database.  
For the steps this method entails, see [Data Pump Conventional Export/Import](#).
- **Data Pump Transportable Tablespace**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database service database are compatible.  
For the steps this method entails, see [Data Pump Transportable Tablespace](#).
- **RMAN Transportable Tablespace with Data Pump**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database service database are compatible.  
For the steps this method entails, see [RMAN Transportable Tablespace with Data Pump](#).
- **RMAN `CONVERT` Transportable Tablespace with Data Pump**

This method can be used only if the database character sets of your on-premises database and the Database service database are compatible.

This method is similar to the Data Pump Transportable Tablespace method, with the addition of the `RMAN CONVERT` command to enable transport between platforms with different endianness. Query `V$TRANSPORTABLE_PLATFORM` to determine if the on-premises database platform supports cross-platform tablespace transport and to determine the endian format of the platform. The Database service platform is little-endian format.

For the steps this method entails, see [RMAN CONVERT Transportable Tablespace with Data Pump](#).

- Data Pump Full Transportable

This method can be used only if the source database release version is 11.2.0.3 or later, and the database character sets of your on-premises database and the Database service database are compatible.

For the steps this method entails, see [Data Pump Full Transportable](#).

## Migrating from Oracle Database 12c CDB to Oracle Database 12c in the Cloud

You can migrate Oracle Database 12c CDB databases from on-premises to Oracle Database 12c databases in the Oracle Cloud Infrastructure Database service using several different methods.

The applicability of some of the migration methods depends on the on-premises database's character set and platform endian format.

If you have not already done so, determine the database character set of your on-premises database, and determine the endian format of the platform your on-premises database resides on. Use this information to help you choose an appropriate method.

- Data Pump Conventional Export/Import

This method can be used regardless of the endian format and database character set of the on-premises database.

For the steps this method entails, see [Data Pump Conventional Export/Import](#).

- **Data Pump Transportable Tablespace**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database database are compatible.  
For the steps this method entails, see [Data Pump Transportable Tablespace](#).
- **RMAN Transportable Tablespace with Data Pump**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database service database are compatible.  
For the steps this method entails, see [RMAN Transportable Tablespace with Data Pump](#).
- **RMAN `CONVERT` Transportable Tablespace with Data Pump**  
This method can be used only if the database character sets of your on-premises database and the Database database are compatible.  
This method is similar to the Data Pump Transportable Tablespace method, with the addition of the RMAN `CONVERT` command to enable transport between platforms with different endianness. Query `V$TRANSPORTABLE_PLATFORM` to determine if the on-premises database platform supports cross-platform tablespace transport and to determine the endian format of the platform. The Database service platform is little-endian format.  
For the steps this method entails, see [RMAN `CONVERT` Transportable Tablespace with Data Pump](#).
- **RMAN Cross-Platform Transportable Tablespace Backup Sets**  
This method can be used only if the database character sets of your on-premises database and the Database service database are compatible.  
For the steps this method entails, see [RMAN Cross-Platform Transportable Tablespace Backup Sets](#).
- **Data Pump Full Transportable**  
This method can be used only if the database character sets of your on-premises database and the Database service database are compatible.  
For the steps this method entails, see [Data Pump Full Transportable](#).

- Unplugging/Plugging (CDB)

This method can be used only if the on-premises platform is little endian, and the on-premises database and Database database have compatible database character sets and national character sets.

For the steps this method entails, see [Unplugging/Plugging a PDB](#).

- Remote Cloning (CDB)

This method can be used only if the on-premises platform is little endian, the on-premises database release is 12.1.0.2 or higher, and the on-premises database and Database service database have compatible database character sets and national character sets.

For the steps this method entails, see [Remote Cloning a PDB](#).

- RMAN Cross-Platform Transportable PDB

This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database service database are compatible.

For the steps this method entails, see [RMAN Cross-Platform Transportable PDB](#).

- SQL Developer and SQL\*Loader to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into your Oracle Database 12c database on the cloud. In this method, you use SQL\*Loader to load the data into your cloud database.

For the steps this method entails, see [SQL Developer and SQL\\*Loader to Migrate Selected Objects](#).

- SQL Developer and `INSERT` Statements to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into your Oracle Database 12c database on the cloud. In this method, you use SQL `INSERT` statements to load the data into your cloud database.

For the steps this method entails, see [SQL Developer and `INSERT` Statements to Migrate Selected Objects](#).

### Migrating from Oracle Database 12c Non-CDB to Oracle Database 12c in the Cloud

You can migrate Oracle Database 12c non-CDB databases from on-premises to Oracle Database 12c databases in Oracle Cloud Infrastructure Database service using several different methods.

The applicability of some of the migration methods depends on the on-premises database's character set and platform endian format.

If you have not already done so, determine the database character set of your on-premises database, and determine the endian format of the platform your on-premises database resides on. Use this information to help you choose an appropriate method.

- **Data Pump Conventional Export/Import**  
This method can be used regardless of the endian format and database character set of the on-premises database.  
For the steps this method entails, see [Data Pump Conventional Export/Import](#).
- **Data Pump Transportable Tablespace**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database database are compatible.  
For the steps this method entails, see [Data Pump Transportable Tablespace](#).
- **RMAN Transportable Tablespace with Data Pump**  
This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database service database are compatible.  
For the steps this method entails, see [RMAN Transportable Tablespace with Data Pump](#).
- **RMAN `CONVERT` Transportable Tablespace with Data Pump**  
This method can be used only if the database character sets of your on-premises database and the Database service database are compatible.  
This method is similar to the Data Pump Transportable Tablespace method, with the addition of the RMAN `CONVERT` command to enable transport between platforms with

different endianness. Query `V$TRANSPORTABLE_PLATFORM` to determine if the on-premises database platform supports cross-platform tablespace transport and to determine the endian format of the platform. The Database service platform is little-endian format.

For the steps this method entails, see [RMAN CONVERT Transportable Tablespace with Data Pump](#).

- RMAN Cross-Platform Transportable Tablespace Backup Sets

This method can be used only if the database character sets of your on-premises database and the Database database are compatible.

For the steps this method entails, see [RMAN Cross-Platform Transportable Tablespace Backup Sets](#).

- Data Pump Full Transportable

This method can be used only if the database character sets of your on-premises database and the Database service database are compatible.

For the steps this method entails, see [Data Pump Full Transportable](#).

- Unplugging/Plugging (non-CDB)

This method can be used only if the on-premises platform is little endian, and the on-premises database and Database service database have compatible database character sets and national character sets.

You can use the unplug/plug method to migrate an Oracle Database 12c non-CDB database to Oracle Database 12c in the cloud. This method provides a way to consolidate several non-CDB databases into a single Oracle Database 12c CDB on the cloud.

For the steps this method entails, see [Unplugging/Plugging Non-CDB](#).

- Remote Cloning (non-CDB)

This method can be used only if the on-premises platform is little endian, the on-premises database release is 12.1.0.2 or higher, and the on-premises database and Database service database have compatible database character sets and national character sets.

You can use the remote cloning method to copy an Oracle Database 12c non-CDB on-premises database to your Oracle Database 12c database in the cloud.

For the steps this method entails, see [Remote Cloning Non-CDB](#).

- SQL Developer and SQL\*Loader to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into your Oracle Database 12c database on the cloud. In this method, you use SQL\*Loader to load the data into your cloud database.

For the steps this method entails, see [SQL Developer and SQL\\*Loader to Migrate Selected Objects](#).

- SQL Developer and INSERT Statements to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into your Oracle Database 12c database on the cloud. In this method, you use SQL INSERT statements to load the data into your cloud database.

For the steps this method entails, see [SQL Developer and INSERT Statements to Migrate Selected Objects](#).

### Data Pump Conventional Export/Import

You can use this method regardless of the endian format and database character set of the on-premises database.

To migrate an on-premises source database, tablespace, schema, or table to the database on a Database service database deployment using Data Pump Export and Import, you perform these tasks:

1. On the on-premises database host, invoke Data Pump Export and export the on-premises database.
2. Use a secure copy utility to transfer the dump file to the Database service compute node.
3. On the Database service compute node, invoke Data Pump Import and import the data into the database.
4. After verifying that the data has been imported successfully, you can delete the dump file.

For information about Data Pump Import and Export, see these topics:

- "Data Pump Export Modes" in *Oracle Database Utilities* for Release [12.2](#), [12.1](#) or [11.2](#).
- "Data Pump Import Modes" in *Oracle Database Utilities* for Release [12.2](#), [12.1](#) or [11.2](#).

### Data Pump Conventional Export/Import: Example

This example provides a step-by-step demonstration of the tasks required to migrate a schema from an on-premises Oracle database to a Database service database.

This example illustrates a schema mode export and import. The same general procedure applies for a full database, tablespace, or table export and import.

In this example, the on-premises database is on a Linux host.

1. On the on-premises database host, invoke Data Pump Export to export the schemas.
  - a. On the on-premises database host, create an operating system directory to use for the on-premises database export files.

```
$ mkdir /u01/app/oracle/admin/orcl/dpdump/for_cloud
```

- b. On the on-premises database host, invoke SQL\*Plus and log in to the on-premises database as the `SYSTEM` user.

```
$ sqlplus system
Enter password: <enter the password for the SYSTEM user>
```

- c. Create a directory object in the on-premises database to reference the operating system directory.

```
SQL> CREATE DIRECTORY dp_for_cloud AS '/u01/app/oracle/admin/orcl/dpdump/for_cloud';
```

- d. Exit from SQL\*Plus.
- e. On the on-premises database host, invoke Data Pump Export as the `SYSTEM` user or another user with the `DATAPUMP_EXP_FULL_DATABASE` role and export the on-premises schemas. Provide the password for the user when prompted.

```
$ expdp system SCHEMAS=fowner DIRECTORY=dp_for_cloud
```

2. Use a secure copy utility to transfer the dump file to the Database service compute

node.

In this example the dump file is copied to the `/u01` directory. Choose the appropriate location based on the size of the file that will be transferred.

- a. On the Database service compute node, create a directory for the dump file.

```
$ mkdir /u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

- b. Before using the `scp` command to copy the export dump file, make sure the SSH private key that provides access to the Database service compute node is available on your on-premises host.
- c. On the on-premises database host, use the SCP utility to transfer the dump file to the Databaseservice compute node.

```
$ scp -i private_key_file \
/u01/app/oracle/admin/orcl/dpdump/for_cloud/expdat.dmp \
oracle@IP_address_DBaaS_VM:/u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

3. On the Database service compute node, invoke Data Pump Import and import the data into the database.

- a. On the Database service compute node, invoke SQL\*Plus and log in to the database as the `SYSTEM` user.

```
$ sqlplus system
Enter password: <enter the password for the SYSTEM user>
```

- b. Create a directory object in the Database service database.

```
SQL> CREATE DIRECTORY dp_from_onprem AS '/u01/app/oracle/admin/ORCL/dpdump/from_onprem';
```

- c. If they do not exist, create the tablespace(s) for the objects that will be imported.
- d. Exit from SQL\*Plus.
- e. On the Database service compute node, invoke Data Pump Import and connect to the database. Import the data into the database.

```
impdp system SCHEMAS=fsowner DIRECTORY=dp_from_onprem
```

4. After verifying that the data has been imported successfully, you can delete the `expdat.dmp` file.

### Data Pump Full Transportable

You can use this method only if the source database release version is 11.2.0.3 or later, and the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database service database are compatible.

You can use the Data Pump full transportable method to copy an entire database from your on-premises host to the database on a Database service database deployment.

To migrate an Oracle Database 11g on-premises database to the Oracle Database 12c database on a Database service database deployment using the Data Pump full transportable method, you perform these tasks:

1. On the on-premises database host, prepare the database for the Data Pump full transportable export by placing the user-defined tablespaces in `READ ONLY` mode.
2. On the on-premises database host, invoke Data Pump Export to perform the full transportable export.
3. Use a secure copy utility to transfer the Data Pump Export dump file and the datafiles for all of the user-defined tablespaces to the Database service compute node.
4. Set the on-premises tablespaces back to `READ WRITE`.
5. On the Database service compute node, prepare the database for the tablespace import.
6. On the Database service compute node, invoke Data Pump Import and connect to the database.
7. After verifying that the data has been imported successfully, you can delete the dump file.

#### **Data Pump Full Transportable: Example**

This example provides a step-by-step demonstration of the tasks required to migrate an Oracle Database 11g database to a Database service 12c database.

In this example, the source database is on a Linux host.

1. On the source database host, prepare the database for the Data Pump full transportable export.

- a. On the source database host, create a directory in the operating system to use for the source export.

```
$ mkdir /u01/app/oracle/admin/orcl/dpdump/for_cloud
```

- b. On the source database host, invoke SQL\*Plus and log in to the source database as the SYSTEM user.

```
$ sqlplus system
Enter password: <enter the password for the SYSTEM user>
```

- c. Create a directory object in the source database to reference the operating system directory.

```
SQL> CREATE DIRECTORY dp_for_cloud AS '/u01/app/oracle/admin/orcl/dpdump/for_cloud';
```

- d. Determine the name(s) of the tablespaces and data files that belong to the user-defined tablespaces by querying DBA\_DATA\_FILES. These files will also be listed in the export output.

```
SQL> SELECT tablespace_name, file_name FROM dba_data_files;
TABLESPACE_NAME FILE_NAME

USERS /u01/app/oracle/oradata/orcl/users01.dbf
UNDOTBS1 /u01/app/oracle/oradata/orcl/undotbs01.dbf
SYSaux /u01/app/oracle/oradata/orcl/sysaux01.dbf
SYSTEM /u01/app/oracle/oradata/orcl/system01.dbf
EXAMPLE /u01/app/oracle/oradata/orcl/example01.dbf
FSDATA /u01/app/oracle/oradata/orcl/fsdata01.dbf
FSINDEX /u01/app/oracle/oradata/orcl/fsindex01.dbf
SQL>
```

- e. On the source database host, set all tablespaces that will be transported (the transportable set) to READ ONLY mode.

```
SQL> ALTER TABLESPACE example READ ONLY;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ ONLY;
Tablespace altered.
```

## CHAPTER 11 Database

```
SQL> ALTER TABLESPACE fsdata READ ONLY;
Tablespace altered.
SQL> ALTER TABLESPACE users READ ONLY;
Tablespace altered.
SQL>
```

- f. Exit from SQL\*Plus.
2. On the source database host, invoke Data Pump Export to perform the full transportable export. Specify `FULL=y` and `TRANSPORTABLE=always`. Because this is an Oracle Database 11g database and full transportable is an Oracle Database 12c feature, specify `VERSION=12`. Provide the password for the `SYSTEM` user when prompted.

```
$ expdp system FULL=y TRANSPORTABLE=always VERSION=12 DUMPFILE=expdat.dmp DIRECTORY=dp_for_cloud
```

3. Use a secure copy utility to transfer the Data Pump Export dump file and the datafiles for all of the user-defined tablespaces to the Database service compute node. In this example the dump file is copied to the `/u01` directory. Choose the appropriate location based on the size of the file that will be transferred.

- a. On the Database service compute node, create a directory for the dump file.

```
$ mkdir /u01/app/oracle/admin/ORCL/dpdump/from_source
```

- b. Before using the `scp` utility to copy files, make sure the SSH private key that provides access to the Database service compute node is available on your source host.
- c. On the source database host, use the `scp` utility to transfer the dump file and all datafiles of the transportable set to the Database service compute node.

```
$ scp -i private_key_file \
/u01/app/oracle/admin/orcl/dpdump/for_cloud/expdat.dmp \
oracle@compute_node_IP_address:/u01/app/oracle/admin/ORCL/dpdump/from_source

$ scp -i private_key_file \
/u01/app/oracle/oradata/orcl/example01.dbf \
oracle@compute_node_IP_address:/u02/app/oracle/oradata/ORCL/PDB2

$ scp -i private_key_file \
/u01/app/oracle/oradata/orcl/fsdata01.dbf \
oracle@compute_node_IP_address:/u02/app/oracle/oradata/ORCL/PDB2
```

```
$ scp -i private_key_file \
/u01/app/oracle/oradata/orcl/fsindex01.dbf \
oracle@compute_node_IP_address:/u02/app/oracle/oradata/ORCL/PDB2

$ scp -i private_key_file \
/u01/app/oracle/oradata/orcl/users01.dbf \
oracle@compute_node_IP_address:/u02/app/oracle/oradata/ORCL/PDB2
```

4. Set the source tablespaces back to READ WRITE.
  - a. Invoke SQL\*Plus and log in as the SYSTEM user.
  - b. Set the user-defined tablespaces back to READ WRITE mode.

```
SQL> ALTER TABLESPACE example READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsdata READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE users READ WRITE;
Tablespace altered.
```

- c. Exit from SQL\*Plus.
5. On the Database service compute node, prepare the PDB for the tablespace import.
  - a. On the Database service compute node, invoke SQL\*Plus and log in to the PDB as the SYSTEM user.
  - b. Create a directory object in the PDB.

```
SQL> CREATE DIRECTORY dp_from_source AS '/u01/app/oracle/admin/ORCL/dpdump/from_source';
```

6. On the Database service compute node, invoke Data Pump Import and connect to the PDB.

Import the data into the database using the `TRANSPORT_DATAFILES` option.

```
$ impdp system@PDB2 FULL=y DIRECTORY=dp_from_source \
TRANSPORT_
DATAFILES='/u02/app/oracle/oradata/ORCL/PDB2/example01.dbf',\
'/u02/app/oracle/oradata/ORCL/PDB2/fsdata01.dbf',\

```

```
'/u02/app/oracle/oradata/ORCL/PDB2/fsindex01.dbf, '\
'/u02/app/oracle/oradata/ORCL/PDB2/users01.dbf'
```

7. After verifying that the data has been imported successfully, you can delete the `expdat.dmp` dump file.

### Data Pump Transportable Tablespace

You can use this method only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Oracle Cloud Infrastructure Database service database are compatible.

The Transportable Tablespace method is generally much faster than a conventional export/import of the same data because the data files containing all of the actual data are simply copied to the destination location. You use Data Pump to transfer only the metadata of the tablespace objects to the new database.

To migrate an on-premises source database to the database deployment on the Database service using the Data Pump Transportable Tablespace method, you perform these tasks:

1. On the on-premises database host, prepare the database for the Data Pump transportable tablespace export.
2. On the on-premises database host, invoke Data Pump Export to perform the transportable tablespace export.
3. Use a secure copy utility to transfer the Data Pump Export dump file and the tablespace datafiles to the Database service compute node.
4. Set the on-premises tablespaces back to `READ WRITE`.
5. On the Databaseservice compute node, prepare the database for the tablespace import.
6. On the Database service compute node, invoke Data Pump Import and connect to the database.
7. Set the tablespaces on the Database service database to `READ WRITE` mode.
8. After verifying that the data has been imported successfully, you can delete the dump file.

### Data Pump Transportable Tablespace: Example

This example provides a step-by-step demonstration of the tasks required to migrate tablespaces in an on-premises Oracle database to a Database service database.

This example performs a migration of the `FSDATA` and `FSINDEX` tablespaces.

In this example, the on-premises database is on a Linux host.

1. On the on-premises database host, prepare the database for the Data Pump transportable tablespace export.
  - a. On the on-premises database host, create a directory in the operating system to use for the on-premises export.

```
mkdir /u01/app/oracle/admin/orcl/dpdump/for_cloud
```

- b. On the on-premises database host, invoke SQL\*Plus and log in to the on-premises database as the `SYSTEM` user.

```
sqlplus system
Enter password: <enter the password for the SYSTEM user>
```

- c. Create a directory object in the on-premises database to reference the operating system directory.

```
SQL> CREATE DIRECTORY dp_for_cloud AS '/u01/app/oracle/admin/orcl/dpdump/for_cloud';
```

- d. Determine the name(s) of the datafiles that belong to the `FSDATA` and `FSINDEX` tablespaces by querying `DBA_DATA_FILES`. These files will also be listed in the export output.

```
SQL> SELECT file_name FROM dba_data_files
2 WHERE tablespace_name = 'FSDATA';
```

```
FILE_NAME
```

```

/u01/app/oracle/oradata/orcl/fsdata01.dbf
```

```
SQL> SELECT file_name FROM dba_data_files
2 WHERE tablespace_name = 'FSINDEX';
```

```
FILE_NAME
```

```

/u01/app/oracle/oradata/orcl/fsindex01.dbf
```

- e. On the on-premises database host, set all tablespaces that will be transported (the transportable set) to READ ONLY mode.

```
SQL> ALTER TABLESPACE fsindex READ ONLY;
Tablespace altered.
SQL> ALTER TABLESPACE fsdata READ ONLY;
Tablespace altered.
```

- f. Exit from SQL\*Plus.
2. On the on-premises database host, invoke Data Pump Export to perform the transportable tablespace export.

On the on-premises database host, invoke Data Pump Export and connect to the on-premises database. Export the on-premises tablespaces using the `TRANSPORT_TABLESPACES` option. Provide the password for the `SYSTEM` user when prompted.

```
expdp system TRANSPORT_TABLESPACES=fsdata,fsindex TRANSPORT_FULL_CHECK=YES DIRECTORY=dp_for_cloud
```

3. Use a secure copy utility to transfer the Data Pump Export dump file and the tablespace datafiles to the Database service compute node.

In this example the dump file is copied to the `/u01` directory. Choose the appropriate location based on the size of the file that will be transferred.

- a. On the Database service compute node, create a directory for the dump file.

```
mkdir /u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

- b. Before using the `scp` utility to copy files, make sure the SSH private key that provides access to the Database service compute node is available on your on-premises host.
- c. On the on-premises database host, use the `scp` utility to transfer the dump file and all datafiles of the transportable set to the Database service compute node.

```
scp -i private_key_file \
/u01/app/oracle/admin/orcl/dpdump/for_cloud/expdat.dmp \
oracle@IP_address_DBaaS_VM:/u01/app/oracle/admin/ORCL/dpdump/from_onprem

$ scp -i private_key_file \
\u01/app/oracle/oradata/orcl/fsdata01.dbf \
oracle@IP_address_DBaaS_VM:/u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

```
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL
$ scp -i private_key_file \u01/app/oracle/oradata/orcl/fsindex01.dbf \
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL
```

4. Set the on-premises tablespaces back to READ WRITE.
  - a. Invoke SQL\*Plus and log in as the SYSTEM user.
  - b. Set the FSDATA and FSINDEX tablespaces back to READ WRITE mode.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ WRITE;
Tablespace altered.
```

- c. Exit from SQL\*Plus.
5. On the Database service compute node, prepare the database for the tablespace import.
  - a. On the Database service compute node, invoke SQL\*Plus and log in to the database as the SYSTEM user.
  - b. Create a directory object in the Database service database.

```
SQL> CREATE DIRECTORY dp_from_onprem AS '/u01/app/oracle/admin/ORCL/dpdump/from_onprem';
```

- c. If the owners of the objects that will be imported do not exist in the database, create them before performing the import. The transportable tablespace mode of import does not create the users.

```
SQL> CREATE USER fsowner
2 PROFILE default
3 IDENTIFIED BY fspass
4 TEMPORARY TABLESPACE temp
5 ACCOUNT UNLOCK;
```

6. On the Database service compute node, invoke Data Pump Import and connect to the database.  
Import the data into the database using the `TRANSPORT_DATAFILES` option.

## CHAPTER 11 Database

---

```
impdp system DIRECTORY=dp_from_onprem \
TRANSPORT_DATAFILES='/u02/app/oracle/oradata/ORCL/fsdata01.dbf', \
'/u02/app/oracle/oradata/ORCL/fsindex01.dbf'
```

7. Set the tablespaces on the Database service database to `READ WRITE` mode.
  - a. Invoke SQL\*Plus and log in as the `SYSTEM` user.
  - b. Set the `FSDATA` and `FSINDEX` tablespaces to `READ WRITE` mode.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ WRITE;
Tablespace altered.
```

- c. Exit from SQL\*Plus.
8. After verifying that the data has been imported successfully, you can delete the `expdat.dmp` dump file.

## Remote Cloning a PDB

You can use this method only if the on-premises platform is little endian, the on-premises database release is 12.1.0.2 or higher, and the on-premises database and Database service database have compatible database character sets and national character sets.

You can use the remote cloning method to copy a PDB from your on-premises Oracle Database 12c database to a PDB in an Oracle Database 12c database on the Database service.

### Migration Tasks

To migrate an Oracle Database 12c PDB to a PDB in a Database service database deployment using the remote cloning method, you perform these tasks:

1. On the on-premises database host, invoke SQL\*Plus and close the on-premises PDB and then reopen it in `READ ONLY` mode.
2. On the Database service compute node, invoke SQL\*Plus and create a database link that enables a connection to the on-premises database.

3. On the Database service compute node, execute the `CREATE PLUGGABLE DATABASE` command to clone the on-premises PDB.
4. On the Database compute node, open the new PDB by executing the `ALTER PLUGGABLE DATABASE OPEN` command.
5. Optionally, on the on-premises database host invoke SQL\*Plus and set the on-premises PDB back to `READ WRITE` mode.

For more information, see "Cloning a Remote PDB or Non-CDB" in *Oracle Database Administrator's Guide* for Release [12.2](#) or [12.1](#).

### Remote Cloning Non-CDB

You can use this method only if the on-premises platform is little endian, the on-premises database release is 12.1.0.2 or higher, and the on-premises database and Database service database have compatible database character sets and national character sets.

You can use the remote cloning method to copy an Oracle Database 12c non-CDB on-premises database to a PDB in an Oracle Database 12c database on the Databaseservice.

### Migration Tasks

To migrate an Oracle Database 12c non-CDB database to a Database service database deployment using the remote cloning method, you perform these tasks:

1. On the on-premises database host, invoke SQL\*Plus and set the on-premises database to `READ ONLY` mode.
2. On the Database service compute node, invoke SQL\*Plus and create a database link that enables a connection to the on-premises database.
3. On the Database service compute node, execute the `CREATE PLUGGABLE DATABASE` command to clone the on-premises non-CDB database.
4. On the Database service compute node, execute the `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` script.
5. On the Database service compute node, open the new PDB by executing the `ALTER`

`PLUGGABLE DATABASE OPEN` command.

6. Optionally, on the on-premises database host invoke SQL\*Plus and set the on-premises database back to `READ WRITE` mode.

For more information, see "Cloning a Remote PDB or Non-CDB" in *Oracle Database Administrator's Guide* for Release [12.2](#) or [12.1](#).

### RMAN Cross-Platform Transportable PDB

This method can be used only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Database service database are compatible.

To migrate an Oracle Database 12c PDB to a PDB in an Oracle Database 12c database on a Database service deployment using the RMAN cross-platform transportable PDB method, you perform these tasks:

1. On the on-premises database host, invoke SQL\*Plus and close the on-premises PDB.
2. On the on-premises database host, execute the `ALTER PLUGGABLE DATABASE UNPLUG` command to generate an XML file containing the list of datafiles that will be plugged in on the cloud database.
3. On the on-premises database host, invoke RMAN and connect to the root. Execute the `BACKUP FOR TRANSPORT PLUGGABLE DATABASE` command.
4. Use a secure copy utility to transfer the XML file and the backup set to the Database service compute node.
5. On the Database service compute node, invoke RMAN and connect to the root. Execute the `RESTORE ALL FOREIGN DATAFILES` command.
6. On the Database service compute node, invoke SQL\*Plus and connect to the root. Execute the `CREATE PLUGGABLE DATABASE` command.
7. On the Database service compute node, execute the `ALTER PLUGGABLE DATABASE OPEN` command.

For more information, see "Performing Cross-Platform Data Transport in CDBs and PDBs" in *Oracle Database Backup and Recovery User's Guide* for Release [12.2](#) or [12.1](#).

### RMAN Cross-Platform Transportable Tablespace Backup Sets

You can use this method only if the database character sets of your on-premises database and the Database service database are compatible.



#### Note

For detailed information on a similar method that enables you to perform a cross-platform transport of an entire database, see the *Oracle Database 12c Backup and Recovery User's Guide* for Release [12.2](#) or [12.1](#). When you transport an entire database to a different platform, the source platform and the destination platform must use the same endian format.

To migrate Oracle Database 12c on-premises tablespaces to an Oracle Database 12c database on a Database service deployment using the RMAN cross-platform transportable backup sets method, you perform these tasks:

1. On the on-premises database host, prepare the database by placing the user-defined tablespaces that you intend to transport in `READ ONLY` mode.
2. On the on-premises database host, invoke RMAN and use the `BACKUP` command with the `TO PLATFORM` or `FOR TRANSPORT` clause and the `DATAPUMP` clause to create a backup set for cross-platform transport. See in "BACKUP" in *Oracle Database Backup and Recovery Reference* for Release [12.2](#) or [12.1](#) for more information on the `BACKUP` command.
3. Use a secure copy utility to transfer the backup sets, including the Data Pump export dump file, to the Database service compute node.
4. Set the on-premises tablespaces back to `READ WRITE`.
5. On the Database service compute node, prepare the database by creating the required

schemas.

6. On the Database service compute node, invoke RMAN and use the `RESTORE` command with the `foreignFileSpec` subclause to restore the cross-platform backup.
7. On the Database service compute node, set the tablespaces on the database to `READ WRITE` mode.

For more information, see "Overview of Cross-Platform Data Transport Using Backup Sets" in *Oracle Database Backup and Recovery User's Guide* for Release [12.2](#) or [12.1](#).

### **RMAN Cross-Platform Transportable Tablespace Backup Sets: Example**

This example provides a step-by-step demonstration of the tasks required to migrate tablespaces in an Oracle Database PDB to a Database service database.

This example performs a migration of the `FSDATA` and `FSINDEX` tablespaces.

In this example, the on-premises database is on a Linux host.

1. On the on-premises database host, prepare the database by creating a directory for the export dump file and placing the user-defined tablespaces that you intend to transport in `READ ONLY` mode..
  - a. On the on-premises database host, create a directory in the operating system to use for the export dump.

```
mkdir /u01/app/oracle/admin/orcl/dpdump/for_cloud
```

- b. On the on-premises data host, invoke SQL\*Plus and log in to the PDB as the `SYSTEM` user..

```
sqlplus system@pdb_servicename
Enter password: enter the password for the SYSTEM user
```

- c. Create a directory object in the on-premises database to reference the operating system directory.

```
SQL> CREATE DIRECTORY dp_for_cloud AS '/u01/app/oracle/admin/orcl/dpdump/for_cloud';
```

- d. On the on-premises database host, set all tablespaces that will be transported

(the transportable set) to READ ONLY mode.

```
SQL> ALTER TABLESPACE fsindex READ ONLY;
SQL> ALTER TABLESPACE fsdata READ ONLY;
```

- e. Exit from SQL\*Plus.
2. On the on-premises database host, invoke RMAN and use the `BACKUP` command with the `TO PLATFORM` or `FOR TRANSPORT` clause and the `DATAPUMP` clause to create a backup set for cross-platform transport.

- a. On the on-premises database host, create an operating system directory for the datafiles.

```
mkdir /u01/app/oracle/admin/orcl/rman_transdest
```

- b. Invoke RMAN and log in as a user that has been granted the `SYSDBA` or `SYSBACKUP` privilege.

```
rman target username@pdb_servicename
```

- c. Execute the `BACKUP` command.

```
RMAN> BACKUP FOR TRANSPORT
2> FORMAT '/u01/app/oracle/admin/orcl/rman_transdest/fs_tbs.bck'
3> TABLESPACE fsdata,fsindex
4> DATAPUMP FORMAT '/u01/app/oracle/admin/orcl/rman_transdest/fs_tbs.dmp';
```

- d. Log out of RMAN.
- e. Optionally, navigate to the directory you specified in the `BACKUP` command to view the files that were created.

```
cd /u01/app/oracle/admin/orcl/rman_transdest
$ ls
fs_tbs.bck fs_tbs.dmp
```

3. Use a secure copy utility to transfer the backup set, including the Data Pump export dump file, to the Database service compute node.

- a. On the Database service compute node, create a directory for the backup set and dump file.

```
mkdir /tmp/from_onprem
```

- b. Before using the `scp` command to copy files, make sure the SSH private key that provides access to the Database service compute node is available on your on-premises host.
- c. On the on-premises database host, use the SCP utility to transfer the backup set and the dump file to the Database service compute node.

```
scp -i private_key_file \
/u01/app/oracle/admin/orcl/rman_transdest/fs_tbs.bck \
oracle@IP_address_DBaaS_VM:/tmp/from_onprem

$ scp -i private_key_file \
/u01/app/oracle/admin/orcl/rman_transdest/fs_tbs.dmp \
oracle@IP_address_DBaaS_VM:/tmp/from_onprem

$
```

4. Set the on-premises tablespaces back to READ WRITE.
  - a. Invoke SQL\*Plus and log in to the PDB as the `SYSTEM` user.
  - b. Set the `FSDATA` and `FSINDEX` tablespaces back to READ WRITE mode.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
SQL> ALTER TABLESPACE fsindex READ WRITE;
```

- c. Exit from SQL\*Plus.
5. On the Database service compute node, prepare the database by creating the required schemas.
  - a. On the Database service compute node, invoke SQL\*Plus and log in to the PDB as the `SYSTEM` user.
  - b. If the owners of the objects that will be imported do not exist in the database, create them before performing the `RESTORE`.

```
SQL> CREATE USER fsowner
2 PROFILE default
3 IDENTIFIED BY fspass
4 TEMPORARY TABLESPACE temp
5 ACCOUNT UNLOCK;
```

6. On the Database service compute node, invoke RMAN and use the `RESTORE` command with the `foreignFileSpec` subclause to restore the cross-platform backup.

- a. Create an operating system directory for the Data Pump Dump file.

```
mkdir /tmp/from_onprem
```

- b. Invoke RMAN and log in to the PDB as a user that has been granted the `SYSDBA` or `SYSBACKUP` privilege.

```
rman target username@pdb_servicename
```

- c. Execute the `RESTORE` command.

```
RMAN> RESTORE FOREIGN TABLESPACE fsdata,fsindex TO NEW
2> FROM BACKUPSET '/tmp/from_onprem/fs_tbs.bck'
3> DUMP FILE DATAPUMP DESTINATION '/tmp/datapump'
4> FROM BACKUPSET '/tmp/from_onprem/fs_tbs.dmp';
```

- d. Exit from RMAN.

7. On the Database service compute node, set the tablespaces to `READ WRITE` mode.

- a. Invoke `SQL*Plus` and log in to the PDB as the `SYSTEM` user.
- b. Set the `FSDATA` and `FSINDEX` tablespaces to `READ WRITE`.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
SQL> ALTER TABLESPACE fsindex READ WRITE;
```

- c. Exit from `SQL*Plus`.

8. After verifying that the data has been imported successfully, you can delete the backup set files that were transported from the on-premises host.

### RMAN Transportable Tablespace with Data Pump

You can use this method only if the on-premises platform is little endian, and the database character sets of your on-premises database and the Databaseservice database are compatible.

You can use this method to eliminate placing the tablespaces in `READ ONLY` mode, as required by the Data Pump Transportable Tablespace method.

To migrate an on-premises source database to a database deployment on the Database service using the RMAN Transportable Tablespace with Data Pump method, you perform these tasks:

1. On the on-premises database host, invoke RMAN and create the transportable tablespace set.
2. Use a secure copy utility to transfer the Data Pump Export dump file and the tablespace datafiles to the Database service compute node.
3. On the Database service compute node, prepare the database for the tablespace import.
4. On the Database service compute node, invoke Data Pump Import and connect to the database. Import the data into the database using the `TRANSPORT_DATAFILES` option.
5. After verifying that the data has been imported successfully, you can delete the dump file.

### RMAN Transportable Tablespace with Data Pump: Example

This example provides a step-by-step demonstration of the tasks required to migrate tablespaces in an on-premises Oracle database to a Database service database.

This example performs a migration of the `FSDATA` and `FSINDEX` tablespaces.

In this example, the on-premises database is on a Linux host.

1. On the on-premises database host, invoke RMAN and create the transportable tablespace set.
  - a. On the on-premises database host, create an operating system directory for the datafiles.

```
mkdir /u01/app/oracle/admin/orcl/rman_transdest
```

- b. On the on-premises data host, create an operating system directory for the RMAN auxiliary instance files.

```
mkdir /u01/app/oracle/admin/orcl/rman_auxdest
```

- c. Invoke RMAN and log in as the SYSTEM user. Enter the password for the SYSTEM user when prompted.

```
rman target system
```

- d. Execute the TRANSPORT TABLESPACE command.

```
RMAN> TRANSPORT TABLESPACE fsdata, fsindex
2> TABLESPACE DESTINATION '/u01/app/oracle/admin/orcl/rman_transdest'
3> AUXILIARY DESTINATION '/u01/app/oracle/admin/orcl/rman_auxdest';
```

- e. Log out of RMAN.

- f. Optionally, navigate to the directory you specified for the TABLESPACE DESTINATION and view the files that were created by the TRANSPORT TABLESPACE operation.

```
cd /u01/app/oracle/admin/orcl/rman_transdest
$ ls
dmpfile.dmp fsdata01.dbf fsindex01.dbf impscript.sql
```

2. Use a secure copy utility to transfer the Data Pump Export dump file and the tablespace datafiles to the Database service compute node.

In this example the dump file is copied to the /u01 directory. Choose the appropriate location based on the size of the file that will be transferred.

- a. On the Database service compute node, create a directory for the dump file.

```
mkdir /u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

- b. Before using the scp command to copy files, make sure the SSH private key that provides access to the Database service compute node is available on your on-premises host.

- c. On the on-premises database host, use the SCP utility to transfer the dump file and all datafiles of the transportable set to the Database service compute node.

```
scp -i private_key_file \
/u01/app/oracle/admin/orcl/rman_transdest/dmpfile.dmp \
oracle@IP_address_DBaaS_VM:/u01/app/oracle/admin/ORCL/dpdump/from_onprem

$ scp -i private_key_file \
/u01/app/oracle/admin/orcl/rman_transdest/fsdata01.dbf \
```

```
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL
$ scp -i private_key_file \
/u01/app/oracle/admin/orcl/rman_transdest/fsindex01.dbf \
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL
```

3. On the Database service compute node, prepare the database for the tablespace import.
  - a. On the Database service compute node, invoke SQL\*Plus and log in to the database as the `SYSTEM` user.
  - b. Create a directory object in the Database service database.

```
SQL> CREATE DIRECTORY dp_from_onprem AS '/u01/app/oracle/admin/ORCL/dpdump/from_onprem';
```

- c. If the owners of the objects that will be imported do not exist in the database, create them before performing the import. The transportable tablespace mode of import does not create the users.

```
SQL> CREATE USER fsowner
2 PROFILE default
3 IDENTIFIED BY fspass
4 TEMPORARY TABLESPACE temp
5 ACCOUNT UNLOCK;
```

4. On the Database service compute node, invoke Data Pump Import and connect to the database.

Import the data into the database using the `TRANSPORT_DATAFILES` option.

```
impdp system DIRECTORY=dp_from_onprem DUMPFILE='dmpfile.dmp' \
TRANSPORT_DATAFILES='/u02/app/oracle/oradata/ORCL/fsdata01.dbf', \
'/u02/app/oracle/oradata/ORCL/fsindex01.dbf'
```

5. After verifying that the data has been imported successfully, you can delete the `dmpfile.dmp` dump file.

## RMAN CONVERT Transportable Tablespace with Data Pump

You can use this method only if the database character sets of your on-premises database and the Database service database are compatible.

This method is similar to the Data Pump Transportable Tablespace method, with the addition of the RMAN `CONVERT` command to enable transport between platforms with different endianness. Query `V$TRANSPORTABLE_PLATFORM` to determine if the on-premises database platform supports cross-platform tablespace transport and to determine the endian format of the platform. The Database service platform is little-endian format.

To migrate tablespaces from your on-premises Oracle database to a database deployment on the Database service using RMAN, you perform these tasks:

1. On the on-premises database host, prepare the database for the Data Pump transportable tablespace export.
2. On the on-premises database host, invoke Data Pump Export to perform the transportable tablespace export.
3. On the on-premises database host, invoke RMAN and use the `CONVERT TABLESPACE` command to convert the tablespace datafile to the Database service platform format. Refer to the Oracle Database Backup and Recovery Reference for more information on the `CONVERT` command.
4. Use a secure copy utility to transfer the Data Pump Export dump file and the converted tablespace datafiles to the Database service compute node.
5. Set the on-premises tablespaces back to `READ WRITE`.
6. On the Database service compute node, prepare the database for the tablespace import.
7. On the Database service compute node, invoke Data Pump Import and connect to the database.
8. On the Database service compute node, set the tablespaces in the database to `READ WRITE` mode.
9. After verifying that the data has been imported successfully, you can delete the dump file.

### **RMAN CONVERT Transportable Tablespace with Data Pump: Example**

This example provides a step-by-step demonstration of the tasks required to migrate tablespaces in an on-premises Oracle database to a Database service database.

In this example, the on-premises database is on a Linux host.

1. On the on-premises database host, prepare the database for the Data Pump transportable tablespace export.

- a. On the on-premises database host, create a directory in the operating system to use for the on-premises export.

```
mkdir /u01/app/oracle/admin/orcl/dpdump/for_cloud
```

- b. On the on-premises database host, invoke SQL\*Plus and log in to the on-premises database as the SYSTEM user.

```
sqlplus system
Enter password: <enter the password for the SYSTEM user>
```

- c. Create a directory object in the on-premises database to reference the operating system directory.

```
SQL> CREATE DIRECTORY dp_for_cloud AS '/u01/app/oracle/admin/orcl/dpdump/for_cloud';
```

- d. On the on-premises database host, set all tablespaces that will be transported (the transportable set) to READ ONLY mode.

```
SQL> ALTER TABLESPACE fsindex READ ONLY;
Tablespace altered.
SQL> ALTER TABLESPACE fsdata READ ONLY;
Tablespace altered.
```

- e. Exit from SQL\*Plus.

2. On the on-premises database host, invoke Data Pump Export to perform the transportable tablespace export.

On the on-premises database host, invoke Data Pump Export and connect to the on-premises database. Export the on-premises tablespaces using the `TRANSPORT_TABLESPACES` option. Provide the password for the SYSTEM user when prompted.

```
expdp system TRANSPORT_TABLESPACES=fsdata,fsindex TRANSPORT_FULL_CHECK=YES DIRECTORY=dp_for_cloud
```

3. On the on-premises database host, invoke RMAN and use the `CONVERT TABLESPACE` command to convert the tablespace datafile to the Database service platform format.

- a. Invoke RMAN.

```
rman target /
```

- b. Execute the RMAN `CONVERT TABLESPACE` command to convert the datafiles and store the converted files in a temporary location on the on-premises database host.

```
RMAN> CONVERT TABLESPACE fsdata, fsindex
 2> TO PLATFORM 'Linux x86 64-bit'
 3> FORMAT '/tmp/%U ';
...
input datafile file number=00006 name=/u01/app/oracle/oradata/orcl/fsdata01.dbf
converted datafile=/tmp/data_D-ORCL_I-1410251631_TS-FSDATA_FNO-6_0aqc9un3
...
input datafile file number=00007 name=/u01/app/oracle/oradata/orcl/fsindex01.dbf
converted datafile=/tmp/data_D-ORCL_I-1410251631_TS-FSINDEX_FNO-7_0bqc9un6
...
```

- c. Take note of the names of the converted files. You will copy these files to the Database service compute node in the next step.
  - d. Exit RMAN.
4. Use a secure copy utility to transfer the Data Pump Export dump file and the converted tablespace datafiles to the Database service compute node.

In this example the dump file is copied to the `/u01` directory. Choose the appropriate location based on the size of the file that will be transferred.

- a. On the Databaseservice compute node, create a directory for the dump file.

```
mkdir /u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

- b. Before using the `scp` command to copy files, make sure the SSH private key that provides access to the Database service compute node is available on your on-premises host.
- c. On the on-premises database host, use the `scp` utility to transfer the dump file and all data files of the transportable set to the Database service compute node.

```
scp -i private_key_file \
/u01/app/oracle/admin/orcl/dpdump/for_cloud/expdat.dmp \
oracle@IP_address_DBaaS_VM:/u01/app/oracle/admin/ORCL/dpdump/from_onprem
```

```
$ scp -i private_key_file \
/tmp/data_D-ORCL_I-1410251631_TS-FSDATA_FNO-6_0aqc9un3 \
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL/fsdata01.dbf

$ scp -i private_key_file \
/tmp/data_D-ORCL_I-1410251631_TS-FSINDEX_FNO-7_0bqc9un6 \
oracle@IP_address_DBaaS_VM:/u02/app/oracle/oradata/ORCL/fsindex01.dbf
```

5. Set the on-premises tablespaces back to READ WRITE.
  - a. Invoke SQL\*Plus and log in as the SYSTEM user.
  - b. Set the FSDATA and FSINDEX tablespaces back to READ WRITE mode.
6. On the Database service compute node, prepare the database for the tablespace import.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ WRITE;
Tablespace altered.
```

- c. Exit from SQL\*Plus.
- a. On the Database service compute node, invoke SQL\*Plus and log in to the database as the SYSTEM user.
- b. Create a directory object in the Database service database.

```
SQL> CREATE DIRECTORY dp_from_onprem AS '/u01/app/oracle/admin/ORCL/dpdump/from_onprem';
```

- c. If the owners of the objects that will be imported do not exist in the database, create them before performing the import. The transportable tablespace mode of import does not create the users.

```
SQL> CREATE USER fsowner
2 PROFILE default
3 IDENTIFIED BY fspass
4 TEMPORARY TABLESPACE temp
5 ACCOUNT UNLOCK;
```

7. On the Database service compute node, invoke Data Pump Import and connect to the database.

Import the data into the Database service database using the `TRANSPORT_DATAFILES` option

```
impdp system DIRECTORY=dp_from_onprem \
TRANSPORT_DATAFILES='/u02/app/oracle/oradata/ORCL/fsdata01.dbf', \
'/u02/app/oracle/oradata/ORCL/fsindex01.dbf'
```

8. On the Database service compute node, set the tablespaces in the database to `READ WRITE` mode.
  - a. Invoke `SQL*Plus` and log in as the `SYSTEM` user.
  - b. Set the `FSDATA` and `FSINDEX` tablespaces to `READ WRITE` mode.

```
SQL> ALTER TABLESPACE fsdata READ WRITE;
Tablespace altered.
SQL> ALTER TABLESPACE fsindex READ WRITE;
Tablespace altered.
```

- c. Exit from `SQL*Plus`.
9. After verifying that the data has been imported successfully, you can delete the `expdat.dmp` dump file.

## RMAN DUPLICATE from an Active Database

This topic explains how to migrate an entire, active container database (CDB) or non-CDB database to Oracle Cloud Infrastructure by using RMAN Active Duplication. The database to be migrated can reside on-premises or in Oracle Cloud Infrastructure Classic. This topic does not cover duplicating a pluggable database, or migrating a pluggable database or non-CDB to a CDB in the cloud.

The following terms are used throughout this topic:

- Source database: The active database to be migrated.
- Target database: The new database (duplicated from the source database) on a DB system in the Oracle Cloud Infrastructure.



### Note

Version 11.2.0.4 databases will be migrated to a DB system using a ACFS storage.

### Prerequisites

For the source database to be migrated, you'll need:

- The source database name, database unique name, listener port, service name, database home patch level, and the password for SYS.
- A copy of the sqlpatch directory from the source database home. This is required for rollback in case the target DB system does not include these patches.
- If the source database is configured with Transparent Data Encryption (TDE), you'll need a backup of the wallet and the wallet password to allow duplication of a database with encrypted data.

When migrating a source database to an existing target database, Oracle recommends that you patch the source environment to the same database bundle patch level as the target database home. If the source environment has an interim patch (previously known as a "one-off" patch) that includes a sqlpatch component, and that sqlpatch is missing from the target environment (or a different cumulative patch is applied), the interim patch should be rolled back in the source environment before the migration, if possible.



### Tip

To check for interim patches installed on the source or target database, use the `$ORACLE_HOME/OPatch/opatch lspatches` command. To roll back SQL changes in the target database, copy the `$ORACLE_HOME/sqlpatch/<patch_number>/postdeinstall.sql` script from the source environment to the cloud environment and execute the `postdeinstall.sql` script.

For the target database, you'll need:

- A target DB system that supports the same database edition as the source database edition. When you launch a DB system, an initial database is created on it. If necessary, you can delete that database and create a new one by using the dbcli command line interface. For more information on creating a DB system, see [Creating Bare Metal and Virtual Machine DB Systems](#). For information about creating a database with the DBCLI, see [Database Commands](#).
- The target database name, database unique name, auxiliary service name, and database home patch level.
- A free TCP port in the target database to setup the auxiliary instance.

If you need to roll back interim patches in the target environment so that the patch level matches that of the source environment, copy the source DB `$ORACLE_HOME/sqlpatch/<patch_number>` directory to the target database home.

### **Migrating Source Databases That Include Patch Set Updates (PSUs)**

In Oracle Cloud Infrastructure DB systems, the database home includes an installation of Database Proactive Bundle Patches. If the source DB uses Patch Set Updates (PSUs), follow the instructions in [MOS Note:1962125.1](#) (Oracle Database - Overview of Database Patch Delivery Methods) for migrating the DB into Oracle Cloud Infrastructure.

### Verifying the Environment

Perform the following steps before you begin the migration:

1. Make sure the source DB system is reachable from the target DB system. You should be able to SSH between the two hosts.
2. On the target host, use the TNSPING utility to make sure the source host listener port works. For example:

```
tnsping <source_host>:1521
```

3. On the target host, use Easy Connect to verify the connection to the source database:

```
<host>:<port>/<service_name>
```

For example:

```
sqlplus system@129.145.0.164:1521/proddb
```

Make sure the connection string does not exceed 64 characters.

4. Copy the required sqlpatch files (for rollback) from the source database home to the target database.
5. Make sure at least one archivelog has been created on the source database, otherwise, the RMAN duplication will fail with an error.
6. If the source database uses wallets, back up the password-based wallet and copy it to the standard location in the DB system:

```
/opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>/
```

7. Make sure the compatibility parameters in the source database are set to at least:
  - 18.0.0.0.0 for an 18.1.0.0 database
  - 12.1.0.2.0 for a 12.1.0.2 or a 12.2.0.1 database
  - 11.2.0.4.0 for an 11.2.0.4 database

## Setting Up Storage on the DB System

1. SSH to the DB System.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile, which will set the PATH to the dbcli directory (`/opt/oracle/dcs/bin`).

```
login as: opc
```

```
[opc@dbsys ~]$ sudo su -
```

3. Use the [dbcli create-dbstorage](#) to set up directories for DATA, RECO, and REDO storage. The following example creates 10GB of ACFS storage for the tdetest database.

```
[root@dbsys ~]# dbcli create-dbstorage --dbname tdetest --dataSize 10 --dbstorage ACFS
```



### Note

When migrating a version 11.2 database, ACFS storage must be specified.

4. Use the [dbcli list-dbstorages](#) command to list the storage ID. You'll need the ID for the next step.

```
[root@dbsys ~]# dbcli list-dbstorages
```

ID	Type	DBUnique Name	Status
9dcdfb8e-e589-4d5f-861a-e5ba981616ed	Acfs	tdetest	Configured

5. Use the [dbcli describe-dbstorage](#) command with the storage ID from the previous step to list the DATA, RECO and REDO locations.

```
[root@dbsys ~]# dbcli describe-dbstorage --id 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
```

```
DBStorage details
```

```

ID: 9dcdfb8e-e589-4d5f-861a-e5ba981616ed
DB Name: tdetest
DBUnique Name: tdetest
```

## CHAPTER 11 Database

---

```
DB Resource ID:
 Storage Type: Acfs
 DATA Location: /u02/app/oracle/oradata/tdetest
 RECO Location: /u03/app/oracle/fast_recovery_area/
 REDO Location: /u03/app/oracle/redo/
 State: ResourceState(status=Configured)
 Created: August 24, 2016 5:25:38 PM UTC
 UpdatedTime: August 24, 2016 5:25:53 PM UTC
```

Note the locations. You'll use them later to set the `db_create_file_dest`, `db_create_online_log_dest`, and `db_recovery_file_dest` parameters for the database.

### Choosing an ORACLE\_HOME

Decide which ORACLE\_HOME to use for the database restore and then switch to that home with the correct ORACLE\_BASE, ORACLE\_HOME, and PATH settings.

To get a list of existing ORACLE\_HOMEs, use the [dbcli list-dbhomes](#) command. To create a new ORACLE\_HOME, use the [dbcli create-dbhome](#) command.

### Copying the Source Database Wallets

Skip this section if the source database is **not** configured with TDE.

1. On the DB system, become the oracle user:

```
sudo su - oracle
```

2. Create the following directory if it does not already exist:

```
mkdir /opt/oracle/dcs/commonstore/wallets/tde/<db_unique_name>
```

3. Copy the ewallet.p12 file from the source database to the directory you created in the previous step.
4. On the target host, make sure that `$ORACLE_HOME/network/admin/sqlnet.ora` contains the following line:

```
ENCRYPTION_WALLET_LOCATION= (SOURCE= (METHOD=FILE) (METHOD_DATA=
(DIRECTORY=/opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

Add the line if it doesn't exist in the file. (The line might not be there if this is a new home and no database has been created yet on this host.)

5. Create the autologin wallet from the password-based wallet to allow auto-open of the wallet during restore and recovery operations.

For version 12c, use the `ADMINISTER KEY MANAGEMENT` command:

```
$cat create_autologin_12.sh

#!/bin/sh
if [$# -lt 2]; then
 echo "Usage: $0 <db_unique_name> <remote_wallet_location>"
 exit 1;
fi

mkdir /opt/oracle/dcs/commonstore/wallets/tde/$1
cp $2/ewallet.p12* /opt/oracle/dcs/commonstore/wallets/tde/$1
rm -f autokey.ora
echo "db_name=$1" > autokey.ora
autokeystoreLog="autologinKeystore_`date +%Y%m%d_%H%M%S_%N`.log"
echo "Enter Keystore Password:"
read -s keystorePassword
echo "Creating AutoLoginKeystore -> "
sqlplus "/as sysdba" <<EOF
spool $autokeystoreLog
set echo on
startup nomount pfile=autokey.ora
ADMINISTER KEY MANAGEMENT CREATE AUTO_LOGIN KEYSTORE
FROM KEYSTORE '/opt/oracle/dcs/commonstore/wallets/tde/$1' -- Keystore location
IDENTIFIED BY "$keystorePassword";
shutdown immediate;
EOF
```

For version 11g, use the `orapki` command:

```
orapki wallet create -wallet wallet_location -auto_login [-pwd <password>]
```

### Setting Up the Static Listener

Set up the static listener for the auxiliary instance for RMAN duplication.

1. On the DB system, create `$ORACLE_HOME/network/admin/listener.ora` and add the following content to it.

```
LISTENER_aux_<db_unique_name>=
 (DESCRIPTION=
 (ADDRESS_LIST=
 (ADDRESS=(PROTOCOL=TCP) (HOST=<hostname> or <ip_address>) (PORT=<available_TCP_port>))
)
)
)
SID_LIST_LISTENER_aux_<db_unique_name>=
 (SID_LIST=
 (SID_DESC=
 (GLOBAL_DBNAME=<auxServiceName_with_domain>
 (ORACLE_HOME=<Oracle_home_for_target_database>
 (SID_NAME=<database_name>
 (ENVS="TNS_ADMIN=<path_to_tnsnames.ora>")
 (ENVS="ORACLE_UNQNAME=<db_unique_name(in lower case)>"))
)
)
)
```

2. Make sure the port specified in `(PORT=<available_TCP_port>)` is open in the DB system's iptables and in the DB system's cloud network Security List.

### Using the RMAN Duplicate Command to Migrate the Database

1. Set the following environment variables for RMAN and SQL Plus sessions for the database:

```
ORACLE_HOME=<path_of_Oracle_home_where_the_database_is_to_be_restored>
ORACLE_SID=<database_name>
ORACLE_UNQNAME=<db_unique_name(in lower case)>
NLS_DATE_FORMAT="mm/dd/yyyy hh24:mi:ss"
```

2. Start the listener:

```
lsnrctl start listener_aux_<db_unique_name>
```

3. Create an `init.ora` file with the minimal required parameters as described in [Creating an Initialization Parameter File and Starting the Auxiliary Instance](#) and use it for the auxiliary instance.
4. Start the auxiliary instance in nomount mode:

```
startup nomount
```

5. Run the following commands to duplicate the database. Note that the example below uses variables to indicate the values to be specified:

```
rman target sys/$sourceSysPassword@$sourceNode:$sourceListenerPort/$sourceDb auxiliary
sys/$auxSysPassword@$targetNode:$targetListenerPort/$auxService<<EOF

spool log to "`date +%Y%m%d_%H%M%S_%N`_duplicate_${targetDbUniqueName}_from_${sourceDb}.log"
set echo on

duplicate target database to $targetDb from active database
password file
spfile
 PARAMETER_VALUE_CONVERT $sourceDb $targetDb $sourceDbUniqueNameCaps $targetDbUniqueNameCaps
set cluster_database='false'
set db_name='$targetDb'
set db_unique_name='$targetDbUniqueName'
set db_create_file_dest='$dataLoc'
set db_create_online_log_dest_1='$redoLoc'
set db_recovery_file_dest='$recoLoc'
set audit_file_dest = '$auditFileDest'
reset control_files
nofilenamecheck
;
EOF
```

### Preparing to Register the Database

Before you register the database:

1. Make sure the database COMPATIBLE parameter value is acceptable.  
For a 11.2 database, the minimum compatibility value is 11.2.0.4.  
For a 12c database, the minimum compatibility value is 12.1.0.2.  
If the value is less than the minimum, the database cannot be registered until you upgrade the database compatibility.
2. Use the following command to verify that the database has registered with the local listener and service name.

```
lsnrctl services
```

3. Use the following command to verify that the password file was restored or created for a new database.

```
ls -ltr $ORACLE_HOME/dbs/orapw<$ORACLE_SID>
```

If the file does not exist, create it using the orapwd command.

```
orapwd file=<$ORACLE_HOME/dbs/orapw<$ORACLE_SID>> password=<sys_password>
```

4. Use the following command to verify that the restored database is open in read write mode.

```
select open_mode from v$database;
```

Read write mode is required to register the database later. Any PDBs must also be in read write mode.

5. From oracle home on the migrated database host, use the following command verify the connection to SYS.

```
conn sys/<password>@<service_name> as sysdba
```

This connection is required to register the database later. Fix any connection issues before continuing.

6. Copy the folder \$ORACLE\_HOME/sqlpatch from source database to the target database. This will enable the dbcli register-database command to rollback any conflicting patches.



### Note

If you are migrating a version 11.2 database, additional steps are required after you register the database. For more information, see [Rolling Back Patches on a Version 11.2 Database](#).

7. Use the following SQL\*Plus command to make sure the database is using the spfile.

```
SHOW PARAMETERS SPFILE
```

### Registering the Database on the DB System

The [dbcli register-database](#) command registers the migrated database to the dcs-agent so it can be managed by the dcs-agent stack.



#### Note

The `dbcli register-database` command is not available on 2-node RAC DB Systems.

As the root user, use the `dbcli register-database` command to register the database on the DB system, for example:

```
[root@dbsys ~]# dbcli register-database --dbclass OLTP --dbshape odb1 --servicename crmdb.example.com --syspassword syspassword
Password for SYS:
{
 "jobId" : "317b430f-ad5f-42ae-bb07-13f053d266e2",
 "status" : "Created",
 "message" : null,
 "reports" : [],
 "createTimestamp" : "August 08, 2016 05:55:49 AM EDT",
 "description" : "Database service registration with db service name: crmdb.example.com",
 "updatedAtTime" : "August 08, 2016 05:55:49 AM EDT"
}
```

### Migrating a Version 12.1 or Later Database That Includes SQL Patch Components

For a 1-node DB system at version 12.1 or higher, the `dbcli register-database` command automates the datapatch execution. Before executing the `dbcli register-database` command, open all PDBs in read-write mode. If you have already run the `dbcli register-database` command and did not open all PDBs, or did not copy the `$ORACLE_HOME/sqlpatch` directory from the source database home, manually rerun the datapatch utility to configure the SQL portion of existing interim patches. This can be done by executing the command `$ORACLE_HOME/OPatch/opatch datapatch`.



### Tip

If the source database includes patch 23170620 and the target database is running with the October 2017 patch or a later one, the `$ORACLE_HOME/sqlpatch` directory does not need to be copied to the target database, because the contents of the patch are already installed in the target database.

### Rolling Back Patches on a Version 11.2 Database

For version 11.2 databases, the `sqlpatch` application is not automated, so any interim patches (previously known as a "one-off" patches) applied to the source database that are not part of the installed PSU must be rolled back manually in the target database. After registering the database, execute the `catbundle.sql` script and then the `postinstall.sql` script with the corresponding PSU patch (or the overlay patch on top of the PSU patch), as described below.



### Tip

Some interim patches may include files written to the `$ORACLE_HOME/rdbms/admin` directory as well as the `$ORACLE_HOME/sqlpatch` directory. Oracle recommends that you roll back these patches in the source database using the instructions in the patch read-me prior to migrating the database to OCI environment. Contact Oracle Support if you need assistance with rolling back these patches.

1. On the DB System, use the `dbcli list-dbhomes` command to find the PSU patch number for the version 11.2 database home. In the following sample command output, the PSU patch number is the second number in the DB Version column:

## CHAPTER 11 Database

```
[root@dbsys ~]# dbcli list-dbhomes
ID Name DB Version
Home Location Status

59d9bc6f-3880-4d4f-b5a6-c140f16f8c64 OraDB11204_home1 11.2.0.4.160719 (23054319, 23054359)
/u01/app/oracle/product/11.2.0.4/dbhome_1 Configured
```

(The first patch number, 23054319 in the example above, is for the OCW component in the database home.)

2. Find the overlay patch, if any, by using the `lsinventory` command. In the following example, patch number **24460960** is the overlay patch on top of the 23054359 PSU patch.

```
$ $ORACLE_HOME/OPatch/opatch lsinventory
...
Installed Top-level Products (1):

Oracle Database 11g 11.2.0.4.0
There are 1 products installed in this Oracle Home.

Interim patches (5) :

Patch 24460960 : applied on Fri Sep 02 15:28:17 UTC 2016
Unique Patch ID: 20539912
 Created on 31 Aug 2016, 02:46:31 hrs PST8PDT
 Bugs fixed:
 23513711, 23065323, 21281607, 24006821, 23315889, 22551446, 21174504
 This patch overlays patches:
 23054359
 This patch needs patches:
 23054359
 as prerequisites
```

3. Start SQL\*Plus and execute the `catbundle.sql` script, for example:

```
SQL> startup
SQL> connect / as sysdba
```

```
SQL> @$ORACLE_HOME/rdbms/admin/catbundle.sql psu apply
exit
```

4. Apply the sqlpatch, using the overlay patch number from the previous step, for example:

```
SQL> connect / as sysdba
SQL> @$ORACLE_HOME/sqlpatch/24460960/postinstall.sql
exit
```

### Creating a Backup Configuration (Optional)

If you would like to manage the database backup with the dbcli command line interface, you can associate a new or existing backup configuration with the migrated database when you register it or after you register it. A backup configuration defines the backup destination and recovery window for the database. As the root user, use the following commands to create, list, and display backup configurations:

- [dbcli update-backupconfig](#)
- [dbcli list-backupconfigs](#)
- [dbcli describe-backupconfig](#)

### Post Migration Checklist

After the database is migrated and registered on the DB system, use the following checklist to verify the results of the migration and perform any post-migration customizations.

1. Make sure the database files were restored in OMF format.
2. Make sure the database is listed in the [dbcli list-databases](#) command output.
3. Check for the following external references in the database and update them if necessary:
  - External tables: If the source database uses external tables, back up that data and migrate it to the target host.
  - Directories: Customize the default directories as needed for the migrated database.

- Database links: Make sure all the required TNS entries are updated in the tnsnames.ora file in ORACLE\_HOME.
  - Email and URLs: Make sure any email addresses and URLs used in the database are still accessible from the DB system.
  - Scheduled jobs: Review the jobs scheduled in source database and schedule similar jobs as needed in the migrated database.
4. If you associated a backup configuration when you registered the database, run a test back up using the [dbcli create-backup](#) command.
  5. Verify that patches have been applied to all PDBs if the migrated database contains CDB and PDBs.
  6. Validate the database performance by using Database Replay and SQL Performance Analyzer for SQL. For more information, see the [Database Testing Guide](#).

### SQL Developer and INSERT Statements to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into an Oracle Database 12c database in the Oracle Cloud Infrastructure Database service.

In this method, you use SQL `INSERT` statements to load the data into your cloud database.

To migrate selected objects to an Oracle Database 12c database in a Database service deployment using SQL Developer and `INSERT` statements, you perform these tasks:

1. Launch SQL Developer, connect to your on-premises database and create a cart containing the objects you want to migrate.
2. In SQL Developer, click the Export Cart icon and select "Insert" in the Format menu.
3. In SQL Developer, open a connection to the Oracle Database 12c database in the Database service and execute the generated script to create the database objects.
4. In SQL Developer, open a connection to the Oracle Database 12c database in the Database service and run the generated script to create the objects and load the data.

### SQL Developer and SQL\*Loader to Migrate Selected Objects

You can use SQL Developer to create a cart into which you add selected objects to be loaded into an Oracle Database 12c database in the Oracle Cloud Infrastructure Database.

In this method, you use SQL\*Loader to load the data into your cloud database.

To migrate selected objects to an Oracle Database 12c database in the Database service deployment using SQL Developer and SQL\*Loader, you perform these tasks:

1. Launch SQL Developer, connect to your on-premises database and create a cart containing the objects you want to load into your cloud database.
2. In SQL Developer, click the Export Cart icon and select "loader" in the Format menu.
3. In SQL Developer, open a connection to the Oracle Database 12c database on the Database service and execute the generated script to create the database objects.
4. Use a secure copy utility to transfer the SQL\*Loader control files and the SQL\*Loader data files to the Database service compute node.
5. On the Database service compute node, invoke SQL\*Loader to load the data using the SQL\*Loader control files and data files for each object.

### Unplugging/Plugging a PDB

You can use this method only if the on-premises platform is little endian, and the on-premises database and the Oracle Cloud Infrastructure Database service database have compatible database character sets and national character sets.

You can use the unplug/plug method to migrate an Oracle Database 12c PDB to a PDB in an Oracle Database 12c database on a Database service database deployment.

To migrate an Oracle Database 12c PDB to a PDB in the Oracle Database 12c database on an Oracle Cloud Infrastructure Database service database deployment using the plug/unplug method, you perform these tasks:

1. On the on-premises database host, invoke SQL\*Plus and close the on-premises PDB.
2. On the on-premises database host, execute the `ALTER PLUGGABLE DATABASE UNPLUG`

command to generate an XML file containing the list of datafiles that will be plugged in to the database on the Database service.

3. Use a secure copy utility to transfer the XML file and the datafiles to the Databaseservice compute node.
4. On the Database service compute node, invoke SQL\*Plus and execute the `CREATE PLUGGABLE DATABASE` command to plug the database into the CDB.
5. On the Database service compute node, open the new PDB by executing the `ALTER PLUGGABLE DATABASE OPEN` command.

For more information, see "Creating a PDB by Plugging an Unplugged PDB into a CDB" in *Oracle Database Administrator's Guide* for Release [12.2](#) or [12.1](#).

### Unplugging/Plugging Non-CDB

You can use this method only if the on-premises platform is little endian, and the on-premises database and the Oracle Cloud Infrastructure Database database have compatible database character sets and national character sets.

You can use the unplug/plug method to migrate an Oracle Database 12c non-CDB database to a PDB in an Oracle Database 12c database on a Database service database deployment. This method provides a way to consolidate several non-CDB databases into a single Oracle Database 12c multitenant database on the Database service.

To migrate an Oracle Database 12c non-CDB database to the Oracle Database 12c database on a Database service deployment using the plug/unplug method, you perform these tasks:

1. On the on-premises database host, invoke SQL\*Plus and set the on-premises database to `READ ONLY` mode.
2. On the on-premises database host, execute the `DBMS_PDB.DESCRIBE` procedure to generate an XML file containing the list of datafiles that will be plugged in on the cloud database.
3. Use a secure copy utility to transfer the XML file and the datafiles to the Database service compute node.

4. On the Database service compute node, invoke SQL\*Plus and execute the `CREATE PLUGGABLE DATABASE` command to plug the database into the CDB.
5. On the Database service compute node, execute the `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` script to delete unnecessary metadata from the `SYSTEM` tablespace of the new PDB.
6. On the Database service compute node, open the new PDB by executing the `ALTER PLUGGABLE DATABASE OPEN` command.
7. Optionally, on the on-premises database host invoke SQL\*Plus and set the on-premises database back to `READ WRITE` mode.

For more information, see "Creating a PDB Using a Non-CDB" in *Oracle Database Administrator's Guide* for Release [12.2](#) or [12.1](#).

## Troubleshooting

These topics cover some common issues you might run into and how to address them.

- [Backup Failures on Bare Metal and Virtual Machine DB Systems](#)
- [Patching Failures on Bare Metal and Virtual Machine DB Systems](#)

### Backup Failures on Bare Metal and Virtual Machine DB Systems

Database backups can fail for various reasons. Typically, a backup fails because either the database host cannot access the object store, or there are problems on the host or with the database configuration.

This topic includes information to help you determine the cause of the failure and fix the problem. The information is organized into several sections, based on the error condition. If you already know the cause, you can skip to the section with the suggested solution. Otherwise, use the procedure in [Determining the Problem](#) to get started.

### Determining the Problem

In the Console, a failed database backup either displays a status of **Failed** or hangs in the **Backup in Progress** or **Creating** state. If the error message does not contain enough information to point you to a solution, you can use the database CLI and log files to gather more data. Then, refer to the applicable section in this topic for a solution.

### To identify the root cause of the backup failure

1. Log on to the host as the root user and navigate to the `/opt/oracle/dcs/bin/` directory.
2. Determine the sequence of operations performed on the database.

```
dbcli list-jobs | grep -i <dbname>
```

Note the last job ID listed with a status other than **Success**.

3. With the job ID you noted from the previous step, use the following command to check the details of that job:

```
dbcli describe-job -i <job_ID> -j
```

Typically, running this command is enough to reveal the root cause of the failure.

4. If you require more information, review the `/opt/oracle/dcs/log/dcs-agent.log` file.

You can find the job ID in this file by using the timestamp returned by the job report in step 2.

5. If the problem details suggest an RMAN issue, review the RMAN logs in the `/opt/oracle/dcs/log/<hostname>/rman/bkup/<db_unique_name>/rman_backup/<yyyy-mm-dd>` directory.



### Note

If the database failure is on a 2-node RAC database, perform steps 3 and 4 on both nodes.

### Database Service Agent Issues

Your Oracle Cloud Infrastructure Database makes use of an agent framework to allow you to manage your database through the cloud platform. Occasionally you might need to restart the dcsagent program if it has the status of **stop/waiting** to resolve a backup failure.

#### To restart the database service agent

1. From a command prompt, check the status of the agent:

```
initctl status initdcsagent
```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

```
initctl start initdcsagent
```

3. Check the status of the agent again to confirm that it has the **start/running** status:

```
initctl status initdcsagent
```

### Oracle Clusterware Issues

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. Occasionally you might need to restart the Clusterware program to resolve a backup failure.

#### To restart the Oracle Clusterware

1. From command prompt, check the status of Oracle Clusterware:

## CHAPTER 11 Database

---

```
crsctl check crs
```

```
crsctl stat res -t
```

2. If Oracle Clusterware is not online, try to restart the program:

```
crsctl start crs
```

3. Check the status of Oracle Clusterware to confirm that it is online:

```
crsctl check crs
```

### Object Store Connectivity Issues

Backing up your database to Oracle Cloud Infrastructure Object Storage requires that the host can connect to the applicable Swift endpoint. You can test this connectivity by using a Swift user.

To ensure your database host can connect to the object store

1. Create a Swift user in your tenancy. See [Working with Auth Tokens](#).
2. With the user you created in the previous step, use the following command to verify the host can access the object store.

```
curl -v -X HEAD -u <user_ID>:<auth_token> https://swiftobjectstorage.<region_name>.oraclecloud.com/v1/<tenant>
```

See [Object Storage FAQ](#) for the correct region to use.

3. If you cannot connect to the object store, refer to [Prerequisites](#) for how to configure object store connectivity.

### Host Issues

One or more of the following conditions on the database host can cause backups to fail:

### INTERACTIVE COMMANDS IN THE ORACLE PROFILE

If an interactive command such as `oraenv`, or any command that might return an error or warning message, was added to the `.bash_profile` file for the `grid` or `oracle` user, Database service operations like automatic backups can be interrupted and fail to complete. Check the `.bash_profile` file for these commands, and remove them.

### THE FILE SYSTEM IS FULL

Backup operations require space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the space available for backups. If the file system has insufficient space, you can remove old log or trace files to free up space.

### INCORRECT VERSION OF THE ORACLE DATABASE CLOUD BACKUP MODULE

Your system might not have the required version of the backup module (`opc_installer.jar`). See [Unable to use Managed Backups in your DB System](#) for details about this known issue. To fix the problem, you can follow the procedure in that section or simply update your DB system and database with the latest bundle patch.

### CHANGES TO THE SITE PROFILE FILE (GLOGIN.SQL)

[Customizing](#) the site profile file (`$ORACLE_HOME/sqlplus/admin/glogin.sql`) can cause managed backups to fail in Oracle Cloud Infrastructure. In particular, interactive commands can lead to backup failures. Oracle recommends that you not modify this file for databases hosted in Oracle Cloud Infrastructure.

## Database Issues

An improper database state or configuration can lead to failed backups.

### DATABASE NOT RUNNING DURING BACKUP

The database must be active and running while the backup is in progress.

### To check that the database is active and running

Use the following command to check the state of your database, and ensure that any

## CHAPTER 11 Database

---

problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```

The system returns a message including the database's instance status. The instance status must be **Open** for the backup to succeed. If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the **Open** status, use the following commands to access the SQL\*Plus command prompt and set the status to **Open**:

```
sqlplus / as sysdba
```

```
alter database open;
```

### ARCHIVING MODE SET TO NOARCHIVELOG

When you provision a new database, the archiving mode is set to `ARCHIVELOG` by default. This is the required archiving mode for backup operations. Check the archiving mode setting for the database and change it to `ARCHIVELOG`, if applicable.

### To check and set the archiving mode

Open an SQL\*Plus command prompt and enter the following command:

```
select log_mode from v$database;
```

If you need to set the archiving mode to `ARCHIVELOG`, start the database in **Mount** status (and not **Open** status), and use the following command at the SQL\*Plus command prompt:

```
alter database archivelog;
```

Be sure to confirm that the `db_recovery_file_dest` parameter points to `+RECO`, and that the `log_archive_dest_1` parameter is set to `USE_DB_RECOVERY_FILE_DEST`.

For RAC databases, one instance must have the **Mount** status when enabling archivelog mode. To enable archivelog mode for a RAC database, perform the following steps:

## CHAPTER 11 Database

---

1. Shutdown all database instances:

```
srvctl stop database -d
```

2. Start one of the database instances in mount state:

```
srvctl start instance -d <db_unique_name> -i <instance_name> -o mount
```

3. Access the SQL\*Plus command prompt:

```
sqlplus / as sysdba
```

4. Enable archive log mode:

```
alter database archivelog;
```

```
exit;
```

5. Stop the database:

```
srvctl stop instance -d <db_unique_name> -i <instance_name>
```

6. Re-start all database instances:

```
srvctl start database -d <db_unique_name>
```

7. At the SQL\*Plus command prompt, confirm the archiving mode is set to ARCHIVELOG:

```
select log_mode from v$database;
```

### STUCK DATABASE ARCHIVER PROCESS AND BACKUP FAILURES

Backups can fail when the database instance has a stuck archiver process. For example, this can happen when the flash recovery area (FRA) is full. You can check for this condition using the `srvctl status database -db <db_unique_name> -v` command. If the command returns the following output, you must resolve the stuck archiver process issue before backups will succeed:

```
Instance <instance_identifier> is running on node *<node_identifier>. Instance status: Stuck Archiver
```

Refer to [ORA-00257:Archiver Error \(Doc ID 2014425.1\)](#) for information on resolving a stuck archiver process.

After resolving the stuck process, the command should return the following output :

## CHAPTER 11 Database

---

```
Instance <instance_identifier> is running on node *<node_identifier>. Instance status: Open
```

If the instance status does not change after you resolve the underlying issue with the device or resource being full or unavailable, try one of the following workarounds:

- Restart the database using the `srvctl` command to update the status of the database in the clusterware
- Upgrade the database to the latest patchset levels

### TEMPORARY TABLESPACE ERRORS

If fixed table statistics are not up to date on the database, backups can fail with errors referencing temporary tablespace present in the `dcS-agent.log` file. For example:

```
select status from v$rman_status where COMMAND_ID=<backup_id>

ERROR at line 1:
ORA-01652: unable to extend temp segment by 128 in tablespace TEMP
```

Gather your fixed table statics as follows to resolve this issue:

```
conn / as sysdba

exec dbms_stats.gather_fixed_objects_stats();
```

### RMAN CONFIGURATION AND BACKUP FAILURES

Editing certain RMAN configuration parameters can lead to backup failures in Oracle Cloud Infrastructure. To check your RMAN configuration, use the `show all` command at the RMAN command line prompt.

See the following list of parameters for details about RMAN the configuration settings that should not be altered for databases in Oracle Cloud Infrastructure.

### RMAN configuration settings that should not be altered

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 30 DAYS;
```

## CHAPTER 11 Database

---

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;

CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 5 BACKUP TYPE TO COMPRESSED BACKUPSET;

CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;

CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' MAXPIECESIZE 2 G FORMAT '%d_%I_%U_%T_%t' PARMS 'SBT_
LIBRARY=/opt/oracle/dcs/commonstore/pkgrepos/oss/odbc/libopc.so ENV=(OPC_
PFILE=/opt/oracle/dcs/commonstore/objectstore/opc_pfile/1578318329/opc_tiger_iad3c8.ora)';

CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO 'SBT_TAPE';

CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;

CONFIGURE ENCRYPTION FOR DATABASE ON;
```

### RMAN RETENTION POLICY AND BACKUP FAILURES

The RMAN retention policy configuration can be the source of backup failures. Using the REDUNDANCY retention policy configuration instead of the RECOVERY WINDOW policy can lead to backup failures. Be sure to use the RECOVERY WINDOW OF 30 DAYS configuration.

#### To configure the RMAN retention policy setting

1. Find the database ID using the following command:

```
dbcli list-databases
```

2. Find the BackupConfigId value for the database using the following command:

```
dbcli describe-database -i <database_id>
```

3. Update the retention policy configuration to RECOVERY WINDOW OF 30 DAYS:

```
dbcli update-backupconfig -i <backup_config_id> --recoverywindow 30
```

### LOSS OF OBJECTSTORE WALLET FILE AND BACKUP FAILURES

RMAN backups fail when an objectstore wallet file is lost. The wallet file is necessary to enable connectivity to the object store.

To confirm that the objectstore wallet file exists and has the correct permissions

1. Find the database ID using the following command:

```
dbcli list-databases
```

2. Find the BackupConfigId value for the database using the following command:

```
dbcli describe-database -i <database_id>
```

3. Find the BackupLocation value for the database using the following command:

```
dbcli describe-backupconfig <backup_config_id>
```

4. Find the file path of the backup config parameter file (opc\_<backup\_location\_value>\_BC.ora) using the following command:

```
locate opc_<backup_location_value>_BC.ora
```

For example:

```
[root@orcl 13aef284-9d6b-4eb6-8751-2988aexample]# locate opc_b9naujWMAXzi9example_BC.ora
/opt/oracle/dcs/commonstore/objectstore/opc_pfile/13aef284-9d6b-4eb6-8751-2988a9example/opc_b9naujWMAXzi9example_BC.ora
```

5. Find the file path to the wallet file in the backup config parameter file by inspecting the value stored in the OPC\_WALLET parameter. To do this, navigate to the directory containing the backup config parameter file and use the following cat command:

```
cat <backup_config_parameter_file>
```

For example:

```
[root@orcl 13aef284-9d6b-4eb6-8751-2988aexample]# cat opc_b9naujWMAXzi9example_BC.ora
OPC_HOST=https://swiftobjectstorage.us-ashburn-1.oraclecloud.com/v1/dbbackupiad
OPC_WALLET='LOCATION=file:/opt/oracle/dcs/commonstore/objectstore/wallets/13aef284-9d6b-4eb6-8751-2988aexample CREDENTIAL_ALIAS=alias_opc'
OPC_CONTAINER=b9naujWMAXzi9example
```

6. Confirm that the `cwallet.sso` file exists in the directory specified in the `OPC_WALLET` parameter, and confirm that the file has the correct permissions. The file permissions should have the octal value of "600" (`-rw-----`). Use the following command:

```
ls -ltr /opt/oracle/dcs/commonstore/objectstore/wallets/<backup_config_id>
```

For example:

```
[root@orcl 13aef284-9d6b-4eb6-8751-2988aexample]# ls -ltr
/opt/oracle/dcs/commonstore/objectstore/wallets/13aef284-9d6b-4eb6-8751-2988aexample

total 4

-rw----- 1 oracle oinstall 0 Apr 20 06:45 cwallet.sso.lck
-rw----- 1 oracle oinstall 1941 Apr 20 06:45 cwallet.sso
```

### TDE Wallet and Backup Failures

#### INCORRECT TDE WALLET LOCATION SPECIFICATION

For backup operations to work, the `$ORACLE_HOME/network/admin/sqlnet.ora` file must contain the `ENCRYPTION_WALLET_LOCATION` parameter formatted exactly as follows:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=
(DIRECTORY=/opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```



#### Important

In this wallet location entry, `$ORACLE_UNQNAME` is an environment variable and should not be replaced with an actual value.

To check the TDE wallet location specification

Use the `cat` command to check the TDE wallet location specification. For example:

## CHAPTER 11 Database

---

```
[oracle@orcl tde]$ cat $ORACLE_HOME/network/admin/sqlnet.ora

ENCRYPTION_WALLET_LOCATION= (SOURCE= (METHOD=FILE) (METHOD_DATA=
(DIRECTORY=/opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

### INCORRECT STATE OF THE TDE WALLET

Database backups fail if the TDE wallet is not in the proper state. The following scenarios can cause this problem:

#### The ORACLE\_UNQNAME environment variable was not set when the database was started using SQL\*Plus

If the database was started using SQL\*Plus, and the ORACLE\_UNQNAME environment variable was not set, the wallet is not opened correctly.

To fix the problem, start the database using the `srvctl` utility:

```
srvctl start database -d <db_unique_name>
```

#### A pluggable database was added with an incorrectly configured master encryption key

In a multitenant environment, each pluggable database (PDB) has its own master encryption key, which is stored in a single keystore used by all containers. After you create or plug in a new PDB, you must create and activate a master encryption key for it. If you do not do so, the STATUS column in the `v$encryption_wallet` view shows the value `OPEN_NO_MASTER_KEY`.

To check the master encryption key status and create a master key, do the following:

1. Review the the STATUS column in the `v$encryption_wallet` view, as shown in the following example:

```
SQL> alter session set container=pdb2;

Session altered.

SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from v$encryption_wallet;
```

```

WRL_TYPE WRL_PARAMETER STATUS
WALLET_TYPE

FILE /opt/oracle/dcs/commonstore/wallets/tde/example_iadxyz/ OPEN_NO_MASTER_KEY
AUTOLOGIN

```

2. Confirm that the PDB is in READ WRITE open mode and is not restricted, as shown in the following example:

```

SQL> show pdbs

CON_ID CON_NAME OPEN MODE RESTRICTED

2 PDB$SEED READ ONLY NO
3 PDB1 READ WRITE NO
4 PDB2 READ WRITE NO

```

The PDB cannot be open in restricted mode (the `RESTRICTED` column must show `NO`). If the PDB is currently in restricted mode, review the information in the `PDB_PLUG_IN_VIOLATIONS` view and resolve the issue before continuing. For more information on the `PDB_PLUG_IN_VIOLATIONS` view and the restricted status, review the [documentation](#) on pluggable database for your Oracle database version.

3. Run the following `DBCLI` commands to change the status to `OPEN`:

```

$ sudo su -
dbcli list-database
dbcli update-tdekey -i <database_ID> -n <PDB_name> -p

```

The `update-tdekey` command shown will prompt you for the admin password.

- a. Confirm that the status of the wallet has changed from `OPEN_NO_MASTER_KEY` to `OPEN` by querying the `v$encryption_wallet` view as shown in step 1.

### INCORRECT CONFIGURATION RELATED TO THE TDE WALLET

Several configuration parameters related to the TDE wallet can cause backups to fail.

#### To check configuration related to the TDE wallet

- Check that the environment's database unique name parameter (ORACLE\_UNQNAME) is set correctly using the following command:

```
srvctl getenv database -d <db_unique_name>
```

For example:

```
[oracle@orcl tde]$ srvctl getenv database -d orclbkp_iadxyz

orclbkp_iadxyz:

ORACLE_UNQNAME=orclbkp_iadxyz

TZ=UTC
```

- Check your `sqlnet.ora` settings to confirm that the file has an `ENCRYPTION_WALLET_LOCATION` parameter with the correct `DIRECTORY` value. Use the following command:  
`cat $ORACLE_HOME/network/admin/sqlnet.ora`

For example:

```
[oracle@orcl tde]$ cat $ORACLE_HOME/network/admin/sqlnet.ora

ENCRYPTION_WALLET_LOCATION= (SOURCE= (METHOD=FILE) (METHOD_DATA=
(DIRECTORY=/opt/oracle/dcs/commonstore/wallets/tde/$ORACLE_UNQNAME)))
```

- Confirm that the wallet status is **open** and the wallet type is **auto login** by checking the `v$encryption_wallet` view. For example:

```
SQL> select status, wrl_parameter, wallet_type from v$encryption_wallet;

STATUS WRL_PARAMETER WALLET_TYPE


```

```
OPEN /opt/oracle/dcs/commonstore/wallets/tde/example_iadxzy/ AUTOLOGIN
```

For pluggable databases (PDBs), be sure that you switch to the appropriate container before querying `v$encryption_wallet` view. For example:

```
[oracle@paulo ~]$ sqlplus / as sysdba

SQL> alter session set container=pdb1;

Session altered.

SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from v$encryption_wallet;

WRL_TYPE WRL_PARAMETER STATUS WALLET_TYPE

--
FILE /opt/oracle/dcs/commonstore/wallets/tde/tiger_iad3c8/ OPEN AUTOLOGIN
```

### MISSING TDE WALLET FILE

The TDE wallet file (`ewallet.p12`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example:

```
[oracle@orcl tde]$ ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/${ORACLE_UNQNAME}/ewallet.p12

-rwx----- 1 oracle oinstall 5680 Apr 18 13:09 /opt/oracle/dcs/commonstore/wallets/tde/orclbkp_
iadxzy/ewallet.p12
```

The TDE wallet file should have file permissions with the octal value "700" (`-rwx-----`), and the owner of this file should be a part of the `oinstall` operating system group.

### MISSING AUTO LOGIN WALLET FILE

The auto login wallet file (`cwallet.sso`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example:

## CHAPTER 11 Database

---

```
[oracle@orcl tde]$ ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/${ORACLE_UNQNAME}/cwallet.sso
-rwx----- 1 oracle oinstall 5725 Apr 18 13:09 /opt/oracle/dcs/commonstore/wallets/tde/orclbkp_
iadxyz/cwallet.sso
```

The auto login wallet file should have file permissions with the octal value "700" (-rwx-----), and the owner of this file should be a part of the `oinstall` operating system group.

### Other Causes of Backup Failures

#### UNMOUNTED COMMONSTORE MOUNT POINT

The mount point `/opt/oracle/dcs/commonstore` must be mounted, or backups will fail.

#### To check the commonstore mount point

Confirm that the mount point `/opt/oracle/dcs/commonstore` is mounted, as shown in the following example:

```
[root@orcl ~]# srvctl config filesystem -volume commonstore -diskgroup data

Volume device: /dev/asm/commonstore-5

Diskgroup name: data

Volume name: commonstore

Canonical volume device: /dev/asm/commonstore-5

Accelerator volume devices:

Mountpoint path: /opt/oracle/dcs/commonstore

Mount point owner: oracle

Mount users:

Type: ACFS
```

### To confirm that ora.data.commonstore.acfs is online

The state for ora.data.commonstore.acfs must be online, or backups will fail. Confirm as shown in the following example:

```
[root@orcl ~]# crsctl stat resource ora.data.commonstore.acfs -v

NAME=ora.data.commonstore.acfs

TYPE=ora.acfs.type

LAST_SERVER=orcl

STATE=OFFLINE

TARGET=OFFLINE

...

STATE_DETAILS=admin unmounted /opt/oracle/dcs/commonstore

...

[root@orcl ~]# ls -ltr /opt/oracle/dcs/commonstore

total 0
```

If the STATE\_DETAILS value is unmounted, mount the file system as shown in the following example:

```
[root@orcl ~]# srvctl start filesystem -volume commonstore -diskgroup data
```

Confirm that the change was successful as shown in the following example:

```
[root@orcl ~]# crsctl stat resource ora.data.commonstore.acfs -v

NAME=ora.data.commonstore.acfs

TYPE=ora.acfs.type

LAST_SERVER=orcl
```

## CHAPTER 11 Database

---

```
STATE=ONLINE on orcl

TARGET=ONLINE

CARDINALITY_ID=ONLINE

...

STATE_DETAILS=mounted on /opt/oracle/dcs/commonstore
```

List the contents of the `commonstore` directory to confirm that it is mounted, as shown in the following example:

```
[root@orcl ~]# ls -ltr /opt/oracle/dcs/commonstore

total 220

drwx----- 2 root root 65536 Apr 18 10:50 lost+found

drwx----- 3 oracle oinstall 20480 Apr 18 11:02 wallets

drwxr-xr-x 3 root root 20480 Apr 20 06:41 pkgrepos

drwxr-xr-x 4 oracle oinstall 20480 Apr 20 06:41 objectstore
```

### THE DATABASE IS NOT PROPERLY REGISTERED

Database backups fail if the database is not registered with the `dc`s-agent. This scenario can occur if you manually migrate the database to Oracle Cloud Infrastructure and do not run the `dbcli register-database` command.

To check whether the database is properly registered, review the information returned by running the `srvctl config database` command and the `dbcli list-databases` command. If either command does not return a record of the database, contact Oracle Support Services.

For instructions on how to register the database, refer to the following topics:

- [Registering the Database on the DB System](#)
- [dbcli register-database](#)

### Obtaining Further Assistance

If you were unable to resolve the problem using the information in this topic, follow the procedures below to collect relevant database and diagnostic information. After you have collected this information, contact [Oracle Support](#).

### To collect database information for use in problem reports

Use the following commands to collect details about your database. Record the output of each command for reference:

```
dbcli list-databases
```

```
dbcli describe-database -i <database_id>
```

```
dbcli describe-component
```

### To collect diagnostic information regarding failed jobs

1. Log on to the host as the root user and navigate to the `/opt/oracle/dcs/bin/` directory.
2. Run the following two commands to generate information about the failed job:

```
dbcli list-jobs |grep -i <dbname>
```

```
dbcli describe-job -i <job_ID> -j
```

The `<job_ID>` in the second command should be the ID of the latest failed job reported from the first command.

3. Run the diagnostics collector script to create a zip file with the diagnostic information for Oracle Support Services.

```
diagcollector.py
```

This command creates a file named `diagLogs-<timestamp>.zip` in the `/tmp` directory.

### To collect DCS agent log files

To collect DCS agent log files, do the following:

1. Log in as opc user.
2. Run the following command:

```
sudo /opt/oracle/dcs/bin/diagcollector.py
```

3. The system returns a message indicating that agent logs are available in a zip file at a specified directory. For example:

```
[opc@prodpr ~]$ sudo /opt/oracle/dcs/bin/diagcollector.py

Log files collected to :/tmp/dcsdiag/diagLogs-1234567890.zip

Logs are being collected to:

/tmp/dcsdiag/diagLogs-1234567890.zip
```

### To collect TDE configuration details

1. Run the `srvctl getenv database -d <db_unique_name>` command and record the output for reference.
2. Record the output of the view `v$encryption_wallet`. For example:

```
SQL> select status, wrl_parameter, wallet_type from v$encryption_wallet;

STATUS WRL_PARAMETER WALLET_TYPE

OPEN /opt/oracle/dcs/commonstore/wallets/tde/example_iadxyz/ AUTOLOGIN
```

3. Record the output of the output of the `ls -ltr <wrl_parameter>` command. For example:

```
[oracle@patchtst ~]$ ls -ltr /opt/oracle/dcs/commonstore/wallets/tde/example_iadxyz/
```

## CHAPTER 11 Database

---

```
total 28
-rw----- 1 oracle asmadmin 2400 May 2 09:42 ewallet_2018050209420381_defaultTag.p12
-rw----- 1 oracle asmadmin 5680 May 2 09:42 ewallet.p12
-rw----- 1 oracle asmadmin 5723 May 2 09:42 cwallet.sso
```

### To collect the RMAN backup report file

Generate RMAN Backup Report File using the following command:

```
dbcli create-rmanbackupreport -i <db_id> -w detailed -rn <report_name>
```

For example:

```
[root@patchtst ~]# dbcli create-rmanbackupreport -i 57fvwxyz-9dc4-45d3-876b-5f850example -w detailed -rn bkpreport1
```

Locate the report file using the `dbcli describe-rmanbackupreport -in <report_name>` command. The location of the report is given in output. For example:

```
[root@patchtst ~]# dbcli describe-rmanbackupreport -in bkpreport1

Backup Report details

ID: b55vwxyz-c49f-4af3-a956-accddexample

Report Type: detailed

Location: Node patchtst: /opt/oracle/dcs/log/patchtst/rman/bkup/example_iadxzy/rman_list_backup_
detail/2018-05-02/rman_list_backup_detail_2018-05-02_11-46-51.0359.log

Database ID: 57fvwxyz-9dc4-45d3-876b-5f850example

CreatedTime: May 2, 2018 11:46:38 AM UTC
```

### Patching Failures on Bare Metal and Virtual Machine DB Systems

Patching operations can fail for various reasons. Typically, an operation fails because a database node is down, there is insufficient space on the file system, or the database host cannot access the object store.

This topic includes information to help you determine the cause of the failure and fix the problem. The information is organized into several sections, based on the error condition. If you already know the cause, you can skip to the section with the suggested solution. Otherwise, use the procedure in [Determining the Problem](#) to get started.

#### Determining the Problem

In the Console, you can identify a failed patching operation by viewing the patch history of a DB system or an individual database. A patch that was not successfully applied displays a status of **Failed** and includes a brief description of the error that caused the failure. If the error message does not contain enough information to point you to a solution, you can use the database CLI and log files to gather more data. Then, refer to the applicable section in this topic for a solution.

#### To identify the root cause of the patching operation failure

1. Log on to the host as the root user and navigate to the `/opt/oracle/dcs/bin/` directory.
2. Determine the sequence of operations performed on the database.

```
dbcli list-jobs
```

Note the last job ID listed with a status other than **Success**.

3. With the job ID you noted from the previous step, use the following command to check the details of that job:

```
dbcli describe-job -i <job_ID> -j
```

Typically, running this command is enough to reveal the root cause of the failure.

4. If you require more information, review the `/opt/oracle/dcs/log/dcs-agent.log` file.

You can find the job ID in this file by using the timestamp returned by the job report in step 2.



### Note

If the patching failure is on a 2-node RAC database, perform steps 3 and 4 on both nodes.

## Database Service Agent Issues

Your Oracle Cloud Infrastructure Database makes use of an agent framework to allow you to manage your database through the cloud platform.

### RESOLVING PATCHING FAILURES CAUSED BY A STOPPED AGENT

Occasionally you might need to restart the `dcsagent` program if it has the status of **stop/waiting** to resolve a patching failure.

### To restart the database service agent

1. From a command prompt, check the status of the agent:

```
initctl status initdcsagent
```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

```
initctl start initdcsagent
```

3. Check the status of the agent again to confirm that it has the **start/running** status:

```
initctl status initdcsagent
```

### RESOLVING PATCHING FAILURES CAUSED BY AN AGENT THAT NEEDS TO BE UPDATED

Patching can also fail if your agent needs to be updated. The system gives the following error message for this failure:

```
Current DcsAgent version is less than or equal to minimum required version.
```

To resolve this issue, perform the steps in [To have Oracle Support update the Oracle Cloud Infrastructure Database service agent](#).

### To have Oracle Support update the Oracle Cloud Infrastructure Database service agent

1. Confirm that the agent (dcsagent) and DCS Admin program (dcsadmin) are running using the following commands:

```
initctl status initdcsagent
```

```
initctl status initdcsadmin
```

If these programs are not running, use the following commands to restart them:

```
initctl start initdcsagent
```

```
initctl start initdcsadmin
```

2. Follow the instructions in [Obtaining Further Assistance](#) to collect your DCS agent log files.
3. Contact [Oracle Support](#) for assistance with updating the agent.

### Object Store Connectivity Issues

Oracle Cloud Infrastructure DB system and database patches are stored in Oracle Cloud Infrastructure Object Storage. Therefore, successful patching operations require connectivity between the DB system host and the Object Storage location from which the patches are downloaded.

### To ensure your database host can connect to Oracle Cloud Infrastructure Object Storage

1. Use the following command to verify the host can access Oracle Cloud Infrastructure Object Storage:

```
dbcli describe-latestpatch
```

Example output indicating success:

```
[root@<host> ~]# dbcli describe-latestpatch
componentType availableVersion

gi 12.2.0.1.180417
gi 12.1.0.2.180417
gi 18.2.0.0.180417
db 11.2.0.4.180417
db 12.2.0.1.180417
db 12.1.0.2.180417
db 18.2.0.0.180417
oak 12.1.2.11.3
oak 12.2.1.1.0
```

Example output indicating failure:

```
[root@<host> ~]# dbcli describe-latestpatch
DCS-10032:Resource patch metadata is not found.Failed to download patchmetadata from
objectstore
```

2. If you cannot connect to the object store, refer to [Prerequisites](#) for how to configure object store connectivity.

### Host and Oracle Clusterware Issues

One or more of the following conditions on the database host can cause patching operations to fail:

#### DATABASE NODE NOT RUNNING DURING THE PATCHING OPERATION

All nodes of the database must be active and running while a patching operation is in progress, whether you are patching the DB system or the database home. Use the Console to check that the status of each node is AVAILABLE, and start the node, if needed.

### THE FILE SYSTEM IS FULL

Patching operations require a minimum of 15 GB of free space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the available space. If the file system has insufficient space, you can remove old log or trace files to free up space.

### THE ORACLE CLUSTERWARE IS NOT RUNNING

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the DB system for patching operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a patching failure.

### To restart the Oracle Clusterware

1. From command prompt, check the status of Oracle Clusterware:

```
crsctl check crs
```

Example output:

```
[grid@<host> ~]$ crsctl check crs
CRS-4638: Oracle High Availability Services is online
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

For more detailed status information, you can run `crsctl stat res -t`.

2. If Oracle Clusterware is not online, try to restart the program:

```
crsctl start crs
```

3. Check the status of Oracle Clusterware to confirm that it is online:

```
crsctl check crs
```

### THE ORACLE GRID INFRASTRUCTURE (GI) IS NOT UPDATED

This problem occurs when you try to patch a database before you patch the DB system of that database. The error description indicates that the Oracle Grid Infrastructure must be updated

## CHAPTER 11 Database

---

first. To resolve this issue, patch the DB system to latest available version. After you patch the DB system, you can retry the database patching operation.

To get the current and latest-available GI versions for the DB system, use the following command:

```
dbcli describe-component
```

### Database Issues

An improper database state can lead to patching failures.

#### DATABASE NOT RUNNING DURING THE PATCHING OPERATION

The database must be active and running for all of the patching tasks to complete. Otherwise, you must run the datapatch task manually.

#### To check that the database is active and running

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```

The system returns a message including the database instance status. The instance status must be **Open** for the patching operation to succeed.

If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the **Open** status, use the following commands to access the SQL\*Plus command prompt and set the status to **Open**:

```
sqlplus / as sysdba
```

```
alter database open;
```

### To run the datapatch task

Before you run the `datapatch` command, ensure that all pluggable databases (PDBs) are open. To open a PDB, you can use SQL\*Plus to execute `ALTER PLUGGABLE DATABASE <pdb_name> OPEN READ WRITE;` against the PDB.

```
$ORACLE_HOME/OPatch/datapatch
```

The `datapatch` command should be run on each database home.

### Obtaining Further Assistance

If you were unable to resolve the problem using the information in this topic, follow the procedures below to collect relevant database and diagnostic information. After you have collected this information, contact [Oracle Support](#).

### To collect diagnostic information regarding failed jobs

1. Log on to the host as the root user and navigate to the `/opt/oracle/dcs/bin/` directory.
2. Run the following two commands to generate information about the failed job:

```
dbcli list-jobs
```

```
dbcli describe-job -i <job_ID> -j
```

The `<job_ID>` in the second command should be the ID of the latest failed job reported from the first command.

3. Run the diagnostics collector script to create a zip file with the diagnostic information for Oracle Support Services.

```
diagcollector.py
```

This command creates a file named `diagLogs-<timestamp>.zip` in the `/tmp` directory.

### To collect DCS agent log files

To collect DCS agent log files, do the following:

1. Log in as opc user.
2. Run the following command:

```
sudo /opt/oracle/dcs/bin/diagcollector.py
```

3. The system returns a message indicating that agent logs are available in a zip file at a specified directory. For example:

```
[opc@prodpr ~]$ sudo /opt/oracle/dcs/bin/diagcollector.py

Log files collected to :/tmp/dcsdiag/diagLogs-1234567890.zip

Logs are being collected to:

/tmp/dcsdiag/diagLogs-1234567890.zip
```

### To collect Oracle Grid Infrastructure and Database log files

If an Oracle Grid Infrastructure or Oracle Database patch failed, you can find log files for these failures in the following locations:

#### Oracle Grid Infrastructure

```
$GI_HOME/cfgtoollogs/
```

#### Oracle Database

```
$ORACLE_HOME/cfgtoollogs/
```

# CHAPTER 12 DNS and Traffic Management

This chapter explains how to create and manage your DNS zones and guide traffic to your endpoints based on various conditions.

## Overview of the DNS Service

The Oracle Cloud Infrastructure Domain Name System (DNS) service lets you [create and manage your DNS zones](#). You can create zones, add records to zones, and allow Oracle Cloud Infrastructure's edge network to handle your domain's DNS queries.

See [Supported Resource Records](#) for additional information.

## DNS Service Components

The following list describes the components used to build a DNS zone and make it accessible from the internet.

### **DOMAIN**

Domain names identify a specific location or group of locations on the Internet as a whole. A common definition of "domain" is the complete portion of the DNS tree that has been delegated to a user's control. For example, *example.com* or *oracle.com*.

### **ZONE**

A zone is a portion of the DNS namespace. A Start of Authority record (SOA) defines a zone. A zone contains all labels underneath itself in the tree, unless otherwise specified.

### **LABEL**

Labels are prepended to the zone name, separated by a period, to form the name of a subdomain. For example, the "www" section of *www.example.com* or the "docs" and "us-ashburn-1" sections of *docs.us-ashburn-1.oraclecloud.com* are labels. Records are associated with these domains.

### CHILD ZONE

Child zones are independent subdomains with their own Start of Authority and Name Server (NS) records. The parent zone of a child zone must contain NS records that refer DNS queries to the name servers responsible for the child zone. Each subsequent child zone creates another link in the delegation chain.

### RESOURCE RECORDS

A record contains specific domain information for a zone. Each record type contains information called record data (RDATA). For example, the RDATA of an A or AAAA record contains an IP address for a domain name, while MX records contain information about the mail server for a domain. OCI normalizes all RDATA into the most machine readable format. The returned presentation of your RDATA may differ from its initial input. For more information about RDATA, please see [Supported DNS Resource Record Types](#).

### DELEGATION

The name servers where your DNS is hosted and managed.

## Ways to Access the DNS Service

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide.

To access the Console, you must use a supported browser. You can use the Console link at the top of this page to go to the sign-in page. Enter your tenancy, user name, and your password.

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see

[Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### DNS Service Capabilities and Limits

The Oracle Cloud Infrastructure DNS service is limited to 1000 zones per account and 25,000 records per zone. Customers with zone and record size needs exceeding these values are encouraged to contact support at [support.oracle.com](https://support.oracle.com). Zone file uploads are limited to 1 megabyte (MB) in size per zone file. If your zone file is larger than 1 MB, you will need to split the zone file into smaller batches to upload all of the zone information.

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for DNS, see [Details for the DNS Service](#).

### Getting Started with DNS

If you're new to Oracle Cloud Infrastructure DNS, this topic gives guidance on how to proceed.

#### **What is DNS?**

The Domain Name System (DNS) translates human-readable domain names to machine-readable IP addresses. A DNS nameserver stores the DNS records for a zone, and responds

with answers to queries against its database. When you type a domain name into your browser, your operating system queries several DNS nameservers until it finds the authoritative nameserver for that domain. The authoritative nameserver then responds with an IP address or other requested record data. The answer is then relayed back to your browser and the DNS record is resolved to the web page.

### Creating a Zone

In this step, you will create a zone. A zone holds the trusted DNS records that will reside on Oracle Cloud Infrastructure's nameservers.

#### To add a zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click **Create Zone**.
3. In the **Create Zone** dialog box, choose one of the following methods:
  - **Manual** - Enter the following:
    - **Zone Name**: Enter the name of a zone you want to create. Avoid entering confidential information.
    - **Zone Type**: If you want to control the zone contents directly within Oracle Cloud Infrastructure, select **Primary**. If you want Oracle Cloud Infrastructure to pull zone contents from an external server, select **Secondary** and enter your **Zone Master Server IP** address.
  - **Import** - Drag and drop, select, or paste a valid zone file into the Import Zone File window. The zone is imported as a primary zone. For information about formatting a zone file, see [Formatting a Zone File](#).
4. Click **Submit**.

The system creates and publishes the zone, complete with the necessary SOA and NS records. For more information on adding a record to your zone, see [To add a zone record](#).

### Delegating Your Zone

In this step, you will delegate your domain with your registrar. Delegating your domain with your domain's registrar makes your Oracle Cloud Infrastructure hosted zone accessible through the internet.

#### To delegate a zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name for the zone you want to delegate. Zone details and a list of records appear.
3. Use the **Type** sort filter to locate the NS records for your zone.
4. Note the name servers in the RDATA field within each NS record.
5. You can use the noted name servers to change your domain's DNS delegation. Refer to your registrar's documentation for instructions.



#### Note

Once delegation has completed, allow 24 hours for your delegation to propagate across the internet.

#### To add a zone record



#### Tip

There are many record types you can add to your zone, depending on your goals for the zone and its DNS management.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the **Zone Name** in which you want to add a record. Zone details and a list of records appear.



### Tip

You can use the Zone Name sort filter to list to sort zone names alphanumerically in ascending or descending order.

3. Click **Add Record**.
4. In the **Add Record** dialog box, select a record type from the drop-down list, and then enter the information for the record. Avoid entering confidential information. For more information about record types, see [Supported Resource Records](#).
5. (Optional) Click the **Add Another Record** check box to add multiple records in succession.
6. Click **Submit**.
7. Once your records have been added, click **Publish Changes**.
8. In the confirmation dialog box, click **Publish Changes**.

### Common DNS Zone Record Types

For a complete list of records supported by Oracle Cloud Infrastructure DNS, see [Supported Resource Records](#).

#### A

An address record used to point a hostname to an IPv4 address. For more information about A records, see [RFC 1035](#).

### AAAA

An address record used point a hostname at an IPv6 address. For more information about AAAA records, see [RFC 3596](#).

### CNAME

A Canonical Name record identifies the canonical name for a domain. For more information about CNAME records, see [RFC 1035](#).



#### Note

Per [RFC 1912](#), CNAMEs cannot be placed at the apex of the zone.

### MX

A Mail Exchanger record defines the mail server accepting mail for a domain. MX records must point to a hostname. MX records must not point to a CNAME or IP address. For more information about MX records, see [RFC 1035](#).

### TXT

A Text record holds descriptive, human readable text, and can also include non-human readable content for specific uses. It is commonly used for SPF records and DKIM records that require non-human readable text items. For more information about TXT records, see [RFC 1035](#).

### Testing DNS Using BIND's dig Tool

Using the Domain Information Groper (dig) command line tool, you can test against the delegation where your domain is hosted, and you will immediately see whether the change took place without accounting for the cache or TTL (Time to Live) that you have configured.

For more information on using dig to test your DNS, see [Testing DNS Using BIND'S dig Tool](#).

### Managing DNS Service Zones

The Oracle Cloud Infrastructure DNS service enables you to manage zones and view zone reports within the Console.

#### Using the Console

##### MANAGING ZONES AND ZONE RECORDS

#### To add a zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click **Create Zone**.
3. In the **Create Zone** dialog box, choose one of the following methods:
  - **Manual** - Enter the following:
    - **Zone Name**: Enter the name of a zone you want to create. Avoid entering confidential information.
    - **Zone Type**: If you want to control the zone contents directly within OCI, select **Primary**. If you want OCI to pull zone contents from an external server, select **Secondary** and enter your **Zone Master Server IP** address.
  - **Import** - Drag and drop, select, or paste a valid zone file into the Import Zone File window. The zone is imported as a primary zone. For information about formatting a zone file or how to amend a zone file exported from GoDaddy.com, please see [Formatting a Zone File](#).
4. Click **Submit**.

The system creates and publishes the zone, complete with the necessary SOA and NS records. For more information on adding a record to your zone, see [To add a zone record](#).

### To update a secondary zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the secondary **Zone Name** you want to update. Zone details and a list of master server IPs appear.



#### Tip

You can use the Zone Type sort filter to sort zone type alphanumerically in ascending or descending order.

3. Select the checkbox for the Master Server IP you want to update, and then select **Edit** from the **Actions** drop-down menu.
4. Make the needed changes, and then click **Submit**.
5. (Optional) Click **Add Master Server** to add another Master Server IP address.
6. Click **Publish Changes**.
7. In the confirmation dialog box, click **Publish Changes**.



#### Tip

For OCI to transfer data from your zone, your nameservers must be able to accept a transfer request from the following IP addresses:  
208.78.68.65, 204.13.249.65, 2600:2001:0:1::65,  
2600:2003:0:1::65

### To delete a zone



#### Warning

Deletion permanently removes a zone from your DNS service.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Select the checkbox for the zone you want to delete.
3. Click **Delete**. The zone is staged for deletion.
4. Click **Publish Changes** to delete the zone.
5. In the confirmation dialog box, click **Publish Changes**.

### To add a zone record



#### Tip

There are many record types you can add to your zone, depending on your goals for the zone and its DNS management.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the **Zone Name** in which you want to add a record. Zone details and a list of records appear.



### Tip

You can use the Zone Name sort filter to list to sort zone names alphanumerically in ascending or descending order.

3. Click **Add Record**.
4. In the **Add Record** dialog box, select a record type from the drop-down list, and then enter the information for the record. Avoid entering confidential information. For more information about record types, see [Supported Resource Records](#).
5. (Optional) Click the **Add Another Record** check box to add multiple records in succession.
6. Click **Submit**.
7. Once your records have been added, click **Publish Changes**.
8. In the confirmation dialog box, click **Publish Changes**.

### To update a zone record



### Note

#### *Protected Records*

You can change various components of the records within your zones, such as time-to-live (TTL) and



relevant RDATA. However, some records contain information that cannot be changed. A lock symbol indicates a protected record. You can attempt changes to such records through the **Actions** menu, but the system might not permit updates to some fields.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the **Zone Name** in which you want to update a record. Zone details and a list of records appear.



### Tip

You can use the Zone Name sort filter to sort zone names alphanumerically in ascending or descending order.

3. To help find a record, you can use the following filter options:
  - Enter the name of the record's domain in the **Search** field.
  - To find unpublished records, select the **Staged** check box.
  - To find published records, select the **Unstaged** check box.
  - Use the **Domain**, **TTL**, or **Type** sort filter to sort records.
4. Select the checkbox for the record you want to update, and select **Edit** from the **Actions** drop-down menu.
5. In the **Edit Record** dialog box, make the needed changes, and then click **Submit**.
6. Click **Publish Changes**.
7. In the confirmation dialog box, click **Publish Changes**.

### REVERTING CHANGES BEFORE PUBLISHING

You can revert records to their current published state before you publish changes. Once a record has been published, it cannot be reverted. Select the checkbox for the record you want to revert, and then select **Revert** from the **Actions** drop-down menu.

### To delete a zone record

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the **Zone Name** in which you want to delete a record. Zone details and a list of records appear.



#### Tip

You can use the Zone Name sort filter to sort zone names alphanumerically in ascending or descending order.

3. Select the checkbox for the record you want to delete, and then select **Delete** from the **Actions** drop-down menu.
4. Click **Publish Changes**.
5. In the confirmation dialog box, click **Publish Changes**.

### To delegate a zone

To make your Oracle Cloud Infrastructure hosted zone accessible through the internet, you must delegate your domain with your domain's registrar.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name for the zone you want to delegate. Zone details and a list of records appear.
3. Use the **Type** sort filter to locate the NS records for your zone.
4. Note the name servers in the RDATA field within each NS record.
5. You can use the noted name servers to change your domain's DNS delegation. Refer to your registrar's documentation for instructions.

### To move a zone to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. In the **List Scope** section, select a compartment.
3. Find the zone in the list, click the the Actions icon (three dots), and then click **Choose New Compartment**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

For more information, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage your DNS zones:

- [GetZone](#)
- [ListZones](#)
- [CreateZone](#)

- [UpdateZone](#)
- [DeleteZone](#)
- [PatchZoneRecords \(add or delete records\)](#)
- [UpdateZoneRecords](#)

### Setting Up Reverse DNS Zones

Reverse DNS, or rDNS, maps an IP address to a hostname. Reverse DNS serves a number of different purposes from email to network troubleshooting. Some of the benefits include:

- Adding a label for network troubleshooting tools such as traceroute.
- Populating the "Received:" header field in an SMTP email.
- Checking for generic reverse DNS such as 1-2-3-4.example.com to identify spammers.
- Verifying a relationship between the owner of a domain name and the owner of the server (IP address).
- Writing a human readable hostname to the log files for system monitoring tools.
- Determining which hostname is affected when maintenance is performed on an IP address.

Before getting started with setting up reverse DNS within your Oracle Cloud Infrastructure account, contact your IP provider and confirm that they allow delegation of your reverse DNS zone. If they do not allow delegation, typically they can host your pointer record (PTR) for you and no reverse DNS configurations are required within your Oracle Cloud Infrastructure account. If they do allow delegation, confirm the exact syntax of the reverse DNS hostname with them, as some providers use slashes and some use dashes. Additionally, if you are delegating a reverse DNS zone, confirm that this zone matches exactly what you configure in your Oracle Cloud Infrastructure account as this is necessary in order for delegation to work properly.

After you create and publish your reverse DNS zone and PTR records, you can update your reverse DNS zone delegation with your IP provider. Delegation changes are not required with your domain registrar with a reverse DNS zone.

Setting up a reverse DNS zone is different for the two types of IP address blocks. Use the following procedures to set up a reverse DNS zone for your IP address block type.

### Using the Console

#### SETTING UP REVERSE DNS FOR CLASSLESS ADDRESS BLOCK (PARTIAL RANGE OF IP ADDRESSES)

To find your reverse DNS zone name using classless address block

1. Make a note of your network IP address. For example, **192.168.15.224/27**.
2. Remove the netmask portion of the address. This is the number after the slash (/). For example, remove the '27' after your IP address, **192.168.15.224/27**.
3. Reverse the order of the remaining octets. For example, **224.15.168.192**.
4. Append **'in-addr.arpa'** to the end of the IP address. For example, **224.15.168.192.in-addr.arpa**.



#### Note

Some assigning authorities require you to use a slash (/) instead of a dash (-) in the reverse address. Ask which character to use when you contact your assigning authority to delegate the reverse address.

5. Add the netmask back into the address. For example, **224-27.15.168.192.in-addr.arpa**.

In this example, **224-27.15.168.192.in-addr.arpa** is your reverse DNS zone name.

To create your DNS zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click

### DNS Zone Management.

2. Click **Create Zone**.
3. In the Create Zone dialog box, choose one of the following methods:
  - **Manual** - Enter the following:
    - a. **Zone Name:** Enter the name of a zone you want to create. Avoid entering confidential information.
    - b. **Zone Type:** If you want to control the zone contents directly within OCI, select **Primary**. If you want OCI to pull zone contents from an external server, select **Secondary** and enter your **Zone Master Server IP** address.
  - **Import** - Drag and drop, select, or paste a valid zone file into the Import Zone File window. The zone is imported as a primary zone. For information about formatting a zone file or how to amend a zone file exported from GoDaddy.com, see [Formatting a Zone File](#).
4. Click **Submit**.

The system creates and publishes the zone, complete with the necessary SOA and NS records.

### To create a pointer record (PTR) for each host address

As part of the process of setting up a reverse DNS zone, you need to add a PTR record for each host address. This is done specifically for reverse DNS zones to ensure requests are properly routed for resolution.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name in which you want to add the PTR record. Zone details and a list of records appear.



### Tip

You can use the Zone Name sort filter to list to sort zone names alphanumerically in ascending or descending order.

3. Click **Add Record**.
4. In the Add Record dialog box, select the **PTR – Pointer** record type from the drop-down list. Enter the following information:
  - a. **Name:** Optional. Name of the subdomain.
  - b. **TTL:** Click the lock icon to unlock this field. All PTR records in the zone will be updated to reflect the last changes to TTL. This value indicates how long you want to allow external nameservers to cache the information about a given DNS record.
  - c. **TTL Unit:** Select the unit of time used for the TTL value.
  - d. **RData Mode:** Select Basic or Advanced format. If you select Advanced, enter the canonical hostname (for example, *example.com*) that the record is going to point to in the RDATA field.
  - e. **Hostname:** The web address of your zone.

For more information about the PTR record type, see [Supported Resource Records](#).
5. Click **Submit**.
6. Once your record has been added, click **Publish Changes**.
7. In the confirmation dialog box, click **Publish Changes**.

### To add CNAME records for each host at your ISP

If your IP provider does not automatically configure the CNAME record on your behalf, you will need to add a CNAME record for each host at your ISP. This is done specifically for

reverse DNS zones to ensure requests are properly routed for resolution.

1. Make a note of the IP address and your desired CNAME for each host in your new reverse DNS zone.
2. Contact your ISP and request that they append a CNAME record for each host in your Oracle Cloud Infrastructure DNS zone to your account with them.
3. Test the reverse DNS path by running the following command:

```
dig -x <insert any regular forward-formatted IP address from the zone> +trace
```

See [Testing DNS Using BIND'S dig Tool](#) for more information.

The returned information should show that your reverse domain is now being resolved.

### To update your zone delegation

To make your Oracle Cloud Infrastructure hosted zone accessible through the internet, you must delegate your domain with your domain's registrar (usually the website where you purchased your domain, such as GoDaddy.com or Bluehost.com).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name for the zone you want to delegate. Zone details and a list of records appear.
3. Use the **Type** sort filter to locate the NS records for your zone.
4. Note the name servers in the RDATA field within each NS record.

You can use the noted name servers to change your domain's DNS delegation. Refer to your registrar's documentation for instructions.

### SETTING UP REVERSE DNS FOR FULL ADDRESS BLOCK

To find your reverse DNS zone name using full address block

1. Make a note of your network IP address. For example, **192.168.15.0**.
2. Remove the netmask portion of the address (the last number in the set of 4). For example, **192.168.15**.
3. Reverse the order of the remaining three octets. For example, **15.168.192**.
4. Append `'in-addr.arpa'` to the end. For example, **15.168.192.in-addr.arpa**

In this example, **15.168.192.in-addr.arpa** is your reverse DNS zone name.

To create your DNS zone

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click **Create Zone**.
3. In the Create Zone dialog box, choose one of the following methods:
  - **Manual** - Enter the following:
    - a. **Zone Name:** Enter the name of a zone you want to create. Avoid entering confidential information.
    - b. **Zone Type:** If you want to control the zone contents directly within OCI, select **Primary**. If you want OCI to pull zone contents from an external server, select **Secondary** and enter your **Zone Master Server IP** address.
  - **Import** - Drag and drop, select, or paste a valid zone file into the Import Zone File window. The zone is imported as a primary zone. For information about formatting a zone file or how to amend a zone file exported from GoDaddy.com, see [Formatting a Zone File](#).
4. Click **Submit**.

The system creates and publishes the zone, complete with the necessary SOA and NS records.

### To create a pointer record (PTR) for each host address

As part of the process of setting up a reverse DNS zone, you need to add a PTR record for each host address. This is done specifically for reverse DNS zones to ensure requests are properly routed for resolution.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name in which you want to add the PTR record. Zone details and a list of records appear.



#### Tip

You can use the Zone Name sort filter to list to sort zone names alphanumerically in ascending or descending order.

3. Click **Add Record**.
4. In the Add Record dialog box, select the **PTR – Pointer** record type from the drop-down list. Enter the following information:
  - a. **Name:** Optional. Name of the subdomain.
  - b. **TTL:** Click the lock icon to unlock this field. All PTR records in the zone will be updated to reflect the last changes to TTL. This value indicates how long you want to allow external nameservers to cache the information about a given DNS record.
  - c. **TTL Unit:** Select the unit of time used for the TTL value.
  - d. **RData Mode:** Select Basic or Advanced format. If you select Advanced, enter the canonical hostname (for example, *example.com*) that the record is going to point

to in the RDATA field.

- e. **Hostname:** The web address of your zone.

For more information about the PTR record type, see [Supported Resource Records](#).

5. Click **Submit**.
6. Once your record has been added, click **Publish Changes**.
7. In the confirmation dialog box, click **Publish Changes**.

### To update your zone delegation

To make your Oracle Cloud Infrastructure hosted zone accessible through the internet, you must delegate your domain with your domain's registrar.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
2. Click the Zone Name for the zone you want to delegate. Zone details and a list of records appear.
3. Use the **Type** sort filter to locate the NS records for your zone.
4. Note the name servers in the RDATA field within each NS record.

You can use the noted name servers to change your domain's DNS delegation. Refer to your registrar's documentation for instructions.

### Supported Resource Records

The Oracle Cloud Infrastructure DNS service supports many resource record types. The following list provides a brief explanation of the purpose of each supported record type. Avoid entering confidential information when entering record data. The RFC links direct you to further information about the record types and data structure.

### Note About RDATA

OCI normalizes all RDATA into the most machine readable format. The returned presentation of your RDATA may differ from its initial input.

### Example:

The RDATA for the ALIAS, CNAME, DNAME, MX, and NS record types may contain one or more absolute domain names. If the specified RDATA for one of these record types does not end in a dot or period to represent the root, the period will be added.

```
www.example.com --> www.example.com.
```

You can use various DNS libraries to normalize your RDATA before input.

Programming Language	Library
Go	<a href="#">DNS Library in Go</a>
Java	<a href="#">dnsjava</a>
Python	<a href="#">dnspython</a>

### DNS Resource Record Types

#### A

An address record used to point a hostname to an IPv4 address. For more information about A records, see [RFC 1035](#).

#### AAAA

An address record used point a hostname at an IPv6 address. For more information about AAAA records, see [RFC 3596](#).

#### ALIAS

A private pseudo-record that allows CNAME functionality at the apex of a zone. You can view and read ALIAS records in Oracle Cloud Infrastructure DNS, but you cannot create them.

### **CAA**

A Certification Authority Authorization record allows a domain name holder to specify one or more Certification Authorities authorized to issue certificates for that domain. For more information about CAA records, see [RFC 6844](#).

### **CDNSKEY**

A Child DNSKEY moves a CDNSSEC key from a child zone to a parent zone. The information provided in this record must match the CDNSKEY information for your domain at your other DNS provider. This record is automatically created if you enable DNSSEC on a primary zone in Oracle Cloud Infrastructure DNS. For more information about CDNSKEY, see [RFC 7344](#).

### **CDS**

A Child Delegation Signer record is a child copy of a DS record, for transfer to a parent zone. For more information about CDS records, see [RFC 7344](#).

### **CERT**

A Certificate record stores public key certificates and related certificate revocation lists in the DNS. For more information about CERT records, see [RFC 2538](#) and [RFC 4398](#).

### **CNAME**

A Canonical Name record identifies the canonical name for a domain. For more information about CNAME records, see [RFC 1035](#).

### **CSYNC**

A Child-to-Parent Synchronization record syncs records from a child zone to a parent zone. For more information about CNAME records, see [RFC 7477](#).

### **DHCID**

A DHCP identifier record provides a way to store DHCP client identifiers in the DNS to eliminate potential hostname conflicts within a zone. For more information about DHCID, see [RFC 4701](#).

### **DKIM**

A Domain Keys Identified Mail is a special TXT record set up specifically to supply a public key used to authenticate arriving mail for a domain. For more information about DKIM records, see [RFC 6376](#).

### **DNAME**

A Delegation Name record has similar behavior to a CNAME record, but allows you to map an entire subtree beneath a label to another domain. For more information about DNAME records, see [RFC 6672](#).

### **DNSKEY**

A DNS Key record documents public keys used for DNSSEC. The information in this record must match the DNSKEY information for your domain at your other DNS provider. For more information about DNSKEY records, see [RFC 4034](#).

### **DS**

A Delegation Signer record resides at the top-level domain and points to a child zone's DNSKEY record. DS records are created when DNSSEC security authentication is added to the zone. For more information about DS records, see [RFC 4034](#).

### **IPSECKEY**

An IPsec Key record stores public keys for a host, network, or application to connect to IP security (IPsec) systems. For more information on IPSECKEY records, see [RFC 4025](#).

### **KEY**

A Key record stores a public key that is associated with a domain name. Currently only used by SIG and TKEY records. IPSECKEY and DNSKEY have replaced key for use in IPsec and DNSSEC, respectively. For more information about KEY records, see [RFC 4025](#).

### **KX**

A Key Exchanger record identifies a key management agent for the associated domain name with some cryptographic systems (not including DNSSEC). For more information about KX records, see [RFC 2230](#).

### **LOC**

A Location record stores geographic location data of computers, subnets, and networks within the DNS. For more information about LOC records, see [RFC 1876](#).

### **MX**

A Mail Exchanger record defines the mail server accepting mail for a domain. MX records must point to a hostname. MX records must not point to a CNAME or IP address. For more information about MX records, see [RFC 1035](#).

### **NS**

A Nameserver record lists the authoritative nameservers for a zone. Oracle Cloud Infrastructure DNS automatically generates NS records at the apex of each new primary zone. For more information about NS records, see [RFC 1035](#).

### **PTR**

A Pointer record reverse maps an IP address to a hostname. This behavior is the opposite of an A Record, which forward maps a hostname to an IP address. PTR records are commonly found in reverse DNS zones. For more information about PTR records, see [RFC 1035](#).

### **PX**

A resource record used in X.400 mapping protocols. For more information about PX records, see [RFC 822](#) and [RFC 2163](#).

### **SOA**

A Start of Authority record specifies authoritative information about a DNS zone, including:

- The primary nameserver.
- The email of the domain administrator.

- The domain serial number.
- Several timers relating to refreshing the zone.

The Oracle Cloud Infrastructure DNS automatically generates an SOA record when a zone is created. For more information about SOA records, see [RFC 1035](#).

### **SPF**

A Sender Policy Framework record is a special TXT record used to store data designed to detect email spoofing. For more information about SPF records, see [RFC 4408](#).

### **SRV**

A Service Locator record allows administrators to use several servers for a single domain. For more information about SRV records, see [RFC 2782](#).

### **SSHFP**

An SSH Public Key Fingerprint record publishes SSH public host key fingerprints using the DNS. For more information about SSHFP records, see [RFC 6594](#).

### **TLSA**

A Transport Layer Security Authentication record associates a TLS server certificate, or public key, with the domain name where the record is found. This relationship is called a TLSA certificate association. For more information about TLSA records, see [RFC 6698](#).

### **TXT**

A Text record holds descriptive, human readable text, and can also include non-human readable content for specific uses. It is commonly used for SPF records and DKIM records that require non-human readable text items. For more information about TXT records, see [RFC 1035](#).

## Formatting a Zone File

A zone file is a text file that describes a DNS zone. The BIND file format is the industry preferred zone file format and has been widely adopted by DNS server software. The format is defined in [RFC 1035](#).

### Example of a Zone File

This is an example of a zone file downloaded from Oracle Cloud Infrastructure DNS.

```
$ORIGIN example.com.
@ 3600 SOA ns1.p30.oraclecloud.net. (
zone-admin.dyndns.com. ; address of responsible party
2016072701 ; serial number
3600 ; refresh period
600 ; retry period
604800 ; expire time
1800) ; minimum ttl
86400 NS ns1.p68.dns.oraclecloud.net.
86400 NS ns2.p68.dns.oraclecloud.net.
86400 NS ns3.p68.dns.oraclecloud.net.
86400 NS ns4.p68.dns.oraclecloud.net.
3600 MX 10 mail.example.com.
3600 MX 20 vpn.example.com.
3600 MX 30 mail.example.com.
60 A 204.13.248.106
3600 TXT "v=spf1 includespf.oraclecloud.net ~all"
mail 14400 A 204.13.248.106
vpn 60 A 216.146.45.240
webapp 60 A 216.146.46.10
webapp 60 A 216.146.46.11
www 43200 CNAME example.com.
```



#### Note

##### *Record Classes*

In the example zone file above, no record classes are displayed. Oracle Cloud Infrastructure DNS only works with Internet (IN) class records but omits the class information in zone files for efficiency purposes.

### Anatomy of a Zone File

`$ORIGIN` indicates a DNS node tree and will typically start a DNS zone file. Any host labels below the origin will append the origin hostname to assemble a fully qualified hostname. Any host label within a record that uses a fully qualified domain terminating with an ending period will not append the origin hostname.

**Example:** With `$ORIGIN example.com.`, any record where the host label field is not followed by a period, `example.com.` will be appended to them.

The "@" symbol is a special label that indicates the `$ORIGIN` should replace the "@" symbol. This is typically used for the apex of a zone.

**SOA Record** – The `$ORIGIN` is followed by the zone's Start Of Authority (SOA) record. An SOA record is required for each zone. It contains the name of the zone, the e-mail address of the party responsible for administering the domain's zone file, the current serial number of the zone, the primary nameserver of the zone, and various timing elements (measured in seconds).

#### SOA RECORD FORMAT

```
@ IN SOA {primary-name-server} {hostmaster-email} (
 {serial-number}
 {time-to-refresh}
 {time-to-retry}
 {time-to-expire}
 {minimum-TTL})
```

- **Primary Name Server** – The nameserver that contains the original zone file and not an AXFR transferred copy.
- **Hostmaster Email** – Address of the party responsible for the zone. A period "." is used in place of an "@" symbol. For email addresses that contain periods, replace the periods with a slash "/".
- **Serial Number** – Version number of the zone. The serial number will increase with each subsequent update to your zone.
- **Time To Refresh** – How long a nameserver should wait prior to checking for a serial number increase within the primary zone file, in seconds. An increased serial number

detected by a secondary DNS nameserver means a transfer is needed to sync your records. Only applies to zones using secondary DNS.

- **Time To Retry** – How long a nameserver should wait prior to retrying to update a zone after a failed attempt, in seconds. Only applies to zones using secondary DNS.
- **Time To Expire** – How long a nameserver should wait prior to considering data from a secondary zone invalid and stop answering queries for that zone, in seconds. Only applies to zones using secondary DNS.
- **Minimum TTL** – Minimum Time To Live (TTL). How long a nameserver or resolver should cache a negative response, in seconds.

### Anatomy of a Record Within a Zone File

A zone file is a collection of resource records with each record entry described in the following sequence:

<b>Format:</b>	Host Label	TTL	Record Class	Record Type	Record Data
<b>Example:</b>	example.com.	60	IN	A	104.255.228.125

- **Host Label** – A host label helps to define the hostname of a record and whether the `$ORIGIN` hostname will be appended to the label. Fully qualified hostnames terminated by a period will not append the origin.
- **TTL** – The Time To Live (TTL) is the amount of time that a DNS record will be cached by an outside DNS server or resolver, in seconds.
- **Record Class** – There are three classes of DNS records: IN (Internet), CH (Chaosnet), and HS (Hesiod). Oracle Cloud Infrastructure DNS only uses the IN class of records.
- **Record Type** – The type of a record, such as CNAME, AAAA, or TXT.
- **Record Data** – The data within a DNS answer, such as an IP address, hostname, or other information. Different record types will contain different types of record data.

### Amending Zone Files Exported from GoDaddy.com for Import

GoDaddy.com exports zone files in a proprietary format. To get the Oracle Cloud Infrastructure DNS service to correctly import a zone file exported from GoDaddy.com, you must directly alter the file. Follow these instructions to update the zone file.

1. Export your zone file from GoDaddy.com. Reference GoDaddy.com's documentation to see how this is done.
2. Open the file in your preferred text editor.
3. Prepend a new line to the file before the SOA record with the following information, including the trailing period: `$ORIGIN [yourdomain].`
4. Once the file has been amended, save the changes to the file and use the zone import function to import the file into your DNS configuration. For more information about zone import, see [Managing DNS Zones](#).



#### Note

If your zone file includes dynamic A records, such as `@ 600 IN A GoCentral Published Site`, you will need to amend these records with the correct IP addresses of your website. Please contact GoDaddy.com for information about how to obtain this information.

**Example:** `@ 600 IN A 192.0.2.255`

#### Example:

This is an example of a zone file exported from GoDaddy.com. The code in bold is the code that needs to be removed from the file for it to be eligible for import into Oracle Cloud Infrastructure DNS.



### Tip

Placing a semi-colon at the beginning of a line is valid comment syntax for a zone file, per [RFC 1035](#), but for ease of use and formatting it is recommended to remove the large section of comments from the beginning of the zone file provided by GoDaddy.com, as shown below.

```
Domain: example.com
; Exported (y-m-d hh:mm:ss): 2019-01-10 13:05:04
;
; This file is intended for use for informational and archival
; purposes ONLY and MUST be edited before use on a production
; DNS server.
;
; In particular, you must update the SOA record with the correct
; authoritative name server and contact e-mail address information,
; and add the correct NS records for the name servers which will
; be authoritative for this domain.
;
; For further information, please consult the BIND documentation
; located on the following website:
;
; http://www.isc.org/
;
; And RFC 1035:
;
; http://www.ietf.org/rfc/rfc1035.txt
;
; Please note that we do NOT offer technical support for any use
; of this zone data, the BIND name server, or any other third-
; party DNS software.
;
; Use at your own risk.
; SOA Record
example.com. 3600 IN SOA ns41.domaincontrol.com. dns.net. (
 2018122702
 28800
```

## CHAPTER 12 DNS and Traffic Management

```

 7200
 604800
 3600
)
; A Records
@ 600 IN A 192.0.2.249
blog 10800 IN A 192.0.2.255
dev 1800 IN A 192.0.2.254
dev01 1800 IN A 192.0.2.253
dev02 1800 IN A 192.0.2.252
dev03 1800 IN A 192.0.2.251
dev04 1800 IN A 192.0.2.250
; CNAME Records
abc123b432dc7785b7ef31f04f25c3e71 1800 IN CNAME verify.bing.com.
akamai 600 IN CNAME www.example.com.edgekey.net.
email 3600 IN CNAME email.secureserver.net.
; MX Records
@ 604800 IN MX 10 amlxe.1.google.com.
@ 604800 IN MX 10 aplxe.1.google.com.
; TXT Records
@ 3600 IN TXT "google-site-verification=3J82-80dbMyCo5Q5C1G11JszeOnZPGCSY1HcPcXg"
@ 3600 IN TXT "google-site-verification=eS_QPYLE_W4nduSrlN-cddxG7ZqOnB743xsbX918"
```

Below is an example of an amended zone file ready to import into Oracle Cloud Infrastructure DNS. The code in bold needs to be prepended to your zone file before import.

```
$ORIGIN example.com.
example.com. 3600 IN SOA ns41.domaincontrol.com. dns.net. (
 2018122702
 28800
 7200
 604800
 3600
)

; A Records
@ 600 IN A 192.0.2.249
blog 10800 IN A 192.0.2.255
dev 1800 IN A 192.0.2.254
dev01 1800 IN A 192.0.2.253
dev02 1800 IN A 192.0.2.252
```

## CHAPTER 12 DNS and Traffic Management

---

```
dev03 1800 IN A 192.0.2.251
dev04 1800 IN A 192.0.2.250abc123b432dc7785b7ef31f04f25c3e71 1800 IN CNAME
verify.bing.com.
; CNAME Records
akamai 600 IN CNAME www.example.edgekey.net.
email 3600 IN CNAME email.secureserver.net.
; MX Records
@ 604800 IN MX 10 amlxe.l.google.com.
@ 604800 IN MX 10 aplxe.l.google.com.
; TXT Records
@ 3600 IN TXT "google-site-verification=3J82-80dbMyCo5Q5C1GM8os1VYVEOnZPGCSY1HcPcXg"
@ 3600 IN TXT "google-site-verification=eS_QPYLE_W4nduSrlN-cddxG7ZqOnB7k7uIG7qrsyu8"
```

### Setting Up HTTP Redirect

The HTTP Redirect service allows you to redirect HTTP traffic to another URL. You can use HTTP Redirect to:

- Redirect all HTTP traffic for an entire zone to another zone. For example, if a company owns example.net and example.com, HTTP Redirect lets the company redirect all HTTP traffic for example.net to example.com. This is a one-to-one mapping; wildcards are not supported.
- Redirect a specific subdomain to an HTTP URL. For example, test.example.com can be redirected to `http://example.net/test/test.php`.
- Redirect a subdomain to a URL with a port number. For example, camera.example.com can be redirected to `http://office.example.com:8080` so a user can view their camera system without typing in the port number each time.
- Permanently redirect a domain name that has been deprecated by displaying a 301 response code. Permanently redirecting a domain name informs search engines and browsers what to do with the information.

### Using the Console

#### To create an HTTP redirect

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **HTTP Redirects**.
2. Select a zone.
3. Click **Create HTTP Redirect**.
4. In the Create Redirect dialog box, enter the following information:
  - **Name:** (Optional) Enter the name of the redirect zone you want to create. Avoid entering confidential information.
  - **Select a Zone:** (Optional) Select a zone from a list of configured zones. If the **Create DNS Record** check box is selected, the zone will be used to build an alias record for the redirect.
  - **Domain:** Enter the domain name from which traffic is redirected.
  - **Target** - Enter the following information for the endpoint where the traffic will be redirected:
    - **Protocol:** The network protocol used to interact with the target.
    - **Host:** The hostname of the target.
    - **Port:** (Optional) The port used to connect to the endpoint. The default is 80 for HTTP and 443 for HTTPS.
    - **Path:** (Optional) The specific path on the target for the redirect. A value of {path} will copy the path from the incoming request.
    - **Query:** (Optional) The query component of the target URL (for example, "?redirected" in "https://target.example.com/path/to/resource?redirected"). Use of the "\" character is not permitted except to escape a following "\", "{", or "}". An empty value results in a redirection target URL with no query component. A static value must begin with a leading "?", optionally followed by other query characters. A request-copying value must exactly match "{query}",

and will be replaced with the query component of the request URL (including a leading "?" if the request URL includes a query component).

- **Response Code:** The response code that is returned with the redirect. If your website was permanently moved to the redirection URL and you want it to be indexed by search engines, select **301 - Moved Permanently**. If you want to indicate that the URL has been temporarily changed to a different address, select **302 - Found**.
  - **Create DNS Record:** Select this check box to create an associated ALIAS record for the redirect in the specified zone. If a record for the zone specified already exists, the DNS record will not be created.
  - **ALIAS TTL in Seconds** - The Time to Live for the ALIAS record before a new ALIAS record is retrieved. The default value is 300.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create**. The redirect zone is added to the redirects list.

### To edit an HTTP redirect

1. Open the navigation menu. **Under Core Infrastructure**, go to **Networking** and click **HTTP Redirects**.
2. Click the name of a redirect zone.
3. Click **Edit**.
4. In the Edit Redirect dialog box, make the needed changes and then click **Save Changes**.

### To delete an HTTP redirect

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **HTTP Redirects**.
2. Select a redirect zone.
3. Click **Delete**.
4. In the Delete Resource dialog box, click **Delete HTTP Redirect**. Any attached records will need to be managed in [DNS Zone Management](#).

### To move an HTTP redirect to another compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **HTTP Redirects**.
2. Click the name of a redirect zone.
3. Click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### Using the API

Use the following operations to manage your HTTP redirect zones:

- [GetHttpRedirect](#)
- [ListHttpRedirects](#)
- [CreateHttpRedirect](#)
- [UpdateHttpRedirect](#)
- [DeleteHttpRedirect](#)
- [ChangeHttpRedirectCompartment](#)

### Testing DNS Using BIND'S dig Tool

Using the Domain Information Groper (dig) command line tool, you can test against the delegation where your domain is hosted, and you will immediately see whether the change took place without accounting for the cache or TTL (Time to Live) that you have configured.



#### Note

Windows users can download the tool from BIND's [website](#). Use Terminal to access dig on Linux and Macintosh systems.

### Using dig

Before using BIND's dig tool, you must access or install dig on your system. Once you have access to dig, you can use dig to test your DNS.

#### To access dig (Mac)

1. From your Applications folder, open the Utilities folder, and then select **Terminal**.
2. When Terminal is open, type a [dig command](#) using a hostname you want to look up.

#### To Install dig (Windows)

1. Go to [BIND's website](#) and download the most current, stable version of BIND.



### Note

BIND supports both 32 and 64 bit Windows systems. Confirm which version of Windows you are using and download the correct version of BIND. View Microsoft's [documentation](#) to determine which version of Windows you are using.

2. Extract the downloaded file and install BIND in the following directory: **C:\Program Files\ISC BIND 9**. Select the **Tools Only** check box.
3. Once BIND is installed, on the Windows menu open the Control Panel, and then open your System properties.
4. On the **Advanced** tab, click **Environment Variables**.
5. Under **System Variables**, select **Path**, and then click **Edit**.
6. At the end of the path in the Edit System Variable window, add **C:\Program Files\ISC BIND 9\bin**, and then click **OK**.
7. In the Edit Variables window, click **OK**. In the System properties window, click **OK**.

### TO OPEN THE COMMAND PROMPT

For Windows versions 8 -10:

1. Click the Windows menu icon.
2. In the **Search** field, type **CMD**.
3. Click **Command Prompt**.

For Windows version 7:

1. On the **Start** menu click **Run**.
2. Enter **CMD**, and then click **OK**.

### To use dig to test your DNS

1. Open Terminal (Mac and Linux) or Command Prompt (Windows).
2. Type `dig <any hostname>`, and then press **Enter**.

The following information is returned:

```

$ dig oracle.com

;<<>> DiG 9.10.6 <<>> oracle.com
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 45619
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 6, ADDITIONAL: 13

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;oracle.com. IN A

;; ANSWER SECTION:
oracle.com. 300 IN A 137.254.16.101

;; AUTHORITY SECTION:
oracle.com. 10800 IN NS dnsmaster5.oracle.com.
oracle.com. 10800 IN NS dnsmaster6.oracle.com.
oracle.com. 10800 IN NS dnsmaster3.oracle.com.
oracle.com. 10800 IN NS dnsmaster1.oracle.com.
oracle.com. 10800 IN NS dnsmaster4.oracle.com.
oracle.com. 10800 IN NS dnsmaster2.oracle.com.

;; ADDITIONAL SECTION:
dnsmaster5.oracle.com. 10800 IN A 192.135.82.70
dnsmaster4.oracle.com. 10800 IN A 192.135.82.52
dnsmaster3.oracle.com. 10800 IN A 192.135.82.36
dnsmaster6.oracle.com. 10800 IN A 192.135.82.84
dnsmaster2.oracle.com. 10800 IN A 10.221.8.13
dnsmaster1.oracle.com. 10800 IN A 192.135.82.4
dnsmaster5.oracle.com. 10800 IN AAAA 2606:b400:1400:4240:4fff:ffff:ffff:9f99
dnsmaster4.oracle.com. 10800 IN AAAA 2606:b400:1400:8140:4fff:ffff:ffff:9f99
dnsmaster3.oracle.com. 10800 IN AAAA 2606:b400:1400:8040:4fff:ffff:ffff:9f99
dnsmaster6.oracle.com. 10800 IN AAAA 2606:b400:1400:4144::41
dnsmaster2.oracle.com. 10800 IN AAAA 2606:b400:1400:280:feed::3
dnsmaster1.oracle.com. 10800 IN AAAA 2606:b400:1400:180:4fff:ffff:ffff:9f99

;; Query time: 163 msec
;; SERVER: 192.135.82.52#53(192.135.82.52)
;; WHEN: Tue Feb 26 14:02:05 EST 2019
;; MSG SIZE rcvd: 469

```

- **Question section:** The query made to the DNS. In this example, we asked for the first available A record for the hostname, oracle.com.

## CHAPTER 12 DNS and Traffic Management

---

- **Answer section:** The first available answer for the query made to the DNS. In this example, we received the A record for the IP address 137.254.16.101.
- **Authority section:** The authoritative nameservers from which the answer to the query was received. These nameservers house the zones for a domain.
- **Additional section:** Additional information the resolver may need but not the answer to the query.

### DIG COMMANDS

Command	Description	Example
<code>dig [hostname]</code>	Returns any A record found within the queried hostname's zone.	<code>dig oracle.com</code>
<code>dig [hostname] [record type]</code>	Returns the records of that type found within the queried hostname's zone.	<code>dig oracle.com MX</code>
<code>dig [hostname] +short</code>	Provides a brief answer, usually just an IP address.	<code>dig oracle.com +short</code>
<code>dig @ [nameserver address] [hostname]</code>	Queries the nameserver directly instead of your ISP's resolver.	<code>dig @dnsmaster6.oracle.com</code>
<code>dig [hostname] +trace</code>	Adding <code>+trace</code> instructs <code>dig</code> to resolve the query from the root nameserver downwards and to report the results from each query step.	<code>dig dyn.com +trace</code>

Command	Description	Example
<code>dig -X [IP address]</code>	Reverse lookup for IP addresses.	<code>dig -X 137.254.16.101</code>
<code>dig [hostname] any</code>	Returns all records for a hostname.	<code>dig oracle.com any</code>

## Overview of the Traffic Management Steering Policies Service

The Oracle Cloud Infrastructure Traffic Management Steering Policies service is a critical component of DNS. Traffic Management Steering Policies enables you to configure policies to serve intelligent responses to DNS queries, meaning different answers (endpoints) may be served for the query depending on the logic the customer defines in the policy. Traffic Management Steering Policies can account for health of answers to provide failover capabilities, provide the ability to load balance traffic across multiple resources, and account for the location where the query was initiated to provide a simple, flexible and powerful mechanism to efficiently steer DNS traffic.

### Traffic Management Steering Policies Service Components

The following list describes the components used to build a traffic management steering policy.

#### **STEERING POLICIES**

A framework to define the traffic management behavior for your zones. Steering policies contain rules that help to intelligently serve DNS answers.

### **ATTACHMENTS**

Allows you to link a steering policy to your zones. An attachment of a steering policy to a zone occludes all records at its domain that are of a covered record type, constructing DNS responses from its steering policy rather than from those domain's records. A domain can have at most one attachment covering any given record type.

### **RULES**

The guidelines steering policies use to filter answers based on the properties of a DNS request, such as the requests geo-location or the health of your endpoints.

### **ANSWERS**

Answers contain the DNS record data and metadata to be processed in a steering policy.

## Ways to Access the Traffic Management Steering Policies Service

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide.

To access the Console, you must use a supported browser. You can use the Console link at the top of this page to go to the sign-in page. Enter your tenancy, user name, and your password.

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Traffic Management Steering Policies Service Capabilities and Limits

The Oracle Cloud Infrastructure Traffic Management Steering Policies service is limited to 100 policies and 1,000 attachments per tenant. See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for Traffic Management Steering Policies, see Details for the Traffic Management Steering Policies Service.

### Traffic Management Steering Policies API Guide

Traffic Management Steering Policies allows you to build and configure traffic management policies using the Oracle Cloud Infrastructure [DNS REST API](#). Use the following guide to learn how policies are constructed using the REST API.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Traffic Management Steering Policy Components

The following list describes the components used to build a Traffic Management Steering Policy.

#### **STEERING POLICIES**

An overall framework to define the traffic management behavior for your zones. Steering policies contain rules that help to intelligently serve DNS answers.

#### **ATTACHMENTS**

Allows you to link a steering policy to your zones. An attachment of a steering policy to a zone occludes all records at its domain that are of a covered record type, constructing DNS responses from its steering policy rather than from those domain's records. A domain can have at most one attachment covering any given record type.

#### **RULES**

The guidelines steering policies use to filter answers based on the properties of a DNS request, such as the requests geo-location or the health of your endpoints.

#### **ANSWERS**

Answers contain the DNS record data and metadata to be processed in a steering policy.

### TEMPLATES

Templates are predefined rule sequences that create a policy type and its intended behavior. Example: The `FAILOVER` template determines answers by checking DNS query against a `FILTER` rule first, then the following rules in succession: `HEALTH`, `PRIORITY`, and `LIMIT`. This gives the domain dynamic failover capability. Policies that define the `template` field with any policy other than `CUSTOM`, must follow the rule sequence outlined for that policy type, otherwise, a `400` status code error will be returned upon policy creation.

### CASES

A rule may optionally include a sequence of cases defining alternate configurations for how it should behave during processing for any given DNS query. When a rule has no sequence of cases, it is always evaluated with the same configuration during processing. When a rule has an empty sequence of cases, it is always ignored during processing. When a rule has a non-empty sequence of cases, its behavior during processing is configured by the first matching case in the sequence. A rule case with no `caseCondition` always matches. A rule case with a `caseCondition` matches only when that expression evaluates to true for the given query.

### Create Steering Policies Using Templates

The following section explains the rule configuration for each type of steering policy template followed by an example POST request ([CreateSteeringPolicy](#)) displaying how to configure each template.

#### FAILOVER

Failover policies allow you to prioritize the order in which you want answers served in a policy (for example, Primary and Secondary). Oracle Cloud Infrastructure Health Checks are leveraged to determine the health of answers in the policy. If the Primary Answer is determined to be unhealthy, DNS traffic will automatically be steered to the Secondary Answer. Each of the following rules must be defined in the order specified below in the `rules` field of your request body when using a `FAILOVER` template:

## CHAPTER 12 DNS and Traffic Management

Order	Rule	Restrictions	Comments
1	FILTER	<ul style="list-style-type: none"><li>No cases are allowed.</li><li>Answer data must be defined in <code>defaultAnswerData</code> using the following JSON:<pre>{   "answerCondition":   "answer.isDisabled != true",   "shouldKeep": true }</pre></li></ul>	
2	HEALTH	<ul style="list-style-type: none"><li>No cases are allowed.</li></ul>	Only included if <code>healthCheckMonitorId</code> is defined for the policy.

Example of a POST `/steeringPolicies` policy using the `FAILOVER` template:

```
{
 "compartmentId": "ocid1...",
 "displayName": "failover between endpoints",
 "ttl": 30,
 "healthCheckMonitorId": "ocid1...",
 "template": "FAILOVER",
 "answers": [
 {
 "name": "server-primary",
 "rtype": "A",
 "rdata": "192.0.2.1",
 "pool": "primary"
 },
 {
 "name": "server-secondary",
 "rtype": "A",
 "rdata": "192.1.2.1",
 "pool": "secondary"
 }
],
]
}
```

```
"rules": [
 {
 "ruleType": "FILTER",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.isDisabled != true",
 "shouldKeep": true
 }
]
 },
 {
 "ruleType": "HEALTH"
 },
 {
 "ruleType": "PRIORITY",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.pool == 'primary'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'secondary'",
 "value": 99
 }
]
 },
 {
 "ruleType": "LIMIT",
 "defaultCount": 1
 }
]
```

### **LOAD\_BALANCE**

Load Balancer policies allow distribution of traffic across multiple endpoints. Endpoints can be assigned equal weights to distribute traffic evenly across the endpoints or custom weights may be assigned for ratio load balancing. Oracle Cloud Infrastructure Health Checks are leveraged to determine the health of the endpoint. DNS traffic will be automatically distributed to the other endpoints, if an endpoint is determined to be

## CHAPTER 12 DNS and Traffic Management

unhealthy. Each of the following rules must be defined in the order specified below in the `rules` field of your request body when using a `LOAD_BALANCE` template:

Order	Rule	Restrictions	Comments
1	FILTER	<ul style="list-style-type: none"><li>No cases are allowed.</li><li>Answer data must be defined in <code>defaultAnswerData</code> using the following JSON:<pre>{  "answerCondition":  "answer.isDisabled != true",  "shouldKeep": true}</pre></li></ul>	
2	HEALTH	<ul style="list-style-type: none"><li>No cases are allowed.</li></ul>	Only included if <code>healthCheckMonitorId</code> is defined for the policy.

Example of a `POST /steeringPolicies` request body using the `LOAD_BALANCE` template:

```
{ "compartmentId": "ocid1...", "displayName": "Weighted load balance for a set of answers with health checks", "ttl": 30, "healthCheckMonitorId": "ocid1...", "template": "LOAD_BALANCE", "answers": [{ "name": "server1", "rtype": "A", "rdata": "192.0.2.1" }, { "name": "server2", "rtype": "A", "rdata": "198.51.100.1" }]}
```

```
],
"rules": [
 {
 "ruleType": "FILTER",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.isDisabled != true",
 "shouldKeep": true
 }
]
 },
 {
 "ruleType": "HEALTH"
 },
 {
 "ruleType": "WEIGHTED",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.name == 'server1'",
 "value": 99
 },
 {
 "answerCondition": "answer.name == 'server2'",
 "value": 1
 }
]
 },
 {
 "ruleType": "LIMIT",
 "defaultCount": 1
 }
]
}
```

### **ROUTE\_BY\_GEO**

Geolocation-based steering policies distribute DNS traffic to different endpoints based on the location of the end user. Customers can define geographic regions composed of originating continent, countries or states/provinces (North America) and define a separate endpoint or set of endpoints for each region. Each of the following rules must be defined in

## CHAPTER 12 DNS and Traffic Management

the order specified below in the `rules` field of your request body when using a `ROUTE_BY_GEO` template:

Order	Rule	Restrictions	Comments
1	FILTER	<ul style="list-style-type: none"><li>No cases are allowed.</li><li>Answer data must be defined in <code>defaultAnswerData</code> using the following JSON:<pre>{  "answerCondition":  "answer.isDisabled != true",  "shouldKeep": true}</pre></li></ul>	
2	HEALTH	<ul style="list-style-type: none"><li>No cases are allowed.</li></ul>	Only included if <code>healthCheckMonitorId</code> is defined for the policy.

Example of a `POST /steeringPolicies` request body using the `ROUTE_BY_GEO` template:

```
{ "compartmentId": "ocid1...", "displayName": "Geolocations mapped to answer pools", "ttl": 30, "healthCheckMonitorId": "ocid1...", "template": "ROUTE_BY_GEO", "answers": [{ "name": "US Server 1", "rtype": "A", "rdata": "10.10.10.10", "pool": "US" }, { "name": "US Server 2", "rtype": "A", "rdata": "10.10.10.11", }
```

## CHAPTER 12 DNS and Traffic Management

---

```
 "pool": "US"
 },
 {
 "name": "EU Server 1",
 "rtype": "A",
 "rdata": "10.10.1.1",
 "pool": "EU"
 },
 {
 "name": "EU Server 2",
 "rtype": "A",
 "rdata": "10.10.1.2",
 "pool": "EU"
 },
 {
 "name": "rest of world 1",
 "rtype": "A",
 "rdata": "203.0.113.1",
 "pool": "Global"
 },
 {
 "name": "rest of world 2",
 "rtype": "A",
 "rdata": "203.0.113.2",
 "pool": "Global"
 }
],
"rules": [
 {
 "ruleType": "FILTER",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.isDisabled != true",
 "shouldKeep": true
 }
]
 },
 {
 "ruleType": "HEALTH"
 }
],
```

## CHAPTER 12 DNS and Traffic Management

---

```
{
 "ruleType": "PRIORITY",
 "cases": [
 {
 "caseCondition": "query.client.geoKey in (geoKey '6255149')",
 "answerData": [
 {
 "answerCondition": "answer.pool == 'US'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'EU'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Global'",
 "value": 3
 }
]
 },
 {
 "caseCondition": "query.client.geokey in (geokey '6255148')",
 "answerdata": [
 {
 "answerCondition": "answer.pool == 'EU'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'US'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Global'",
 "value": 3
 }
]
 },
 {
 "answerData": [
 {
```

## CHAPTER 12 DNS and Traffic Management

---

```
 "answerCondition": "answer.pool == 'Global'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'US'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'EU'",
 "value": 3
 }
]
}
]
},
{
 "ruleType": "LIMIT",
 "defaultCount": 1
}
]
```

For a list of geographic keys to use in the `geokey` field, see [Traffic Management Steering Policy geokeys](#).

### **ROUTE\_BY\_ASN**

ASN-based steering policies enable you to steer DNS traffic based on Autonomous System Numbers (ASN). DNS queries originating from a specific ASN or set of ASNs can be steered to a specified endpoint. Each of the following rules must be defined in the order specified below in the `rules` field of your request body when using a `ROUTE_BY_ASN` template:

## CHAPTER 12 DNS and Traffic Management

Order	Rule	Restrictions	Comments
1	FILTER	<ul style="list-style-type: none"><li>No cases are allowed.</li><li>Answer data must be defined in <code>defaultAnswerData</code> using the following JSON:<pre>{  "answerCondition":  "answer.isDisabled != true",  "shouldKeep": true}</pre></li></ul>	
2	HEALTH	<ul style="list-style-type: none"><li>No cases are allowed.</li></ul>	Only included if <code>healthCheckMonitorId</code> is defined for the policy.

Example of a POST `/steeringPolicies` request body using the `ROUTE_BY_ASN` template:

```
{ "compartmentId": "ocidl...", "displayName": "ASNs mapped to pools", "ttl": 30, "template": "ROUTE_BY_ASN", "answers": [{ "name": "MIT Server", "rtype": "A", "rdata": "10.10.10.10", "pool": "MIT" }, { "name": "Google Fiber Server", "rtype": "A", "rdata": "10.10.1.1", "pool": "Google Fiber" }, { "name": "Other",
```

## CHAPTER 12 DNS and Traffic Management

---

```
"rtype": "A",
"rdata": "203.0.113.1",
"pool": "Other"
}
],
"rules": [
{
"ruleType": "FILTER",
"defaultAnswerData": [
{
"answerCondition": "answer.isDisabled != true",
"shouldKeep": true
}
]
},
{
"ruleType": "PRIORITY",
"cases": [
{
"caseCondition": "query.client.asn == 3",
"answerData": [
{
"answerCondition": "answer.pool == 'MIT'",
"value": 1
},
{
"answerCondition": "answer.pool == 'Google Fiber'",
"value": 2
},
{
"answerCondition": "answer.pool == 'Other'",
"value": 3
}
]
}
],
},
{
"caseCondition": "query.client.asn == 16591",
"answerdata": [
{
"answerCondition": "answer.pool == 'Google Fiber'",
```

```
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'MIT'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Other'",
 "value": 3
 }
]
 },
 {
 "answerData": [
 {
 "answerCondition": "answer.pool == 'Other'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'MIT'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Google Fiber'",
 "value": 3
 }
]
 }
],
{
 "ruleType": "LIMIT",
 "defaultCount": 1
}
]
```

### **ROUTE\_BY\_IP**

IP Prefix-based steering policies enable customers to steer DNS traffic based on the IP Prefix of the originating query. Each of the following rules must be defined in the order

## CHAPTER 12 DNS and Traffic Management

specified below in the `rules` field of your request body when using a `ROUTE_BY_IP` template:

Order	Rule	Restrictions	Comments
1	FILTER	<ul style="list-style-type: none"><li>No cases are allowed.</li><li>Answer data must be defined in <code>defaultAnswerData</code> using the following JSON: <pre>"answer.isDisabled != true",</pre></li></ul>	<pre>{   "answerCondition":   "shouldKeep": true }</pre>
2	HEALTH	<ul style="list-style-type: none"><li>No cases are allowed.</li></ul>	Only included if <code>healthCheckMonitorId</code> is defined for the policy.

Example of a `POST /steeringPolicies` request body using the `ROUTE_BY_IP` template:

```
{
 "compartmentId": "ocidl...",
 "displayName": "IP subnets mapped to answer pools",
 "ttl": 30,
 "template": "ROUTE_BY_IP",
 "answers": [
 {
 "name": "MIT Server",
 "rtype": "A",
 "rdata": "10.10.10.10",
 "pool": "MIT"
 },
 {
 "name": "Google Fiber Server",
 "rtype": "A",
 "rdata": "10.10.1.1",
```

## CHAPTER 12 DNS and Traffic Management

---

```
 "pool": "Google Fiber"
 },
 {
 "name": "Other",
 "rtype": "A",
 "rdata": "203.0.113.1",
 "pool": "Other"
 }
],
"rules": [
 {
 "ruleType": "FILTER",
 "defaultAnswerData": [
 {
 "answerCondition": "answer.isDisabled != true",
 "shouldKeep": true
 }
]
 },
 {
 "ruleType": "PRIORITY",
 "cases": [
 {
 "caseCondition": "query.client.address in (subnet '18.0.0.0/9')",
 "answerData": [
 {
 "answerCondition": "answer.pool == 'MIT'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'Google Fiber'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Other'",
 "value": 3
 }
]
 }
]
 }
],
{
```

## CHAPTER 12 DNS and Traffic Management

---

```
"caseCondition": "query.client.address in (subnet '136.32.0.0/11')",
"answerdata": [
 {
 "answerCondition": "answer.pool == 'Google Fiber'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'MIT'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Other'",
 "value": 3
 }
]
},
{
 "answerData": [
 {
 "answerCondition": "answer.pool == 'Other'",
 "value": 1
 },
 {
 "answerCondition": "answer.pool == 'MIT'",
 "value": 2
 },
 {
 "answerCondition": "answer.pool == 'Google Fiber'",
 "value": 3
 }
]
}
],
{
 "ruleType": "LIMIT",
 "defaultCount": 1
}
]
```

### **CUSTOM**

Custom policies allow you to create complex policies combining the capabilities of failover, load balancing, geolocation, ASN and IP prefix steering. Custom templates do not require a regimented sequence of rules and it is recommended to contact Oracle Cloud Infrastructure support before creating a custom policy.

### **Rule Types**

#### **FILTER**

Uses boolean data associated with answers, keeping answers only if the rule's `shouldKeep` value is `true`.

#### **HEALTH**

Utilizes Oracle Cloud Health Check monitors to determine the health of your endpoints and add and remove answers from your policy as needed. A health check monitor must be referenced in a health rule to have an effect on the policy. For more information about Health Checks, see Health Checks.

#### **WEIGHTED**

Uses a number between 0 and 255 used to determine how often an answer will be served in relation to other answers. Answers with higher values are more likely to be returned.

#### **PRIORITY**

Uses an integer associated with each answer to sort answers from lowest to highest value. Example: An answer with a priority value of 1 would be returned before an answer with a priority value of 10 in the list of answers. Answers that do not have a priority value assigned to them will be moved to the end of the list of answers.

#### **LIMIT**

Uses a count property to filter away all but the first answers in the list.

# Managing Traffic Management Steering Policies

## Policy Types

### **FAILOVER**

Failover policies allow you to prioritize the order in which you want answers served in a policy (for example, Primary and Secondary). Oracle Cloud Infrastructure Health Checks are leveraged to determine the health of answers in the policy. If the Primary Answer is determined to be unhealthy, DNS traffic will automatically be steered to the Secondary Answer.

### **LOAD BALANCER**

Load Balancer policies allow distribution of traffic across multiple endpoints. Endpoints can be assigned equal weights to distribute traffic evenly across the endpoints or custom weights may be assigned for ratio load balancing. Oracle Cloud Infrastructure Health Checks are leveraged to determine the health of the endpoint. DNS traffic will be automatically distributed to the other endpoints, if an endpoint is determined to be unhealthy.

### **GEOLOCATION STEERING**

Geolocation steering policies distribute DNS traffic to different endpoints based on the location of the end user. Customers can define geographic regions composed of originating continent, countries or states/provinces (North America) and define a separate endpoint or set of endpoints for each region.

### **ASN STEERING**

ASN steering policies enable you to steer DNS traffic based on Autonomous System Numbers (ASN). DNS queries originating from a specific ASN or set of ASNs can be steered to a specified endpoint.

### **IP PREFIX STEERING**

IP Prefix steering policies enable customers to steer DNS traffic based on the IP Prefix of the originating query.

### Typical Traffic Steering Scenarios

This section describes several typical scenarios for using Traffic Management Steering Policies.

#### **BASIC FAILOVER**

You can leverage Traffic Management Steering Policies to provide automated failover between primary and secondary servers.

#### **CLOUD MIGRATION**

Weighted load balancing supports controlled migration from your data center to Oracle Cloud Infrastructure servers. You can steer a small amount of traffic (1%) to your new resources in the cloud to verify everything is working as expected. You can then increase the ratios until you are comfortable with fully migrating all DNS traffic to the cloud.

#### **LOAD BALANCING ACROSS MULTIPLE SERVERS FOR SCALE**

You can configure load balancing pools of multiple servers. Traffic Management Steering Policies can automatically distribute DNS traffic across the set of servers. Health Checks may also be used and traffic will be automatically redirected to healthy servers, if a server is determined to be unhealthy.

#### **HYBRID ENVIRONMENTS**

Since Traffic Management Steering Policies is an agnostic service, it may be used to not only steer traffic to Oracle Cloud Infrastructure resources, but can also be used to steer traffic to any publicly exposed (internet resolvable) resources, including other cloud providers and enterprise data centers.

#### **WORLDWIDE GEOLOCATION TREATMENT**

You can divide your global users into geographically defined regions (for example, state/province level in NA, country level for rest of world) and steer customers to specified resources based on their location. This helps to ensure global, high performing internet resolution, and supports functions such as ring fencing. For example, keeping traffic from China in China and block traffic outside of China into China.

### CANARY TESTING

Leveraging IP Prefix steering, you can configure policies to serve different responses for your internal users versus external users.

### ZERO-RATING SERVICES

ASN steering conditional steering based on the originating enterprise, mobile operator or other communications provider in support of various commercial agreements that may be in place. Essentially, preferred ASNs can be directed to free resources, while all other traffic can be directed to paid resources.

### Using the Console

#### MANAGING TRAFFIC MANAGEMENT STEERING POLICIES

To create a Load Balancer policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click **Create Traffic Management Steering Policy**.
3. In the **Create Traffic Management Steering Policy** dialog box, select **Load Balancer**.
4. Enter the following information:
  - **Policy Name:** The unique name that identifies policy.
  - **Policy TTL:** The Time to Live for responses from the steering policy. If not specified, the system will set this value on the steering policy.
  - **Maximum Answer Count:** The maximum number of answers returned for the policy.
  - **Answer(s):** Answer pools contain the group of answers that will be served in response to DNS queries.

- **Name:** A unique name to identify the answer. Avoid entering confidential information.
  - **Type:** The record type that will be provided as the answer.
  - **RDATA:** A valid domain name or IP address to add as an answer.
  - **Weight:** A number between 0 and 255 used to determine how often an answer is served in relation to other answers. Answers with higher values are more likely to be served.
  - **Eligible:** Select the check box to indicate that the answer is available within the pool to be used in response to queries. Alternatively, select **Mark pool answers eligible** or **Mark pool answers ineligible** from the **Actions** drop-down menu.
- **Attach Health Check:** Select an existing Health Check to be included as part of the policy, add a new one, or select **None**.
  - **Attach Domain(s):** (Optional) The domain name and domain OCID you want to attach to the policy. Additional domains can be added in this section.
5. Click **Create Policy**.

The system creates and publishes the policy.

### To create a Failover policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click **Create Traffic Management Steering Policy**.
3. In the **Create Traffic Management Steering Policy** dialog box, select **Failover**.
4. Enter the following information:
  - **Policy Name:** The unique name that identifies policy. Avoid entering confidential information.

- **Policy TTL:** The Time to Live for responses from the steering policy. If not specified, the system will set this value on the steering policy.
  - **Maximum Answer Count:** The maximum number of answers returned for the policy. For priority-based policies, the first valid answer is returned.
  - **Answer Pool(s):** Answer pools contain the group of answers that will be served in response to DNS queries.
    - **Answer Pool Name:** A user-friendly name for the answer pool, unique within the steering policy. Avoid entering confidential information.
    - **Name:** A unique name to identify the answer. Avoid entering confidential information.
    - **Type:** The record type that will be provided as the answer.
    - **RDATA:** A valid domain name or IP address to add as an answer.
    - **Weight:** A number between 0 and 255 used to determine how often an answer is served in relation to other answers. Answers with higher values are more likely to be served.
    - **Eligible:** Select the check box to indicate that the answer is available within the pool to be used in response to queries. Alternatively, select **Mark pool answers eligible** or **Mark pool answers ineligible** from the **Actions** drop-down menu.
  - **Pool Priority:** Failover priority rules specify the priority of answers that are served in a policy. If the primary answer is unavailable, traffic is steered to the next answer in the list.
    - **Pool:** Select the priority in which the answers are served.
  - **Attach Health Check:** Select an existing Health Check to be included as part of the policy, add a new one, or select **None**.
  - **Attach Domain(s):** The domain name and domain OCID you want to attach to the policy. Additional domains can be added in this section.
5. Click **Create Policy**.

The system creates and publishes the policy.

### To create a Geolocation Steering policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click **Create Traffic Management Steering Policy**.
3. In the **Create Traffic Management Steering Policy** dialog box, select **Geolocation Steering**.
4. Enter the following information:
  - **Policy Name:** The unique name that identifies policy. Avoid entering confidential information.
  - **Policy TTL:** The Time to Live for responses from the steering policy. If not specified, the system will set this value on the steering policy.
  - **Maximum Answer Count:** The maximum number of answers returned for the policy. For priority-based policies, the first valid answer is returned.
  - **Answer Pool(s):** Answer pools contain the group of answers that will be served in response to DNS queries.
    - **Answer Pool Name:** A user-friendly name for the answer pool, unique within the steering policy. Avoid entering confidential information.
    - **Name:** A unique name to identify the answer. Avoid entering confidential information.
    - **Type:** The record type that will be provided as the answer.
    - **RDATA:** A valid domain name or IP address to add as an answer.
    - **Eligible:** Select the check box to indicate that the answer is available within the pool to be used in response to queries. Alternatively, select **Mark pool answers eligible** or **Mark pool answers ineligible** from the **Actions** drop-down menu.

- **Geolocation Steering Rules:** Geolocation steering rules specify the priority of answers that are served in a policy. If the primary answer is unavailable, traffic is steered to the next answer in the list. Additional rules and priorities can be added in this section.
    - **Geolocation:** Select a location that will be used to distribute DNS traffic.
    - **Pool Priority:** Select the priority in which the answers are served.
    - **Global Catch-all:** Adding a global catch-all allows you to specify answer pools for queries that do not match any of the specified rules you have added. Click **Add Global Catch-all** and select the pool priorities.
  - **Attach Health Check:** Select an existing Health Check to be included as part of the policy, add a new one, or select **None**.
  - **Attach Domain(s):** The domain name and domain OCID you want to attach to the policy. Additional domains can be added in this section.
5. Click **Create Policy**.

The system creates and publishes the policy.

### To create an ASN Steering policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click **Create Traffic Management Steering Policy**.
3. In the **Create Traffic Management Steering Policy** dialog box, select **ASN Steering**.
4. Enter the following information:
  - **Policy Name:** The unique name that identifies policy. Avoid entering confidential information.
  - **Policy TTL:** The Time to Live for responses from the steering policy. If not specified, the system will set this value on the steering policy.

- **Maximum Answer Count:** The maximum number of answers returned for the policy. For priority-based policies, the first valid answer is returned.
  - **Answer Pool(s):** Answer pools contain the group of answers that will be served in response to DNS queries.
    - **Answer Pool Name:** A user-friendly name for the answer pool, unique within the steering policy. Avoid entering confidential information.
    - **Name:** A unique name to identify the answer. Avoid entering confidential information.
    - **Type:** The record type that will be provided as the answer.
    - **RDATA:** A valid domain name or IP address to add as an answer.
    - **Eligible:** Select the check box to indicate that the answer is available within the pool to be used in response to queries. Alternatively, select **Mark pool answers eligible** or **Mark pool answers ineligible** from the **Actions** drop-down menu.
  - **ASN Steering Rules:** ASN steering rules specify the priority of answers that are served in a policy. If the primary answer is unavailable, traffic is steered to the next answer in the list.
    - **ASN:** Enter an Autonomous System Number (ASN) that will be used to distribute DNS traffic.
    - **Pool Priority:** Select the priority in which the answers are served.
    - **Global Catch-all:** Adding a global catch-all allows you to specify answer pools for queries that do not match any of the specified rules you have added. Click **Add Global Catch-all** and select the pool priorities.
  - **Attach Health Check:** Select an existing Health Check to be included as part of the policy, add a new one, or select **None**.
  - **Attach Domain(s):** The domain name and domain OCID you want to attach to the policy. Additional domains can be added in this section.
5. Click **Create Policy**.

The system creates and publishes the policy.

### To create an IP Prefix Steering policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click **Create Traffic Management Steering Policy**.
3. In the **Create Traffic Management Steering Policy** dialog box, select **IP Prefix Steering**.
4. Enter the following information:
  - **Policy Name:** The unique name that identifies policy. Avoid entering confidential information.
  - **Policy TTL:** The Time to Live for responses from the steering policy. If not specified, the system will set this value on the steering policy.
  - **Maximum Answer Count:** The maximum number of answers returned for the policy. For priority-based policies, the first valid answer is returned.
  - **Answer Pool(s):** Answer pools contain the group of answers that will be served in response to DNS queries.
    - **Answer Pool Name:** A user-friendly name for the answer pool, unique within the steering policy. Avoid entering confidential information.
    - **Name:** A unique name to identify the answer. Avoid entering confidential information.
    - **Type:** The record type that will be provided as the answer.
    - **RDATA:** A valid domain name or IP address to add as an answer.
    - **Eligible:** Select the check box to indicate that the answer is available within the pool to be used in response to queries. Alternatively, select **Mark pool answers eligible** or **Mark pool answers ineligible** from the **Actions** drop-down menu.

- **IP Prefix Steering Rules:** IP prefix steering rules specify the priority of answers that are served in a policy. If the primary answer is unavailable, traffic is steered to the next answer in the list.
    - **Subnet Address:** Enter a subnet address that will be used to distribute DNS traffic.
    - **Pool Priority:** Select the priority in which the answers are served.
    - **Global Catch-all:** Adding a global catch-all allows you to specify answer pools for queries that do not match any of the specified rules you have added. Click **Add Global Catch-all** and select the pool priorities.
  - **Attach Health Check:** Select an existing Health Check to be included as part of the policy, add a new one, or select **None**.
  - **Attach Domain(s):** The domain name and domain OCID you want to attach to the policy. Additional domains can be added in this section.
5. Click **Create Policy**.

The system creates and publishes the policy.

### To update a policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click the **Policy Name** you want to update. Policy information and a list of attached domains appear.



#### Tip

You can use search for a policy by name in the **Search** field. You can also use the **Time Created** sort filter to sort the policies chronologically in ascending or descending order.

3. Click **Edit**.
4. Make the needed changes, and then click **Save**.

### To attach a domain to an existing policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click the **Policy Name** you want to update. Policy information and a list of attached domains appear.



#### Tip

You can use search for a policy by name in the **Search** field. You can also use the **Time Created** sort filter to sort the policies chronologically in ascending or descending order.

3. Click **Add Attached Domain(s)**.
4. In the Add Attached Domain(s) dialog box, enter the domain and select a zone.
5. Click **Submit**.

### To edit an attached domain

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Click the **Policy Name** you want to update. Policy information and a list of attached domains appear.



### Tip

You can use search for a policy by name in the **Search** field. You can also use the **Time Created** sort filter to sort the policies chronologically in ascending or descending order.

3. For the attached domain you want to edit, click the Actions icon (three dots), and then click **Edit Attached Domain**.
4. In the Attached Domain(s) dialog box, enter the domain and select a zone.
5. Click **Save**.

### To delete a policy

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Traffic Management Steering Policies**.
2. Select the check box for the policy you want to delete.
3. Click **Delete**. The policy is staged for deletion.
4. Click **Publish Changes** to delete the policy.
5. In the confirmation dialog box, click **Publish Changes**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For more information about managing steering policies using the API, see [Traffic Management Steering Policies API Guide](#).

Use the following operations to manage your steering policies:

## CHAPTER 12 DNS and Traffic Management

---

- [CreateSteeringPolicy](#)
- [ListSteeringPolicies](#)
- [GetSteeringPolicy](#)
- [UpdateSteeringPolicy](#)
- [DeleteSteeringPolicy](#)

Use the following operations to manage your steering policy attachments:

- [CreateSteeringPolicyAttachment](#)
- [ListSteeringPolicyAttachments](#)
- [GetSteeringPolicyAttachment](#)
- [UpdateSteeringPolicyAttachment](#)
- [DeleteSteeringPolicyAttachment](#)

### Traffic Management Steering Policy geokeys

Use these keys as values for the `geokey` fields of `caseConditions` in `ROUTE_BY_GEO` steering policies.

#### Continent geokeys

Continent Name	geoKey
Africa	6255146
Antarctica	6255152
Asia	6255147
Europe	6255148
North America	6255149

## CHAPTER 12 DNS and Traffic Management

---

Continent Name	geoKey
Oceania	6255151
South America	6255150

### Country geokeys

Country Name	geoKey
Afghanistan(AF)	1149361
Aland Islands(AX)	661882
Albania(AL)	783754
Algeria(DZ)	2589581
American Samoa(AS)	5880801
Andorra(AD)	3041565
Angola(AO)	3351879
Anguilla(AI)	3573511
Antarctica(AQ)	6697173
Antigua and Barbuda(AG)	3576396
Argentina(AR)	3865483
Armenia(AM)	174982
Aruba(AW)	3577279
Australia(AU)	2077456

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Austria(AT)	2782113
Azerbaijan(AZ)	587116
Bahamas(BS)	3572887
Bahrain(BH)	290291
Bangladesh(BD)	1210997
Barbados(BB)	3374084
Belarus(BY)	630336
Belgium(BE)	2802361
Belize(BZ)	3582678
Benin(BJ)	2395170
Bermuda(BM)	3573345
Bhutan(BT)	1252634
Bolivia(BO)	3923057
Bonaire, Saint Eustatius and Saba(BQ)	7626844
Bosnia and Herzegovina(BA)	3277605
Botswana(BW)	933860
Bouvet Island(BV)	3371123
Brazil(BR)	3469034
British Indian Ocean Territory(IO)	1282588

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
British Virgin Islands(VG)	3577718
Brunei(BN)	1820814
Bulgaria(BG)	732800
Burkina Faso(BF)	2361809
Burundi(BI)	433561
Cambodia(KH)	1831722
Cameroon(CM)	2233387
Canada(CA)	6251999
Cape Verde(CV)	3374766
Cayman Islands(KY)	3580718
Central African Republic(CF)	239880
Chad(TD)	2434508
Chile(CL)	3895114
China(CN)	1814991
Christmas Island(CX)	2078138
Cocos (Keeling) Islands(CC)	1547376
Colombia(CO)	3686110
Comoros(KM)	921929
Congo(CG)	2260494

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Cook Islands(CK)	1899402
Costa Rica(CR)	3624060
Croatia(HR)	3202326
Cuba(CU)	3562981
Curacao(CW)	7626836
Cyprus(CY)	146669
Czech Republic(CZ)	3077311
Democratic Republic of the Congo(CD)	203312
Denmark(DK)	2623032
Djibouti(DJ)	223816
Dominica(DM)	3575830
Dominican Republic(DO)	3508796
East Timor(TL)	1966436
Ecuador(EC)	3658394
Egypt(EG)	357994
El Salvador(SV)	3585968
Equatorial Guinea(GQ)	2309096
Eritrea(ER)	338010
Estonia(EE)	453733

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
Ethiopia(ET)	337996
Falkland Islands(FK)	3474414
Faroe Islands(FO)	2622320
Fiji(FJ)	2205218
Finland(FI)	660013
France(FR)	3017382
French Guiana(GF)	3381670
French Polynesia(PF)	4030656
French Southern Territories(TF)	1546748
Gabon(GA)	2400553
Gambia(GM)	2413451
Georgia(GE)	614540
Germany(DE)	2921044
Ghana(GH)	2300660
Gibraltar(GI)	2411586
Greece(GR)	390903
Greenland(GL)	3425505
Grenada(GD)	3580239
Guadeloupe(GP)	3579143

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Guam(GU)	4043988
Guatemala(GT)	3595528
Guernsey(GG)	3042362
Guinea(GN)	2420477
Guinea-Bissau(GW)	2372248
Guyana(GY)	3378535
Haiti(HT)	3723988
Heard Island and McDonald Islands(HM)	1547314
Honduras(HN)	3608932
Hong Kong(HK)	1819730
Hungary(HU)	719819
Iceland(IS)	2629691
India(IN)	1269750
Indonesia(ID)	1643084
Iran(IR)	130758
Iraq(IQ)	99237
Ireland(IE)	2963597
Isle of Man(IM)	3042225
Israel(IL)	294640

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
Italy(IT)	3175395
Ivory Coast(CI)	2287781
Jamaica(JM)	3489940
Japan(JP)	1861060
Jersey(JE)	3042142
Jordan(JO)	248816
Kazakhstan(KZ)	1522867
Kenya(KE)	192950
Kiribati(KI)	4030945
Kuwait(KW)	285570
Kyrgyzstan(KG)	1527747
Laos(LA)	1655842
Latvia(LV)	458258
Lebanon(LB)	272103
Lesotho(LS)	932692
Liberia(LR)	2275384
Libya(LY)	2215636
Liechtenstein(LI)	3042058
Lithuania(LT)	597427

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Luxembourg(LU)	2960313
Macau(MO)	1821275
Macedonia(MK)	718075
Madagascar(MG)	1062947
Malawi(MW)	927384
Malaysia(MY)	1733045
Maldives(MV)	1282028
Mali(ML)	2453866
Malta(MT)	2562770
Marshall Islands(MH)	2080185
Martinique(MQ)	3570311
Mauritania(MR)	2378080
Mauritius(MU)	934292
Mayotte(YT)	1024031
Mexico(MX)	3996063
Micronesia(FM)	2081918
Moldova(MD)	617790
Monaco(MC)	2993457
Mongolia(MN)	2029969

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Montenegro(ME)	3194884
Montserrat(MS)	3578097
Morocco(MA)	2542007
Mozambique(MZ)	1036973
Myanmar(MM)	1327865
Namibia(NA)	3355338
Nauru(NR)	2110425
Nepal(NP)	1282988
Netherlands(NL)	2750405
New Caledonia(NC)	2139685
New Zealand(NZ)	2186224
Nicaragua(NI)	3617476
Niger(NE)	2440476
Nigeria(NG)	2328926
Niue(NU)	4036232
Norfolk Island(NF)	2155115
North Korea(KP)	1873107
Northern Mariana Islands(MP)	4041468
Norway(NO)	3144096

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Oman(OM)	286963
Pakistan(PK)	1168579
Palau(PW)	1559582
Palestinian territories(PS)	6254930
Panama(PA)	3703430
Papua New Guinea(PG)	2088628
Paraguay(PY)	3437598
Peru(PE)	3932488
Philippines(PH)	1694008
Pitcairn(PN)	4030699
Poland(PL)	798544
Portugal(PT)	2264397
Puerto Rico(PR)	4566966
Qatar(QA)	289688
Reunion(RE)	935317
Romania(RO)	798549
Russia(RU)	2017370
Rwanda(RW)	49518
Saint Barthelemy(BL)	3578476

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
Saint Helena(SH)	3370751
Saint Kitts and Nevis(KN)	3575174
Saint Lucia(LC)	3576468
Saint Martin(MF)	3578421
Saint Pierre and Miquelon(PM)	3424932
Saint Vincent and the Grenadines(VC)	3577815
Samoa(WS)	4034894
San Marino(SM)	3168068
Sao Tome and Principe(ST)	2410758
Saudi Arabia(SA)	102358
Senegal(SN)	2245662
Serbia(RS)	6290252
Seychelles(SC)	241170
Sierra Leone(SL)	2403846
Singapore(SG)	1880251
Sint Maarten(SX)	7609695
Slovakia(SK)	3057568
Slovenia(SI)	3190538
Solomon Islands(SB)	2103350

## CHAPTER 12 DNS and Traffic Management

---

<b>Country Name</b>	<b>geoKey</b>
Somalia(SO)	51537
South Africa(ZA)	953987
South Georgia and the South Sandwich Islands(GS)	3474415
South Korea(KR)	1835841
South Sudan(SS)	7909807
Spain(ES)	2510769
Sri Lanka(LK)	1227603
Sudan(SD)	366755
Suriname(SR)	3382998
Svalbard and Jan Mayen(SJ)	607072
Swaziland(SZ)	934841
Sweden(SE)	2661886
Switzerland(CH)	2658434
Syria(SY)	163843
Taiwan(TW)	1668284
Tajikistan(TJ)	1220409
Tanzania(TZ)	149590
Thailand(TH)	1605651

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
Togo(TG)	2363686
Tokelau(TK)	4031074
Tonga(TO)	4032283
Trinidad and Tobago(TT)	3573591
Tunisia(TN)	2464461
Turkey(TR)	298795
Turkmenistan(TM)	1218197
Turks and Caicos Islands(TC)	3576916
Tuvalu(TV)	2110297
U.S. Virgin Islands(VI)	4796775
Uganda(UG)	226074
Ukraine(UA)	690791
United Arab Emirates(AE)	290557
United Kingdom(GB)	2635167
United States(US)	6252001
United States Minor Outlying Islands(UM)	5854968
Uruguay(UY)	3439705
Uzbekistan(UZ)	1512440
Vanuatu(VU)	2134431

## CHAPTER 12 DNS and Traffic Management

---

Country Name	geoKey
Vatican City(VA)	3164670
Venezuela(VE)	3625428
Vietnam(VN)	1562822
Wallis and Futuna(WF)	4034749
Western Sahara(EH)	2461445
Yemen(YE)	69543
Zambia(ZM)	895949
Zimbabwe(ZW)	878675

### United States geokeys

State Name	geoKey
Alabama	4829764
Alaska	5879092
Arizona	5551752
Arkansas	4099753
California	5332921
Colorado	5417618
Connecticut	4831725
Delaware	4142224

## CHAPTER 12 DNS and Traffic Management

---

<b>State Name</b>	<b>geoKey</b>
District of Columbia	4138106
Florida	4155751
Georgia	4197000
Hawaii	5855797
Idaho	5596512
Illinois	4896861
Indiana	4921868
Iowa	4862182
Kansas	4273857
Kentucky	6254925
Louisiana	4331987
Maine	4971068
Maryland	4361885
Massachusetts	6254926
Michigan	5001836
Minnesota	5037779
Mississippi	4436296
Missouri	4398678
Montana	5667009

## CHAPTER 12 DNS and Traffic Management

---

<b>State Name</b>	<b>geoKey</b>
Nebraska	5073708
Nevada	5509151
New Hampshire	5090174
New Jersey	5101760
New Mexico	5481136
New York	5128638
North Carolina	4482348
North Dakota	5690763
Ohio	5165418
Oklahoma	4544379
Oregon	5744337
Pennsylvania	6254927
Rhode Island	5224323
South Carolina	4597040
South Dakota	5769223
Tennessee	4662168
Texas	4736286
Utah	5549030
Vermont	5242283

## CHAPTER 12 DNS and Traffic Management

---

State Name	geoKey
Virginia	6254928
Washington	5815135
West Virginia	4826850
Wisconsin	5279468
Wyoming	5843591

### Canada Provinces geokeys

Province Name	geoKey
Alberta	5883102
British Columbia	5909050
Manitoba	6065171
New Brunswick	6087430
Newfoundland and Labrador	6354959
Northwest Territories	6091069
Nova Scotia	6091530
Nunavut	6091732
Ontario	6093943
Prince Edward Island	6113358
Quebec	6115047

## CHAPTER 12 DNS and Traffic Management

---

Province Name	geoKey
Saskatchewan	6141242
Yukon	6185811

# CHAPTER 13 Email Delivery

This chapter explains how to send large volume email.

## Overview of the Email Delivery Service

Oracle Cloud Infrastructure Email Delivery is an email sending service that provides a fast and reliable managed solution for sending high-volume emails that need to reach your recipients' inbox. Email Delivery provides the tools necessary to send application-generated email for mission-critical communications such as receipts, fraud detection alerts, multi-factor identity verification, and password resets.

Oracle Cloud Infrastructure's Email Deliverability team manages the platform using key deliverability metrics to ensure the best sending reputation possible for your emails.

The following items are provided to you when you send email using the Email Delivery service:

- Unique mailbox provider SMTP configurations on our Mail Transfer Agents (MTA)
- Bounce collection
- User complaint collection
- Email authentication standards
- Deliverability performance

## Email Delivery Service Components

Email Delivery uses the components described in this section.

### **APPROVED SENDERS**

An Approved Sender is a resource that equates to the "From" address. An approved sender is associated with a compartment and only exists in the region where the approved sender was configured. If you need to have the same approved sender in another region,

it must be created in the other region. For example, if you create an approved sender in the US West (Phoenix) region, you cannot send email through the US East (Ashburn) region.

### **SUPPRESSION LIST**

The Suppression List is included on your Email Delivery console user interface and from the API. Email Delivery automatically adds email addresses with bounce codes showing permanent failures or user complaints to the suppression list to protect your sender reputation. Email Delivery will not send any messages to these recipients in the future.

Reasons for suppression currently include:

- Complaints
- Hard bounces
- Repetitive soft bounces
- Manual entries
- List-unsubscribe requests

### **SPF AUTHENTICATION**

Sender Policy Framework (SPF) is used by email receivers to detect email spoofing. Using SPF, an email receiver can check if the Internet Protocol (IP) is explicitly authorized to send for that domain.

SPF is implemented by publishing a special TXT record to a domain's DNS records. The TXT record declares which hosts are allowed to send mail on behalf of this domain.

Receiving mail servers check the SPF records of sending domains to verify that the email's source IP address is authorized to send from that domain. Without SPF, a spam or phishing email can be "spoofed" to appear that the email comes from a legitimate domain. Domains that implement SPF are much more likely to block emails attempting to spoof your domain.

For an overview of how SPF works, see [Sender Policy Framework](#). For details on SPF record syntax, see [SPF Record Syntax](#).

### Regions and Availability Domains

Email Delivery is available in the US West (Phoenix), US East (Ashburn), and UK South (London) regions. For more information, see [Regions and Availability Domains](#).

The sending application is not required to be located in the region where email is sent. For example, if your sending application is located in a region where Email Delivery is not currently available, you would configure email from one of the regions where it is available. In the Console, change your region to US West (Phoenix), US East (Ashburn), or UK South (London) and create an approved sender. When creating SMTP credentials, any region can be used, as identities are global assets. Configure your application to send email to the region where you created the approved sender (US West (Phoenix), US East (Ashburn), or UK South (London) endpoint) using the SMTP credentials.

When Email Delivery is available in more regions, you can configure Email Delivery in the same region as the sending application to improve performance.

### Configuring a New Region

If you want to start sending email from a new region, keep the following in mind:

- An approved sender must be created in the new region.
- SMTP credentials are global, however, it is recommended that you generate SMTP credentials for a new user (without console access) in the new region so that the credentials are not shared with other regions. Ensure that the user has the correct privileges.
- Email must be sent to the new regional SMTP connection endpoint.
- The suppression list and approved senders are regional Email Delivery assets. For example, if an email sent from the US West (Phoenix) region bounces, the recipient email address will be added to the US West (Phoenix) region suppression list. This recipient would not be added to other region suppression lists. If you are sending email from different regions, approved senders must be created in each region.
- SPF must be set up on each subdomain. For example, in your DNS setup, create a TXT record for `notification.eu-frankfurt-1.oraclecloud.com` and paste the following

```
information from the dialog box into the record: v=spf1
include:spf.oracleemaildelivery.com -all
```

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [SDKs and Other Tools](#).

To access the Console, you must use a supported browser. You can use the Console link at the top of this page to go to the sign-in page. You are prompted to enter your cloud tenant, your user name, and your password. For general information about using the API, see [About the API](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Email Delivery supports the following authentication types for control plane operations (management endpoint):

- **Instance Authorization:** The IAM service feature that enables instances to be authorized actors (or principals) to perform actions on service resources. Each compute instance has its own identity, and it authenticates using the certificates that are added to it. These certificates are automatically created, assigned to instances and rotated, preventing the need for you to distribute credentials to your hosts and rotate them.
- **Cross-Tenancy:** Cross-tenancy authorization allows customers to share resources between tenancies. To authorize a cross-tenancy request, the request must be endorsed by the requester's tenancy and permitted by the target tenancy.
- **Federated:** Federated authentication enables an administrator to configure a relationship between an identity provider and a service provider. When you federate Oracle Cloud Infrastructure with an identity provider, you manage users and groups in the identity provider. You manage authorization in Oracle Cloud Infrastructure's IAM service. Oracle Cloud Infrastructure tenancies are federated with Oracle Identity Cloud Service by default.



### Note

Instance authorization, cross-tenancy, and federated authentication types do not apply to SMTP email sending. An approved sender and SMTP credentials are required and must be associated with the same tenancy for SMTP email sending.

## SMTP Authentication and Connection Endpoints

Email Delivery only supports the AUTH PLAIN command when using SMTP authentication. If the sending application is not flexible with the AUTH command, an SMTP proxy/relay can be used. For more information about the AUTH command, see [AUTH Command and its Mechanisms](#).

Use the following regional endpoints for establishing SMTP connections for sending.

- US West (Phoenix): smtp.us-phoenix-1.oraclecloud.com
- US East (Ashburn): smtp.us-ashburn-1.oraclecloud.com
- UK South (London): smtp.email.uk-london-1.oci.oraclecloud.com

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about available Email Delivery service metrics and how to view them, see [Email Delivery Metrics](#).

### Email Delivery Service Capabilities and Limits

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

Customers that sign up for a free Oracle Cloud trial are limited to:

- A volume of 200 emails a day.
- Five approved senders.
- Each user is limited to a maximum of two SMTP credentials.
- Sending rates are limited to ten emails per minute.
- Inline attachments.

Enterprise accounts are limited to:

- A volume of 50,000 emails a day.
- 10,000 approved senders.
- Sending rates are limited to 18,000 emails per minute.
- Inline attachments.



### Note

The Email Delivery platform supports higher volumes. Limits are set as a safeguard for our customers' reputation. To file a service request to increase the email sending limit, open the navigation menu. Under **Governance and Administration**, go to **Service Limits**. Click **Request a service limit increase**.



### Note

Currently, Email Delivery supports messages up to 2 MB, inclusive of message headers, body, and attachments. This is not a limit set per tenant. Larger message sizes will be available in the future.

## Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for Email Delivery, see [Details for the Email Service](#).

Permissions are required for managing and using approved senders and the suppression list. For example:

## CHAPTER 13 Email Delivery

---

- To enable all operations on approved senders for a specific user group:

```
Allow group <Your Group Name> to manage approved-senders in tenancy
```

- To enable all operations on suppressions for a specific user group:

```
Allow group <Your Group Name> to manage suppressions in tenancy
```

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Email Delivery supports applying tags to approved senders.

### Integration with Oracle Cloud Infrastructure Services

Email Delivery audits the following events:

- Creating a sender (CreateSender)
- Deleting a sender (DeleteSender)
- Retrieving details about a sender (ListSenders)

To view logs for events in the Email Delivery service, your user must be in a group with the ability to view all of the Audit event logs in the tenancy. For more information, see [Viewing Audit Log Events](#).

### Getting Started with Email Delivery

You can set up the Email Delivery service within the Console. To begin sending email with Email Delivery, complete the following steps:

1. [Generate SMTP credentials for a user.](#)
2. [Set up permissions.](#)
3. [Create an approved sender.](#)

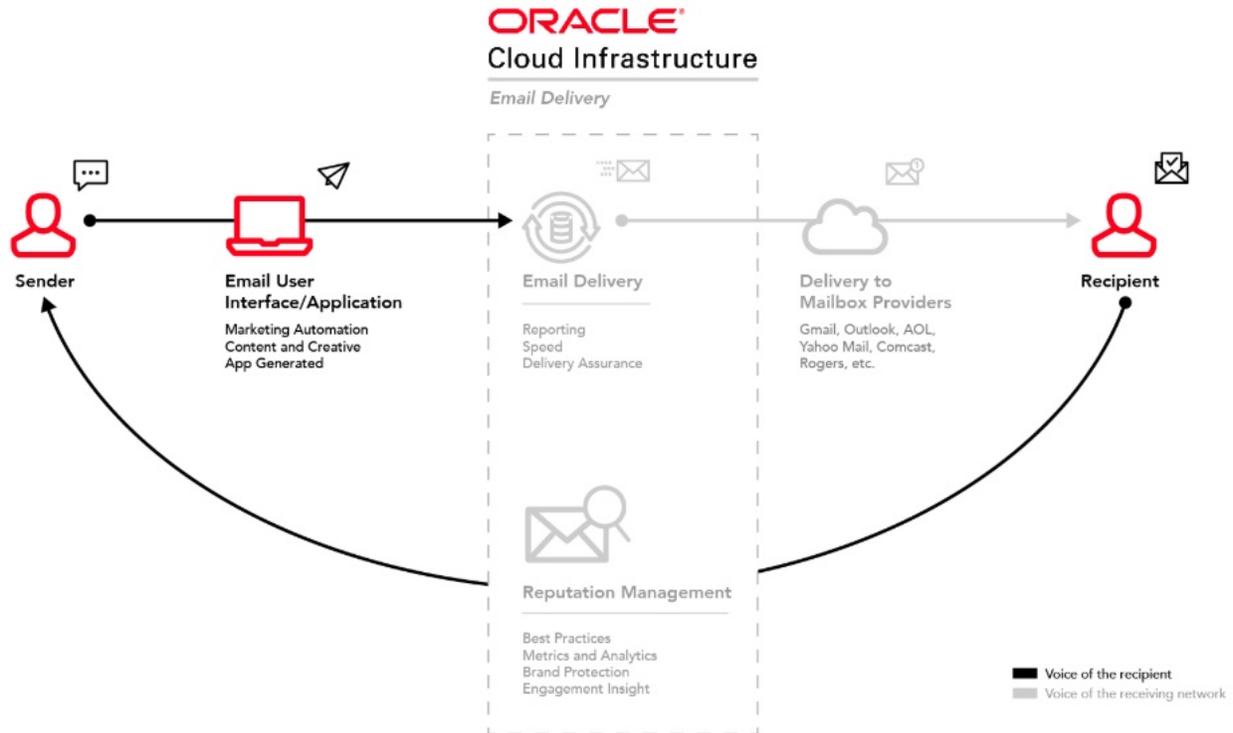
4. [Configure SPF on the approved sender domain.](#)
5. [Configure the SMTP connection.](#)
6. [Begin sending email.](#)

For more information, see [Getting Started with Email Delivery](#).

### Getting Started with Email Delivery

Email Delivery provides a highly scalable, cost effective, and reliable way to send email from your applications. Email Delivery includes developer-friendly tools to quickly send application-generated email for mission-critical communications such as receipts, programmatic notifications, or password reset emails.

## Email Delivery Basics



When you use Email Delivery, we become your outbound email server. If you have an existing email server, you can keep it and configure it to send through Email Delivery. The Email Delivery service will take care of the feedback loops and platform reputation automatically.

## Getting Started

This topic gives guidance on how to get started with Email Delivery. For complete details about the service and its components, see [Overview of the Email Delivery Service](#).

### Email Configuration Options

You can configure Oracle Cloud Infrastructure using the Console (a browser-based interface), [REST API](#), [SDKs](#), [CLI](#) or [Terraform](#).

### Using the Email Delivery SDK

The Email Delivery SDK is available in several programming languages. For information on installing and configuring the Oracle Cloud Infrastructure SDKs, see [Developer Tools](#).

Examples of SDK usage can be found on GitHub, including:

- [Example: Email Delivery SDK for Java](#)
- [Example: Email Delivery SDK for Python](#)
- [Example: Email Delivery SDK for Ruby](#)
- [Example: Email Delivery SDK for Go](#)

### Configuring Third-Party Applications

The following information describes how you can configure third-party applications to send email through Email Delivery:

- [Integrating Oracle Application Express with Email Delivery](#)
- [Integrating Postfix with Email Delivery](#)
- [Integrating Oracle Enterprise Manager with Email Delivery](#)
- [Integrating Mailx with Email Delivery](#)
- [Integrating Swaks with Email Delivery](#)
- [Integrating Sendmail with Email Delivery](#)
- [Integrating JavaMail with Email Delivery](#)
- [Integrating PeopleSoft with Email Delivery](#)
- [Integrating Python with Email Delivery](#)

### Sending Email

To begin sending email with Email Delivery, complete the following steps:

#### Generate SMTP credentials for a user.

Simple Mail Transfer Protocol (SMTP) credentials are necessary to send email through Email Delivery. Each user is limited to a maximum of two SMTP credentials. If more than two are required, SMTP credentials must be generated that are associated with another existing user or more users must be created.

**Best Practice:** A security best practice is to generate SMTP credentials for a new user instead of your Console user that already has permissions assigned to it. For detailed instructions on creating a user, see [Adding Users](#).

#### To generate SMTP credentials for a user.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. Locate the user in the list that has permissions to manage email, and then click the user's name to view the details.



#### Tip

If your user does not have permissions to view or create users, you can create SMTP credentials under your user. Open the **Profile** menu () and click **User Settings**.

2. Click **SMTP Credentials**.
3. Click **Generate SMTP Credentials**.
4. Enter a **Description** of the SMTP Credentials in the dialog box.
5. Click **Generate SMTP Credentials**. A user name and password is displayed.

6. Copy the user name and password for your records and click **Close**.

### Set up permissions.

The new user must be assigned to a group with permissions to manage `approved-senders` and `suppressions`.

### To create a policy to allow a group to manage approved senders and suppressions

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.

A list of the policies in the compartment you're viewing is displayed.

2. If you want to attach the policy to a compartment other than the one you're viewing, select the desired compartment from the list on the left. Where the policy is attached controls who can later modify or delete it (see [Policy Attachment](#)).

3. Click **Create Policy**.

4. Enter the following:

- **Name:** A unique name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later.
- **Description:** A friendly description. You can change this later if you want to.
- **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in format YYYY-MM-DD format. For more information, see [Policy Language Version](#).
- **Statement:** Enter the following policy statement:

```
Allow group <group name> to use approved-senders in compartment <compartment name>
```

For more information about policies and policy syntax, see [Policy Basics](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

The new policy will go into effect typically within 10 seconds.

### To add the new user to the group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Locate the user in the list.
3. Click the user.  
Its details are displayed.
4. Click **Groups**.
5. Click **Add User to Group**.
6. Select the group from the drop-down list, and then click **Add**.

Make sure to let the user know which compartment(s) they have access to.

### Create an approved sender.

You must set up an approved sender for all "From:" addresses sending mail via Oracle Cloud Infrastructure or mail will be rejected. An approved sender is associated with a compartment and only exists in the region where the approved sender was configured. That is, if you create

an approved sender in the Phoenix (PHX) region, you cannot send email through the Ashburn (IAD) region.

**Best Practice:** Approved senders should not be created in the root compartment. If approved senders exist in the root compartment, you are required to create a policy to manage approved senders in the entire tenant. Creating approved senders in a compartment other than the root allows the policy to be specific to that compartment.

### To create an approved sender using the Console

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**. Ensure that you are in the correct compartment. Your user must be in a group with permissions to manage `approved-senders` in this compartment.
2. Click **Create Approved Sender** within the **Approved Senders** view.
3. Enter the email address you want to list as an approved sender in the **Add Sender** dialog box.
4. Click **Add**. The email address is added to your Approved Senders list.



#### Tip

Approved senders are unique to tenancies. If an attempt is made to create a duplicate approved sender within a tenancy, the service will return a 409 Conflict error.

### To create an approved sender using the API

The following example shows how to create an approved sender. For more information about creating an approved sender, see [CreateSender](#).

## CHAPTER 13 Email Delivery

---

```
POST /20170907/senders

 {
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaaat7uqcb6zoxvzoga4d4vh4dtweciavepacd3skz56atf3qp73d7fx",
 "emailAddress": "user@example.com",
 }
}
```

### Configure SPF on the approved sender domain.

Sender Policy Framework (SPF) is used by email receivers to detect email spoofing. Using SPF, an email receiver can check if the Internet Protocol (IP) is explicitly authorized to send for that domain. SPF is implemented by publishing a special TXT record to a domain's DNS records. The TXT record declares which hosts are allowed to send mail on behalf of this domain. Receiving mail servers check the SPF records of sending domains to verify that the email's source IP address is authorized to send from that domain. Without SPF, a spam or phishing email can be "spoofed" to appear that the email comes from a legitimate domain. Domains that implement SPF are much more likely to block emails attempting to spoof your domain. For an overview of how SPF works, see [Sender Policy Framework](#). For details on SPF record syntax, see [SPF Record Syntax](#).

The Approved Senders section within the Console provides validation of an SPF record for each of your approved senders.

### To configure SPF

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**.
2. Select the checkbox for the approved sender you want to view SPF details for and click **View SPF**.



### Tip

You can search for an approved sender by using the Search field. Addresses can be sorted alphanumerically or by creation date in ascending or descending order.

3. The Manage SPF dialog box appears indicating whether an SPF record for the approved sender exists.
  - If your domain does not currently have an SPF record, the information necessary to add an SPF record in your DNS setup is displayed. See [Managing DNS Service Zones](#) for instructions on adding a zone record in Oracle Cloud Infrastructure. If your DNS setup resides with another provider, please reference their documentation for adding a TXT record to your domain.
    - In your DNS setup, create a TXT record and paste the following information from the dialog box into the record: `v=spf1 include:spf.oracleemaildelivery.com -all`
  - If your domain currently has an SPF record, add the following information to the record to add Oracle Cloud Infrastructure Email Delivery:  
`include:spf.oracleemaildelivery.com`

### Configure the SMTP connection.

Set up and test your SMTP connection using an SMTP library or product such, as [Postfix](#) or [Sendmail](#), to send email through Oracle Cloud Infrastructure Email Delivery.

### SMTP Connection Endpoints

Use the following regional endpoints for establishing SMTP connections for sending.

- PHX: smtp.us-phoenix-1.oraclecloud.com
- IAD: smtp.us-ashburn-1.oraclecloud.com
- LHR: smtp.email.uk-london-1.oci.oraclecloud.com

### TLS Requirements

Oracle maintains strict security policies and only accepts email traffic using Transport Layer Security (TLS). Use of TLS 1.2 is mandatory to send email using Oracle Cloud Infrastructure.

The approved TLS 1.2 ciphers are:

- TLS\_DHE\_RSA\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_DHE\_DSS\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_DHE\_DSS\_WITH\_AES\_128\_CBC\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_CBC\_SHA256
- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256

### To access SMTP sending information to configure the connection in your system

Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Configuration**. The following information is displayed:

- **Manage SMTP Credentials:** Access your user credentials. Use the SMTP user credentials (in plain text) when validating your connection.
- **Server Name:** Regional SMTP endpoint
- **Port:** Email Delivery supports TLS on port 25 or 587.
- **Use Transport Layer Security (TLS):** This field indicates if TLS, the standard means of performing encryption in transit for email, is being used. Customers must encrypt

email while it is in transit to the Oracle Cloud Infrastructure Email Delivery service. Encrypted emails are protected from being read during transit.



### Tip

Java applications (including JavaMail) must be updated to the latest version to ensure the latest protocols, ciphers, and security patches are in compliance with Oracle's supported security policies and ciphers.

### Begin sending email.

Use Email Delivery to begin sending email.

### Suppression List

As you begin to send email, Email Delivery automatically adds email addresses with bounce codes showing permanent failures or user complaints to the suppression list to protect your sender reputation. Email Delivery will not send any messages to these recipients in the future. Reasons for suppression currently include:

- Complaints
- Hard bounces
- Repetitive soft bounces
- Manual entries
- List-unsubscribe requests

### To manually add an email address to the suppression list using the Console

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and

## CHAPTER 13 Email Delivery

---

click **Email Suppression List**.

2. Click **Add Suppression**.
3. In the Add Suppression dialog box, enter the email address.
4. Click **Add**. The email address is added to the suppression list.

For more information, see [Managing the Suppression List](#).

### To manually add an email address to the suppression list using the API

The following example shows how to add an email address to the suppression list. For more information about managing the suppressions list, see [GetSuppression](#) and [DeleteSuppression](#).

```
POST /20170907/suppressions
```

```
{
 "compartmentId":
"ocidl.compartment.ocl1..aaaaaaaaat7uqcb6zoxvzoga4d4vh4dtweciavepacd3skz56atf3qp73d7fx",
 "emailAddress": "user@example.com",
}
```

### Using the API

You can access Oracle Cloud Infrastructure using the [REST API](#). Instructions for the API are included in topics throughout this guide. For a list of available SDKs, see [SDKs and Other Tools](#).

### Regions

See [Regions and Availability Domains](#) for information on regions Email Delivery is available in.

### Limits

See [Email Delivery Service Capabilities and Limits](#) for information on new account and enterprise account limits.

### Best Practices

This section describes best practices for using Email Delivery.

**Volume Testing** - In order to maintain our sender reputation and yours, testing at volume needs to be done using the following best practice.

- Use a recipient address at the [email-blackhole.com](#) domain, such as [example@email-blackhole.com](#). Email Delivery will accept the mail but will not deliver it to an inbox.
- If large volume emails are sent to valid email addresses, these will get rejected by receivers and will result in a large amount of hard bounces. This will negatively affect IP reputation. For testing bounce processing, send small amounts of emails to a domain that does not have an MX record, in other words, the domain does not exist.

**Deliverability** - To help you learn and manage the habits that affect your sending reputation, see [Deliverability Best Practices](#).

**Sending to Email Aliases** - When sending email to an alias, the alias is considered one recipient. When sending email to a distribution group or list set up in an email client such as Apple Mail or Outlook, a separate email is sent for each recipient in the group.

## Generate SMTP Credentials for a User

Simple Mail Transfer Protocol (SMTP) credentials are necessary to send email through Email Delivery. Each user is limited to a maximum of two SMTP credentials. If more than two are required, SMTP credentials must be generated on other existing users or more users must be created.

A security best practice is to generate SMTP credentials for a new user instead of your Console user that already has permissions assigned to it. For detailed instructions on creating

a user, see [Adding Users](#). The new user must be assigned to a group with permissions to manage approved-senders and suppressions. For example:

```
Allow group <group name> to use approved-senders in compartment <compartment name>
```

### Using the Console

#### To generate SMTP credentials for a user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. Locate the user in the list that has permissions to manage email, and then click the user's name to view the details.



#### Tip

If your user does not have permissions to view or create users, you can create SMTP credentials under your user. Open the **Profile** menu (👤) and click **User Settings**.

2. Click **SMTP Credentials**.
3. Click **Generate SMTP Credentials**.
4. Enter a **Description** of the SMTP Credentials in the dialog box.
5. Click **Generate SMTP Credentials**. A user name and password is displayed.
6. Copy the user name and password for your records and click **Close**.

### Managing Approved Senders

You must set up an approved sender for all "From:" addresses sending mail via Oracle Cloud Infrastructure or mail will be rejected. An approved sender is associated with a compartment and only exists in the region where the approved sender was configured. That is, if you create

an approved sender in the US West (Phoenix) region, you cannot send email through the US East (Ashburn) region.

Approved senders should not be created in the root compartment. If approved senders exist in the root compartment, you are required to create a policy to manage approved senders in the entire tenant. Creating approved senders in a compartment other than the root allows the policy to be specific to that compartment.

### Moving Approved Senders to a Different Compartment

You can move approved senders from one compartment to another. To manage approved senders and use approved senders to send mail, user groups must have an associated identity policy in the new compartment. For more information, see [Managing Compartments](#).

### Using the Console

#### To create an approved sender

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**. Ensure that you are in the correct compartment. Your user must be in a group with permissions to manage `approved-senders` in this compartment.
2. Click **Create Approved Sender** within the **Approved Senders** view.
3. Enter the email address you want to list as an approved sender in the **Create Approved Sender** dialog box.

**Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

4. Click **Create Approved Sender**. The email address is added to your Approved Senders list.



### Tip

Approved senders are unique to tenancies. If an attempt is made to create a duplicate approved sender within a tenancy, the service will return a 409 Conflict error.

### To delete an approved sender

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**.
2. Find the approved sender you're interested in, click the Actions icon (three dots), and then click **Delete**.
3. In the confirmation dialog box, click **Confirm**. The email address is removed from the Approved Senders list.

### To move an approved sender to a different compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**.
2. In the **List Scope** section, select a compartment.
3. Find the approved sender in the list, click the the Actions icon (three dots), and then click **Choose New Compartment**.
4. Choose the destination compartment from the list.
5. Click **Move Approved Sender**.

For more information, see [Managing Compartments](#).

### To manage tags for an approved sender

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**.
2. Find the approved sender you're interested in, click the Actions icon (three dots), and then click **View Tags** to view or edit existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage your approved senders:

- [CreateSender](#)
- [GetSender](#)
- [ListSenders](#)
- [DeleteSender](#)

### Configure SPF

The Approved Senders section within the Console provides validation of an SPF record for each of your approved senders.

### Using the Console

#### To configure SPF

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Approved Senders**.
2. Select the checkbox for the approved sender you want to view SPF details for and click **View SPF**.



#### Tip

You can search for an approved sender by using the Search field. Addresses can be sorted alphanumerically or by creation date in ascending or descending order.

3. The Manage SPF dialog box appears indicating whether an SPF record for the approved sender exists.
  - If your domain does not currently have an SPF record, the information necessary to add an SPF record in your DNS setup is displayed. See [Managing DNS Service Zones](#) for instructions on adding a zone record in Oracle Cloud Infrastructure. If your DNS setup resides with another provider, please reference their documentation for adding a TXT record to your domain.
    - In your DNS setup, create a TXT record and paste the following information from the dialog box into the record: `v=spf1 include:spf.oracleemaildelivery.com -all`
  - If your domain currently has an SPF record, add the following information to the record to add Oracle Cloud Infrastructure Email Delivery:  
`include:spf.oracleemaildelivery.com`

## Configure SMTP Connection

Set up and test your SMTP connection using an SMTP library or product, such as [Postfix](#) or [Sendmail](#), to send email through Oracle Cloud Infrastructure Email Delivery.

To access SMTP sending information to configure the connection in your system, open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Configuration**. The following information is displayed:

- **Manage SMTP Credentials:** Access your user credentials. Use the SMTP user credentials (in plain text) when validating your connection.
- **Server Name:** The Email Delivery service hostname.
- **Port:** Email Delivery supports TLS on port 25 or 587.
- **Use Transport Layer Security (TLS):** This field indicates if TLS, the standard means of performing encryption in transit for email, is being used. Customers must encrypt email while it is in transit to the Oracle Cloud Infrastructure Email Delivery service. Encrypted emails are protected from being read during transit.



### Important

Java applications (including JavaMail) must be updated to the latest version to ensure the latest protocols, ciphers, and security patches are in compliance with Oracle's supported security policies and ciphers.

## SMTP Connection Endpoints

Use the following regional endpoints for establishing SMTP connections for sending.

- PHX: smtp.us-phoenix-1.oraclecloud.com
- IAD: smtp.us-ashburn-1.oraclecloud.com

- LHR: smtp.email.uk-london-1.oci.oraclecloud.com

### TLS Requirements

Oracle maintains strict security policies and only accepts email traffic using Transport Layer Security (TLS). Use of TLS 1.2 is mandatory to send email using Oracle Cloud Infrastructure.

The approved TLS 1.2 ciphers are:

- TLS\_DHE\_RSA\_WITH\_AES\_256\_CBC\_SHA
- TLS\_DHE\_RSA\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_DHE\_DSS\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA
- TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA256
- TLS\_DHE\_RSA\_WITH\_AES\_128\_CBC\_SHA
- TLS\_DHE\_DSS\_WITH\_AES\_128\_CBC\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_256\_GCM\_SHA384
- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256
- TLS\_ECDHE\_ECDSA\_WITH\_AES\_128\_CBC\_SHA256
- TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256

### Managing the Suppression List

Manually add an email address to the suppression list to prevent it from being part of your sending list.

Users are required to have correct permissions to manage the suppression list. Currently, identity policies for suppression must be at the tenant level (not at the compartment level). The following is an example of the permission policy statement.

```
Allow group <group name> to manage suppressions in tenancy
```

Suppressions are stored at the tenancy level. Therefore any request requiring a `compartmentId` must provide the `tenancyId` as the `compartmentId`. For example:

```
Allow group <ordinary users> to inspect approved-senders in tenancy
Allow group <power users> to read approved-senders in tenancy
Allow group <sender admins> to manage approved-senders in tenancy
Allow user <mail user> to use approved-senders in tenancy where target.approved-sender.senderId =
<senderId>
```

```
Allow group <ordinary users> to inspect suppressions in tenancy
Allow group <power users> to read suppressions in tenancy
Allow group <sender admins> to manage suppressions in tenancy
```

### Using the Console

To manually add an email address to the suppression list

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Suppression List**.
2. Click **Add Suppression**.
3. In the Add Suppression dialog box, enter the email address.
4. Click **Add**. The email address is added to the Suppression List.

To delete an email address from the suppression list

1. Open the navigation menu. Under **Solutions and Platform**, go to **Email Delivery** and click **Email Suppression List**.
2. Select the checkbox for the email address you want to delete and then click **Delete**.



### Tip

You can search for an email address by using the Search field. Addresses can be sorted alphanumerically or by creation date in ascending or descending order.

3. In the confirmation dialog box, click **OK**. The email address is removed from the Suppression List.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage your suppressions:

- [CreateSuppression](#)
- [GetSuppression](#)
- [ListSuppressions](#)
- [DeleteSuppression](#)

## Email Delivery Metrics

You can monitor the health, capacity, and performance of your Email Delivery by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_emaildelivery` (the Email Delivery service).

### Overview of the Email Delivery Service Metrics

Oracle Cloud Infrastructure Email Delivery (Email Delivery) is an email sending service that provides a fast and reliable managed solution for sending high-volume emails that need to reach your recipients' inbox. The Email Delivery service metrics help you measure counts for accepted mail, which consists of the unique emails accepted by the Email Delivery service to send. Emails are defined by the number of unique emails, as well as the number of unique recipients per message attempted to be delivered, resulting in successful delivery and blocked email. For example, sending an email with 10 recipients means 10 emails accepted.



#### Note

Email Delivery is billed for every 1,000 emails accepted.

For more information, see [Overview of the Email Delivery Service](#).

### Prerequisites

- IAM policies: To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_emaildelivery

The metrics listed in the following table are automatically available for any policies you create. You do not need to enable monitoring on the resource to get these metrics. However,

## CHAPTER 13 Email Delivery

---

your tenancy must have Email Delivery configured and must send mail to make the `oci_emaildelivery` metric space available in the Metrics Explorer feature.

Each metric includes the following dimensions:

### **RESOURCEID**

The OCID of the policy to which the metric applies.

<b>Metric</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Description</b>	<b>Dimensions</b>
EmailsAccepted	<b>EmailsAccepted</b>	count	The number of unique emails accepted by the Email Delivery service.	resourceID

## Using the Console

Email Delivery service metrics are currently only available using the Metrics Explorer feature in the Console. For more information about metrics, see [Viewing Metric Charts](#).

### To view Email Delivery metric charts

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.  
For **Metric Namespace**, select `oci_emaildelivery`.
2. Select a metric to view from the **Metric Name** field.
3. Select a qualifier specified in the **Dimension Name** field. For example, the dimension `resourceId` is specified in the metric definition for `EmailsAccepted`.
4. Select the value you want to use for the specified dimension in the **Dimension Value** field. For example, the resource identifier for your instance of interest.
5. Click **Update Chart**.

The chart will be updated with the metrics that have been requested.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

## Integrating Oracle Application Express with Email Delivery

### Configure Oracle Application Express to Send Email Through Email Delivery

You can use the `APEX_MAIL` package to send emails from Oracle Application Express applications deployed in Autonomous Transaction Processing. See [Creating an Autonomous Database](#) and [Autonomous Transaction Processing](#) for more information.

Before you use `APEX_MAIL` you must configure Oracle Cloud Infrastructure Email Delivery in your Application Express instance.

To enable `APEX_MAIL` functionality in your Application Express instance in Autonomous Transaction Processing:

1. Identify the SMTP connection endpoint for Email Delivery. You configure the endpoint as the SMTP Host in your Application Express instance in Step 4. See [Configure SMTP Connection](#) for more information.
2. Generate SMTP credentials for Email Delivery. Your Application Express instance uses credentials to authenticate with Email Delivery servers when you send email. See [Generate SMTP Credentials for a User](#) for more information.

3. Create an approved sender for Email Delivery. You need to complete this step for all email addresses you use as the "From" with `APEX_MAIL.SEND` calls, as the Application Email From Address in your apps, or in the `SMTP_FROM` instance parameter. See [Managing Approved Senders](#) for more information.
4. Connect to your Autonomous Transaction Processing as ADMIN user using a SQL client and configure the following SMTP parameters using `APEX_INSTANCE_ADMIN.SET_PARAMETER`:
  - `SMTP_HOST_ADDRESS`: Specifies the SMTP connection endpoint from Step 1.
  - `SMTP_USERNAME` Specifies the SMTP credential user name from Step 2.
  - `SMTP_PASSWORD` Specifies the SMTP credential password from Step 2.

For example:

```
BEGIN
APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_HOST_ADDRESS', 'smtp.us-phoenix-1.oraclecloud.com');
APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_USERNAME', 'ocid1.user.oc1.username');
APEX_INSTANCE_ADMIN.SET_PARAMETER('SMTP_PASSWORD', 'password');
COMMIT;
END;
/
```

5. Send a test email using APEX SQL Workshop, SQL Commands specifying one of the approved senders from Step 3 as "From". For example:

```
BEGIN
APEX_MAIL.SEND(p_from => 'alice@example.com',
 p_to => 'bob@example.com',
 p_subj => 'Email from Oracle Autonomous Database',
 p_body => 'Sent using APEX_MAIL');
END;
/
```

6. To monitor email delivery in your Application Express instance:
  - a. Sign in to APEX Administration Services.
  - b. Open the Manage Instance page.
  - c. Click the Mail Queue link in the Manage Meta Data section.

Alternatively, query `APEX_MAIL_QUEUE` and `APEX_MAIL_LOG` views using a SQL client.

### More Information

- [Creating Applications with Oracle Application Express in Autonomous Database](#)
- [APEX\\_MAIL](#) in Oracle Application Express API Reference
- [APEX\\_INSTANCE\\_ADMIN](#) in Oracle Application Express API Reference

## Integrating Postfix with Email Delivery

### Configure Postfix to Send Email Through Email Delivery

You can use Postfix to send emails through Email Delivery. Before you use Postfix you must configure Oracle Cloud Infrastructure Email Delivery in your Postfix application.



#### Note

The paths and commands used below for specifying file locations are specific to Ubuntu/Debian; your file paths or editing commands may differ depending on the operating system you are using. The changes to the configuration files are the same.

To enable Postfix to integrate with Email Delivery:

1. Make sure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

The SMTP credentials are required to configure Postfix to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Update the Postfix `main.cf` file.

To open the `main.cf` file, run the following command:

```
sudo vi /etc/postfix/main.cf
```

Add the following information to the end of the file:

```
smtp_tls_security_level = may
smtp_sasl_auth_enable = yes
smtp_sasl_password_maps = hash:/etc/postfix/sasl_passwd
smtp_sasl_security_options =
```

If the following line is present, either remove the line or turn it off:

```
smtpd_use_tls = yes
```

3. Update `relayhost` to include your SMTP connection endpoint and port. For example:

```
relayhost = smtp.us-ashburn-1.oraclecloud.com:587
```

4. Create the `sasl_passwd` file in the same directory as `main.cf`.

Run the following command:

```
sudo vi /etc/postfix/sasl_passwd
```

5. Add your relay host and port by entering:

```
server:port user:pass
```

where:

- `server` is your relay host and `port` is 25 or 587.
- `user` is the user name and `pass` is the password you received when you generated your SMTP credentials.

6. Enter the permissions in the password file.

Run the following command:

```
sudo chown root:root /etc/postfix/sasl_passwd && sudo chmod 600 /etc/postfix/sasl_passwd
```

7. Generate the password hash.

Run the following command:

```
sudo postmap hash:/etc/postfix/sasl_passwd
```

8. Reload Postfix.

Run the following command:

```
sudo postfix reload
```

9. Test the configuration by sending a test email.

Run the following command:

```
echo "This is a test message" | mail -s "Test" -r "<approved sender email address>" <recipient email address>
```

If you want to monitor the log while you send the test email, open a separate Terminal window and run the following command before running the test command:

```
log stream --predicate '(process == "smtpd") || (process == "smtp")' --info
```

A `status=sent (250 Ok)` message in the log indicates the email was sent successfully.



### Note

If you are using SASL authentication, you must use the following RPM package: `cyrus-sasl-plain`. See the [PostFix website](#) for further documentation on configuring SASL authentication.

### More Information

- See the [Postfix website](#) for more information on Postfix configuration.
- See [TLS errors when integrating with Postfix](#) for troubleshooting techniques related to Email Delivery.

## Integrating Oracle Enterprise Manager with Email Delivery

### Configure Oracle Enterprise Manager to Send Email Through Email Delivery

You can use Oracle Enterprise Manager to send emails through Email Delivery. Before you use Oracle Enterprise Manager, you must configure Oracle Cloud Infrastructure Email Delivery in your Oracle Enterprise Manager application.



#### Note

For information on installing Oracle Enterprise Manager, see [Setting Up Oracle Enterprise Manager on Oracle Cloud Infrastructure](#).

To enable Oracle Enterprise Manager to integrate with Email Delivery:

1. Make sure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

The SMTP credentials are required to configure Oracle Enterprise Manager to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. In Oracle Enterprise Manager, go to the **Setup** menu and click **Initial Setup Console**.
3. In the **Initial Setup Console** section, click **Configure Mail Servers** in the navigation pane.
4. In the **Sender Identify** section, click **Edit**.
5. Enter the name of the administrator or system that should send the email notifications and the email address from which the notifications should be sent, and then click **OK**.
6. In the **Outgoing Mail (SMTP) Servers** section, click **Create**.
7. Enter the mail server host name, the mail server credentials, and the encryption method to be used, and then click **OK**.
8. Select the outgoing mail server you wish to test and select **Test Mail Server**. Note the confirmation message in the console and verify that you received the test email in your inbox.



### Note

If you configure multiple outgoing mail servers, automatic failover and load balancing is performed in round robin fashion.

## Integrating Mailx with Email Delivery

### Configure Mailx to Send Email Through Email Delivery

You can use Mailx to send emails through Email Delivery. Before you use Mailx you must configure Oracle Cloud Infrastructure Email Delivery in your Mailx application.



#### Note

These steps assume you are logged into an Oracle Linux instance. Other distributions of Linux may have different commands and file locations.

To enable Mailx to integrate with Email Delivery:

1. Make sure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



#### Note

The SMTP credentials are required to configure Mailx to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Update the Mailx `mail.rc` file.

To open the `mail.rc` file, run the following command:

```
sudo vi ~/etc/mail.rc
```

Add the following information to the end of the file:

```
#smtp config

set nss-config-dir=/etc/pki/nssdb/
set smtp-use-starttls
```

## CHAPTER 13 Email Delivery

---

```
set smtp-auth=plain
set smtp=<SMTP connection endpoint>:25
set from=<from_email_address>
set smtp-auth-user=<OCID from smtp credentials>
set smtp-auth-password=<password from smtp credentials>

#write and quit file
:wq!
```

### 3. Test the configuration by sending a test email.

Run the following command:

```
echo "Test Email" | mail -v -s "Send an email via mailx" -r "from_name<from_email_address>" -S
replyto="from_name<>from_email_address>" -S smtp="SMTP connection endpoint:25" -S smtp-use-
starttls -S smtp-auth=plain -S smtp-auth-user='<ocid from smtp credentials>' -S smtp-auth-
password='<password from smtp credentials>' -S ssl-verify=ignore <recipient_email_address>
```

## More Information

- For network security services, see the [Mailx](#) documentation.

## Integrating Swaks with Email Delivery

Swaks (Swiss Army Knife SMTP) is a transaction-based tool you can use to test SMTP configurations in Email Delivery. Before you use Swaks, you must configure Email Delivery and take note of your SMTP sending information and SMTP credentials.



### Note

Many options and parameters can be used to test various scenarios with Swaks. When Swaks evaluates an option (that is, a flag with parameters), it does so in three steps:

- First, it looks for a configuration file (default location or specified with `--config`).
- Next, it looks for options in environment variables.
- Finally, it looks at command line options. At each step, any options set earlier are overridden.

## Assumptions

The following procedures assume the following:

- The following example supplies options to Swaks via the command line in long form, for example, `--server` as opposed to the short form, `-s`.
- The following example assumes the default behavior to connect through network sockets.
- A local certificate is not required for a TLS connection to be negotiated. The following example assumes the default behavior where Swaks does not attempt certificate verification.
- Swaks is primarily intended for use on UNIX-like operating systems with functionality based on known standards so it should work on most modern mail servers.

## Configure Swaks to Send Email Through Email Delivery

To enable Swaks to test the configuration of Email Delivery:

1. Ensure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

The SMTP credentials are required to configure Swaks to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Ensure Swaks is installed. The installation process differs depending on which operating system you are using. For example, run the following command to install Swaks on Oracle Linux:

```
sudo yum install swaks -y
```

3. To send a test email with Swaks, run the following command:

```
swaks --pipeline -tls --server <smtp.region.oraclecloud.com> --port <587 or 25> --auth-user '<username OCID from SMTP credentials>' --auth-pass '<password>' --from '<sender email address>' --to '<recipient email address>' --data '<email message>'
```

For example:

```
swaks --pipeline -tls --server smtp.us-ashburn-1.oraclecloud.com --port 25 --auth-user 'ocid1.user.oc1..<unique_ID>' --auth-pass '<password>' --from 'sender@example.com' --to 'recipient@example.com' --data 'From: sender@example.com\nDate: Thu, 13 Sep 2019\nSubject: Test Send\n\nTest email'
```

Note the following when sending email with Swaks:

- The `-tls` parameter is required.
- The `--pipeline` parameter is supported to make use of SMTP pipelining.
- The `--port <number>` parameter or `:<port number>` syntax can be used to specify the port.

## More Information

- See the [Swaks documentation](#) for more information.

## Integrating JavaMail with Email Delivery

JavaMail provides a platform-independent and protocol-independent framework to build mail and messaging applications. Before you use JavaMail, you must configure Email Delivery and take note of your SMTP sending information and SMTP credentials. This guide uses the Eclipse IDE and the JavaMail API to send email through Email Delivery.



### Important

Java applications (including JavaMail) must be updated to the latest version to ensure that the latest protocols, ciphers, and security patches are in compliance with Oracle's supported security policies and ciphers.

## Configure JavaMail to Send Email Through Email Delivery

To enable JavaMail to test the configuration of Email Delivery:

1. Ensure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

The SMTP credentials are required to configure JavaMail to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Open a browser and go to <https://github.com/javaee/javamail/releases>.
3. Under **Downloads**, select **javax.mail.jar** to download the latest version of JavaMail.

4. Create a project in Eclipse by performing the following steps:
  - a. In Eclipse, open the **File** menu. Select **New**, and then click **Java Project**.
  - b. In the Create a Java Project dialog box, enter a project name, and then click **Next**.
  - c. In the Java Settings dialog box, select the **Libraries** tab.
  - d. Click **Add External JARs**.
  - e. In the JAR Selection dialog box, browse to the folder in which you downloaded JavaMail. Select the **javax.mail.jar** file, and then click **Open**.
  - f. In the Java Settings dialog box, click **Finish**.
5. In Eclipse, in the Package Explorer window, expand your project.
6. Under your project, right-click the src directory, select **New**, and then click **Class**.
7. In the New Java Class dialog box, enter "OCIemail" in the **Name** field and then click **Finish**.
8. Enter the following code in **OCIemail.java** to send a test email with JavaMail:

```
import java.util.Properties;
import javax.mail.Message;
import javax.mail.Session;
import javax.mail.Transport;
import javax.mail.internet.InternetAddress;
import javax.mail.internet.MimeMessage;

public class OCIemail {

 // Replace FROM with your "From" address.
 // This address must be added to Approved Senders in the console.
 static final String FROM = "<sender_email_address>";
 static final String FROMNAME = "<sender_name>";

 // Replace TO with a recipient address.
 static final String TO = "<recipient_email_address>";

 // Replace smtp_username with your Oracle Cloud Infrastructure SMTP username generated in
 console.
```

## CHAPTER 13 Email Delivery

---

```
static final String SMTP_USERNAME = "<username OCID from SMTP credentials>";

// Replace smtp_password with your Oracle Cloud Infrastructure SMTP password generated in
console.
static final String SMTP_PASSWORD = "<SMTP password>";

// Oracle Cloud Infrastructure Email Delivery hostname.
static final String HOST = "<SMTP endpoint>";

// The port you will connect to on the SMTP endpoint. Port 25 or 587 is allowed.
static final int PORT = 587;

static final String SUBJECT = "<subject of your email>";
static final String BODY = String.join(

 System.getProperty("line.separator"),
 "<h1>OCI Email Delivery test</h1>",
 "<p>This email was sent with OCI Email Delivery using the ",
 "https://github.com/javaee/javamail'>Javamail Package",
 " for Java."

);

public static void main(String[] args) throws Exception {

 // Create a Properties object to contain connection configuration information.

 Properties props = System.getProperties();
 props.put("mail.transport.protocol", "smtp");
 props.put("mail.smtp.port", PORT);

 //props.put("mail.smtp.ssl.enable", "true"); //the default value is false if not set
 props.put("mail.smtp.auth", "true");
 props.put("mail.smtp.auth.login.disable", "true"); //the default authorization order is
"LOGIN PLAIN DIGEST-MD5 NTLM". 'LOGIN' must be disabled since Email Delivery authorizes as
'PLAIN'
 props.put("mail.smtp.starttls.enable", "true"); //TLSv1.2 is required
 props.put("mail.smtp.starttls.required", "true"); //Oracle Cloud Infrastructure required

 // Create a Session object to represent a mail session with the specified properties.
```

## CHAPTER 13 Email Delivery

---

```
Session session = Session.getDefaultInstance(props);

// Create a message with the specified information.
MimeMessage msg = new MimeMessage(session);
msg.setFrom(new InternetAddress(FROM, FROMNAME));
msg.setRecipient(Message.RecipientType.TO, new InternetAddress(TO));
msg.setSubject(SUBJECT);
msg.setContent(BODY, "text/html");

// Create a transport.
Transport transport = session.getTransport();

// Send the message.

try
{
 System.out.println("Sending Email now...standby...");

 // Connect to OCI Email Delivery using the SMTP credentials specified.
 transport.connect(HOST, SMTP_USERNAME, SMTP_PASSWORD);

 // Send email.
 transport.sendMessage(msg, msg.getAllRecipients());
 System.out.println("Email sent!");
}

catch (Exception ex) {

 System.out.println("The email was not sent.");
 System.out.println("Error message: " + ex.getMessage());

}

finally

{
```

```
// Close & terminate the connection.
transport.close();

}

}

}
```

9. In the **OCIemail.java** file, replace the following with your own values:



### Note

Email addresses are case-sensitive. Ensure that the addresses are the same as the ones you entered in Approved Senders in the console.

- **FROM** - Replace with your sender email address. This email address must be added to the Approved Senders list in Email Delivery first.
  - **TO** - Replace with your recipient email address.
  - **SMTP credentials** - Replace smtp\_username and smtp\_password with your Oracle Cloud Infrastructure SMTP username and password generated in the console.
  - **HOST** - Replace with the Email Delivery SMTP endpoint. For example, smtp.us-ashburn-1.oraclecloud.com.
10. Refer to the requirements for [configuring an SMTP connection](#) with Email Delivery. TLSv1.2 is required for Email Delivery. Some default settings of [Javamail](#) need to be disabled. For example, JavaMail authorizes in a certain order. The default authorization order is "LOGIN PLAIN DIGEST-MD5 NTLM". Since Email Delivery authorizes as "PLAIN", "LOGIN" needs to be disabled. For example, the following code is entered in **OCIemail.java** file to configure the SMTP connection:

```
//props.put("mail.smtp.ssl.enable", "true"); //default is false if not set
props.put("mail.smtp.auth", "true");
props.put("mail.smtp.auth.login.disable", "true");
props.put("mail.smtp.starttls.enable", "true");
props.put("mail.smtp.starttls.required", "true");
```

11. Open the **File** menu and click **Save**.
12. To build the project, open the **Project** menu and then select **Build Project**. If this option is disabled, you may have automatic building enabled.
13. To start the program and send the email, open the **Run** menu and then click **Run**.
14. Review the output. If the email was successfully sent, the console displays "Email sent successfully!" Otherwise, it displays an error message.
15. Log into the recipient inbox to verify receipt of the email.

### More Information

- See the [JavaMail](#) documentation for more information.

## Integrating Sendmail with Email Delivery

### Configure Sendmail to Send Email Through Email Delivery

You can use Sendmail to send emails through Email Delivery. Before you use Sendmail you must configure Oracle Cloud Infrastructure Email Delivery in your Sendmail application.



#### Note

The steps below are for configuring Sendmail to send email via Oracle Cloud Infrastructure Email Delivery. These steps were tested on an Ubuntu 18.04 compute instance.

To enable Sendmail to integrate with Email Delivery:

1. Make sure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

SMTP credentials are required to configure Sendmail to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Run the following update and install commands:

```
sudo apt update
sudo apt install sendmail
sudo apt install m4
```

3. In a file editor such as `vi`, update `/etc/mail/authinfo`.  
Run the following command:

```
sudo vi /etc/mail/authinfo
```



### Note

If `/etc/mail/authinfo` doesn't exist, you can create it by running the command `sudo vi /etc/mail/authinfo`.

Add the following line:

```
AuthInfo:<SMTP connection endpoint> "U:root" "I:<username from smtp credentials>" "P:<password
from smtp credentials>" "M:PLAIN"

#write and quit file
:wq!
```

4. Generate the `/etc/mail/authinfo.db` file.

## CHAPTER 13 Email Delivery

---

Run the following command:

```
sudo sh -c 'makemap hash /etc/mail/authinfo.db < /etc/mail/authinfo'
```

5. Add support for relaying to the Oracle Cloud Infrastructure Email Delivery SMTP endpoint.

Run the following command:

```
sudo sh -c 'echo "Connect:<SMTP connection endpoint> RELAY" >> /etc/mail/access'
```

6. Regenerate `/etc/mail/access.db`.

Run the following command:

```
sudo sh -c 'makemap hash /etc/mail/access.db < /etc/mail/access'
```

7. Create a backup of the `sendmail.cf` and `sendmail.mc` files.

Run the following command:

```
sudo sh -c 'cp /etc/mail/sendmail.cf /etc/mail/sendmail_cf.backup && cp /etc/mail/sendmail.mc /etc/mail/sendmail_mc.backup'
```

8. Update the `/etc/mail/sendmail.mc` file.

Run the following command:

```
sudo vi /etc/mail/sendmail.mc
```

Find the `MAILER()` definitions.

Type `/MAILER` and press `ENTER`.

In Insert mode, add the following settings before any `MAILER()` definitions:

```
define(`SMART_HOST', `<SMTP connection endpoint>')dnl
define(`RELAY_MAILER_ARGS', `TCP $h 25')dnl
define(`confAUTH_MECHANISMS', `LOGIN PLAIN')dnl
FEATURE(`authinfo', `hash -o /etc/mail/authinfo.db')dnl
MASQUERADE_AS(`<sending_domain>')dnl
FEATURE(masquerade_envelope)dnl
FEATURE(masquerade_entire_domain)dnl
```

Disable Insert mode.

Run the following command:

```
#write and quit file
:wq!
```

## CHAPTER 13 Email Delivery

---

9. Make Sendmail writeable.

Run the following command:

```
sudo chmod 666 /etc/mail/sendmail.cf
```

10. Regenerate `sendmail.cf`.

Run the following command:

```
sudo sh -c 'm4 /etc/mail/sendmail.mc > /etc/mail/sendmail.cf'
```



### Note

If you receive an error, such as "Command not found" or "No such file or directory," confirm that the `m4` and `sendmail` packages are installed on your system.

11. Reset permissions for `sendmail.cf` to read only.

Run the following command:

```
sudo chmod 644 /etc/mail/sendmail.cf
```

12. Restart Sendmail.

Run the following command:

```
sudo /etc/init.d/sendmail restart
```

13. Test the configuration by sending a test email.

Run the following command:

```
/usr/sbin/sendmail -vf <from_email_address> <recipient_email_address>
```

Enter the details of the email. After each line press `Enter`.

For example:

```
From: <from_email_address>
To: <recipient_email_address>
Subject: OCI Email Delivery test email
This is a test message sent from OCI Email Delivery using Sendmail.
```

Press `Ctrl + D` to send the email.

14. Verify receipt of the test email.



### Note

You can troubleshoot an issue by reviewing the Sendmail log on your mail server, located at `/var/log/mail.log`.

## More Information

- For more information, see the [Sendmail Installation and Operation Guide](#).

## Integrating PeopleSoft with Email Delivery

### Configure PeopleSoft to Send Email Through Email Delivery

You can use PeopleSoft to send emails through Email Delivery. Before you use PeopleSoft, you must configure Oracle Cloud Infrastructure Email Delivery in your PeopleSoft application.



### Note

The following steps require familiarity with [PeopleSoft documentation](#). (This link points to PeopleTools 8.57, which is the latest version at the time this article was published. Please refer to the documentation for your specific PeopleTools version.)

To enable PeopleSoft to integrate with Email Delivery:

1. Make sure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

SMTP credentials are required to configure PeopleSoft to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Open a Chrome browser and navigate to the SMTP connection endpoint (for example, <https://smtp.us-phoenix-1.oraclecloud.com>).
  - a. Click the certificate, and then click the **Certification Path** tab.
  - b. Select the **DigiCert** root certificate.
  - c. Click **View Certificate**.
  - d. Click the **Details** tab for the Digitrust certificate, and then select **Copy to File**.
  - e. Select **Base-64 X-509**.
  - f. Save the certificate.
3. Repeat the steps above for the **Digicert SHA2 Secure Server CA** intermediate certificate.
4. Add the certificates to the PeopleSoft application.

Log into the Pure Internet Architecture (PIA) as user "PS" and import the certificates into the target environment. See [Installing Application Server-Based Digital Certificates](#) and refer to the **Adding CA Authorities and Installing Root Certificates** section.
5. Encrypt the SMTP password in the config file.

You can encrypt the SMTP password using the PIA or the PSCipher utility.

### Using the PIA

- a. Open the navigation menu on the PeopleSoft dashboard. Go to **PeopleTools**, and then select **Integration Broker**.

- b. Select **Configuration**, and then click **Gateways**. Select the default LOCAL gateway.
- c. Click **Gateway Setup Properties**. The default user ID is administrator and the default password is the password selected during setup.
- d. Click the **Advanced Properties Page** link.
- e. Click **Password Encryption** at the bottom of the page. This is where you will encrypt your password.

### Using PSCipher

The PSCipher utility can be found under `$PS_CFG_HOME/webserv/<DOMAIN>/piabin` where `<DOMAIN>` is your web server domain.

Run the following command:

```
./PSCipher.sh <password>
```



#### Note

The password can have special characters so you will need to enclose the password in single quotes. For example:

```
./PSCipher.sh '#rpassword$) {'
```

6. Update the SMTP settings on the PeopleSoft Application server. For more information, see [SMTP Settings](#) in the PeopleSoft documentation.  
Establish an ssh connection to the PeopleSoft Application server machine (as username "opc") and do the following:
  - a. Switch user to `psadm2` (for example, `sudo su - psadm2`).



### Note

psadm2 is the PeopleTools domain user who creates and configures the Application Server domain.

- b. Navigate to the Appserver configuration directory.

Run the following command:

```
$ cd $PS_CFG_HOME/appserv/APPDOM
```

- c. Back up the original psappsrv.cfg file.
- d. Add the following information to the psappsrv.cfg file:

```
SMTPServer=<SMTP connection endpoint>
SMTPUserName=<username from SMTP credentials>
SMTPUserPassword=<encrypted SMTP password>
SMTPPort=587
SMTPUseSSL=N
SMTPSSLPort=587
SMTPTLSEnable=true
SMTPTLSRequired=true
```



### Note

Do not include a space between the "=" and the values because the space could be counted in the value for the password, causing an authentication failure.

7. Add the primary email address for the PeopleSoft application user who is trying to send notification from within the application. In this example, the user is "PS".  
Log in as "PS" and do the following:
  - a. Open the navigation menu on the PeopleSoft dashboard. Go to **PeopleTools**, and then select **Security**.
  - b. Select **User Profiles**, and then click **user Profiles**. Find the profile for "PS".

- c. On the **General** tab, click **Edit Email Addresses**.
  - d. Enter the approved sender email address as the primary email address.
8. Log out of the PeopleSoft application.
  9. Reboot the application server using the `PSADMIN` utility. See [Using the Application Server Administration Menu](#).
  10. Test the email notification delivery.  
Log into the PIA as "PS", and select **Notify** anywhere in the console. For example, you can do the following:
    - a. Go to **Peopletools**, and then select **Web Profile**.
    - b. Select **Web Profile Configuration**.
    - c. Click **Search**, and then click **PROD** in the search results.
    - d. Click **Notify**, enter the notification details, and then click **OK**.

Confirm receipt of the test email.

To debug SMTP errors (optional):

1. You can add the following parameter to help with SMTP debugging: `SMTPTrace=1`  
LogFence should be set to 5 to use this parameter. The system writes the log information to `SMTP<DDMM>.log` in `%PS_SERVDIR%/LOGS` by default, or the custom value set for Log Directory.

For example:

```
$PS_CFG_HOME\APPSERV\domain\LOGS\SMTP6_27.log
```

2. After you set this parameter, you will need to reboot the Application server. Once this parameter is set, you can monitor the SMTP log.
3. Type `ls` and find the SMTP file for the date you sent the email.
4. Run the following command:

```
tail -f <smtp log file and date>
```

For example,

## CHAPTER 13 Email Delivery

---

```
tail -f SMTP9_17.log
```

Search for any errors in the output.

### More Information

- [SMTP Settings \(PeopleSoft\)](#)
- [Encrypting Passwords in the PeopleSoft Pure Internet Architecture](#)
- My Oracle Support: [Is There a Way to Both Authenticate And Secure Emails From PeopleSoft?](#)

## Integrating Python with Email Delivery

You can use Python to send emails through Email Delivery. Before you can send email you must configure Email Delivery in Python.



### Note

These steps assume you are logged into an Oracle Linux instance. Other distributions of Linux may have different commands and file locations.

### Configure Python to Send Email Through Email Delivery

To enable Python to test the configuration of Email Delivery:

1. Ensure Email Delivery is configured to send email. See [Getting Started with Email Delivery](#).



### Note

The SMTP credentials are required to configure Python to use Email Delivery. Be sure to note the user name and password when you generate the SMTP credentials.

2. Ensure Python is installed. The installation process differs depending on which operating system you are using. For example, run the following command to install Python on Oracle Linux:

```
sudo yum install python3 -y
```

3. In a file editor such as vi, create a python script to test Email Delivery. Run the following command:

```
sudo vi ociemail.py
```

4. In the **ociemail.py** file, replace the variables with your own values. For example:

```

#python script for sending SMTP configuration with Oracle Cloud Infrastructure Email Delivery
import smtplib
import email.utils
from email.mime.multipart import MIMEMultipart
from email.mime.text import MIMEText

Replace sender@example.com with your "From" address.
This address must be verified.
SENDER = 'sender@example.com'
SENDERNAME = 'Sender Name'

Replace recipient@example.com with a "To" address. If your account
is still in the sandbox, this address must be verified.
RECIPIENT = 'recipient@example.com'

Replace the USERNAME_SMTP value with your Email Delivery SMTP username.
USERNAME_SMTP = 'ocidl.user.oc1..<unique_ID>@ocidl.tenancy.oc1..<unique_ID>.vf.com'
```

## CHAPTER 13 Email Delivery

---

```
Replace the PASSWORD_SMTP value with your Email Delivery SMTP password.
PASSWORD_SMTP = '<password>'

If you're using Email Delivery in a different region, replace the HOST value with an SMTP
endpoint. Use port 25 or 587 to connect to the SMTP endpoint.
HOST = "smtp.us-ashburn-1.oraclecloud.com"
PORT = 587

The subject line of the email.
SUBJECT = 'Email Delivery Test (Python smtplib)'

The email body for recipients with non-HTML email clients.
BODY_TEXT = ("Email Delivery Test\r\n"
 "This email was sent through the Email Delivery SMTP "
 "Interface using the Python smtplib package."
)

The HTML body of the email.
BODY_HTML = """<html>
<head></head>
<body>
 <h1>Email Delivery SMTP Email Test</h1>
 <p>This email was sent with Email Delivery using the
 Python

 smtplib library.</p>
</body>
</html>"""

Create message container - the correct MIME type is multipart/alternative.
msg = MIMEMultipart('alternative')
msg['Subject'] = SUBJECT
msg['From'] = email.utils.formataddr((SENDERNAME, SENDER))
msg['To'] = RECIPIENT

Record the MIME types of both parts - text/plain and text/html.
part1 = MIMEText(BODY_TEXT, 'plain')
part2 = MIMEText(BODY_HTML, 'html')
```

## CHAPTER 13 Email Delivery

---

```
Attach parts into message container.
According to RFC 2046, the last part of a multipart message, in this case
the HTML message, is best and preferred.
msg.attach(part1)
msg.attach(part2)

Try to send the message.
try:
 server = smtplib.SMTP(HOST, PORT)
 server.ehlo()
 server.starttls()
 #smtplib docs recommend calling ehlo() before & after starttls()
 server.ehlo()
 server.login(USERNAME_SMTP, PASSWORD_SMTP)
 server.sendmail(SENDER, RECIPIENT, msg.as_string())
 server.close()
Display an error message if something goes wrong.
except Exception as e:
 print ("Error: ", e)
else:
 print ("Email successfully sent!")
```

5. To send a test email with Python, run the following command from the directory the script is located in:

```
python3 ociemail.py
```

### More Information

- More Python script examples can be found on [GitHub](#).

## Troubleshooting Email Delivery

This topic provides troubleshooting solutions for problems you might encounter using Email Delivery.

### TLS errors when integrating with Postfix

- If you are encountering TLS errors when attempting to integrate Postfix with Email Delivery, ensure that the following setting is removed from the Postfix `main.cf` file, as it has been deprecated:

```
smtp_use_tls = yes
```

- Use the following setting instead to turn on TLS:

```
smtp_tls_security_level = may
```

Using this setting, the Postfix SMTP server announces STARTTLS support to remote SMTP clients, but does not require that clients use TLS encryption.

- If you want to enforce the use of TLS, so that the Postfix SMTP server announces STARTTLS and accepts no mail without TLS encryption, use the following setting:

```
smtp_tls_security_level = encrypt
```

For more information, see [Postfix TLS Support](#).

### Connectivity Issues



#### Note

Email Delivery does not prohibit connectivity from any source IP range. Any IP that attempts to connect to Email Delivery will be accepted.

Refer to [SMTP Connection Endpoints](#) for a list of regional endpoints to establish SMTP connections for sending.

### To troubleshoot a problem connecting to endpoint network ports

- Ensure that you have the correct endpoint DNS name or IP address for the region and

that you have been whitelisted to use the endpoint.

- Ping the endpoint to ensure that you can reach the endpoint.
  1. Open a command prompt.
  2. Use the following command to ensure you can reach the endpoint.

```
ping <SMTP endpoint>
```

For example, `ping smtp.us-ashburn-1.oraclecloud.com`

If you are unable to ping the endpoint successfully, you are experiencing a network connectivity issue. If you are able to ping the SMTP endpoint, you will now need to test connectivity on ports 25 or 587.

- Test connectivity to the endpoint using port 25 or 587. Use a utility such as Telnet or netcat to attempt to connect to the port manually.
  1. Open a command prompt.
  2. Use the following command to test the network connection.

```
telnet <SMTP endpoint> <port>
```

For example, `telnet smtp.us-ashburn-1.oraclecloud.com 25`

The port is open and the test is successful if a blank screen appears. If you are unable to connect to the ports using telnet, you are experiencing a network connectivity issue.

### To troubleshoot a problem connecting to an external mail transfer agent (MTA)

Use the following steps to determine whether you are able to communicate with an external service on the required ports 25 or 587. If you are unable to connect successfully, you are experiencing a network connectivity issue. If you are able to connect to an external MTA, the network connectivity issue is within Oracle Cloud Infrastructure.

- Connect to an external MTA such as Google's mail exchangers.
  1. Open a command prompt.
  2. Use the following command to retrieve one of Google's MX server records.

```
dig MX google.com
```

3. Use the following command to test connectivity to the endpoint port 25 or 587 against Google's MX servers.

```
telnet <IP address> <port>
```

If you are unable to connect to Google's MX servers, this confirms that you are having issues connecting to mail servers (port 25 or 587). It is possible that your egress rules are filtering traffic at the VCN.

If you can connect to an external MTA (that is, you are able to communicate with a public SMTP endpoint on the correct ports) but you cannot connect to Email Delivery public SMTP endpoints on those ports, create a service request with My Oracle Support with this information.

## Common Errors Returned by Email Delivery

### API Errors

For a complete list of common errors returned by all the services for Oracle Cloud Infrastructure, see [API Errors](#).

### Common SMTP Errors Returned by Email Delivery

The following table lists the common errors returned by the Email Delivery SMTP service.

## CHAPTER 13 Email Delivery

HTTP Status Code	Error Code	Description
451	Server error	An unexpected error has occurred during the SMTP conversation.
451	Error in Processing	An unexpected error has occurred during the SMTP conversation.
452	System storage error	The server is unable to persist the message in its delivery queue.
455	Maximum messages sent per minute reached : limit is <limit>	The SMTP send burst rate (of messages accepted per minute period) has been exceeded.
455	Maximum messages sent per day reached : limit is <limit>	The SMTP daily send rate (of messages accepted per 24 hour period) has been exceeded.
471	Authorization failed: address <address> not authorized	Authorization of the address (either in the envelope or message) has failed for the SMTP user.
501	Invalid command argument, not a valid Base64 string	The base64 encoded AUTH (PLAIN) secret is invalid.
501	Invalid command argument, does not contain NUL	The base64 encoded AUTH (PLAIN) secret does not contain NUL field separator(s).

HTTP Status Code	Error Code	Description
501	Invalid command argument, does not contain the second NUL	The base64 encoded AUTH (PLAIN) secret does not contain NUL field separator(s).
504	Method not supported	The client has attempted to use an unsupported AUTH mechanism with our service.
504	AUTH mechanism mismatch	The client has sent an invalid AUTH command to our service.
523	Exceeds byte limit	The message has exceeded the size limit enforced by the service (see server response to EHLO for size restriction).
535	Authentication credentials invalid	Authentication of the SMTP user has failed.
535	Authentication required	The client has sent commands that require SMTP authentication succeeded before the service is able to process (that is, commands are being sent out of order).
553	<address> Invalid email address	The RFC-822 Internet Address sent by the client is invalid.
554	Message parse error	The RFC-2822 Internet Message is invalid (and unable to be parsed by the server).

# Deliverability Best Practices

Deliverability Best Practices help you to learn and manage the habits that affect your sending reputation. These six recommendations can help lower your email bounce rate, stay off blacklists, lower your complaint rate, and improve your email sender reputation.

## Implement an Opt-in Process

An opt-in process is a method for your users to subscribe to your mailing list, which gives you permission to send messages. Only send messages to subscribers who have opted-in to your mailing list. There are two types of opt-in procedures.

- **Single opt-in (unconfirmed):** A user provides their email address and gives permission to receive relevant messages. Once the address is provided, messages can be sent without confirming the email address belongs to the user who provided it.
- **Double opt-in (confirmed):** A user provides their email address, but before the first mailing, a confirmation email is sent to the account owner. The email requires action from the account owner to confirm that future messages are wanted. An account can be verified by having the owner click a link for reply to the email. The confirmation email ensures that the address was not added to a third-party mailing list without consent.

## Purge Unengaged Users

Remove unengaged users by implementing a process. If a recipient is not engaging with your mail by either opening or clicking the email, this might be an indication that the email account is not in use or that the recipient is no longer interested in your content. If the recipient does not use the email account, eventually the mailbox provider terminates the account or transforms the account into a spam trap. Remove recipients who have not engaged with your email in a time frame defined by your business model. Purging unengaged users helps your deliverability by increasing your user engagement rate.

## Review Your Subscriber List

When reviewing your subscriber list, keep these things in mind:

- Eliminate duplicate addresses before sending. If addresses that do not exist are mailed to multiple times, your hard bounce rate could be inflated.
- Ensure that a previous suppression list (possibly from another email service provider) was not accidentally included.
- Verify that subscribers have opted-in. Do not send to an old list that you found.
- Restrict users from uploading their email client's contact list in a "select all" fashion. Forcing users to select addresses individually prevents users from accidentally including potentially out of date or expired addresses.

### Evaluate Your Sending Frequency

Sending too many emails in a short time might aggravate recipients, causing the recipients to mark your messages as spam. This is called list fatigue. Ensure that your message cadence aligns with the expected frequency of your content. Reducing frequency might reduce spam complaints. Ensure that your content is relevant to your subscribers. Keep your email messages consistent to your audience. A person who subscribed to a list for coupon updates might not want regular emails about auto loan finance rates. These unexpected messages are likely to be marked as spam, which decreases your sender reputation.

### Easily Accessible Unsubscribe URL

Unsubscribing helps your inbox success by sending only to recipients that engage by opening or clicking. When people complain, your sending reputation is harmed. Make it easy for recipients to be removed from the list. Do not hide the unsubscribe URL at the bottom of the message. A small percentage of users scroll to the bottom of the email and search for a small URL. Most users mark the email as spam.

### Canadian Anti-Spam Law (CASL) Guide

Canada's Anti-Spam Law (CASL) is one of the best guides to ensuring your compliance with the law, users' desire, and the intended filtering that most mailbox providers use. If you are a Canadian email sender or you send email to Canadian residents, you must comply with CASL. The following information is intended to help provide you with some guidance for complying

with CASL. This article does not constitute legal advice, nor is it intended supplement or otherwise affect your rights or obligations under your service agreement with Oracle, including your obligations under Oracle's Acceptable Use Policy. If you have questions about CASL or the legality of your sending practices, we encourage you to speak with an attorney who specializes in that subject matter.

### **What is covered by CASL?**

CASL and its related regulations apply to any "commercial electronic message" sent from or to Canadian computers and devices in Canada. Electronic messages that are merely routed through Canadian computer systems are not subject to CASL.

A "commercial electronic message" is any message that:

- Is in an electronic format, including emails, instant messages, text messages, and some social media communications.
- Is sent to an electronic address, including email addresses, instant message accounts, phone accounts, and social media accounts; and
- Contains a message encouraging recipients to take part in some type of commercial activity, including the promotion of products, services, people/personas, companies, or organizations.

### **Are there any types of messages that are exempt from CASL?**

These types of electronic messages are exempt from CASL for various reasons.

- Messages to family or a person with established personal relationship.
- Messages to an employee, consultant, or person associated with your business.
- Responses to a current customer, or someone who has inquired in the last six months.
- Messages that will be opened or accessed in a foreign country, including the U.S., China, and most of Europe.
- Messages sent on behalf of a charity or political organization for the purposes of raising funds or soliciting contributions.
- Messages attempting to enforce a legal right or court order.

- Messages that provide warranty, recall, safety, or security information about a product or service purchased by the recipient.
- Messages that provide information about a purchase, subscription, membership, account, loan, or other ongoing relationship, including delivery of product updates or upgrades.
- A single message to a recipient without an existing relationship based on a referral. The full name of the referring person must be disclosed in the message. The referrer might be family or have another relationship with the person to whom you are sending.

If your message does not meet one of these criteria, consent is required under CASL. Not all of the previous messages listed are permitted under the [Oracle Cloud Hosting and Delivery Policy](#).

### **What is “express consent”?**

Under CASL, “express consent” means a written or oral agreement to receive specific types of messages. For example, “You want to receive monthly newsletters and weekly discount notifications from Oracle”.

Express consent is only valid if your request for consent clearly and simply describes the following information:

- Your purpose in obtaining consent.
- A description of messages you will be sending.
- The name and contact information (physical mailing address and telephone number, email address, or website URL) of the requestor.
- A statement that the recipient can unsubscribe at any time.

The requestor can be you or someone for whom you are asking. If you are requesting consent on behalf of a client, the name and contact information of the client must be included with the consent request.

### **What is “implied consent”?**

Under CASL, you can only obtain implied consent when certain circumstances exist, including when:

## CHAPTER 13 Email Delivery

---

- A recipient has purchased a product, service or made another business deal, contract, or membership with your organization in the last 24 months.
- You are a registered charity or political organization, and the recipient has made a donation or gift, has volunteered, or attended a meeting organized by you.
- A professional message is sent to someone whose email address was given to you, or is conspicuously published, and who has not published or told you that unsolicited messages are not wanted.

### **What type of consent is required?**

After July 1, 2017, you can only send to recipients with express consent or whose implied consent is valid under CASL.

### **Some additional requirements**

In addition to understanding what qualifies as CASL-regulated message, and what type of consent is needed, there are a few other details to keep in mind.

- Retention of a record of consent confirmations is required.
- When requesting consent, checkboxes cannot be pre-filled to suggest consent. Each subscriber must check the box themselves for consent to be valid.
- All messages sent must include the following:
  - your name
  - the person on whose behalf you are sending (if any)
  - your physical mailing address and telephone number
  - your email address or website URL
- All messages sent after consent must also include an unsubscribe mechanism, and unsubscribes must be processed within ten days.

### **Where can I find more information on CASL?**

The full text of the law can be found on the website for the [Canadian Justice Department](#). The Canadian Radio and Telecommunications Commission has also set up an [FAQ](#) page and some guidelines for obtaining consent. If you have any questions, we encourage you to contact an attorney who is familiar with the law.

### **Oracle Cloud Hosting and Delivery Policy**

Often, the [Oracle Cloud Hosting and Delivery Policy](#) is more stringent than CASL requirements. It is important that you review Oracle policies before using the service.

### Troubleshooting Undelivered Emails

The following issues can cause an email to be undelivered:

- The recipient is on the Suppression List.
- An authentication failure or an issue with the format of the email message occurred. For example, if the SMTP "From" address is not the same as the "From" address in the email body, the email is rejected. The addresses must match and be an Approved Sender. Refer to your sending application's logs to review any issues.

If you are unable to resolve the issue, you can go to My Oracle Support and create a service request. See [Creating a Service Request](#) for more information.

# CHAPTER 14 Events

This chapter explains how to create automation in your tenancy.

## Overview of Events

Oracle Cloud Infrastructure Events enables you to create automation based on the state changes of resources throughout your tenancy. Use Events to allow your development teams to automatically respond when a resource changes its state.

Here are some examples of how you might use Events:

- Send a notification to a DevOps team when a database backup completes.
- Convert files of one format to another when files are uploaded to an Object Storage bucket.



### Note

Events is not available in Oracle Cloud Infrastructure Government Cloud.

## How Events Works

Oracle Cloud Infrastructure services emit *events*, which are structured messages that indicate changes in resources. Events (the messages, not the service) follow the [CloudEvents](#) industry standard format hosted by the [Cloud Native Computing Foundation \(CNCF\)](#). This standard allows for interoperability between various cloud providers or on-premises systems and cloud providers. An event could be a create, read, update, or delete (CRUD) operation, a resource lifecycle state change, or a system event impacting a resource. For example, an event can be emitted when a backup completes or fails, or a file in an Object Storage bucket is added, updated, or deleted.

Services emit events for resources or data. For example, Object Storage emits events for buckets and objects. Services emit different types of events for resources, which are distinguished as *event types*. Buckets and objects have event types of create, update, and delete, for example. Event types are the changes that produce events by a given resource. For a list of services that produce events and the event types that those services track, see [Services that Produce Events](#).

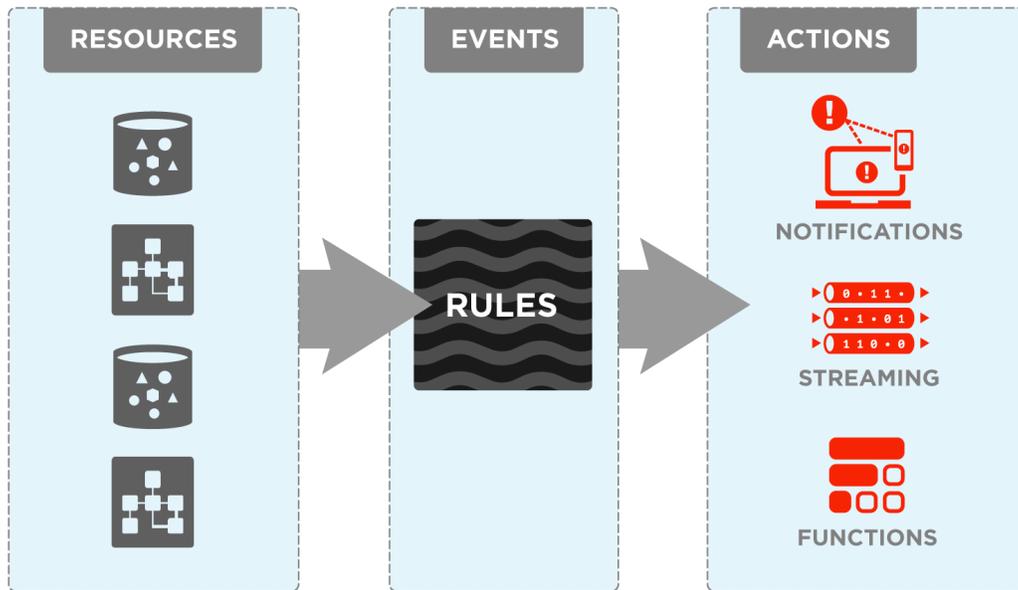
You work with events by creating *rules*. Rules include a filter you define to specify events produced by the resources in your tenancy. The filter is flexible:

- You can define filters that match only certain events or all events.
- You can define filters based on the way resources are tagged or the presence of specific values in attributes from the event itself.

Rules must also specify an *action* to trigger when the filter finds a matching event. Actions are responses you define for event matches. You set up select Oracle Cloud Infrastructure services that the Events service has established as actions (more on these select services follows). The resources for these services act as destinations for matching events. When the filter in the rule finds a match, the Events service delivers the matching event to one or more of the destinations you identified in the rule. The destination service that receives the event then processes the event in whatever manner you defined. This delivery provides the automation in your environment.

You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions:

- [Notifications](#)
- [Streaming](#)
- [Functions](#)



## Events Concepts

The following concepts are essential to working with Events.

### EVENTS

An automatic notification of a state change as reported by an event-emitting Oracle Cloud Infrastructure resource. For example, a database resource emits a `backup.begin` event when a backup begins.

### EVENT TYPES

A distinction between the different types of events. For more information, see [Services that Produce Events](#).

### RULES

A JSON object you create to subscribe to an event type and trigger an action should that event occur. For example, a rule might specify that `backup.end` event types from databases trigger the Notifications service to send an email to a particular DevOps engineer. For more information, see [Matching Events with Filters](#).

### ACTIONS

Rules must also specify an *action* to trigger when the filter finds a matching event. Actions are responses you define for event matches. You set up select Oracle Cloud Infrastructure services that the Events service has established as actions. The resources for these services act as destinations for matching events. When the filter in the rule finds a match, the Events service delivers the matching event to one or more of the destinations you identified in the rule. The destination service that receives the event then processes the event in whatever manner you defined. This delivery provides the automation in your environment.

You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions:

- [Notifications](#)
- [Streaming](#)
- [Functions](#)

## Region Availability

Events is currently available in all regions of the [commercial realm](#). Events is currently not available in regions within the Government Cloud realm.

## Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Administrators: You must write IAM policy that authorize users to work with rules. For more information, see [Events and IAM Policies](#).

### Limits on Events Resources

The Events service has a limitation of 50 rules per tenancy.

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

### Service Gateway and Events

The Events service also supports private access from Oracle Cloud Infrastructure resources in a VCN through a service gateway. A service gateway allows connectivity to the Events public endpoints from private IP addresses in private subnets. For example, you can manage rules over the Oracle Cloud Infrastructure backbone instead of over the internet. You can optionally use IAM policies to control which VCNs or ranges of IP addresses can access Events. See [Access to Oracle Services: Service Gateway](#) for details.

### Getting Started with Events

This topic introduces you to creating automation with Events. You create a simple rule that sends a notification whenever someone creates a bucket in a particular compartment in your tenancy.



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Setting Up for Events

To try out the Events service for this tutorial, you must have these things set up first:

- Create IAM policy for Events
- Create a topic and subscription to use as an action



### Important

A tenancy administrator must configure your tenancy for Events. These configurations give you access to an Oracle Cloud Infrastructure tenancy with the necessary IAM policy and a resource to use as an action.

### Create Users, Groups, and Compartments

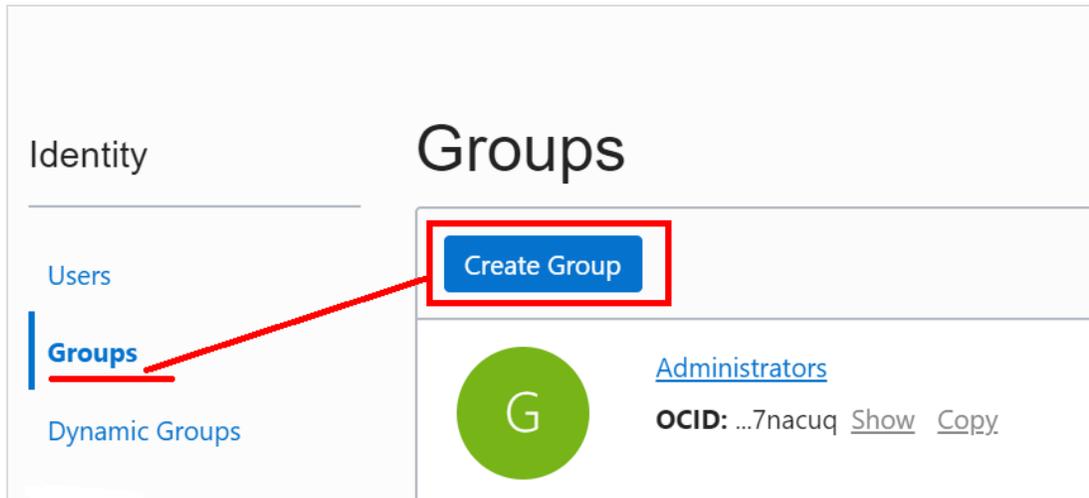
You can use existing users, groups, and compartments or make new ones.

#### To create groups and users

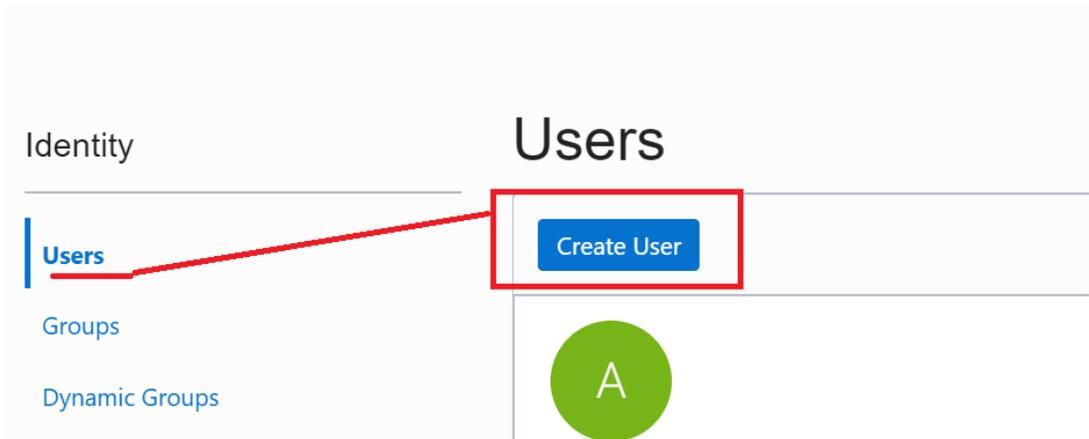
If suitable users and groups for assigning users permissions to work with rules don't already exist, log in to the Console as a tenancy administrator and create them.

1. Log in to the Console as a tenancy administrator.
2. If you need a group for Events, perform these steps:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**. A list of the groups in your tenancy is displayed.
  - b. Click **Create Group** and create a new group (see [To create a group](#)). Give the group a meaningful name and description. Avoid entering confidential

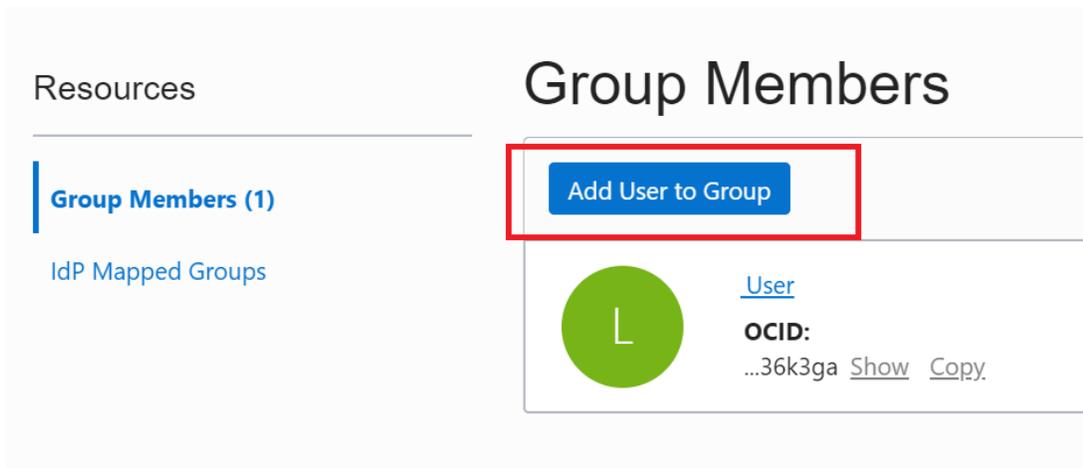
information.



3. If you need user accounts for Events, perform these steps:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. A list of the users in your tenancy is displayed.
  - b. Click **Create User** and create one or more new users (see [To create a user](#)).



4. If users haven't been added to groups already, perform these steps:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**. A list of the groups in your tenancy is displayed.
  - b. Click the group you want to use for Events.
  - c. Click **Add User to Group**.



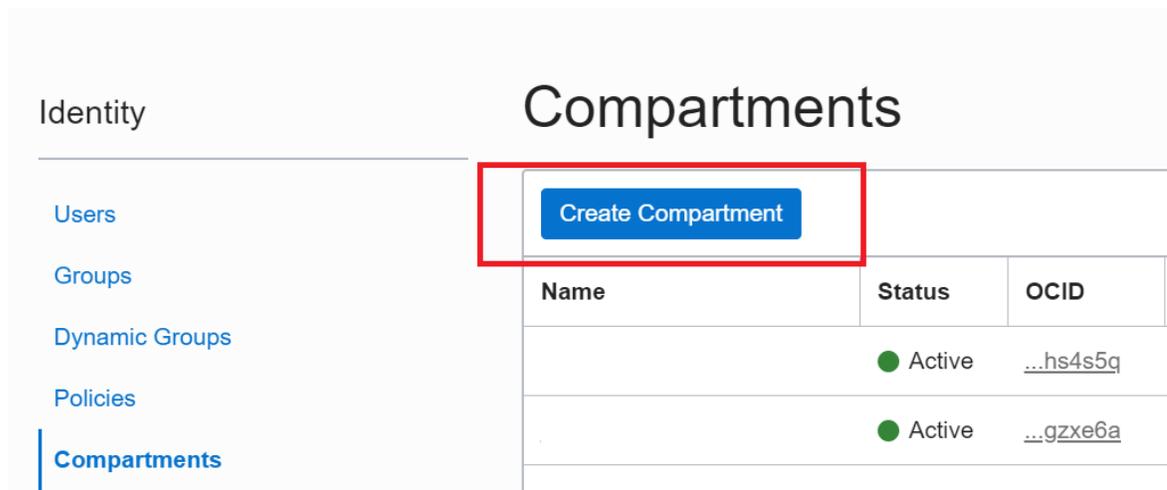
- d. Select the users from the drop-down list, and then click **Add**.

### To create a compartment

If suitable compartment for rules and the resources that emit events doesn't already exist, log in to the Console as a tenancy administrator and create it.

1. Log in to the Console as a tenancy administrator.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**. A list of the compartments in your tenancy is displayed.
3. Click **Create Compartment** and create a new compartment (see [To create a compartment](#)). Give the compartment a meaningful name and description. Avoid

entering confidential information.



The screenshot shows the Oracle Cloud IAM console interface. On the left, there is a navigation menu under the heading 'Identity' with options: Users, Groups, Dynamic Groups, Policies, and Compartments (which is selected). The main content area is titled 'Compartment' and features a blue 'Create Compartment' button, which is highlighted with a red rectangular box. Below the button is a table listing existing compartments.

Name	Status	OCID
	● Active	...hs4s5q
	● Active	...gzxe6a

### Create IAM Policy for Events

Before users can start using Events to create automation, as a tenancy administrator you must create IAM policy:

To create a policy that allows users to create and manage rules

1. Log in to the Console as a tenancy administrator.
2. In the Console, open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**. A list of the policies in the compartment you're viewing is displayed.
3. Select the root compartment.
4. Click **Create Policy**.
5. Enter the following:
  - **Name:** A meaningful name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later. Avoid entering confidential information.

- **Description:** A meaningful description. You can change this later if you want to. Avoid entering confidential information.
- **Statement:** Enter the following policy statements to give users in the group the ability to manage and create rules:  
This line gives the user inspect access to resources in compartments to select actions.

```
allow group <RuleAdmins> to inspect compartments in tenancy
```

This line gives the user access to defined tags to apply filter tags to rules.

```
allow group <RuleAdmins> to use tag-namespaces in tenancy
```

These lines give the user access to Streaming resources for actions

```
allow group <RuleAdmins> to inspect streams in tenancy
allow group <RuleAdmins> to use stream-push in tenancy
allow group <RuleAdmins> to use stream-pull in tenancy
```

These lines give the user access to Functions resources for actions.

```
allow group <RuleAdmins> to use virtual-network-family in tenancy
allow group <RuleAdmins> to manage function-family in tenancy
```

This line give the user access to Notifications topics for actions.

```
allow group <RuleAdmins> to use ons-topic in tenancy
```

This line gives the user manage access to rules for Events.

```
allow group <RuleAdmins> to manage cloudevents-rules in tenancy
```

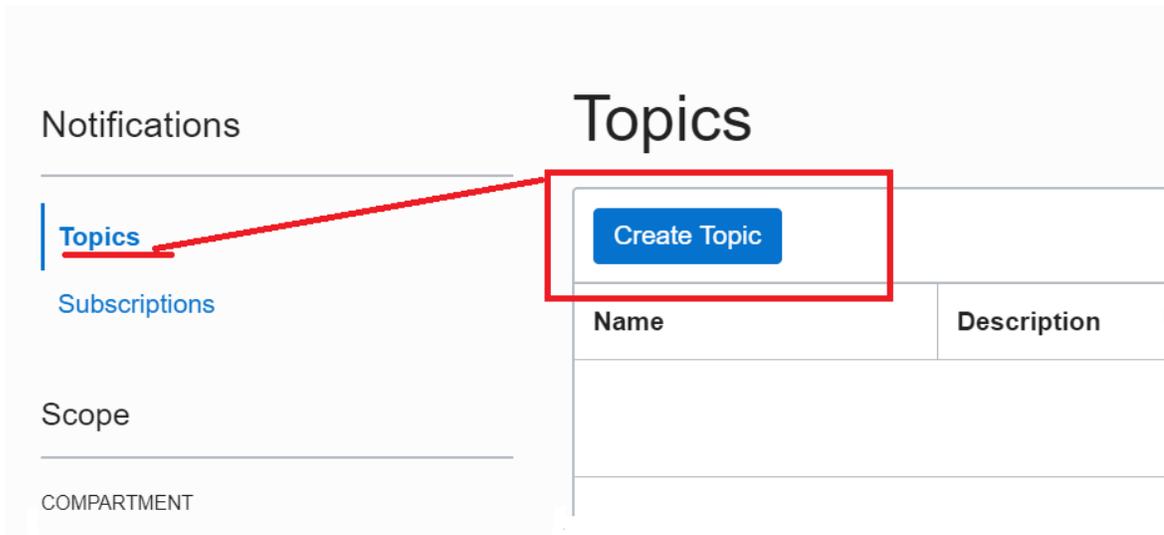
### 6. Click **Create**.

### Create Notifications Topic and Subscription

If a suitable Notifications topic doesn't already exist, then you must log in to the Console as a tenancy administrator and create it. Whether you use an existing topic or create a new one, add an email address as a subscription so that you can monitor that email account for notifications.

### To create a topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Click **Create Topic** at the top of the topic list.



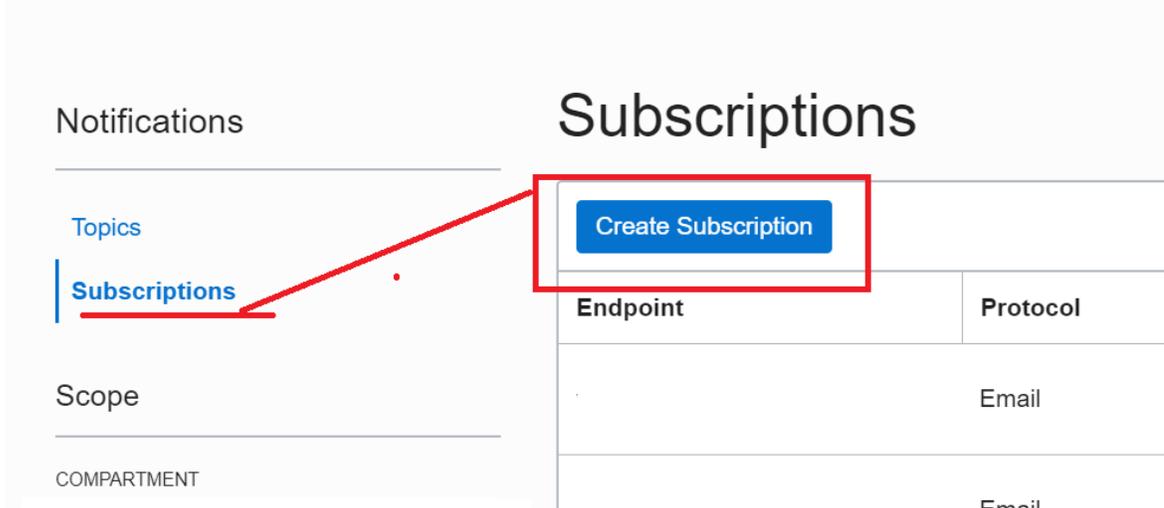
3. In the **Create Topic** dialog box, configure your topic.
  - **Name:** Required. Specify a friendly name for the topic. It must be unique; validation is case-sensitive. Avoid entering confidential information.
  - **Description:** Optional. Enter a description for the topic. Avoid entering confidential information.
4. Click **Create**.

### To create a subscription

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.

## CHAPTER 14 Events

2. Click the name of the topic that you created in the previous step or the topic you intend to use for this tutorial.
3. On the topic detail page, click **Create Subscription**.



4. In the **Create Subscription** dialog box, select **Email**, and then type an email address.

The screenshot shows the 'Create Subscription' dialog box. At the top, it says 'Create Subscription' with 'help' and 'cancel' links. Below this, there are three sections: 'TOPIC', 'PROTOCOL', and 'EMAIL'. The 'TOPIC' section has a dropdown menu with 'Events-Hello-World' selected. The 'PROTOCOL' section has a dropdown menu with 'Email' selected, and this section is highlighted with a red rectangular box. The 'EMAIL' section has an empty text input field, which is also highlighted with a red rectangular box.

5. Click **Create**.

The subscription has been created and a subscription confirmation URL will be sent. The subscription remains in "Pending" status until it has been confirmed.

### To confirm a subscription

- In the confirmation email sent to the address you specified in the previous procedure, click the confirmation URL.

## Using the Console to Create a Rule

Use the Console to create a rule with a pattern that matches bucket creation events emitted by Object Storage. Specify the Notifications topic you created as an action to deliver matching events. To test your rule, create a bucket. Object Storage emits an event which triggers the action. Check the email specified in the subscription to receive your notification.

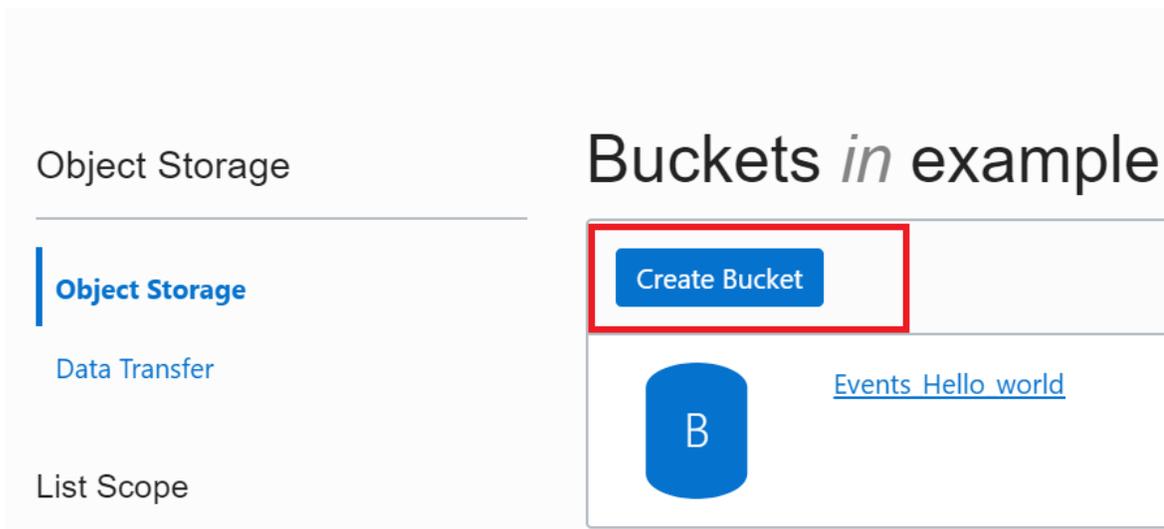
### To create a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click **Create Rule**. Events compares the rules you create in this compartment to event messages emitted from resources in this compartment and any child compartments.
3. Enter the following.
  - **Display Name:** Specify a friendly name for the rule. You can change this name later. Avoid entering confidential information.
  - **Description:** Specify a description of what the rule does. You can change this description later. Avoid entering confidential information.
4. In **Event Matching**, select **Event Type**.

- a. In **Service Name**, select **Object Storage**.
  - b. In **Event Type**, select **Object Storage - Create Bucket**.
5. In **Actions**, specify the actions to trigger when the filter finds a match:
  - a. In **Action Type**, select **Notifications**.
  - b. In **Notifications Compartment**, select the compartment that contains the topic.
  - c. In **Topic**, select the topic.
6. Click **Create Rule**.

### To create a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Select the compartment where you created your rule (or any of its subordinate compartments).
3. Click **Create Bucket**.



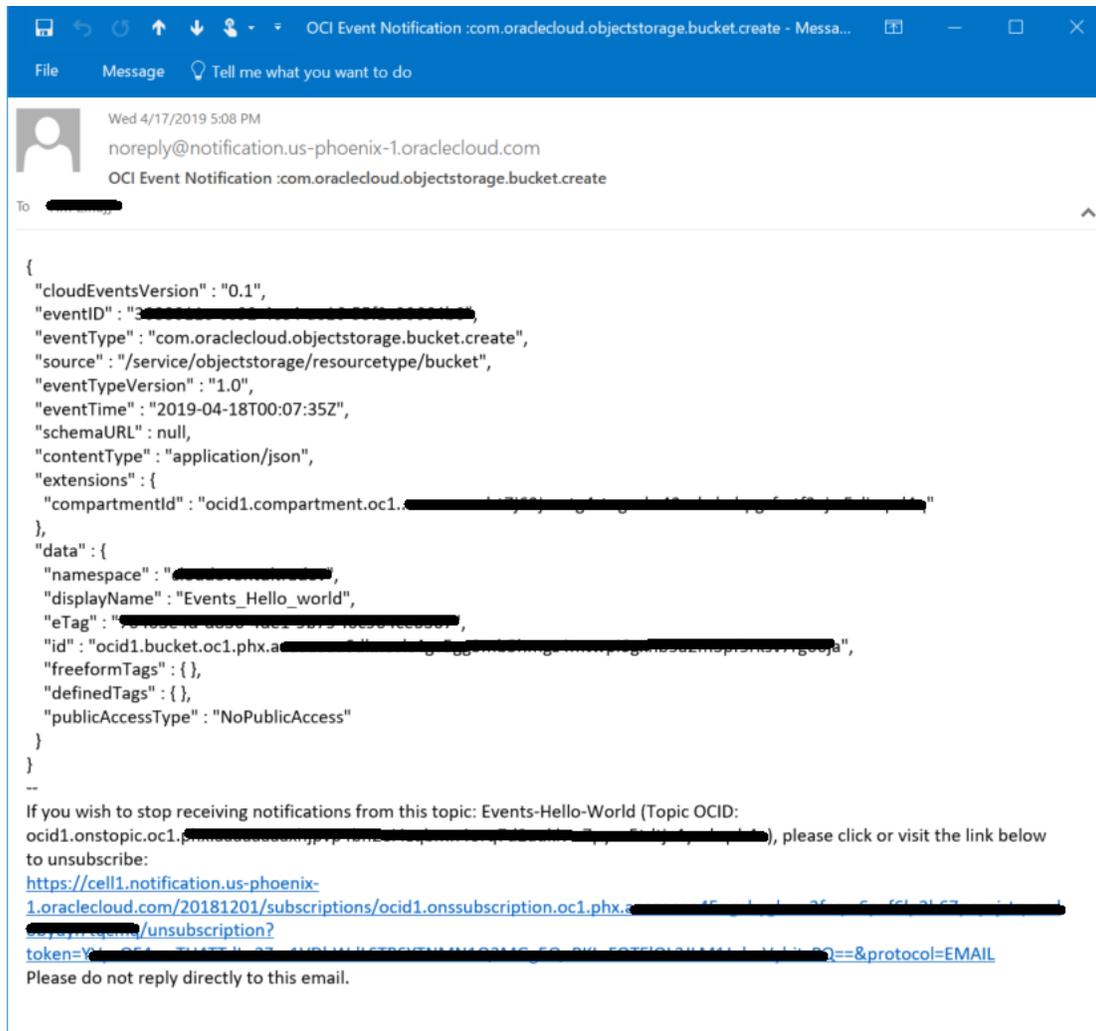
4. In the **Create Bucket** dialog, specify the attributes of the bucket:

- **Name:** Required. A user-friendly name or description. Avoid entering confidential information.
  - **Storage Tier:** Select the tier in which you want to store your data. Available tiers include:
    - **Standard** is the primary default Object Storage tier for storing data that is accessed frequently and requires fast and immediate access.
    - **Archive** is a special tier for storing data that is accessed infrequently and requires long retention periods. Access to data in the **Archive** tier is not immediate. You must restore archived data before it's accessible.
5. Click **Create Bucket**.

### To receive your notification

- Log in to the email account you specified in the previous procedure to receive the notification about the bucket being created.

## CHAPTER 14 Events





### Tip

You will receive notifications each time a bucket is created in the compartment (or any of its sub compartments) until you disable the rule.

## Using the CLI to Create a Rule

When you use the CLI to create a rule, you work a little differently than using the Console.

- To specify the actions for your rule, use a JSON formatted file. You create this file before you create the rule, and the file simplifies the amount of information you must type at the command line.
- To specify an event to match, use a JSON formatted string. You type this right into the console as you create the rule.

## To create an action file

1. Create a file and add the following content.

```
{
 "actions": [
 {
 "actionType": "ONS",
 "description": "string",
 "isEnabled": true,
 "topicId": "<topic_OCID>"
 }
]
}
```

Tip: You can specify functions, streams, or topics as an action.

## Example action file template

```
{
```

```
"actions": [
 {
 "actionType": "FAAS",
 "description": "string",
 "functionId": "<function_OCID>",
 "isEnabled": true
 },
 {
 "actionType": "ONS",
 "description": "string",
 "isEnabled": true,
 "topicId": "<topic_OCID>"
 },
 {
 "actionType": "OSS",
 "description": "string",
 "isEnabled": true,
 "streamId": "<stream_OCID>"
 }
]
```

2. Fill in *<topic\_OCID>* with actual topic OCID value from your tenancy.
3. Add a description.
4. Save the file with *action.json* as the file name.

### To create a rule

Open a command prompt and run `oci events rule create` to create a rule.

Use the following options:

- `display-name` indicates the name of the rule in the Console
- `is-enabled` indicates whether the rule is evaluated.
- `condition` a JSON formatted string used to indicate a pattern for event matching (see Examples for usage).

### Examples

The following example shows how to pass a simple condition that matches all events. Everything between the double quotes ( " ") is a string, while the brackets { } indicate JSON:

## CHAPTER 14 Events

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "{}" --
compartment-id <compartment_OCID> --actions file://action.json --wait-for-state=ACTIVE
```

To pass complex input to the CLI as a JSON string, you must enclose the entire block in double quotes. Inside the block, each double quote for the key and value strings must be escaped with a backslash (\) character.

For example:

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "
{"eventType":["com.oraclecloud.objectstorage.createbucket"]}" --compartment-id <compartment_
OCID> --actions file://action.json --wait-for-state=ACTIVE
```

In PowerShell, to escape double quotes, you must use two characters: The backslash (\) and the back tick (`).

For example (Windows PowerShell):

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "
{"eventType\`":["com.oraclecloud.objectstorage.createbucket\`"]}" --compartment-id
<compartment_OCID> --actions file://action.json --wait-for-state=ACTIVE
```



### Tip

The `condition` option does not support using a file to pass the JSON formatted string.

- `compartment-id` indicates the compartment where the rule applies. Events evaluates messages from resources in this compartment and any child compartments.
- `actions` indicates the location in the local file system of the JSON formatted file you created to specify the actions for a rule.
- `wait-for-state=` when used with `ACTIVE` indicates that the CLI should wait for the service to create the rule, do another GET operation, and then display the rule in the active state. Without the option, the CLI displays the rule immediately in the creating state.

For example:

## CHAPTER 14 Events

```
oci events rule create --display-name CLI-created_rule --is-enabled true --condition "{\"eventType\": [\"com.oraclecloud.objectstorage.createbucket\"]}" --compartment-id <compartment_OCID> --actions <path_to_json_formatted_actions_file> --wait-for-state=ACTIVE
```



### Note

Replace the values in *<compartment\_OCID>* and *<path\_to\_json\_formatted\_actions\_file>* with the actual values from your tenancy and local file system.

When you run the preceding command, the CLI prompts you about the rule and its display:

```
Action completed. Waiting until the resource has entered state: ACTIVE
{
 "data": {
 "actions": {
 "actions": [
 {
 "action-type": "ONS",
 "description": "Notifications action",
 "id": "ocidl.eventaction.oc1.phx.<unique_ID>",
 "lifecycle-message": null,
 "lifecycle-state": "ACTIVE",
 "topic-id": "ocidl.onstopic.oc1.phx.<unique_ID>"
 }
]
 },
 "compartment-id": "ocidl.compartment.oc1..<unique_ID>",
 "condition": "{\"eventType\": [\"com.oraclecloud.objectstorage.createbucket\"]}",
 "defined-tags": {},
 "description": null,
 "display-name": "CLI-created_rule",
 "freeform-tags": {},
 "id": "ocidl.eventrule.oc1.phx.<unique_ID>",
 "is-enabled": true,
 "lifecycle-message": null,
 "lifecycle-state": "ACTIVE",
 "time-created": "2019-04-25T01:32:56.855000+00:00"
 },
}
```

```
"etag": "<unique_ID>--gzip"
}
```

### To create a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Select the compartment where you created your rule (or any of its subordinate compartments).
3. Click **Create Bucket**.
4. In the **Create Bucket** dialog, specify the attributes of the bucket:
  - **Name**: Required. A user-friendly name or description. Avoid entering confidential information.
  - **Storage Tier**: Select the tier in which you want to store your data. Available tiers include:
    - **Standard** is the primary default Object Storage tier for storing data that is accessed frequently and requires fast and immediate access.
    - **Archive** is a special tier for storing data that is accessed infrequently and requires long retention periods. Access to data in the **Archive** tier is not immediate. You must restore archived data before it's accessible.
5. Click **Create Bucket**.

### To receive your notification

- Log in to the email account you specified in the previous procedure to receive the notification about the bucket being created.



### Tip

You receive notifications each time a bucket is created in the compartment (or any of its sub compartments) until you disable the rule.

## Matching Events with Filters

This topic describes how to match events with pattern filters in rules to build automation.

### Background

To understand filtering, it's helpful to review the structure of an actual event message. Events uses JSON objects to represent events. This is an event:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.objectstorage.deletebucket",
 "source": "objectstorage",
 "eventTypeVersion": "1.0",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_bucket",
 "resourceId": "ocidl.compartment.oc1..<unique_ID>",
 "availabilityDomain": "NfHZ:PHX-AD-2",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
```

```
 "Operations": {
 "CostCenter": "42"
 },
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess",
 "eTag": "f8fffb6e9-f602-460f-a6c0-00b5abfa24c7"
 }
 }
}
```

Two key points to remember about all events:

- Events all have the same set of top-level attributes, which are known as the event envelope. With one exception, most of these top-level attributes are not that useful for creating filters. The exception is `eventType`, which identifies the type of event included in the payload.
- The payload of the event appears within the `data` attribute. The information in this field depends on which service produced the event and the event type requested. The information in the payload is useful for isolating one event from another with a filter.

For more information about the envelope, see [Contents of an Event Message](#). For a list of all the services that produce events, see [Services that Produce Events](#).

### Event Matching with Filters

Rules use filters to select events and route them for delivery to action resources. A rule is represented as a JSON object, similar to an event. The filter is an attribute of the rule, and the attribute is named `condition`. A filter either matches an event or it does not.

A few important things to remember about filters:

- Fields not mentioned in a filter are ignored. You can create a valid filter that matches all event messages with two curly brackets.

- For a filter to match an event, the event must contain all the field names listed in the filter. Field names must appear in the filter with the same nesting structure used in the event.
- Rules apply to events in the compartment in which you create them and any child compartments. This means that a filter specified by a rule only matches events emitted from resources in the same compartment or any child compartments.
- Wildcard matching is supported with the asterisk (\*) character. See [Examples of Wildcard Matching in Filters](#).

### Examples of Simple Filters

The following filter matches every event in the compartment and any child compartments where you create the rule.

```
{
...

 "condition": "{ }"
}
```

When you add fields to the filter, you limit the events that the filter can match. For example, the following filter matches only `deletebucket` events.

```
{
...

 "condition": "{
 "eventType": "com.oraclecloud.objectstorage.deletebucket"
 }"
}
```

To create a filter for more than one event type, use an array in `eventType`. The following filter matches `deletebucket` and `createbucket` events.

```
{
...

 "condition": "{
```

```
 "eventType": [
 "com.oraclecloud.objectstorage.deletebucket",
 "com.oraclecloud.objectstorage.createbucket"
]
 }"
}
```

### Examples of Filters with Event Payload Attributes

Both of the following filters would match the event at the top of the page. The first because filter specifies two fields and both fields appear in the event, the second because the "NoPublicAccess" type appears in the event.

The important thing to note is how the field names in the filter match the nesting structure of the event.

```
{
 ...

 "condition": "{
 "data": {
 "compartmentName": "example_name",
 "resourceName": "my_bucket"
 }
 }"
}
```

```
{
 ...

 "condition": "{
 "data": {
 "additionalDetails": {
 "publicAccessType": "NoPublicAccess"
 }
 }
 }"
}
```

## CHAPTER 14 Events

---

Neither of the following filters would match the event at the top of this page. The first because the filter specifies a `PublicAccessType` not found in the event. The second because the event specifies a name for different bucket.

```
{
...

 "condition": "{
 "data": {
 "compartmentName": "example_name",
 "resourceName": "my_bucket",
 "additionalDetails": {
 "publicAccessType": "PublicAccess"
 }
 }
 }"
```

```
{
...

 "condition": "{
 "data": {
 "additionalDetails": {
 "publicAccessType": "NoPublicAccess"
 }
 }
 }"
```

### Examples of Arrays in Filters

Arrays in filters match events if any of the values in the filter match a value in an event. The following filter would match the event at the top of the page because the name of the bucket in the event is included in an array in the filter.

```
{
...

 "condition": "{
 "data": {
```

```
 "resourceName": [
 "my_bucket_2",
 "my_bucket_1",
 "my_bucket"
],
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess"
 }
 }
}"
}
```

You can use an array in `eventType` (or any of the top-level fields), the event payload as shown in the preceding example, or both the event payload and a top-level field.

```
{
...
 "condition": "{
 "eventType": [
 "com.oraclecloud.objectstorage.deletebucket",
 "com.oraclecloud.objectstorage.createbucket"
],
 "data": {
 "resourceName": [
 "my_bucket_2",
 "my_bucket_1",
 "my_bucket"
],
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess"
 }
 }
 }"
}
```

### Examples of Wildcard Matching in Filters

The following are a few things to consider about wildcard matching with filters.

## CHAPTER 14 Events

---

- Use the wildcard only in attribute values. You cannot use the asterisk for matching in keys.
- An attribute value with only an asterisk matches all values for the associated attribute name, but *not* null.
- The period character has no special meaning in a filter.

You can add the asterisk at the start of a string, in the middle, or at the end. All of the filters that follow match the event at the top of the page.

- The first matches because the wildcard in `displayName` matches the bucket naming pattern.
- The second one matches because the `publicAccessType` uses a wildcard. Because of the use of the wildcard, these first two filters would also match events from buckets with a similar naming pattern and would include events from buckets with or without public access.
- The third one matches because the event type includes all types of bucket events.

```
{
...
 "condition": "{
 "data": {
 "resourceName": "my_bucket*",
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess"
 }
 }
 }"
```

```
{
...
 "condition": "{
 "data": {
 "resourceName": [
```

```
 "my_bucket_2",
 "my_bucket_1",
 "my_bucket"
],
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "*"
 }
}
}"
}
```

```
{
...

"condition": "{
 "eventType": "com.oraclecloud.objectstorage.*bucket",

 "data": {
 "resourceName": [
 "my_bucket_2",
 "my_bucket_1",
 "my_bucket"
],
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess"
 }
 }
}"
}
```

## Events and IAM Policies

This topic describes how an administrator must write IAM policy for the Events service. If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about how to write IAM policy for Events, see [Details for the Events Service](#).

### Allow Users to Work with Rules

These IAM policies allow users to manage or list rules.

### Let users list rules in a compartment

**Type of access:** Ability to list Events rules.

**Where to create the policy:** In the tenancy.

```
Allow group RuleReaders to read cloudevents-rules in tenancy
```

The preceding policy allows RuleReaders to list rules in the tenancy.

### Let admins manage rules in a compartment

**Type of access:** Ability to manage Events rules, including creating, deleting, updating or moving rules to a different compartment.

**Where to create the policy:** In the tenancy.

This line gives the user inspect access to resources in compartments to select actions.

```
allow group <RuleAdmins> to inspect compartments in tenancy
```

This line gives the user access to defined tags to apply filter tags to rules.

```
allow group <RuleAdmins> to use tag-namespaces in tenancy
```

These lines give the user access to Streaming resources for actions

```
allow group <RuleAdmins> to inspect streams in tenancy
allow group <RuleAdmins> to use stream-push in tenancy
allow group <RuleAdmins> to use stream-pull in tenancy
```

These lines give the user access to Functions resources for actions.

```
allow group <RuleAdmins> to use virtual-network-family in tenancy
allow group <RuleAdmins> to manage function-family in tenancy
```

This line give the user access to Notifications topics for actions.

```
allow group <RuleAdmins> to use ons-topic in tenancy
```

This line gives the user manage access to rules for Events.

```
allow group <RuleAdmins> to manage cloudevents-rules in tenancy
```

# Managing Rules for Events

This topic describes how to manage rules for the Events service. For more information about Events, see [Overview of Events](#).

## Prerequisites for Creating Rules

- Action resources: You must have resources already set up to specify as an action. The Events service invokes the action specified in the rule by delivering the event message to action resources, which can include topics, streams, or functions. Every rule must have at least one action. The Events service can invoke any of the following services by delivering an event message for processing:
  - [Notifications](#)
  - [Streaming](#)
  - [Functions](#)
- IAM policies: To manage or list rules, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform a task and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information, see [Events and IAM Policies](#).
- Event messages: To create rules, the resources you want to monitor with the rule must emit events. For more information, see [Services that Produce Events](#).

### Working with Rules



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

A typical workflow for setting up rule might follow this pattern:

- 1. Identify action resources**

Set up or identify whatever action resources you intend to use with the rule. For example, you might set up a Notifications topic and create subscriptions for the DevOps team so that they are notified when backups complete. If a topic already exists, you can use it instead of creating a topic. The resources you specify for actions do not have to be in the same compartment as the rule.

- 2. Plan filtering**

Ensure the resources that you want to monitor emit events to the Events service and plan your pattern matching strategy. For example, you might want to monitor backups on Autonomous Data Warehouse instances in the ABC compartment. Ensure Autonomous Data Warehouse instances emit an event type you can use to create the automation you require. Review the example JSON event to determine the best way to identify those resources in filters. See [Matching Events with Filters](#) and [Services that Produce Events](#).

- 3. Create the rule**

Rules apply to events in the compartment in which you create them and any child compartments. Create a rule in the compartment with the resource you want to monitor and specify where to deliver matching events. For example, in the ABC compartment, you might create a rule that filters for Autonomous Data Warehouse backup events.

Since Events has no requirement about the location of action resources, you could specify a topic in the XYZ compartment as the resource to deliver any matching events.

### Managing Tags for Rules

You can add tags to your resources to help you organize them according to your business needs. You can add tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

#### Tags and Event Filtering

With Events, you can also use tags to target resources in your tenancy. You target resources by adding the tag to a filter in a rule. A *filter tag* helps you hone automation by targeting only resources that contain a particular tag. For example, let's say you have dozens of Database instances in your tenancy, but only a few of the most critical of these instances have the tag "Operations." You could create a rule that triggers a particular action for resources that only contain the "Operations" tag.

Policy for working with filter tags is no different from policy for working with tags.

#### To manage filter tags

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click rule's name.
3. In the **Resources** menu, click **Event Matching**.
4. In the **Filter Tags** section, you can view or edit existing filter tags, or click **Add Filter Tag** to add new ones.

### To manage tags for rules

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click rule's name.
3. Click the **Tags** tab to view or edit existing tags, or click **Add Tags** to add new ones.

For more information, see [Resource Tags](#).

### Move Rules to a Different Compartment

You can move rules from one compartment to another. When you move a rule to a new compartment, you stop monitoring events from resources in the current compartment and begin monitoring events in the new compartment (and any child compartments). After you move the rule to the new compartment, inherent policies apply immediately and affect access to the rules through the Console. Moving rules doesn't affect access by the Events service to actions defined in rules. For more information, see [Moving Resources to a Different Compartment](#).

### Monitoring Rules

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For more information about monitoring the rules you create, see [Events Metrics](#).

### Object Events and the Events Service

Events for objects are handled differently than other resources. Objects do not emit events by default. Use the [Console](#), [CLI](#), or [API](#) to enable a bucket to emit events for object state changes. You can enable events for object state changes during or after bucket creation.

### Using the Console

#### To create a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click **Create Rule**. Events compares the rules you create in this compartment to event messages emitted from resources in this compartment and any child compartments.
3. Enter the following.
  - **Display Name:** Specify a friendly name for the rule. You can change this name later. Avoid entering confidential information.
  - **Description:** Specify a description of what the rule does. You can change this description later. Avoid entering confidential information.
4. In **Rule Conditions**, create a filter:

#### To add an event type

- a. Select **Event Type**.
- b. Select a **Service Name**.
- c. In **Event type**, select one or more event types for this service.
- d. Click **+ Another Condition** and select **Event Type** to add event types for a different service.

This filter will match events of the event types you specify.

#### To add an attribute

You must first select an event type to add an attribute.

- a. Select **Attribute**.
- b. Select an **Attribute Name**.
- c. Enter an **Attribute Value**. Attribute values are optional.
- d. Click **+ Another Condition** and select **Attribute** to add another attribute.  
This filter will match events of the events types with the attributes you specify.

### To add a filter tag

- a. Select **Filter Tag**
- b. Select a **Tag Namespace**.  
To specify a free-form tag, select **None (apply a free-form tag)**.
- c. Select a **Tag Key**.
- d. Enter a **Tag Value**. Tag values are optional.
- e. Click **+ Another Condition** and select **Filter Tag** to add another filter tag.  
This filter will match events with the tags you specify.

Filter tags help you to hone automation by targeting only resources that contain a particular tag. If you want to use tags to organize your rules, use resource tags instead. For more information, see [Tags and Event Filtering](#).



#### Tip

You can leave this field entirely blank to match all events. See [Matching Events with Filters](#).

### To validate this rule

You can only evaluate a rule against one event type at a time. To test different event

types, repeat these steps as necessary.

- a. Click **Validate Rule**.

The **Test Rule** panel opens.

- b. In **Service Name**, select a service if necessary.

- c. In **Event Type**, select an event type, if necessary.

A example event appears based on the selections you made. Edit the values in the event to match the values for any attributes and tags you added to your rule. For more information, see [Contents of an Event Message](#).

- d. Click **Check if Example Event Matches Rule**.

If the rule doesn't match, use the rule editor to modify any of the following:

- Add or remove event types
- Add or remove values or attributes
- Add or remove tags
- Insert wildcards

For more information, see [Matching Events with Filters](#).

- e. Click **Close**.

### To view reference events

- a. Click **View example events (JSON)**.

The **View Example Events** panel opens.

- b. In **Service Name**, select a service if necessary.

- c. In **Event Type**, select an event type, if necessary.

A example event appears based on the selections you made. Use the events viewer to browse reference events.

- d. Click **Done**.

For more information, see [Contents of an Event Message](#) and [Matching Events with Filters](#).

5. In **Actions**, specify the actions resources to trigger when the filter finds a match:

### To select a topic

- a. Select **Notifications**.
- b. Select the **Notifications Compartment**.
- c. Select the **Topic**.
- d. Click **+ Another Action** and select **Notifications** to add another topic.

### To select a stream

- a. Select **Streaming**.
- b. Select the **Stream Compartment**.
- c. Select the **Stream**.
- d. Click **+ Another Action** and select **Streaming** to add another stream.

### To select a function

- a. Select **Functions**.
- b. Select the **Functions Compartment**.
- c. Select a **Functions Application**.
- d. Select the **Functions ID**.
- e. Click **+ Another Action** and select **Functions** to add another function.

6. Click **Create Rule**.

### To edit a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** that has the rule you want to edit.
3. For the rule you want to edit, click the Actions icon (three dots), and then click **Edit**.
4. Make your changes and click **Save Changes**.

### To disable or enable a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** that has the rule you want to work with.
3. For the rule, you want change, click the Actions icon (three dots), and then take one of the following actions:
  - Click **Disable**
  - Click **Enable**
4. Confirm when prompted.

### To move a rule to a different compartment

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. In the **Scope** section, select a compartment.
3. Find the rule in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### To validate a rule

You can only evaluate a rule against one event type at a time. Repeat as necessary to test different event types.

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to test.
3. Click **Validate Rule**.
4. Take one or more of the following actions:
  - If there are no event types in the rule, select the service and event type you want to test.
  - If you want to test a different event type than the one selected by default, select the service and event type you want to test.
  - If you added attribute values or filter tags to the rule, edit the example data in the event to match the values in your rule.
5. Click **Check if Example Event Matches Rule**.

For more information, see [Matching Events with Filters](#) and [Contents of an Event Message](#).

### To delete a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose the **Compartment** that has rule you want to delete.
3. For the rule you want to delete, click the Actions icon (three dots) , and then click **Delete**.
4. Confirm when prompted.

### To add an action to a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Actions**.
4. Click **Add**.
  - The **Add Action** box appears. Configure the action resources:

#### To select a topic

- a. Select **Notifications**.
- b. Select the **Notifications Compartment**.
- c. Select the **Topic**.

#### To select a stream

- a. Select **Streaming**.
- b. Select the **Stream Compartment**.
- c. Select the **Stream**.

#### To select a function

- a. Select **Functions**.
  - b. Select the **Functions Compartment**.
  - c. Select a **Functions Application**.
  - d. Select the **Functions ID**.
- **Action State:** Select to enable the action. Clear to disable.

5. Click **Add Action**.

### To edit an action

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Actions**.
4. Select an action.
5. Go to **Actions** and click **Edit**.  
The **Edit Action** box appears.
6. Make your changes and click **Save Changes**.

### To enable or disable an action

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Actions**.
4. Select an action.
5. Go to **Actions** and specify **Enable** or **Disable**.
6. Confirm when prompted.

### To remove an action

1. Open the navigation menu. Under the **Solutions and Platform** group, go to

**Application Integration** and click **Events Service**.

2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Actions**.
4. Select an action.
5. Go to **Actions** and click **Remove**.
6. Confirm when prompted.  
Each rule must have one action.

### To add event types to a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Click **Add Event Type**.
5. In **Service Name**, select a service.
6. In **Event Type**, select an event type for this service.
7. Click **Add Event Type**.

### To edit event types for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.

4. Select an event type.
5. Click **Edit**.  
The **Edit Event Type** box appears.
6. Make your changes and click **Save Changes**.

### To remove event types for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Select the check box next to the event types you want to remove.  
**Tip:** To select the entire list, select the check box in the header row.
5. Click **Remove**.
6. Confirm when prompted.

### To add attributes to a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Click **Add Attribute**.  
The **Add Attribute** box appears. Configure the attribute:

- **Attribute Name:** Specify an attribute or tag to narrow matching results.
  - Select an attribute name. The list of attribute names is based on the event types you selected. If you select no event types, you cannot add an attribute.
  - If you specify an attribute here, you limit the events that match this rule.
- **Attribute Values:** Specify one or more values for the attribute name.
  - a. Enter a value. As you type, the value appears under the field with (New) appended. Select the value with (New) appended to add the value to **Attribute Values**.



The screenshot shows a dialog box titled "ATTRIBUTE VALUES". Inside, there is a text input field containing "Value 1". Below the input field, a dropdown menu is open, showing "Value 1 (New)" as the selected option. The dropdown menu is highlighted in light blue.

- b. Enter more values for attribute name in the same manner as before.



The screenshot shows the "ATTRIBUTE VALUES" dialog box with two values added. The input field now contains "Value 1 Value 2". Each value has a small 'x' icon to its left, indicating it can be removed. A dropdown menu is open below the input field, showing "Value 2 (New)" as the selected option. The dropdown menu is highlighted in light blue.

Here are some things to consider about attribute values:

- Use an asterisk to create a wildcard. See [Examples of Wildcard Matching in Filters](#).
  - Multiple values for an attribute name broaden your results. If any of the values you enter here match a value in an event, the rule matches. See [Examples of Arrays in Filters](#).
5. Click **Add attribute**.

### To edit attributes for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Select an attribute.
5. Click **Edit**.  
The **Edit Attribute** box appears.
6. Make your changes and click **Save Changes**.

### To remove attributes for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Select the check box next to the attributes you want to remove.  
**Tip:** To select the entire list, select the check box in the header row.
5. Click **Remove**.
6. Confirm when prompted.

### To add filter tags to a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.

2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Click **Add Filter Tag**.
5. In **Tag Namespace**, do one of the following:
  - Select a namespace to add a defined tag as a filter.
  - Select **None (apply a free-form tag)** to add a free-form tag as a filter.
6. In **Tag Key**, do one of the following:
  - Select the tag key for the defined tag.
  - Enter the tag key for the free-form tag.
7. Enter a **Value**.
8. Click **Add Filter Tag**.

### To edit filter tags for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Select a filter tag.
5. Click **Edit**.  
The **Edit Attribute** box appears.
6. Make your changes and click **Save Changes**.

### To remove filter tags for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.
2. Choose a **Compartment** you have permission to work in, and then click the **Name** of the rule you want to update.
3. In the **Resources** menu, click **Event Matching**.
4. Select the check box next to the filter tags you want to remove.  
**Tip:** To select the entire list, select the check box in the header row.
5. Click **Remove**.
6. Confirm when prompted.

### Using the Command Line Interface (CLI)

When you use the CLI to create a rule, you work a little differently than using the Console.

- To specify the actions for your rule, use a JSON formatted file. You create this file before you create the rule, and the file simplifies the amount of information you must type at the command line.
- To specify an event to match, use a JSON formatted string. You type this right into the console as you create the rule.

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To create an action JSON file

To specify the actions for your rule, use a JSON formatted file. For more information, see [Using a JSON File for Complex Input](#).

1. Create a file and add the following content. This content doesn't have to be escaped or on a single line, it just has to contain valid JSON.

```
{
 "actions": [
 {
 "actionType": "FAAS",
 "description": "string",
 "functionId": "<function_OCID>",
 "isEnabled": true
 },
 {
 "actionType": "ONS",
 "description": "string",
 "isEnabled": true,
 "topicId": "<topic_OCID>"
 },
 {
 "actionType": "OSS",
 "description": "string",
 "isEnabled": true,
 "streamId": "<stream_OCID>"
 }
]
}
```

2. Edit the file and remove any objects you don't want to use as an action. For example, if you wanted to only use Notifications as an action, then you would delete all the other objects.

```
{
 "actions": [
 {
 "actionType": "ONS",
 "description": "string",
 "isEnabled": true,
 "topicId": "<topic_OCID>"
 }
]
}
```

3. Edit the file and fill in any *variables* with actual values from your tenancy, as shown in the following example.

```
{
 "actions": [
 {
 "actionType": "ONS",
 "description": "string",
 "isEnabled": true,
 "topicId": "<topic_OCID>"
 }
]
}
```

```
]
}
```

4. Add a description.
5. Save the file as `action.json`
6. To create a rule and specify Notifications as an action, run the following command.

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "{}" --
compartment-id <compartment_OCID> --actions file://action.json
```

### To create a rule

Open a command prompt and run `oci events rule create` to create a rule.

Use the following options:

- `display-name` indicates the name of the rule in the Console
- `is-enabled` indicates whether Events should evaluate the rule.
- `condition` a JSON formatted string used to indicate a pattern for event matching (see Examples for usage).

### Examples

The following example shows how to pass a simple condition that matches all events. Everything between the double quotes ( " ") is a string, while the brackets { } indicate JSON:

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "{}" --
compartment-id <compartment_OCID> --actions file://action.json --wait-for-state=ACTIVE
```

To pass complex input to the CLI as a JSON string, you must enclose the entire block in double quotes. Inside the block, each double quote for the key and value strings must be escaped with a backslash (\) character.

For example:

## CHAPTER 14 Events

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "{\n\"eventType\":[\n\"com.oraclecloud.objectstorage.createobject\"]}" --compartment-id <compartment_OCID> --actions file://action.json --wait-for-state=ACTIVE
```

In PowerShell, to escape double quotes, you must use two characters: The backslash (\) and the back tick (`).

For example, in Windows PowerShell:

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition "{\`\"eventType\`\":[\`\"com.oraclecloud.objectstorage.createobject\`\"]}" --compartment-id <compartment_OCID> --actions file://action.json --wait-for-state=ACTIVE
```



### Tip

The `condition` option does not support using a file to pass the JSON formatted string.

For information on creating filters, see [Matching Events with Filters](#).

- `compartment-id` indicates the compartment where the rule applies. Events evaluates messages from resources in this compartment *and any subordinate compartments*.
- `actions` indicates the location in the local file system of the JSON formatted file you created to specify the actions for a rule.
- `wait-for-state=` when used with `ACTIVE` indicates that the CLI should wait for the service to create the rule, do another `GET` operation, and then display the rule in the active state. Without the option, the CLI displays the rule immediately in the creating state.

For example:

```
oci events rule create --display-name <friendly_name> --is-enabled true --condition <json_formatted_string> --compartment-id <compartment_OCID> --actions <json_formatted_file> --wait-for-state=ACTIVE
```



### Note

Replace the values in *<compartment\_OCID>* and *<json\_formatted\_file>* with the actual values from your tenancy and the local file system.

### To delete a rule

Open a command prompt and run `oci events rule delete` to delete a single rule. For example:

```
oci events rule delete --rule-id <rule_OCID>
```

The command returns a prompt, asking for confirmation. Type `y` to delete the rule.

### To get rule metadata

You can get rule metadata using the CLI. The Console displays this metadata in the **Rule Details** tab.

Open a command prompt and run `oci events rule get` to get information about a single rule. For example:

```
oci events rule get --rule-id <rule_OCID>
```

The command returns the following information:

```
{
 "data": {
 "actions": {
 "actions": [
 {
 "action-type": "ONS",
 "description": null,
 "id": "ocidl.eventaction.ocl.phx.<unique_ID>",
 "lifecycle-message": null,

```

## CHAPTER 14 Events

---

```
 "lifecycle-state": "ACTIVE",
 "topic-id": "ocid1.onstopic.oc1.phx.<unique_ID>"
 }
],
"compartment-id": "ocid1.compartment.oc1..<unique_ID>",
"condition": "{\n \"eventType\": [\n
\n\"com.oraclecloud.databaseservice.autonomous.datawarehouse.backup.end\", \n \"CustomEventType\"\n
]\n}",
"defined-tags": null,
"description": null,
"display-name": "rule_name",
"freeform-tags": null,
"id": "ocid1.eventrule.oc1.phx.<unique_ID>",
"is-enabled": true,
"lifecycle-message": null,
"lifecycle-state": "ACTIVE",
"time-created": "2019-01-23T00:48:20.155000+00:00"
},
"etag": "<unique_ID>--gzip"
}
```

### To get a list of rules

Open a command prompt and run `oci events rule list` to list the rules in a compartment. For example:

```
oci events rule list --compartment-id <compartment_OCID>
```

The command returns the following information:

```
{
 "data": [
 {
 "compartment-id": "ocid1.compartment.oc1..<unique_ID>",
 "condition": "{}",
 "description": "Example_Rule",
 "display-name": "rule_1",
 "id": "ocid1.eventrule.oc1.phx.<unique_ID>",
 "is-enabled": true,
 "lifecycle-state": "ACTIVE",
```

## CHAPTER 14 Events

---

```
 "time-created": "2019-01-22T20:10:53.562000+00:00"
 },
 {
 "compartment-id": "ocidl.compartment.oc1..<unique_ID>",
 "condition": "{}",
 "description": null,
 "display-name": "rule_2",
 "id": "ocidl.eventrule.oc1.phx.<unique_ID>",
 "is-enabled": true,
 "lifecycle-state": "ACTIVE",
 "time-created": "2019-01-22T20:27:25.099000+00:00"
 },
 ...

 {
 "compartment-id": "ocidl.compartment.oc1..<unique_ID>",
 "condition": "{\"eventType\": [\"com.oraclecloud.objectstorage.createobject\"]}",
 "description": null,
 "display-name": "rule_75",
 "id": "ocidl.eventrule.oc1.phx.<unique_ID>",
 "is-enabled": true,
 "lifecycle-state": "ACTIVE",
 "time-created": "2019-01-22T23:08:12.379000+00:00"
 }
]
}
```

### To update a rule

Open a command prompt and run `oci events rule update` to update a rule.

To update the condition for a rule:

```
oci events rule update --rule-id <rule_OCID> --condition <json_formatted_string>
```

For example:

```
oci events rule update --rule-id ocidl.eventrule.oc1.phx.<unique_ID> --condition "{}" --wait-for-state=ACTIVE
```

## CHAPTER 14 Events

---

The previous command would update the condition of the rule to use an empty JSON string. The CLI updates the rule, waits for the rule to update and change to the active state (only if you used the `--wait-for-state` option), then displays the updated rule.

Use the following options to update a rule:

- `display-name`
- `description`
- `is-enabled`
- `condition`
- `actions`
- `freeform-tags`
- `defined-tags`

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage rules:

- [ChangeRuleCompartment](#)
- [CreateRule](#)
- [DeleteRule](#)
- [GetRule](#)
- [UpdateRule](#)
- [ListRules](#)

## Contents of an Event Message

This topic describes the contents of an event message. Every event message includes two main parts:

- Envelope: a container for all event messages
- Payload: the data from the resource emitting the event message

### Event Envelope

These attributes for an event envelope are the same for all events. The structure of the envelope follows the [CloudEvents](#) industry standard format hosted by the [Cloud Native Computing Foundation \(CNCF\)](#).

Property	Description
<code>cloudEventsVersion</code>	<p>The version of the CloudEvents specification.</p> <div style="border: 1px solid #0070c0; background-color: #e6f2ff; padding: 10px; margin-top: 10px;">  <p><b>Note</b></p> <p>Events uses version 0.1 specification of the CloudEvents event envelope.</p> </div>
<code>contentType</code>	Set to <code>application/json</code> . The content type of the data contained in the <code>data</code> attribute.
<code>data</code>	The payload of the event. All of the information within <code>data</code> comes from the resource emitting the event. See the following table for more detail on the structure of the payload.
<code>eventID</code>	The UUID of the event. This identifier is not an OCID, but just a unique ID for the event.
<code>eventTime</code>	The time of the event, expressed in <a href="#">RFC 3339</a> timestamp format.

## CHAPTER 14 Events

Property	Description
<code>eventType</code>	<p>The type of event that happened. For a list of all services that produce events and the even types that those services track, see <a href="#">_ Services that Produce Events</a>.</p> <div data-bbox="656 575 1281 919"><p><b>Note</b></p><p>The service that produces the event can also add, remove, or change the meaning of a field by publishing a new version of an <code>eventType</code> and revising the <code>eventTypeVersion</code> field.</p></div>
<code>eventTypeVersion</code>	The version of the event type.
<code>extensions</code>	The OCID of the compartment from which the event originates. If the event originates from the root compartment of the tenancy, then this attribute specifies a tenancy OCID. This attribute is mandatory in the Oracle Cloud Infrastructure implementation of the <a href="#">CloudEvents</a> specification.
<code>source</code>	The resource that produced the event. For example, an Autonomous Database or an Object Storage bucket.

### Payload

The data in these fields depends on which service produced the event and the event type it defines.

## CHAPTER 14 Events

Property	Description
compartmentId	The OCID of the compartment of the resource emitting the event.
compartmentName	The name of the compartment of the resource emitting the event.
resourceName	The name of the resource emitting the event.
resourceId	An OCID or an ID for the resource emitting the event.
availabilityDomain	The availability domain of the resource emitting the event.
freeFormTags	Free-form tags added to the resource emitting the event.
definedTags	Defined tags added to the resource emitting the event.
additionalDetails	<p>A container for attributes unique to the resource emitting the event. In the example bucket event that follows, the payload includes three Object Storage attributes:</p> <ul style="list-style-type: none"><li>• namespace</li><li>• publicAccessType</li><li>• eTag</li></ul> <p>To determine what attributes are included for other resources, retrieve an event or consult the reference samples listed on <a href="#">_ Services that Produce Events</a>.</p>

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### An Example Event

The following is an example bucket event emitted by Object Storage.

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.objectstorage.deletebucket",
 "source": "objectstorage",
 "eventTypeVersion": "1.0",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_bucket",
 "resourceId": "ocidl.compartment.oc1..<unique_ID>",
 "availabilityDomain": "NfHZ:PHX-AD-2",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "namespace": "example_namespace",
 "publicAccessType": "NoPublicAccess",
 "eTag": "f8ffb6e9-f602-460f-a6c0-00b5abfa24c7"
 }
 }
}
```

### Services that Produce Events

This topic lists the Oracle Cloud Infrastructure services that emit events:

- [Analytics Cloud](#)
- [Block Volume](#)
- [Compute](#)
- [Database](#)
- [File Storage](#)
- [Functions](#)
- [IAM](#)
- [Integration](#)
- [Networking](#)
- [Notifications](#)
- [Object Storage](#)
- [Resource Manager](#)

### About Event Types and Example Reference Events

Services emit event messages by resource type. Event messages use a combination of an event type and a data payload (from the resource) to identify state changes.

In this section:

- Event types are organized by service, then by resource type
- There is one reference example per resource type if the payload contains the same attributes for all event types

See [Matching Events with Filters](#) and [Contents of an Event Message](#).

### Analytics Cloud

For details about events emitted by Analytics Cloud, see [Service Events](#).

## Block Volume

Block Volume resources that emit events:

- [Block Volumes](#) and [Block Volume Backups](#)
- [Boot Volumes](#) and [Boot Volume Backups](#)
- [Volume Groups and Volume Group Backups](#)

### Block Volume Event Types

These are the event types that block volumes emit:

Friendly Name	Event Type
Change Volume Compartment Begin	com.oraclecloud.blockvolumes.changevolumecompartment.begin
Change Volume Compartment End	com.oraclecloud.blockvolumes.changevolumecompartment.end
Create Volume Begin	com.oraclecloud.blockvolumes.createvolume.begin
Create Volume End	com.oraclecloud.blockvolumes.createvolume.end
Delete Volume Begin	com.oraclecloud.blockvolumes.deletevolume.begin
Delete Volume End	com.oraclecloud.blockvolumes.deletevolume.end
Delete Volume Kms Key Begin	com.oraclecloud.blockvolumes.deletevolumekmskey.begin
Update Volume	com.oraclecloud.blockvolumes.updatevolume
Update Volume Begin	com.oraclecloud.blockvolumes.updatevolume.begin
Update Volume End	com.oraclecloud.blockvolumes.updatevolume.end
Update Volume Kms Key Begin	com.oraclecloud.blockvolumes.updatevolumekmskey.begin
Update Volume Kms Key End	com.oraclecloud.blockvolumes.updatevolumekmskey.end

### Block Volume Example

This is a reference event for block volumes:

```
{
 "eventType": "com.oraclecloud.blockvolumes.createvolume.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "BlockVolumes",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_volume",
 "resourceId": "ocidl.volume.oc1..<unique_ID>",
 "availabilityDomain": "<availability_domain>",
 }
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Block Volume Backup Event Types

These are the event types that block volume backups emit:

Friendly Name	Event Type
Change Volume Backup Compartment	com.oraclecloud.blockvolumes.changevolumebackupcompartment
Copy Volume Backup Begin	com.oraclecloud.blockvolumes.copyvolumebackup.begin
Copy Volume Backup End	com.oraclecloud.blockvolumes.copyvolumebackup.end
Create Volume Backup Begin	com.oraclecloud.blockvolumes.createvolumebackup.begin

## CHAPTER 14 Events

Friendly Name	Event Type
Create Volume Backup End	com.oraclecloud.blockvolumes.createvolumebackup.end
Create Volume Backup Policy Assignment	com.oraclecloud.blockvolumes.createvolumebackuppolicyassignment
Delete Volume Backup Begin	com.oraclecloud.blockvolumes.deletevolumebackup.begin
Delete Volume Backup End	com.oraclecloud.blockvolumes.deletevolumebackup.end
Delete Volume Backup Policy Assignment	com.oraclecloud.blockvolumes.deletevolumebackuppolicyassignment
Update Volume Backup	com.oraclecloud.blockvolumes.updatevolumebackup
Update Volume Backup Policy	com.oraclecloud.blockvolumes.updatevolumebackuppolicy

### Block Volume Backup Example

This is a reference event for block volume backups:

```
{
 "eventType": "com.oraclecloud.blockvolumes.createvolumebackup.end",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "BlockVolumes",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_volumebackup via policy:gold",
 "resourceId": "ocidl.volumebackup.oc1..<unique_ID>",
 "additionalDetails": {
 "sourceType": "SCHEDULED",
 "volumeId": "ocidl.volume.oc1..<unique_ID>"
 }
 }
}
"eventID": "<unique_ID>",
"extensions": {
```

## CHAPTER 14 Events

```
"compartmentId": "ocidl.compartment.ocl..<unique_ID>"
}
```

### Boot Volume Event Types

These are the event types that boot volumes emit:

Friendly Name	Event Type
Change Boot Volume Compartment Begin	com.oraclecloud.blockvolumes.changebootvolumecompartment.begin
Change Boot Volume Compartment End	com.oraclecloud.blockvolumes.changebootvolumecompartment.end
Create Boot Volume Begin	com.oraclecloud.blockvolumes.createbootvolume.begin
Create Boot Volume End	com.oraclecloud.blockvolumes.createbootvolume.end
Delete Boot Volume Begin	com.oraclecloud.blockvolumes.deletebootvolume.begin
Delete Boot Volume End	com.oraclecloud.blockvolumes.deletebootvolume.end
Delete Boot Volume Kms Key Begin	com.oraclecloud.blockvolumes.deletebootvolumekmskey.begin
Update Boot Volume	com.oraclecloud.blockvolumes.updatebootvolume
Update Boot Volume Begin	com.oraclecloud.blockvolumes.updatebootvolume.begin
Update Boot Volume End	com.oraclecloud.blockvolumes.updatebootvolume.end
Update Boot Volume Kms Key Begin	com.oraclecloud.blockvolumes.updatebootvolumekmskey.begin
Update Boot Volume Kms Key End	com.oraclecloud.blockvolumes.updatebootvolumekmskey.end

## Boot Volume Example

This is a reference event for boot volumes:

```
{
 "eventType": "com.oraclecloud.blockvolumes.createbootvolume.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "BlockVolumes",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_volume",
 "resourceId": "ocidl.volume.oc1..<unique_ID>",
 "availabilityDomain": "<availability_domain>",
 }
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

## Boot Volume Backup Event Types

These are the event types that boot volume backups emit:

Friendly Name	Event Type
Change Boot Volume Backup Compartment	com.oraclecloud.blockvolumes.changebootvolumebackupcompartment
Create Boot Volume Backup Begin	com.oraclecloud.blockvolumes.createbootvolumebackup.begin
Create Boot Volume Backup End	com.oraclecloud.blockvolumes.createbootvolumebackup.end

## CHAPTER 14 Events

Friendly Name	Event Type
Delete Boot Volume Backup Begin	com.oraclecloud.blockvolumes.deletebootvolumebackup.begin
Delete Boot Volume Backup End	com.oraclecloud.blockvolumes.deletebootvolumebackup.end
Update Boot Volume Backup	com.oraclecloud.blockvolumes.updatebootvolumebackup

### Boot Volume Backup Example

This is a reference event for boot volume backups:

```
{
 "eventType": "com.oraclecloud.blockvolumes.createbootvolume.end",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "BlockVolumes",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_volumebackup via policy:gold",
 "resourceId": "ocid1.volumebackup.oc1..<unique_ID>",
 "additionalDetails": {
 "sourceType": "SCHEDULED",
 "volumeId": "ocid1.volume.oc1..<unique_ID>"
 }
 }
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 }
}
```

## CHAPTER 14 Events

### Volume Groups and Volume Group Backups

These are the event types that volume groups and volume group backups emit:

Friendly Name	Event Type
Change Volume Group Compartment	<code>com.oraclecloud.blockvolumes.changevolumegroupcompartment</code>
Change Volume Group Backup Compartment	<code>com.oraclecloud.blockvolumes.changevolumegroupbackupcompartment</code>
Create Volume Group	<code>com.oraclecloud.blockvolumes.createvolumegroup</code>
Create Volume Group Begin	<code>com.oraclecloud.blockvolumes.createvolumegroup.begin</code>
Create Volume Group End	<code>com.oraclecloud.blockvolumes.createvolumegroup.end</code>
Create Volume Group Backup Begin	<code>com.oraclecloud.blockvolumes.createvolumegroupbackup.begin</code>
Create Volume Group Backup End	<code>com.oraclecloud.blockvolumes.createvolumegroupbackup.end</code>
Delete Volume Group Begin	<code>com.oraclecloud.blockvolumes.deletevolumegroup.begin</code>
Delete Volume Group End	<code>com.oraclecloud.blockvolumes.deletevolumegroup.end</code>
Delete Volume Group Backup Begin	<code>com.oraclecloud.blockvolumes.deletevolumegroupbackup.begin</code>
Delete Volume Group Backup End	<code>com.oraclecloud.blockvolumes.deletevolumegroupbackup.end</code>
Update Volume Group	<code>com.oraclecloud.blockvolumes.updatevolumegroup</code>
Update Volume Group Backup	<code>com.oraclecloud.blockvolumes.updatevolumegroupbackup</code>

### Volume Group Example

This is a reference event for volume groups:

```
{
 "eventType": "com.oraclecloud.blockvolumes.createvolumegroup",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "BlockVolumes",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "resourceName": "my_volumegroup",
 "resourceId": "ocidl.volumegroup.oc1..<unique_ID>",
 "availabilityDomain": "<availability_domain>",
 }
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Compute

Compute resources that emit events:

- [Autoscaling configurations and autoscaling policies](#)
- [Cluster networks](#)
- [Console histories](#)
- [Images](#)
- [Instances and instance attachments](#)
- [Instance configurations](#)
- [Instance console connections](#)
- [Instance pools](#)

## Autoscaling Event Types

These are the event types that autoscaling configurations and autoscaling policies emit:

Friendly Name	Event Type
Change Autoscaling Configuration Compartment	<code>com.oraclecloud.autoscaling.changeautoscalingconfigurationcompartment</code>
Create Autoscaling Configuration	<code>com.oraclecloud.autoscaling.createautoscalingconfiguration</code>
Delete Autoscaling Configuration	<code>com.oraclecloud.autoscaling.deleteautoscalingconfiguration</code>
Scaling Action	<code>com.oraclecloud.autoscaling.scalingaction</code>
Update Autoscaling Configuration	<code>com.oraclecloud.autoscaling.updateautoscalingconfiguration</code>
Update Autoscaling Policy	<code>com.oraclecloud.autoscaling.updateautoscalingpolicy</code>

## Autoscaling Example

This is a reference event for autoscaling:

```
{
 "eventType": "com.oraclecloud.autoscaling.scalingaction",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "autoscaling",
 "eventTime": "2019-08-21T04:00:10.046Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "example autoscaling configuration",
 "resourceId": "ocid1.autoscalingconfiguration.oc1.phx.<unique_ID>",
 }
}
```

## CHAPTER 14 Events

```
"additionalDetails": {
 "policyName": "my_policy_name",
 "ruleName": "my_scale_up_condition",
 "actionType": "SCALE_OUT",
 "previousSize": 1,
 "newSize": 2
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
```

### Cluster Network Event Types

These are the event types that cluster networks emit:

Friendly Name	Event Type
Change Cluster Network Compartment	com.oraclecloud.computemanagement.changeclusternetworkcompartment
Create Cluster Network Begin	com.oraclecloud.computemanagement.createclusternetwork.begin
Create Cluster Network End	com.oraclecloud.computemanagement.createclusternetwork.end
Terminate Cluster Network Begin	com.oraclecloud.computemanagement.terminateclusternetwork.begin
Terminate Cluster Network End	com.oraclecloud.computemanagement.terminateclusternetwork.end

### Cluster Networks Example

This is a reference event for most cluster network events:

```
{
 "eventType": "com.oraclecloud.computemanagement.createclusternetwork.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
}
```

## CHAPTER 14 Events

```
"source": "ComputeManagement",
"eventTime": "2019-09-12T21:45:09.036Z",
"contentType": "application/json",
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_cluster_network",
 "resourceId": "ocidl.clusternetwork.oc1.uk-london-1.<unique_ID>",
 "availabilityDomain": "<availability_domain>"
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

Create cluster network end and terminate cluster network end don't include the availability domain.

### Console History Event Types

These are the event types that console histories emit:

Friendly Name	Event Type
Capture Console History Begin	com.oraclecloud.computeapi.captureconsolehistory.begin
Capture Console History End	com.oraclecloud.computeapi.captureconsolehistory.end
Delete Console History	com.oraclecloud.computeapi.deleteconsolehistory
Update Console History	com.oraclecloud.computeapi.updateconsolehistory

### Console History Example

This is a reference event for console histories:

```
{
 "eventType": "com.oraclecloud.computeapi.captureconsolehistory.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
}
```

## CHAPTER 14 Events

```
"source": "ComputeApi",
"eventTime": "2019-08-20T21:58:13.554Z",
"contentType": "application/json",
"data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceId": "ocid1.consolehistory.oc1.iad.<unique_ID>",
 "availabilityDomain": "SoSC:PHX-AD-3"
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
```

### Image Event Types

These are the event types that images emit:

Friendly Name	Event Type
Add Image Shape Compatibility	com.oraclecloud.computeapi.addimageshapecompatibility
Change Image Compartment	com.oraclecloud.computeapi.moveimage
Create Image Begin	com.oraclecloud.computeapi.createimage.begin
Create Image End	com.oraclecloud.computeapi.createimage.end
Delete Image	com.oraclecloud.computeapi.deleteimage
Export Image Begin	com.oraclecloud.computeapi.exportimage.begin
Export Image End	com.oraclecloud.computeapi.exportimage.end
Remove Image Shape Compatibility	com.oraclecloud.computeapi.removeimageshapecompatibility
Update Image	com.oraclecloud.computeapi.updateimage

### Image Example

## CHAPTER 14 Events

This is a reference event for most image events:

```
{
 "eventType": "com.oraclecloud.computeapi.exportimage.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeApi",
 "eventTime": "2019-08-27T04:12:37.397Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.tenancy.ocl..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_image",
 "resourceId": "ocidl.image.ocl.iad.<unique_ID>",
 "availabilityDomain": "SoSC:PHX-AD-3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.tenancy.ocl..<unique_ID>"
 }
}
```

Change image compartment doesn't include the resource name or availability domain.

### Instance Event Types

These are the event types that Compute instances and instance attachments emit:

Friendly Name	Event Type
Attach Boot Volume Begin	com.oraclecloud.computeapi.attachbootvolume.begin
Attach Boot Volume End	com.oraclecloud.computeapi.attachbootvolume.end
Attach Secondary VNIC Begin	com.oraclecloud.computeapi.attachvnic.begin
Attach Secondary VNIC End	com.oraclecloud.computeapi.attachvnic.end
Attach Volume Begin	com.oraclecloud.computeapi.attachvolume.begin

## CHAPTER 14 Events

Friendly Name	Event Type
Attach Volume End	com.oraclecloud.computeapi.attachvolume.end
Change Instance Compartment Begin	com.oraclecloud.computeapi.changeinstancecompartment.begin
Change Instance Compartment End	com.oraclecloud.computeapi.changeinstancecompartment.end
Detach Boot Volume Begin	com.oraclecloud.computeapi.detachbootvolume.begin
Detach Boot Volume End	com.oraclecloud.computeapi.detachbootvolume.end
Detach Secondary VNIC Begin	com.oraclecloud.computeapi.detachvnic.begin
Detach Secondary VNIC End	com.oraclecloud.computeapi.detachvnic.end
Detach Volume Begin	com.oraclecloud.computeapi.detachvolume.begin
Detach Volume End	com.oraclecloud.computeapi.detachvolume.end
Instance Action Begin	com.oraclecloud.computeapi.instanceaction.begin
Instance Action End	com.oraclecloud.computeapi.instanceaction.end
Launch Instance Begin	com.oraclecloud.computeapi.launchinstance.begin
Launch Instance End	com.oraclecloud.computeapi.launchinstance.end
Terminate Instance Begin	com.oraclecloud.computeapi.terminateinstance.begin
Terminate Instance End	com.oraclecloud.computeapi.terminateinstance.end
Update Instance	com.oraclecloud.computeapi.updateinstance

### Compute Instance Example

## CHAPTER 14 Events

---

This is a reference event for most instance events (attach/detach volume and boot volume events don't include additional details):

```
{
 "eventType": "com.oraclecloud.computeapi.launchinstance.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeApi",
 "eventTime": "2019-08-15T21:21:48.586Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_instance",
 "resourceId": "ocid1.instance.oc1.phx.<unique_ID>",
 "availabilityDomain": "SoSC:PHX-AD-3",
 "additionalDetails": {
 "imageId": "ocid1.image.oc1.phx.<unique_ID>",
 "shape": "VM.Standard2.1",
 "type": "CustomerVmi"
 }
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 }
}
```

This is a reference event for attach/detach VNIC events:

```
{
 "eventType": "com.oraclecloud.computeapi.attachvnic.end",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeApi",
 "eventTime": "2019-08-15T21:21:48.586Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_instance",
 "resourceId": "ocid1.instance.oc1.phx.<unique_ID>",
 }
}
```

## CHAPTER 14 Events

```
"availabilityDomain": "SoSC:PHX-AD-3",
"additionalDetails": {
 "subnetId": "ocidl.subnet.oc1.phx.<unique_ID>"
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

### Instance Configuration Event Types

These are the event types that Compute instance configurations emit:

Friendly Name	Event Type
Change Instance Configuration Compartment	com.oraclecloud.computemanagement.changeinstanceconfigurationcompartment
Create Instance Configuration	com.oraclecloud.computemanagement.createinstanceconfiguration
Delete Instance Configuration	com.oraclecloud.computemanagement.deleteinstanceconfiguration
Launch Instance Configuration Begin	com.oraclecloud.computemanagement.launchinstanceconfiguration.begin
Launch Instance Configuration End	com.oraclecloud.computemanagement.launchinstanceconfiguration.end
Update Instance Configuration	com.oraclecloud.computemanagement.updateinstanceconfiguration

### Compute Instance Configuration Example

## CHAPTER 14 Events

This is a reference event for most instance configuration events:

```
{
 "eventType": "com.oraclecloud.computemanagement.createinstanceconfiguration",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeManagement",
 "eventTime": "2019-08-12T22:52:01.062Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_instance_configuration",
 "resourceId": "ocidl.instanceconfiguration.oc1.phx..<unique_ID>",
 "availabilityDomain": "<availability_domain>"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

Launch instance configuration end doesn't include the availability domain.

### Instance Console Connection Event Types

These are the event types that Compute instance console connections emit:

Friendly Name	Event Type
Create Instance Console Connection Begin	com.oraclecloud.computeapi.createinstanceconsoleconnection.begin
Create Instance Console Connection End	com.oraclecloud.computeapi.createinstanceconsoleconnection.end
Delete Instance Console Connection Begin	com.oraclecloud.computeapi.deleteinstanceconsoleconnection.begin
Delete Instance Console Connection End	com.oraclecloud.computeapi.deleteinstanceconsoleconnection.end

### Compute Instance Console Connection Example

This is a reference event for instance console connections:

```
{
 "eventType": "com.oraclecloud.computeapi.createinstanceconsoleconnection.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeApi",
 "eventTime": "2019-08-12T14:47:35.762Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceId": "ocidl.instanceconsoleconnection.oc1.phx.<unique_ID>",
 "availabilityDomain": "SoSC:PHX-AD-3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Instance Pool Event Types

These are the event types that Compute instance pools emit:

Friendly Name	Event Type
Attach Load Balancer Begin	com.oraclecloud.computemanagement.attachloadbalancer.begin
Attach Load Balancer End	com.oraclecloud.computemanagement.attachloadbalancer.end
Change Instance Pool Compartment	com.oraclecloud.computemanagement.changeinstancepoolcompartment
Create Instance Pool Begin	com.oraclecloud.computemanagement.createinstancepool.begin
Create Instance Pool End	com.oraclecloud.computemanagement.createinstancepool.end

## CHAPTER 14 Events

Friendly Name	Event Type
Detach Load Balancer Begin	com.oraclecloud.computemanagement.detachloadbalancer.begin
Detach Load Balancer End	com.oraclecloud.computemanagement.detachloadbalancer.end
Reset Instance Pool Begin	com.oraclecloud.computemanagement.resetinstancepool.begin
Reset Instance Pool End	com.oraclecloud.computemanagement.resetinstancepool.end
Soft Reset Instance Pool Begin	com.oraclecloud.computemanagement.softresetinstancepool.begin
Soft Reset Instance Pool End	com.oraclecloud.computemanagement.softresetinstancepool.end
Start Instance Pool Begin	com.oraclecloud.computemanagement.startinstancepool.begin
Start Instance Pool End	com.oraclecloud.computemanagement.startinstancepool.end
Stop Instance Pool Begin	com.oraclecloud.computemanagement.stopinstancepool.begin
Stop Instance Pool End	com.oraclecloud.computemanagement.stopinstancepool.end
Terminate Instance Pool Begin	com.oraclecloud.computemanagement.terminateinstancepool.begin
Terminate Instance Pool End	com.oraclecloud.computemanagement.terminateinstancepool.end
Update Instance Pool Begin	com.oraclecloud.computemanagement.updateinstancepool.begin
Update Instance Pool End	com.oraclecloud.computemanagement.updateinstancepool.end

### Compute Instance Pools Example

This is a reference event for most instance pool events:

```
{
 "eventType": "com.oraclecloud.computemanagement.createinstancepool.begin",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "ComputeManagement",
 "eventTime": "2019-08-12T22:52:01.343Z",
 "contentType": "application/json",
```

## CHAPTER 14 Events

---

```
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "my_instance_pool",
 "resourceId": "ocidl.instancepool.oc1.phx.<unique_ID>",
 "availabilityDomain": "<availability_domain>"
},
"eventID": "<unique_id>",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

These instance pool events don't include the availability domain: create instance pool end, detach load balancer end, reset instance pool end, soft reset instance pool end, start instance pool end, stop instance pool end, terminate instance pool end, and update instance pool end.

## Database

Database resources that emit events:

- **Autonomous Database Resources:**
  - [Autonomous Databases](#)
  - [Autonomous Container Databases](#)
  - [Autonomous Exadata Infrastructure instances](#)
- **Exadata Cloud at Customer Resources:**
  - [Exadata Infrastructure](#)
  - [VM cluster networks](#)
  - [VM clusters](#)
  - [Backup destinations](#)
  - [Database nodes](#)
  - [Database Homes](#)
  - [Databases](#)

• **Bare metal, virtual machine, and Exadata DB system resources**

- [DB systems](#)
- [Database nodes](#)
- [Database Homes](#)
- [Databases](#)
- [Data Guard associations](#)

**Autonomous Database Event Types**

These are the event types that Autonomous Databases emit:

Friendly Name	Event Type
Change Compartment Begin	com.oraclecloud.databaseservice.changeautonomousdatabasecompartment.begin
Change Compartment End	com.oraclecloud.databaseservice.changeautonomousdatabasecompartment.end
Create Backup Begin	com.oraclecloud.databaseservice.autonomous.database.backup.begin
Create Backup End	com.oraclecloud.databaseservice.autonomous.database.backup.end
Create Begin	com.oraclecloud.databaseservice.autonomous.database.instance.create.begin
Create End	com.oraclecloud.databaseservice.autonomous.database.instance.create.end
Restore Begin	com.oraclecloud.databaseservice.autonomous.database.restore.begin
Restore End	com.oraclecloud.databaseservice.autonomous.database.restore.end
Start Begin	com.oraclecloud.databaseservice.startautonomousdatabase.begin
Start End	com.oraclecloud.databaseservice.startautonomousdatabase.end

## CHAPTER 14 Events

Friendly Name	Event Type
Stop Begin	com.oraclecloud.databaseservice.stopautonomousdatabase.begin
Stop End	com.oraclecloud.databaseservice.stopautonomousdatabase.end
Terminate Begin	com.oraclecloud.databaseservice.deleteautonomousdatabase.begin
Terminate End	com.oraclecloud.databaseservice.deleteautonomousdatabase.end

### Autonomous Database Example

This is a reference event for Autonomous Databases:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.databaseservice.autonomous.database.backup.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-07-10T14:06:23Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_database",
 "resourceId": "ocid1.autonomousdatabase.oc1.phx.<unique_ID>",
 "availabilityDomain": "SoSC:PHX-AD-3",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "cpuCoreCount": 1,
 "lifecycleState": "PROVISIONING",
 "dataStorageSizeInTBs": 1,
 "timeCreated": "2019-07-10T14:06:10.905Z",
 "timeUpdated": "2019-07-10T14:06:10.905Z",
 "serviceConsoleUrl": null,
 "licenseType": null,
 }
 }
}
```

## CHAPTER 14 Events

```
"workloadType": "<Data Warehouse | Transaction Processing>",
"autonomousDatabaseType": "<Dedicated Infrastructure | Severless>"
}
}
}
```

### Autonomous Container Database Event Types

These are the event types that Autonomous Container Databases emit:

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changeautonomouscontainerdatabasecompartment
Create Backup Begin	com.oraclecloud.databaseservice.autonomous.container.database.backup.begin
Create Backup End	com.oraclecloud.databaseservice.autonomous.container.database.backup.end
Create Begin	com.oraclecloud.databaseservice.autonomous.container.database.instance.create.begin
Create End	com.oraclecloud.databaseservice.autonomous.container.database.instance.create.end
Maintenance Begin	com.oraclecloud.databaseservice.autonomous.container.database.maintenance.begin
Maintenance End	com.oraclecloud.databaseservice.autonomous.container.database.maintenance.end
Maintenance Reminder	com.oraclecloud.databaseservice.autonomous.container.database.maintenance.reminder
Maintenance Scheduled	com.oraclecloud.databaseservice.autonomous.container.database.maintenance.scheduled

## CHAPTER 14 Events

Friendly Name	Event Type
Restart Begin	com.oraclecloud.databaseservice.restartautonomouscontainerdatabase.begin
Restart End	com.oraclecloud.databaseservice.restartautonomouscontainerdatabase.end
Restore Begin	com.oraclecloud.databaseservice.autonomous.container.database.restore.begin
Restore End	com.oraclecloud.databaseservice.autonomous.container.database.restore.end
Terminate Begin	com.oraclecloud.databaseservice.terminateautonomouscontainerdatabase.begin
Terminate End	com.oraclecloud.databaseservice.terminateautonomouscontainerdatabase.end
Update Begin	com.oraclecloud.databaseservice.autonomous.container.database.instance.update.begin
Update End	com.oraclecloud.databaseservice.autonomous.container.database.instance.update.begin

### Autonomous Container Database Example

This is a reference event for Autonomous Container Databases:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.databaseservice.autonomous.container.database.backup.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-06-27T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_container_database",
 }
}
```

## CHAPTER 14 Events

```
"resourceId": "<unique_ID>",
"availabilityDomain": "all",
"freeFormTags": {},
"definedTags": {},
"additionalDetails": {
 "cpuCoreCount": null,
 "lifecycleState": "ACTIVE",
 "dataStorageSizeInTBs": null,
 "timeCreated": "2019-06-27T21:15:59.000Z",
 "timeUpdated": "2019-06-27T21:16:04.389Z",
 "dbUniqueName": "dwrrdtsr_phx289",
 "dbHomeId": "ocid1.autonomoushome.oc1.phx.<unique_ID>",
 "dbName": "dwrrdtsr"
 "autonomousContainerDatabaseId": "ocid1.autonomouscontainerdatabase.oc1.phx.<unique_ID>"
}
}
```

### Autonomous Exadata Infrastructure Event Types

These are the event types that Autonomous Exadata Infrastructure instances emit:

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changeautonomousexadatainfrastructurecompartment
Create Begin	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.instance.create.begin
Create End	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.instance.create.end
Maintenance Begin	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.maintenance.begin

## CHAPTER 14 Events

Friendly Name	Event Type
Maintenance End	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.maintenance.end
Maintenance Reminder	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.maintenance.reminder
Maintenance Scheduled	com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.maintenance.scheduled
Terminate Begin	com.oraclecloud.databaseservice.terminateautonomousexadatainfrastructure.begin
Terminate End	com.oraclecloud.databaseservice.terminateautonomousexadatainfrastructure.end

### Autonomous Exadata Infrastructure Example

This is a reference event for Autonomous Exadata Infrastructure instances:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType":
"com.oraclecloud.databaseservice.autonomous.exadata.infrastructure.instance.create.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-07-10T23:28:12Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_exadata_instance",
 "resourceId": "ocidl.autonomousexainfrastructure.oc1.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-3",
```

## CHAPTER 14 Events

```
"freeFormTags": {},
"definedTags": {},
"additionalDetails": {
 "cpuCoreCount": 92,
 "lifecycleState": "TERMINATED",
 "dataStorageSizeInTBs": null,
 "timeCreated": "2019-07-10T23:13:43.136Z",
 "timeUpdated": "2019-07-10T23:28:12.390Z",
 "serviceConsoleUrl": null,
 "licenseType": null,
 "dbName": null
}
}
```

### Exadata Infrastructure Event Types

These are the event types that Exadata Infrastructure instances emit:

Friendly Name	Event Type
Activate Begin	com.oraclecloud.databaseservice.activateexadatainfrastructure.begin
Activate End	com.oraclecloud.databaseservice.activateexadatainfrastructure.end
Change Compartment	com.oraclecloud.databaseservice.changeexadatainfrastructurecompartment
Configuration File Download	com.oraclecloud.databaseservice.downloadexadatainfrastructureconfigfile
Create Begin	com.oraclecloud.databaseservice.createexadatainfrastructure.begin
Create End	com.oraclecloud.databaseservice.createexadatainfrastructure.end
Delete Begin	com.oraclecloud.databaseservice.deleteexadatainfrastructure.begin
Delete End	com.oraclecloud.databaseservice.deleteexadatainfrastructure.end

## CHAPTER 14 Events

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.updateexadatainfrastructure.begin
Update End	com.oraclecloud.databaseservice.updateexadatainfrastructure.end

### Exadata Infrastructure Example

This is a reference event for Exadata Infrastructure instances:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createexadatainfrastructure.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.ocl..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.ocl..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_exadata_infra",
 "resourceId": "ExadataInfra-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "id": "ocidl.id.ocl...<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-29T12:00:00.000Z",
 "timeUpdated": "2019-08-29T12:30:00.000Z",
 "lifecycleDetails": "detail message",
 "shape": "ExadataCC.Base3.48",
 "timeZone": "US/Pacific",
 "displayName": "testDisplayName"
 }
 }
}
```

## CHAPTER 14 Events

### VM Cluster Network Event Types

These are the event types that VM cluster networks emit:

Friendly Name	Event Type
Create Begin	<code>com.oraclecloud.databaseservice.createvmclusternetwork.begin</code>
Create End	<code>com.oraclecloud.databaseservice.createvmclusternetwork.end</code>
Network Validation File Download	<code>com.oraclecloud.databaseservice.downloadvmclusternetworkconfigfile</code>
Terminate Begin	<code>com.oraclecloud.databaseservice.deletevmclusternetwork.begin</code>
Terminate End	<code>com.oraclecloud.databaseservice.deletevmclusternetwork.end</code>
Update Begin	<code>com.oraclecloud.databaseservice.createvmclusternetwork.begin</code>
Update End	<code>com.oraclecloud.databaseservice.createvmclusternetwork.end</code>
Validate Begin	<code>com.oraclecloud.databaseservice.validatevmclusternetwork.begin</code>
Validate End	<code>com.oraclecloud.databaseservice.validatevmclusternetwork.end</code>

### VM Cluster Network Example

This is a reference event for VM cluster networks:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createvmclusternetwork.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
```

## CHAPTER 14 Events

```
"compartmentId": "ocidl.compartment.oc1..<unique_ID>",
"compartmentName": "example_name",
"resourceName": "my_vmcluster_network",
"resourceId": "VmClusterNetwork-unique_ID",
"availabilityDomain": "all",
"freeFormTags": {},
"definedTags": {},
"additionalDetails": {
 "id": "ocidl.id.oc1..<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-29T12:00:00.000Z",
 "timeUpdated": "2019-08-29T12:30:00.000Z",
 "lifecycleDetails": "detail message",
 "exadataInfrastructureId": "ExadataInfra-unique_ID",
 "displayName": "testDisplayName"
}
}
```

### VM Cluster Event Types

These are the event types that VM clusters emit:

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changevmclustercompartment
Create Begin	com.oraclecloud.databaseservice.createvmcluster.begin
Create End	com.oraclecloud.databaseservice.createvmcluster.end
Terminate Begin	com.oraclecloud.databaseservice.deletevmcluster.begin
Terminate End	com.oraclecloud.databaseservice.deletevmcluster.end
Update Begin	com.oraclecloud.databaseservice.updatevmcluster.begin
Update End	com.oraclecloud.databaseservice.updatevmcluster.end

### VM Cluster Example

This is a reference event for VM clusters:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createvmclusternetwork.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_vmcluster_network",
 "resourceId": "VmClusterNetwork-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "id": "ocid1.id..oc1...<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-29T12:00:00.000Z",
 "timeUpdated": "2019-08-29T12:30:00.000Z",
 "lifecycleDetails": "detail message",
 "exadataInfrastructureId": "ExadataInfra-unique_ID",
 "displayName": "testDisplayName"
 }
 }
}
```

### Backup Destination Event Types

These are the event types that backup destinations emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Change Compartment	com.oraclecloud.databaseservice.changebackupdestinationcompartment
Create	com.oraclecloud.databaseservice.createbackupdestination
Terminate	com.oraclecloud.databaseservice.deletebackupdestination
Update	com.oraclecloud.databaseservice.updatebackupdestination

### Backup Destination Example

This is a reference event for backup destinations:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createbackupdestination",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_backupdestination",
 "resourceId": "BackupDestination-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {}
 }
}
```

### Database Node Event Types (Cloud at Customer)

These are the event types that database nodes emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.dbnodeaction.begin
Update End	com.oraclecloud.databaseservice.dbnodeaction.end

### Database Node Example

This is a reference event for database nodes:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.dbnodeaction.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-06-27T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_dbnode",
 "resourceId": "DbNode-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "id": "ocidl.id.oc1...<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-26T12:00:00.000Z",
 "timeUpdated": "2019-08-26T12:30:00.000Z",
 "dbSystemId": "ocidl.dbsystem.oc1.phx.<unique_ID>",
 "lifecycleDetails": "detail message",
 "vmClusterId": "VmCluster-unique_ID",
 "dbHostId": "dbHost-unique_ID",
 "nodeNumber": 2,
 "powerAction": "HardReset",
 "hostName": "testHostName"
 }
 }
}
```

## CHAPTER 14 Events

```
}
}
```

### Database Home Event Types (Cloud at Customer)

These are the event types that Database Homes emit:

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createdbhome.begin
Create End	com.oraclecloud.databaseservice.createdbhome.end
Terminate Begin	com.oraclecloud.databaseservice.deletedbhome.begin
Terminate End	com.oraclecloud.databaseservice.deletedbhome.end

### Database Home Example

This is a reference event for Database Homes:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createdbhome.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_dbhome",
 "resourceId": "DbHome-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
```

## CHAPTER 14 Events

```
"id": "ocid1.id.oc1...<unique_ID>",
"lifecycleState": "AVAILABLE",
"timeCreated": "2019-08-29T12:00:00.000Z",
"timeUpdated": "2019-08-29T12:30:00.000Z",
"lifecycleDetails": "detail message",
"dbSystemId": "DbSystem-unique_ID",
"dbVersion": "19.0.0.0",
"recordVersion": 4,
"displayName": "testDisplayName"
}
}
}
```

### Database Event Types (Cloud at Customer)

These are the event types that databases emit:

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createdatabase.begin
Create End	com.oraclecloud.databaseservice.createdatabase.end
Restore Begin	com.oraclecloud.databaseservice.restoredatabase.begin
Restore End	com.oraclecloud.databaseservice.restoredatabase.end
Terminate Begin	com.oraclecloud.databaseservice.deletedatabase.begin
Terminate End	com.oraclecloud.databaseservice.deletedatabase.end
Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Update End	com.oraclecloud.databaseservice.updatedatabase.end

### Database Example

This is a reference event for databases:

```
{
 "cloudEventsVersion": "0.1",
```

## CHAPTER 14 Events

```
"eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
"eventType": "com.oraclecloud.databaseservice.restoredatabase.begin",
"source": "databaseservice",
"eventTypeVersion": "<version>",
"eventTime": "2019-06-27T21:16:04Z",
"contentType": "application/json",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
},
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_database",
 "resourceId": "Database-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "id": "ocidl.id..oc1...<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-26T12:00:00.000Z",
 "timeUpdated": "2019-08-26T12:30:00.000Z",
 "dbSystemId": "dbSystem-unique_ID",
 "displayName": "testDisplayName",
 "lifecycleDetails": "detail message",
 "vmClusterId": "VmCluster-<unique_ID>",
 "backupType": "FULL",
 "dbHomeId": "dbHome-<unique_ID>",
 "dbVersion": "19.0.0.0",
 "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
 "autoBackupsEnabled": "true",
 "recoveryWindow": 30,
 "backupDestinationId": "backupDestination-<unique_ID>",
 "backupDestinationType": "OBJECT_STORAGE",
 "backupDestinationName": "my_backup_destination_name",
 "exadataInfrastructureId": "ExadataInfrastructure-<unique_ID>",
 "dbUniqueName": "akv_tgh_unqna"
 }
}
}
```

## CHAPTER 14 Events

### DB System Event Types

These are the event types that DB systems emit:

Friendly Name	Event Type
Change Compartment Begin	com.oraclecloud.databaseservice.changedbssystemcompartment.begin
Change Compartment End	com.oraclecloud.databaseservice.changedbssystemcompartment.end
Create Begin	com.oraclecloud.databaseservice.launchdbssystem.begin
Create End	com.oraclecloud.databaseservice.launchdbssystem.end
Terminate Begin	com.oraclecloud.databaseservice.terminatedbssystem.begin
Terminate End	com.oraclecloud.databaseservice.terminatedbssystem.end
Update IORM Begin	com.oraclecloud.databaseservice.updateiormconfig.begin
Update IORM End	com.oraclecloud.databaseservice.updateiormconfig.end

### DB System Example

This is a reference event for DB Systems:

```
{
 "cloudEventsVersion": "0.1",
 "contentType": "application/json",
 "data": {
 "additionalDetails": {
 "cpuCoreCount": 1,
 "dataStoragePercentage": 80,
 "dataStorageSizeInGBs": 256,
 "exadataIormConfig": "null",
 "licenseType": "LICENSE_INCLUDED",
 "lifecycleMessage": null,
 "lifecycleState": "PROVISIONING",
 "nsgIds": "null",
 "patchHistoryEntries": "null",
 "sshPublicKeys": "...",
 "version": null
 }
 }
}
```

## CHAPTER 14 Events

```
 },
 "availabilityDomain": "XXIT:US-ASHBURN-AD-1",
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "18_1",
 "resourceId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
 "resourceName": "myDBsystem"
 },
 "eventID": "0c1f15b1-4bf2-4f27-8a78-a48d446aeb6f",
 "eventTime": "2019-10-25T20:30:46.836Z",
 "eventType": "com.oraclecloud.databaseservice.launchdbsystem.begin",
 "eventTypeVersion": "2.0",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "source": "DatabaseService"
}
```

### Database Node Event Types (DB Systems)

These are the event types that database nodes emit:

Friendly Name	Event Type
Update Begin	com.oraclecloud.databaseservice.dbnodeaction.end
Update End	com.oraclecloud.databaseservice.dbnodeaction.end

### Database Node Example

This is a reference event for database nodes:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.databaseservice.db.node.reboot.begin",
 "source": "databaseservice",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-07-29T04:43:24Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 }
}
```

## CHAPTER 14 Events

```
},
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "Demo",
 "resourceName": "",
 "resourceId": "ocidl.dbnode.oc1.phx.<unique_ID>",
 "availabilityDomain": "TGjA:PHX-AD-2",
 "freeFormTags": null,
 "definedTags": null,
 "additionalDetails": {
 "cpuCoreCount": null,
 "lifecycleState": "STARTING",
 "dataStorageSizeInTBs": null,
 "timeCreated": "2019-06-13T04:31:05.190Z",
 "timeUpdated": "2019-07-29T04:43:06.455Z",
 "hostName": "ora18c",
 "lifecycleDetails": null,
 "dbSystemId": "ocidl.dbsystem.oc1.phx.<unique_ID>",
 "dbHostId": "DbHost-<unique_ID>",
 "nodeNumber": null
 }
}
}
```

### Database Home Types (DB Systems)

These are the event types that Database Homes emit:

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createdbhome.begin
Create End	com.oraclecloud.databaseservice.createdbhome.end
Terminate Begin	com.oraclecloud.databaseservice.deletedbhome.begin
Terminate End	com.oraclecloud.databaseservice.deletedbhome.end

### Database Home Example

This is a reference event for Database Homes:

## CHAPTER 14 Events

---

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
 "eventType": "com.oraclecloud.databaseservice.createdbhome.begin",
 "source": "databaseservice",
 "eventTypeVersion": "<version>",
 "eventTime": "2019-08-29T21:16:04Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_dbhome",
 "resourceId": "DbHome-unique_ID",
 "availabilityDomain": "all",
 "freeFormTags": {},
 "definedTags": {},
 "additionalDetails": {
 "id": "ocid1.id.oc1...<unique_ID>",
 "lifecycleState": "AVAILABLE",
 "timeCreated": "2019-08-29T12:00:00.000Z",
 "timeUpdated": "2019-08-29T12:30:00.000Z",
 "lifecycleDetails": "detail message",
 "dbSystemId": "DbSystem-unique_ID",
 "dbVersion": "19.0.0.0",
 "recordVersion": 4,
 "displayName": "testDisplayName"
 }
 }
}
```

### Database Event Types (DB Systems)

These are the event types that databases emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Automatic Backup Begin	com.oraclecloud.databaseservice.automaticbackupdatabase.begin
Automatic Backup End	com.oraclecloud.databaseservice.automaticbackupdatabase.end
Create Backup Begin	com.oraclecloud.databaseservice.backupdatabase.begin
Create Backup End	com.oraclecloud.databaseservice.backupdatabase.end
Restore Begin	com.oraclecloud.databaseservice.restoredatabase.begin
Restore End	com.oraclecloud.databaseservice.restoredatabase.end
Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Update End	com.oraclecloud.databaseservice.updatedatabase.end

### Database Example

This is a reference event for databases:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.databaseservice.database.backup.begin",
 "source": "databaseservice",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-07-29T03:43:44Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "sic-dbaas",
 "resourceName": "autoBack",
 "resourceId": "ocidl.database.oc1.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {},
 "definedTags": {}
 }
}
```

## CHAPTER 14 Events

```
"additionalDetails": {
 "cpuCoreCount": null,
 "lifecycleState": "AVAILABLE",
 "dataStorageSizeInTBs": null,
 "timeCreated": "2019-07-29T00:36:22.701Z",
 "timeUpdated": "2019-07-29T03:43:44.171Z",
 "lifecycleDetails": null,
 "vmClusterId": null,
 "dbHomeId": "ocid1.dbhome.oc1.phx.<unique_ID>",
 "dbUniqueName": "autoBack_phx1w7",
 "dbVersion": "18.6.0.0.190416",
 "databaseEdition": "ENTERPRISE_EDITION_EXTREME",
 "workloadType": null,
 "autoBackupsEnabled": false,
 "recoveryWindow": "30",
 "backupDestinationId": null,
 "backupDestinationType": null,
 "backupDestinationName": null,
 "databaseId": null,
 "exadataInfrastructureId": null
}
}
```

### Data Guard Association Event Types

These are the event types that Data Guard associations emit:

Friendly Name	Event Type
Create Begin	com.oraclecloud.databaseservice.createdataguardassociation.begin
Create End	com.oraclecloud.databaseservice.createdataguardassociation.end
Failover Begin	com.oraclecloud.databaseservice.failoverdataguardassociation.begin
Failover End	com.oraclecloud.databaseservice.failoverdataguardassociation.end
Reinstate Begin	com.oraclecloud.databaseservice.reinstatedataguardassociation.begin

## CHAPTER 14 Events

Friendly Name	Event Type
Reinstate End	com.oraclecloud.databaseservice.reinstatedataguardassociation.end
Switchover Begin	com.oraclecloud.databaseservice.switchoverdataguardassociation.begin
Switchover End	com.oraclecloud.databaseservice.switchoverdataguardassociation.end

### Data Guard Association Example

This is a reference event for Data Guard associations:

```
{
 "cloudEventsVersion": "0.1",
 "contentType": "application/json",
 "data": {
 "additionalDetails": {
 "ApplyLag": null,
 "DGConfigId": "7e8eff2b-a4cd-474a-abd5-940b05c0b1fd",
 "DGConfigState": "null",
 "DatabaseId": "ocid1.database.oc1.iad.<unique_ID>",
 "DbHomeId": "ocid1.dbhome.oc1.iad.<unique_ID>",
 "DbSystemId": "ocid1.dbsystem.oc1.iad.<unique_ID>",
 "LastSyncedTime": null,
 "SyncState": "null",
 "dcsDgUpdateTimestamp": null,
 "lastUpdatedIdentifier": null,
 "lifeCycleMessage": null,
 "lifecycleState": "PROVISIONING",
 "timeCreated": "2019-10-25T21:42:19.041Z",
 "timeUpdated": "2019-10-25T21:42:19.041Z"
 },
 "availabilityDomain": "XXIT:US-ASHBURN-AD-1",
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "18_1",
 "resourceId": "ocid1.dgassociation.oc1.iad.<unique_ID>"
 },
 "eventID": "5b8b7fbf-2e9a-4730-9761-e52715b7bc79",
 "eventTime": "2019-10-25T21:42:16.579Z",
 "eventType": "com.oraclecloud.databaseservice.createdataguardassociation.begin",
 "eventTypeVersion": "2.0",
 "extensions": {
```

## CHAPTER 14 Events

```
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "source": "DatabaseService"
}
```

### File Storage

File Storage resources that emit events:

- [File Systems](#) and [Snapshots](#)
- [Mount Targets](#)
- [Exports](#) and [Export Sets](#)

### File System Event Types

These are the event types that file systems emit:

Friendly Name	Event Type
Change File System Compartment	com.oraclecloud.filestorage.changefilesystemcompartment
Create File System	com.oraclecloud.filestorage.createfilesystem
Delete File System	com.oraclecloud.filestorage.deletefilesystem
Update File System	com.oraclecloud.filestorage.updatefilesystem

### File System Example

This is a reference event for file systems:

```
{
 "eventType": "com.oraclecloud.filestorage.createfilesystem",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "filestorage",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
```

## CHAPTER 14 Events

```
"compartmentId": "ocid1.compartment.oc1..<unique_id>",
"compartmentName": "example_name",
"resourceName": "my_filesystem",
"resourceId": "ocid1.filesystem.oc1..<unique_id>",
"availabilityDomain": "availability_domain",
"freeFormTags": {
 "Department": "Finance"
},
"definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
}
},
"eventID": "unique_ID",
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_id>"
}
}
```

### Snapshot Event Types

These are the event types that snapshots emit:

Friendly Name	Event Type
Create Snapshot	com.oraclecloud.filestorage.createsnapshot
Delete Snapshot	com.oraclecloud.filestorage.deletesnapshot

### Snapshot Example

This is a reference event for snapshots:

```
{
 "eventType": "com.oraclecloud.filestorage.createsnapshot",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "filestorage",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
}
```

## CHAPTER 14 Events

```
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>",
 "compartmentName": "example_name",
 "resourceName": "my_snapshot",
 "resourceId": "ocidl.snapshot.oc1..<unique_id>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
},
"eventID": "unique_ID",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>"
}
}
```

### Mount Target Event Types

These are the event types that mount targets emit:

Friendly Name	Event Type
Change Mount Target Compartment	com.oraclecloud.filestorage.changemounttargetcompartment
Create Mount Target	com.oraclecloud.filestorage.createmounttarget
Delete Mount Target	com.oraclecloud.filestorage.deletemounttarget
Update Mount Target	com.oraclecloud.filestorage.updatemounttarget

### Mount Target Example

This is a reference event for mount targets:

```
{
 "eventType": "com.oraclecloud.filestorage.createmounttarget",
}
```

## CHAPTER 14 Events

```
"cloudEventsVersion": "0.1",
"eventTypeVersion": "2.0",
"source": "filestorage",
"eventTime": "2019-08-12T17:51:42.789Z",
"contentType": "application/json",
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>",
 "compartmentName": "example_name",
 "resourceName": "my_mounttarget",
 "resourceId": "ocidl.mounttarget.oc1..<unique_id>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
},
"eventID": "unique_ID",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>"
}
}
```

### Export Event Types

These are the event types that exports emit:

Friendly Name	Event Type
Create Export	com.oraclecloud.filestorage.createexport
Delete Export	com.oraclecloud.filestorage.deleteexport
Update Export	com.oraclecloud.filestorage.updateexport

### Export Example

This is a reference event for exports:

## CHAPTER 14 Events

```
{
 "eventType": "com.oraclecloud.filestorage.createexport",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "filestorage",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>",
 "compartmentName": "example_name",
 "resourceName": "my_export",
 "resourceId": "ocidl.export.oc1..<unique_id>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 },
 "eventID": "unique_ID",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>"
 }
}
```

### Export Set Event Types

These are the event types that export sets emit:

Friendly Name	Event Type
Delete Export Set	com.oraclecloud.filestorage.deleteexportset
Update Export Set	com.oraclecloud.filestorage.updateexportset

### Export Set Example

This is a reference event for export sets :

## CHAPTER 14 Events

---

```
{
 "eventType": "com.oraclecloud.filestorage.updateexportset",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "filestorage",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>",
 "compartmentName": "example_name",
 "resourceName": "my_exportset",
 "resourceId": "ocidl.exportset.oc1..<unique_id>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 },
 "eventID": "unique_ID",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_id>"
 }
}
```

## Functions

Functions resources that emit events:

- [Application Event Types](#)
- [Function Event Types](#)

### Application Event Types

These are the event types that applications emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Change Application Compartment	com.oraclecloud.functions.changeapplicationcompartment
Create Application	com.oraclecloud.functions.createapplication
Delete Application	com.oraclecloud.functions.deleteapplication
Update Application	com.oraclecloud.functions.updateapplication

### Application Example

This is an example event for applications:

```
{
 "eventType": "com.oraclecloud.functions.createapplication",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "functions",
 "eventTime": "2019-07-22T09:33:44.754Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "my_compartment",
 "resourceName": "my-application",
 "resourceId": "ocidl.fnapp.oc1.phx.<unique_ID>",
 "availabilityDomain": "AD3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Function Event Types

These are the event types that functions emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Create Function	com.oraclecloud.functions.createfunction
Delete Function	com.oraclecloud.functions.deletefunction
Update Function	com.oraclecloud.functions.updatefunction

### Function Example

This is an example event for functions:

```
{
 "eventType": "com.oraclecloud.functions.createfunction",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "functions",
 "eventTime": "2019-07-22T09:33:44.754Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "my_compartment",
 "resourceName": "my-function",
 "resourceId": "ocidl.fnfunc.oc1.phx.<unique_ID>",
 "availabilityDomain": "AD3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

## IAM

IAM resources that emit events:

- [Authentication Policy Event Types](#)
- [Credentials Event Types](#)
- [Dynamic Group Event Types](#)

## CHAPTER 14 Events

---

- [Group Event Types](#)
- [Identity Provider Event Types](#)
- [Multi-Factor Authentication TOTP Device Event Types](#)
- [Policy Event Types](#)
- [User Event Types](#)

### Authentication Policy Event Types

This is the event type that authentication policies emit:

Friendly Name	Event Type
Update Authentication Policy	com.oraclecloud.identityControlPlane.UpdateAuthenticationPolicy

### Authentication Policy Example

This is a reference event for authentication policy events:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.UpdateAuthenticationPolicy",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_compartment",
 "resourceId": "ocidl.compartment.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 }
}
```

## CHAPTER 14 Events

```
}
},
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
}
```

### Credentials Event Types

These are the event types that credentials emit.

Friendly Name	Event Type
Create Auth Token	com.oraclecloud.identityControlPlane.CreateAuthToken
Create Customer Secret Key	com.oraclecloud.identityControlPlane.CreateCustomerSecretKey
Create or Reset Password	com.oraclecloud.identityControlPlane.CreateOrResetPassword
Create SMTP Credential	com.oraclecloud.identityControlPlane.CreateSmtplibCredential
Create Swift Password	com.oraclecloud.identityControlPlane.CreateSwiftPassword
Delete API Key	com.oraclecloud.identityControlPlane.DeleteApiKey
Delete Auth Token	com.oraclecloud.identityControlPlane.DeleteAuthToken
Delete Customer Secret Key	com.oraclecloud.identityControlPlane.DeleteCustomerSecretKey
Delete SMTP Credential	com.oraclecloud.identityControlPlane.DeleteSmtplibCredential
Delete Swift Password	com.oraclecloud.identityControlPlane.DeleteSwiftPassword
Update Auth Token	com.oraclecloud.identityControlPlane.UpdateAuthToken
Update Authentication Policy	com.oraclecloud.identityControlPlane.UpdateAuthenticationPolicy
Update Customer Secret Key	com.oraclecloud.identityControlPlane.UpdateCustomerSecretKey

## CHAPTER 14 Events

Friendly Name	Event Type
Update SMTP Credential	com.oraclecloud.identityControlPlane.UpdateSmtCredential
UpdateSwift Password	com.oraclecloud.identityControlPlane.UpdateSwiftPassword
Upload API KEY	com.oraclecloud.identityControlPlane.UploadApiKey

### Credentials Example

This is a reference event for most credential events (create or reset password don't include additional details):

```
{
 "eventType": "com.oraclecloud.identityControlPlane.DeleteApiKey",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_user",
 "resourceId": "<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "userId": "ocidl.user.oc1..<unique_ID>"
 }
 },
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

## CHAPTER 14 Events

---

```
}
}
```

### Dynamic Group Event Types

These are the event types that dynamic groups emit.

Friendly Name	Event Type
Create Dynamic Group	<code>com.oraclecloud.identityControlPlane.CreateDynamicGroup</code>
Delete Dynamic Group	<code>com.oraclecloud.identityControlPlane.DeleteDynamicGroup</code>
Update Dynamic Group	<code>com.oraclecloud.identityControlPlane.UpdateDynamicGroup</code>

### Dynamic Group Example

This is a reference event for dynamic groups:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreateDynamicGroup",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_dynamicgroup",
 "resourceId": "ocid1.dynamicgroup.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 }
}
```

## CHAPTER 14 Events

```
},
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
}
```

### Group Event Types

These are the event types that groups emit.

Friendly Name	Event Type
Add User to Group	com.oraclecloud.identityControlPlane.AddUserToGroup
Create Group	com.oraclecloud.identityControlPlane.CreateGroup
Delete Group	com.oraclecloud.identityControlPlane.DeleteGroup
Remove User From Group	com.oraclecloud.identityControlPlane.RemoveUserFromGroup
Update Group	com.oraclecloud.identityControlPlane.UpdateGroup

### Group Example

This is a reference event for some groups (create, delete, and update events don't include additional details):

```
{
 "eventType": "com.oraclecloud.identityControlPlane.AddUserToGroup",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_group",
 "resourceId": "ocid1.groupmembership.oc1.<unique_ID>",
 "availabilityDomain": "availability_domain",
```

## CHAPTER 14 Events

```
"freeFormTags": {
 "Department": "Finance"
},
"definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
},
"additionalDetails": {
 "userId": "ocidl.user.oc1..<unique_ID>",
 "groupId": "ocidl.group.oc1..<unique_ID>"
}
},
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

### Identity Provider Event Types

These are the event types that identity providers emit.

Friendly Name	Event Type
Add User to IdP Group	com.oraclecloud.identityControlPlane.AddUserToIdpGroup
Create Identity Provider	com.oraclecloud.identityControlPlane.CreateIdentityProvider
Create Identity Provider Group	com.oraclecloud.identityControlPlane.CreateIdentityProviderGroup
Create IdP Group Mapping	com.oraclecloud.identityControlPlane.CreateIdpGroupMapping
Create IdP User	com.oraclecloud.identityControlPlane.CreateIdpUser
Delete Identity Provider	com.oraclecloud.identityControlPlane.DeleteIdentityProvider
Delete Identity Provider Group	com.oraclecloud.identityControlPlane.DeleteIdentityProviderGroup

## CHAPTER 14 Events

Friendly Name	Event Type
Delete IdP Group Mapping	com.oraclecloud.identityControlPlane.DeleteIdpGroupMapping
Delete IdP User	com.oraclecloud.identityControlPlane.DeleteIdpUser
Remove User From IdP Group	com.oraclecloud.identityControlPlane.RemoveUserFromIdpGroup
Reset IdP SCIM Client	com.oraclecloud.identityControlPlane.ResetIdpScimClient
Update Identity Provider	com.oraclecloud.identityControlPlane.UpdateIdentityProvider
Update IdP Group Mapping	com.oraclecloud.identityControlPlane.UpdateIdpGroupMapping

### Identity Provider Example

The following reference events are for identity provider events that include additional details. Some identity providers events do not include additional details. These events are create, delete, and update identity providers, as well as delete identity provider group, delete IdP user, and reset IdP SCIM.

This is a reference event for adding and removing users from IdP groups:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.AddUserToIdpGroup",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_group",
 "resourceId": "ocid1.idpgroup.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
```

## CHAPTER 14 Events

---

```
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "userId": "ocidl.user.oc1..<unique_ID>"
 }
},
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

This is a reference event for create, update, and delete IdP group mapping:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreateIdpGroupMapping",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_identityprovider",
 "resourceId": "ocidl.idpgroupmapping.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "idpGroupName": "my_group",
 "groupId": "ocidl.group.oc1..<unique_ID>"
 }
 },
 "extensions": {
```

## CHAPTER 14 Events

---

```
"compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

This is a reference event for create IdP user and create IdP group:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreateIdentityProviderGroup",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_idpgroup",
 "resourceId": "ocidl.idpgroup.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "externalIdentifier": "my_externalidentifier"
 }
 },
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Multi-Factor Authentication TOTP Device Event Types

These are the event types that MFA TOTP devices emit.

## CHAPTER 14 Events

Friendly Name	Event Type
Activate MFA TOTP Device	com.oraclecloud.identityControlPlane.ActivateMfaTotpDevice
Create MFA TOTP Device	com.oraclecloud.identityControlPlane.CreateMfaTotpDevice
Delete MFA TOTP Device	com.oraclecloud.identityControlPlane.DeleteMfaTotpDevice
Generate MFA TOTP Device Seed	com.oraclecloud.identityControlPlane.GenerateTotpSeed

### Multi-Factor Authentication TOTP Devices Example

This is a reference event for MFA TOTP Devices:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreateMfaTotpDevice",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_user",
 "resourceId": "ocidl.credential.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 },
 "additionalDetails": {
 "userId": "ocidl.user.oc1..<unique_ID>"
 }
 },
 "extensions": {
```

## CHAPTER 14 Events

```
"compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
}
```

### Policy Event Types

These are the event types that policies emit.

Friendly Name	Event Type
Create Policy	com.oraclecloud.identityControlPlane.CreatePolicy
Delete Policy	com.oraclecloud.identityControlPlane.DeletePolicy
Update Policy	com.oraclecloud.identityControlPlane.UpdatePolicy

### Policy Example

This is a reference event for policies:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreatePolicy",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_policy",
 "resourceId": "ocid1.policy.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 }
}
```

## CHAPTER 14 Events

```
}
},
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

### User Event Types

These are the event types that users emit.

Friendly Name	Event Type
Create User	com.oraclecloud.identityControlPlane.CreateUser
Delete User	com.oraclecloud.identityControlPlane.DeleteUser
Update User	com.oraclecloud.identityControlPlane.UpdateUser
Update User Capabilities	com.oraclecloud.identityControlPlane.UpdateUserCapabilities
Update User State	com.oraclecloud.identityControlPlane.UpdateUserState

### User Example

This is a reference event for users:

```
{
 "eventType": "com.oraclecloud.identityControlPlane.CreateUser",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "identityControlPlane",
 "eventID": "<unique_ID>",
 "eventTime": "2019-10-21T17:23:54.095Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_user",
 "resourceId": "ocidl.user.oc1..<unique_ID>",
 "availabilityDomain": "availability_domain",
```

## CHAPTER 14 Events

---

```
"freeFormTags": {
 "Department": "Finance"
},
"definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
}
},
"extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
}
}
```

### Integration

For details about events emitted by Oracle Integration, see [Automating with Events](#).

### Networking

Networking resources that emit events:

- [NAT gateways](#)
- [Route tables](#)
- [Security lists](#)
- [Service gateways](#)
- [Virtual cloud networks \(VCNs\)](#)

## NAT Gateway Event Types

These are the event types that NAT gateways emit:

Friendly Name	Event Type
Create NAT Gateway	com.oraclecloud.natgateway.createnatgateway
Delete NAT Gateway	com.oraclecloud.natgateway.deletenatgateway
Update NAT Gateway	com.oraclecloud.natgateway.updatenatgateway
Change NAT Gateway Compartment	com.oraclecloud.natgateway.changenatgatewaycompartment

## NAT Gateway Example

This is a reference event for NAT gateways:

```
{
 "eventType": "com.oraclecloud.natgateway.createnatgateway",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "natgateway",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "example_name",
 "resourceId": "ocidl.natgateway.oci.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 },
 "eventID": "<unique_ID>",
 "extensions": {
```

## CHAPTER 14 Events

```
"compartmentId": "ocid1.compartment.oci..<unique_ID>"
}
}
```

### Route Table Event Types

These are the event types that route tables emit:

Friendly Name	Event Type
Create Route Table	com.oraclecloud.virtualnetwork.createroutetable
Delete Route Table	com.oraclecloud.virtualnetwork.deleteroutetable
Update Route Table	com.oraclecloud.virtualnetwork.updateroutetable
Change Route Table Compartment	com.oraclecloud.virtualnetwork.changeroutetablecompartment

### Route Table Example

This is a reference event for route tables:

```
{
 "eventType": "com.oraclecloud.virtualnetwork.createroutetable",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "virtualNetwork",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocid1.compartment.oci..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "example_name",
 "resourceId": "ocid1.routetable.oci.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 }
}
```

## CHAPTER 14 Events

```
 }
 }
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>"
}
}
```

### Security List Event Types

These are the event types that security lists emit:

Friendly Name	Event Type
Create Security List	com.oraclecloud.virtualnetwork.createsecuritylist
Delete Security List	com.oraclecloud.virtualnetwork.deletesecuritylist
Update Security List	com.oraclecloud.virtualnetwork.updatesecuritylist
Change Security List Compartment	com.oraclecloud.virtualnetwork.changesecuritylistcompartment

### Security List Example

This is a reference event for security lists:

```
{
 "eventType": "com.oraclecloud.virtualnetwork.createsecuritylist",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "virtualNetwork",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "example_name",
 "resourceId": "ocidl.securitylist.oci.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {
```

## CHAPTER 14 Events

```
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocid1.compartment.oci..<unique_ID>"
}
}
```

### Service Gateway Event Types

These are the event types that service gateways emit:

Friendly Name	Event Type
Create Service Gateway	com.oraclecloud.servicegateway.createservicegateway
Delete Service Gateway Start	com.oraclecloud.servicegateway.deleteservicegateway.begin
Delete Service Gateway End	com.oraclecloud.servicegateway.deleteservicegateway.end
Update Service Gateway	com.oraclecloud.servicegateway.updateservicegateway
Change Service Gateway Compartment	com.oraclecloud.servicegateway.changeservicegatewaycompartment
Attach Service	com.oraclecloud.servicegateway.attachserviceid
Detach Service	com.oraclecloud.servicegateway.detachserviceid

### Service Gateway Example

This is a reference event for service gateways:

```
{
 "eventType": "com.oraclecloud.servicegateway.createservicegateway",
```

## CHAPTER 14 Events

```
"cloudEventsVersion": "0.1",
"eventTypeVersion": "2.0",
"source": "servicegateway",
"eventTime": "2019-08-12T17:51:42.789Z",
"contentType": "application/json",
"data": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "example_name",
 "resourceId": "ocidl.servicegateway.oci.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>"
}
}
```

### VCN Event Types

These are the event types that VCNs emit:

Friendly Name	Event Type
Create VCN	com.oraclecloud.virtualnetwork.createvcn
Delete VCN	com.oraclecloud.virtualnetwork.deletevcn
Update VCN	com.oraclecloud.virtualnetwork.updatevcn

### VCN Example

This is a reference event for VCNs:

```
{
 "eventType": "com.oraclecloud.virtualnetwork.createvcn",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "virtualNetwork",
 "eventTime": "2019-08-12T17:51:42.789Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "example_name",
 "resourceId": "ocidl.vcn.oci.phx.<unique_ID>",
 "availabilityDomain": "XXIT:PHX-AD-1",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oci..<unique_ID>"
 }
}
```

## Notifications

Notifications resources that emit events:

- [Subscriptions](#)
- [Topics](#)

### Subscriptions Event Types

These are the event types that subscriptions emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Create Subscription	com.oraclecloud.notification.createsubscription
Delete Subscription	com.oraclecloud.notification.deletesubscription
Move Subscription	com.oraclecloud.notification.movesubscription
Resend Subscription Confirmation	com.oraclecloud.notification.resendsubscriptionconfirmation
Update Subscription	com.oraclecloud.notification.updatesubscription

### Subscription Example

This is a reference event for subscriptions:

```
{
 "eventType": "com.oraclecloud.notification.createsubscription",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "notification",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "my_compartment",
 "resourceName": "ons-subscription",
 "resourceId": "ocidl.onssubscription.oc1..<unique_ID>",
 "availabilityDomain": "AD3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

### Topics Event Types

These are the event types that topics emit:

## CHAPTER 14 Events

Friendly Name	Event Type
Create Topic	com.oraclecloud.notification.createtopic
Delete Topic	com.oraclecloud.notification.deletetopic
Move Topic	com.oraclecloud.notification.movetopic
Update Topic	com.oraclecloud.notification.updatetopic

### Topic Example

This is a reference event for topics:

```
{
 "eventType": "com.oraclecloud.notification.createtopic",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "notification",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "my_compartment",
 "resourceName": "my_topic",
 "resourceId": "ocidl.onstopic.oc1..<unique_ID>",
 "availabilityDomain": "AD3"
 },
 "eventID": "<unique_ID>",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 }
}
```

## Object Storage

Object Storage resources that emit events:

- [Buckets](#)
- [Objects](#)

## Buckets Event Types

These are the event types that buckets emit:

Friendly Name	Event Type
Create Bucket	com.oraclecloud.objectstorage.createbucket
Delete Bucket	com.oraclecloud.objectstorage.deletebucket
Update Bucket	com.oraclecloud.objectstorage.updatebucket

## Bucket Example

This is an example event for buckets:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.objectstorage.createbucket",
 "source": "objectstorage",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-01-10T21:19:24Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "compartmentName": "example_name",
 "resourceName": "my_bucket",
 "resourceId": "ocid1.compartment.oc1..<unique_ID>",
 "availabilityDomain": "all",
 "freeFormTags": {
 "Department": "Finance"
 },
 "definedTags": {
 "Operations": {
 "CostCenter": "42"
 }
 }
 },
 "additionalDetails": {
```

## CHAPTER 14 Events

```
"namespace": "example_namespace",
"publicAccessType": "NoPublicAccess",
"eTag": "f8ffb6e9-f602-460f-a6c0-00b5abfa24c7"
}
}
}
```

### Objects Event Types

Events for objects are handled differently than other resources. Objects do not emit events by default. Use the [Console](#), [CLI](#), or [API](#) to enable a bucket to emit events for object state changes. You can enable events for object state changes during or after bucket creation.

These are the event types that objects emit:

Friendly Name	Event Type
Create Object	com.oraclecloud.objectstorage.createobject
Delete Object	com.oraclecloud.objectstorage.deleteobject
Update Object	com.oraclecloud.objectstorage.updateobject

### Object Example

This is an example event for objects:

```
{
 "cloudEventsVersion": "0.1",
 "eventID": "<unique_ID>",
 "eventType": "com.oraclecloud.objectstorage.createobject",
 "source": "objectstorage",
 "eventTypeVersion": "2.0",
 "eventTime": "2019-07-10T13:37:11Z",
 "contentType": "application/json",
 "extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
 },
 "data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "Example_Compartment,"
 }
}
```

## CHAPTER 14 Events

```
"resourceName": "v1/log/10.0.6.166",
"resourceId": "",
"availabilityDomain": "all",
"additionalDetails": {
 "eTag": "8162db5b-50d7-4947-a576-4401798ed2fa",
 "namespace": "my_namespace",
 "archivalState": null,
 "bucketName": "my_bucket",
 "bucketId": "ocid1.bucket.oc1.<unique_ID>"
}
}
```

### Resource Manager

Resource Manager resources that emit events:

- [Jobs](#)
- [Stacks](#)

### Job Event Types

These are the event types that jobs emit:

Friendly Name	Event Type
Cancel Job	com.oraclecloud.oracleresourcemanager.canceljob
Create Job Begin	com.oraclecloud.oracleresourcemanager.createjob.begin
Create Job End	com.oraclecloud.oracleresourcemanager.createjob.end
Update Job	com.oraclecloud.oracleresourcemanager.updatejob

This is a reference event for jobs:

```
{
 "eventType": "com.oraclecloud.oracleresourcemanager.updateJob",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
```

## CHAPTER 14 Events

```
"source": "OracleResourceManager",
"eventTime": "2019-07-23T01:46:37.606Z",
"contentType": "application/json",
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "example_name",
 "resourceId": "ocidl.ormjob.oc1.phx.<unique_ID>",
 "availabilityDomain": "availability_domain"
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

### Stack Event Types

These are the event types that stacks emit:

Friendly Name	Event Type
Change Compartment Begin	com.oraclecloud.oracleresourcemanager.changestackcompartment.begin
Change Compartment End	com.oraclecloud.oracleresourcemanager.changestackcompartment.end
Create Stack	com.oraclecloud.oracleresourcemanager.createstack
Delete Stack	com.oraclecloud.oracleresourcemanager.deletestack
Update Stack	com.oraclecloud.oracleresourcemanager.updatestack

This is a reference event for stacks:

```
{
 "eventType": "com.oraclecloud.oracleresourcemanager.createstack",
 "cloudEventsVersion": "0.1",
 "eventTypeVersion": "2.0",
 "source": "OracleResourceManager",
 "eventTime": "2019-07-23T01:32:10.866Z",
}
```

```
"contentType": "application/json",
"data": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "compartmentName": "example_compartment",
 "resourceName": "example_name",
 "resourceId": "ocidl.ormstack.oc1.phx.<unique_ID>",
 "availabilityDomain": "availability_domain"
},
"eventID": "<unique_ID>",
"extensions": {
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>"
}
}
```

## Events Metrics

You can monitor performance of your rules by using [metrics](#), [alarms](#), and [notifications](#). This topic describes the metrics emitted by the metric namespace `oci_cloudevents` (the Events service).

Resources: rules. Also measures data for events, which are not resources.

### Prerequisites

**IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Overview of the Events Service Metrics

You create rules that specify which events should be delivered to other services for processing. This delivery creates the automation in your tenancy. A rule identifies an event

pattern to match and specifies other services to deliver matching events to. Metrics help you measure the success of the rules you create (in terms of pattern matching and delivery) and the quality and scope of the emitted events in your tenancy. For more information, see [Overview of Events](#).

### Available Metrics: oci\_cloudevents

The metrics listed in the following table are automatically available for rules you create. You do not need to enable monitoring to get these metrics.

Each metric includes one or more of the following dimensions:

**RESOURCEID**

The OCID of the rule or compartment to which the metric applies.

**EVENTTYPE**

The type of event emitted by a resource.

**RESOURCEDISPLAYNAME**

The name of the rule.

**ACTIONTYPE**

One or more of the following types of resources that receives an event from the Events service.

- [Notifications](#)
- [Streaming](#)
- [Functions](#)

## CHAPTER 14 Events

---

<b>Metric</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Description</b>	<b>Dimensions</b>
PublishedEvents	Events Emitted	count	Total number of events emitted by resources in a compartment.	eventType resourceId
MatchedEvents	Events Matched	count	If you view the default chart from a rule, this metric provides the total number of events matched for the rule. If you view the chart from the Service Metrics page, this metric gives a total number of matched events for all the rules in a compartment.	resourceDisplayName resourceId

## CHAPTER 14 Events

---

Metric	Metric Display Name	Unit	Description	Dimensions
DeliverySucceedEvents	Events Delivered	count	If you view the default chart from a rule, this metric provides the total number of successful deliveries to actions for the rule. If you view the chart from the Service Metrics page, this metric gives a total number of successful deliveries to actions for all the rules in a compartment.	actionType resourceDisplayName resourceId

## CHAPTER 14 Events

---

Metric	Metric Display Name	Unit	Description	Dimensions
DeliveryFailedEvents	Delivery Failure	count	If you view the default chart from a rule, this metric provides the total number of unsuccessful deliveries to actions for the rule. If you view the chart from the Service Metrics page, this metric gives a total number of unsuccessful deliveries to actions for all the rules in a compartment.	

### Using the Console

#### To view default metric charts for a rule

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.

2. Choose the **Compartment** that contains the rule you want to view, and then click the rule's name.

3. Click **Metrics**.

The **Metrics** page displays a default set of charts for the current rule.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for a compartment

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Application Integration** and click **Events Service**.

2. Choose the **Compartment** that contains the rules you want to monitor.

3. Click **Metrics**.

The **Metrics** page displays a default set of charts for the current compartment.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 15 File Storage

This chapter explains how to create file systems, how to manage them, and how to mount them to write files.

## Overview of File Storage

Oracle Cloud Infrastructure File Storage service provides a durable, scalable, secure, enterprise-grade network file system. You can connect to a File Storage service file system from any bare metal, virtual machine, or container instance in your Virtual Cloud Network (VCN). You can also access a file system from outside the VCN using Oracle Cloud Infrastructure FastConnect and Internet Protocol security (IPSec) virtual private network (VPN).

Large Compute clusters of thousands of instances can use the File Storage service for high-performance shared storage. Storage provisioning is fully managed and automatic as your use scales from a single byte to exabytes without upfront provisioning. You have redundant storage for resilient data protection.

The File Storage service supports the Network File System version 3.0 (NFSv3) protocol. The service supports the Network Lock Manager (NLM) protocol for file locking functionality.

Use the File Storage service when your application or workload includes big data and analytics, media processing, or content management, and you require Portable Operating System Interface (POSIX)-compliant file system access semantics and concurrently accessible storage. The File Storage service is designed to meet the needs of applications and users that need an enterprise file system across a wide range of use cases, including the following:

- **General Purpose File Storage:** Access to an unlimited pool of file systems to manage growth of structured and unstructured data.
- **Big Data and Analytics:** Run analytic workloads and use shared file systems to store persistent data.

- **Lift and Shift of Enterprise Applications:** Migrate existing Oracle applications that need NFS storage, such as Oracle E-Business Suite and PeopleSoft.
- **Databases and Transactional Applications:** Run test and development workloads with Oracle, MySQL, or other databases.
- **Backups, Business Continuity, and Disaster Recovery:** Host a secondary copy of relevant file systems from on premises to the cloud for backup and disaster recovery purposes.
- **MicroServices and Docker:** Deliver stateful persistence for containers. Easily scale as your container-based environments grow.



### Tip

Watch a [video introduction](#) to the service and its capabilities.



### Note

File Storage is not available in Oracle Cloud Infrastructure Government Cloud realms.

## File Systems Concepts

Using the File Storage service requires an understanding of the following concepts, including some that pertain to Oracle Cloud Infrastructure Networking:

### **MOUNT TARGET**

An NFS endpoint that lives in a subnet of your choice and is highly available. The mount target provides the IP address or DNS name that is used in the mount command when connecting NFS clients to a file system. A single mount target can export many file systems. By default, you can create two mount targets per account per availability domain, but you can request an increase. See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. See [Managing Mount Targets](#) for more information about working with this resource.

### **EXPORT**

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). For more information, see [Managing Mount Targets](#).

### **EXPORT SET**

Collection of one or more exports that control what file systems the mount target exports using NFSv3 protocol and how those file systems are found using the NFS mount protocol. Each mount target has an export set. Each file system associated with the mount target has at least one export in the export set.

### **EXPORT PATH**

A path that is specified when an export is created. It uniquely identifies the file system within the mount target, letting you associate up to 100 file systems to a single mount target. This path is unrelated to any path within the file system itself, or the client mount point path.

The File Storage service adds an export that pairs the file system's Oracle Cloud Identifier (OCID) and path.

See [Paths in File Systems](#) for more information.

### EXPORT OPTIONS

NFS export options are a set of parameters within the export that specify the level of access granted to NFS clients when they connect to a mount target. An NFS export options entry within an export defines access for a single IP address or CIDR block range. For more information, see [Working with NFS Export Options](#).

### VIRTUAL CLOUD NETWORK (VCN)

A private network that you set up in the Oracle data centers, with firewall rules and specific types of communication gateways that you can choose to use. A VCN covers a single, contiguous IPv4 CIDR block of your choice. For more information about VCNs, see [VCNs and Subnets](#) in the Oracle Cloud Infrastructure Networking documentation.

You can set up a service gateway and give your VCN private access to the File Storage service. A service gateway can be used only by resources in the gateway's own VCN. Traffic to the service will not travel through the internet. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the File Storage service. Be sure to update route tables for any subnets that need to access File Storage through the service gateway.

For more information and detailed instructions, see [Setting Up a Service Gateway in the Console](#)

### SUBNETS

Subdivisions you define in a VCN (for example, 10.0.0.0/24 and 10.0.1.0/24). Subnets contain virtual network interface cards (VNICs), which attach to instances. A subnet can span a region or exist in a single availability domain . A subnet consists of a contiguous range of IP addresses that do not overlap with other subnets in the VCN. For each subnet, you specify the routing rules and security lists that apply to it. For more information about subnets, see [VCNs and Subnets](#) in the Oracle Cloud Infrastructure Networking documentation.

### SECURITY LISTS

Virtual firewall rules for your VCN. Your VCN comes with a default security list, and you can add more. These security lists provide ingress and egress rules that specify the types of traffic allowed in and out of the instances. You can choose whether a given rule is stateful or stateless. Security list rules must be set up so that clients can connect to file system mount targets. For more information about how security lists work in Oracle Cloud Infrastructure, see [Security Lists](#) in the Networking documentation. For information about setting up specific security list rules required for mount target traffic, see [Configuring VCN Security List Rules for File Storage](#). [About Security](#) explains how security lists interact with other types of security in your file system.

### SNAPSHOTS

Snapshots provide a consistent, point-in-time view of your file system, and you can take as many snapshots as you need. You pay only for the storage used by your data and metadata, including storage capacity used by snapshots. Each snapshot reflects only data that changed from the previous snapshot. For more information, see [Managing Snapshots](#).

## Encryption

The File Storage service encrypts all data at rest. By default all file systems are encrypted using Oracle-managed encryption keys. You have the option to encrypt all of your file systems using the keys that you own and manage using the Key Management service. For more information, see [Overview of Key Management](#).

For how to use your own key for new file systems, see [Creating File Systems](#). See [To assign a key to a file system](#) for how to assign or change the key for an existing file system.

## Data Transfers

FastConnect offers you the ability to accelerate data transfers. You can leverage the integration between FastConnect and the File Storage service to perform initial data migration, workflow data transfers for large files, and disaster recovery scenarios between two regions, among other things.

### File Storage Space Allocation

The File Storage service allocates space in blocks of variable size in a way that is fine-tuned to minimize total customer cost and optimize performance for modern workloads. The minimum block size used is 8192 bytes. For example, if you create a 1-byte file, we allocate 8192 bytes. We use larger blocks to store larger files. This method of allocation might cause a different block count for files than expected after they are copied from another storage device to your Oracle Cloud Infrastructure file system.

### How File Storage Permissions Work

File Storage service resources include file systems, mount targets, and export sets. The AUTH\_UNIX style of authentication and permission checking is supported for remote NFS client requests. You use Oracle Cloud Infrastructure Identity and Access Management (IAM) policy language to define access to Oracle Cloud Infrastructure resources. You can consider exports and snapshots subsidiary resources of export sets and file systems, respectively. As such, they do not need their own permissions. Related resources include Oracle Cloud Infrastructure Compute instances and Oracle Cloud Infrastructure Networking virtual cloud networks (VCNs).

Oracle Cloud Infrastructure users require resource permissions to create, delete, and manage resources. Without the appropriate IAM permissions, you cannot export a file system through a mount target. Until a file system has been exported, Compute instances cannot mount it. For more information about creating an IAM policy, see [Let users create, manage, and delete file systems](#).

If you have successfully exported a file system on a subnet, then you use Networking security lists to control traffic to and from the subnet and, therefore, the mount target. Security lists act as a virtual firewall, allowing only the network traffic you specify to and from the IP addresses and port ranges configured in your ingress and egress rules. The security list you create for the subnet lets hosts send and receive packets and mount the file system. If you have firewalls on individual instances, use FastConnect, or use a virtual private network (VPN), the settings for those might also impact security at the networking layer. For more information about creating a security list for the File Storage service, see [Creating File](#)

[Systems](#). See [About Security](#) for more information on how different types of security work together in your file system.

### Regions and Availability Domains

You can use the File Storage service in all regions. For a list of supported regions, see [Regions and Availability Domains](#).

When you create file systems and mount targets, you specify the availability domain they are created in. All file system data is then stored entirely within the availability domain the file system resides in. Within an availability domain, the File Storage service uses synchronous replication and high availability failover to keep your data safe and available.

You cannot move a file system to a different availability domain or region. However, you can take a snapshot of your data and use a tool such as `rsync` to copy your data to a different availability domain or region. See [Managing Snapshots](#) for more information on using snapshots to protect your data.

While it is possible to access mount targets from any availability domain in a region, for optimal performance, place File Storage resources in the same availability domain as the Compute instances that access them.

Subnets can be either AD-specific or regional. You can create File Storage resources in either type of subnet. Regional subnets allow Compute instances to connect to any mount target in the subnet regardless of AD, with no additional routing configuration. However, to minimize latency, place mount targets in the same AD as Compute instances just as you would in an AD-specific subnet. For more information, see [About Regional Subnets](#).

### Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

The following File Storage resources emit events:

- File systems
- Snapshots
- Mount targets
- Exports
- Export sets

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see

[Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Limits on Your File Storage Components

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

You can use the "Classic View" of the **Service Limits** page in the Console to see File Storage service **Limit** and **Usage** data for your tenancy.

#### To view Limit and Usage data for File Storage

1. Open the navigation menu. Under Governance and Administration, click Governance, then click **Limits, Quotas and Usage**.
2. Click **Switch Back to Classic View**.
3. In Service Limits, click on **File Storage** to expand.

### About Security

This topic discusses different methods you can use to secure your file systems.



#### Tip

Watch a [video](#) about security in File Storage.

## Access Control

**Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM)** [uses policies](#) to control what users can do within Oracle Cloud Infrastructure, such as creating instances, a VCN and its security rules, mount targets, and file systems.

**Network security** controls which instance IP addresses or CIDR blocks can connect to a host file system. It uses VCN [security list rules](#) to allow or deny traffic to the mount target, and therefore access to any associated file system.

**NFS export options** apply access control on each file system export based on source IP address.

**NFS v.3 Unix security** controls what users can do on the instance, such as installing applications, creating directories, mounting external file systems by a local mount point, and reading and writing files.

This security layer...	Uses these...	To control actions like...
Oracle Cloud Infrastructure (Oracle Cloud Infrastructure)	OCI Users and policies	Creating instances and VCNs. Creating, listing, and associating file systems and mount targets.
Network security	IP addresses, CIDR blocks, security lists	Connecting the client instance to the mount target.

This security layer...	Uses these...	To control actions like...
NFS export options	File system exports, IP addresses, Unix users	Privileged source port connection, reading and writing files, and limiting root user access on a per-file system basis.
NFS v.3 Unix security	Unix users, file mode bits	Mounting file systems, reading and writing files.

You create users and groups in Oracle Cloud Infrastructure. Then, you can use policies to specify which users and groups can create, access, or modify resources such as file systems, mount targets, and export options.

The network security layer allows you to use VCN security lists to block the appropriate ports from specific IP addresses and CIDR blocks and restrict host access. However, it's on an 'all or nothing' basis - the client either can or cannot access the mount target, and therefore all file systems associated with it. See [Working with NFS Export Options](#) to specify granular controls on a per-file system basis.

File Storage service supports the AUTH\_UNIX style of authentication and permission checking for remote NFS client requests. When mounting file systems, we recommend that you use the `-nosuid` option. This option disables set-user-identifier or set-group-identifier bits. Remote users are prevented from gaining higher privileges using a `setuid` program. For more information, see [Mounting File Systems](#).

Remember that users in UNIX aren't the same as users in Oracle Cloud Infrastructure - they're not linked or associated in any way. The Oracle Cloud Infrastructure policy layer doesn't govern anything that happens inside the file system, the UNIX security layer does. Conversely, the UNIX security layer doesn't govern creating file systems or mount targets in Oracle Cloud Infrastructure.

NFS export options are a method of applying access control at the network security layer and the NFS v.3 Unix security layer. You can use NFS export options to limit access levels by IP addresses or CIDR blocks connecting to multiple file systems through exports of an associated mount target. Access can be restricted so that each client's file system is inaccessible and invisible to the other, allowing for managed hosted environment security. Moreover, you can set permissions for read-only, read/write, or root-squash for your file systems. See [Working with NFS Export Options](#) for more information.

### Encryption

The Oracle Cloud Infrastructure File Storage service always encrypts all file systems at rest. By default all file systems are encrypted using the Oracle-provided encryption keys.

You have the option to encrypt all of your file systems using the keys that you own and manage using the Key Management service. For more information, see [Overview of Key Management](#). If you do not configure a file system to use the Key Management service or you later unassign a key from the file system, the File Storage service uses the Oracle-provided encryption key instead. For how to use your own key for new file systems, see [Creating File Systems](#). See [To assign a key to a file system](#) for instructions about how to assign or change the key for an existing file system.

### Configuring VCN Security List Rules for File Storage

When you create a VCN, a default security list is also created. Rules in the security list are used to allow or deny traffic to a subnet. Before you can mount a file system, you must configure security list rules to allow traffic to the mount target subnet. File Storage requires stateful ingress to TCP ports 111, 2048, 2049, and 2050 and stateful ingress to UDP ports 111 and 2048. File storage also requires stateful egress from TCP ports 111, 2048, 2049, and 2050 and stateful egress from UDP port 111.

## CHAPTER 15 File Storage

Ingress Rules

Add Ingress Rules						
Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
No	10.0.0.0/16	TCP	All	2048-2050		TCP traffic for ports: 2048-2050
No	10.0.0.0/16	TCP	All	111		TCP traffic for ports: 111
No	10.0.0.0/16	UDP	All	2048		UDP traffic for ports: 2048
No	10.0.0.0/16	UDP	All	111		UDP traffic for ports: 111

Showing 4 Item(s) < Page 1 >

Egress Rules

Add Egress Rules						
Stateless	Destination	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
No	10.0.0.0/16	TCP	2048-2050	All		TCP traffic for ports: All
No	10.0.0.0/16	TCP	111	All		TCP traffic for ports: All
No	10.0.0.0/16	UDP	111	All		UDP traffic for ports: All

Showing 3 Item(s) < Page 1 >

See [Security Lists](#) for more information about how security lists work in Oracle Cloud Infrastructure. See [About Security](#) for information about how security lists work with other types of security in File Storage.

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let network admins manage a cloud network](#) covers management of all networking components, including security lists. See the [Policy Reference](#) for more information.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Using the Console

#### To configure security list rules for mount target traffic

Security list rules allow ingress and egress for the following:

- Open Network Computing Remote Procedure Call (ONC RPC) rpcbind utility protocol
  - Network File System (NFS) protocol
  - Network File System (MOUNT) protocol
  - Network Lock Manager (NLM) protocol
1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  2. In the **Scope** section, select the compartment that contains the subnet associated with your file system.
  3. Click the name of the cloud network associated with your file system.
  4. On the details page for the cloud network, in **Resources**, and then click **Security Lists**.
  5. Click the name of the security list used by the subnet associated with your file system.
  6. In **Resources**, click **Ingress Rules**.
  7. Click **Add Ingress Rules** and add the following ingress rule allowing **TCP** traffic.
    - Specify that it's a stateful rule by leaving the check box clear. (For more information about stateful and stateless rules, see [Stateful Versus Stateless Rules](#)). By default, rules are stateful unless you specify otherwise.
    - To allow traffic from the subnet of the cloud network, click **Source Type**, choose **CIDR**, and then enter the CIDR block for the subnet.
    - Click **IP Protocol**, and then click **TCP**.
    - In **Source Port Range**, specify the range of ports that you want to allow traffic from. Alternatively, accept the default of **All** to allow traffic from any source port.



### Important

We recommend that NFS clients be limited to reserved ports. To do this, set the **Source Port** range to **1-1023**. You can also set export options for a file system to require clients to connect from a privileged source port. For more information, see [Working with NFS Export Options](#).

- Click **Destination Port Range**, and then enter **2048-2050**.
8. Click **+ Additional Ingress Rule** and create a second stateful ingress rule allowing **TCP** traffic to a **Destination Port Range** of **111**.
  9. Click **+ Additional Ingress Rule** and create a third stateful ingress rule allowing **UDP** traffic to a **Destination Port Range** of **2048**.
  10. Click **+ Additional Ingress Rule** and create a fourth stateful ingress rule allowing **UDP** traffic to a **Destination Port Range** of **111**.
  11. When you're done, click **Add Ingress Rules**.
  12. Next, create the egress rules. In **Resources**, click **Egress Rules**.
  13. Click **Add Egress Rules** and add the following egress rule allowing **TCP** traffic:
    - Specify that it's a stateful rule by leaving the check box clear.
    - Click **Destination Type**, choose **CIDR**, and then enter the CIDR block for the subnet.
    - Click **IP Protocol**, and then click **TCP**.
    - In **Source Port Range**, enter **2048-2050**.
    - In **Destination Port Range**, accept the default of **All** to allow traffic to any destination port.

14. Click **+ Additional Egress Rule** and add a second stateful egress rule allowing **TCP** traffic from a **Source Port Range** of **111**.
15. Click **+ Additional Egress Rule** and add a third stateful egress rule allowing **UDP** traffic from a **Source Port Range** of **111**.
16. When you're done, click **Add Egress Rules**.

Next steps:

- [Create a mount target and associated file system](#)
- [Mount a file system](#)

## Working with NFS Export Options

This topic describes the basic features of NFS export options, and how to control client access to your file system.

### Overview

NFS export options enable you to create more granular access control than is possible using just security list rules to limit VCN access. You can use NFS export options to specify access levels for IP addresses or CIDR blocks connecting to file systems through exports in a mount target. Access can be restricted so that each client's file system is inaccessible and invisible to the other, providing better security controls in multi-tenant environments.

Using NFS export option access controls, you can limit clients' ability to connect to the file system and view or write data. For example, if you want to allow clients to consume but not update resources in your file system, you can set access to Read Only. You can also reduce client root access to your file systems and map specified User IDs (UIDs) and Group IDs (GIDs) to an anonymous UID/GID of your choice. For more information about how NFS export options work with other security layers, see [About Security](#).



### Tip

Watch a [video](#) about working with NFS export options in File Storage.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to manage NFS export options.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Exports

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

See [To create an export for a file system](#) for more information.

### **NFS Export Options**

NFS export options are a set of parameters within the export that specify the level of access granted to NFS clients when they connect to a mount target. An NFS export options entry within an export defines access for a single IP address or CIDR block range.

Each separate client IP address or CIDR block you want to define access for needs a separate export options entry in the export. For example, if you want to set options for NFS client IP addresses 10.0.0.6, 10.0.0.8, and 10.0.0.10, you need to create three separate entries, one for each IP address.

File Storage service considers the listed order of each export options entry for the export. During an NFS request by a client, File Storage service applies the first set of options that matches the client Source IP address. Only the first set is applied; the rest are ignored.

For example, consider the following two export options entries specifying access for an export:

Entry 1: Source: 10.0.0.0/16, Access: Read Only

Entry 2: Source: 10.0.0.8, Access: Read/Write

In this case, clients who connect to the export from IP address 10.0.0.8 have Read Only access. The request Source IP address is contained in the CIDR block specified in the first entry, and File Storage Service applies the options in the first match.



### Important

File systems can be associated with one or more exports, contained within one or more mount targets. If the client **source** IP address does not match any entry on the list for a single export, then that export is not visible to the client. However, the file system could be accessed through other exports on the same or other mount targets. **To completely deny client access to a file system, be sure that the client source IP address or CIDR block is not included in any export for any mount target associated with the file system.**

The following options can be set to control export access:

- **Source:** The IP address or CIDR block of a connecting NFS client.
- **Require Privileged Source Port (true/false):** This setting determines whether the NFS clients specified in **source** are required to connect from a privileged source port. Privileged ports are any port including 1-1023. On Unix-like systems, only the root user can open privileged ports. Setting this value to **true** disallows requests from unprivileged ports. The default for this setting is different depending on how the export is created. Creating an export without an explicit `ClientOption` array sets the `requirePrivilegedSourcePort` attribute of the client option to **false**. When you create a `ClientOption` array explicitly, `requirePrivilegedSourcePort` defaults to **true**. For example, creating an export in the Console using the default selections sets `requirePrivilegedSourcePort` to **false**. Creating an export in the API along with a `ClientOption` array sets `requirePrivilegedSourcePort` to **true**.



### Important

When **Require Privileged Source Port** is set to **true**, you also have to follow these additional configuration steps:

1. When mounting the file system from a Unix-like system, include the `resvport` option in your mount command when mounting. For example:

```
sudo mount -o resvport 10.x.x.x:/fs-export-path
/mnt/yourmountpoint
```

For more information, see [Mounting File Systems From Unix-Style Instances](#).

2. When mounting the file system from a Windows system, be sure the **UseReserverdPorts** registry key value is set to **1**.

For more information, see [Mounting File Systems From Windows Instances](#).

- **Access (Read\_Only, Read\_Write):** This setting specifies the **source** NFS client access. If unspecified, defaults to **Read\_Write**.
- **Identity Squash: (All, Root, None):** This setting determines whether the **source** clients accessing the file system have their User ID (UID) and Group ID (GID) remapped to **anonymousUid** and **anonymousGid**. If you choose **All**, all users and groups are remapped. If **Root**, only the root user UID/GID combination 0/0 is remapped. If **None**, no users are remapped. If unspecified, defaults to **None**.
- **anonymousUid:** This setting is used along with the **Identity Squash** option. When remapping users, you can use this setting to change the default anonymousUid of **65534** to any user ID of your choice.

- **anonymousGid:** This setting is used along with the **Identity Squash** option. When remapping groups, you can use this setting to change the default anonymousGid of **65534** to any group ID of your choice.

### Typical Access Control Scenarios

When you create file system and export, the NFS export options for that file system are set to the following defaults, which allow full access for all NFS client source connections. These defaults must be changed if you want to restrict access:

- **Source:** 0.0.0.0/0 (All)
- **Require Privileged Source Port:** False
- **Access:** Read\_Write
- **Identity Squash:** None

#### SCENARIO A: CONTROL HOST BASED ACCESS

Provide a managed hosted environment for two clients. The clients share a mount target, but each has their own file system, and cannot access each other's data. For example:

- Client A, who is assigned to CIDR block 10.0.0.0/24, requires Read/Write access to file system A, but not file system B.
- Client B, who is assigned to CIDR block 10.1.1.0/24, requires Read/Write access to file system B, but not file system A.
- Client C, who is assigned to CIDR block 10.2.2.0/24, has no access of any kind to file system A or file system B.
- Both file systems A and B are associated to a single mount target, MT1. Each file system has an export contained in the export set of MT1.

Since Client A and Client B access the mount target from different CIDR blocks, you can set the client options for both file system exports to allow access to only a single CIDR block. Client C is denied access by not including its IP address or CIDR block in the NFS export options for any export of either file system.

### Console Example

Set the export options for file system A to allow Read/Write access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. ⓘ

Source	Ports	Access	Squash	UID	GID
10.0.0.0/24	Privileged ↕	Read/Write ↕	None ↕	Not used	Not used

+ Another Option

Update

Set the export options for file system B to allow Read/Write access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. (i)

Source	Ports	Access	Squash	UID	GID
10.1.1.0/24	Privileged <span style="font-size: 0.8em;">⇅</span>	Read/Write <span style="font-size: 0.8em;">⇅</span>	None <span style="font-size: 0.8em;">⇅</span>	Not used	Not used
					<span style="font-size: 0.8em;">⋮</span>
					<span style="font-size: 0.8em;">+ Another Option</span>

Update

## CLI Example

Set the export options for file system A to allow Read\_Write access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

```
oci fs export update --export-id <File_system_A_export_ID> --export-options '[{"source":"10.0.0.0/24","require-privileged-source-port":"true","access":"READ_WRITE","identity-squash":"NONE","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

Set the export options for file system B to allow Read\_Write access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

```
oci fs export update --export-id <File_system_B_export_ID> --export-options '[{"source":"10.1.1.0/24","require-privileged-source-port":"true","access":"READ_WRITE","identity-squash":"NONE","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

## API Example

Set the export options for file system A to allow READ\_WRITE access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

## CHAPTER 15 File Storage

---

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.0.0.0/24",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "NONE",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```

Set the export options for file system B to allow READ\_WRITE access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

```
PUT /<Current_API_Version>/exports/<File_System_B_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.1.1.0/24",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "NONE",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```

### SCENARIO B: LIMIT THE ABILITY TO WRITE DATA

Provide data to customers for consumption, but don't allow them to update the data.

For example, you'd like to publish a set of resources in file system A for an application to consume, but not change. The application connects from IP address 10.0.0.8.

## Console Example

Set the source IP address 10.0.0.8 to Read Only in the export for file system A:

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. (i)

Source	Ports	Access	Squash	UID	GID
10.0.0.8	Privileged ▾	Read Only ▾	None ▾	Not used	Not used

+ Another Option

Update

## CLI Example

Set the source IP address 10.0.0.8 to READ\_ONLY in the export for file system A:

```
oci fs export update --export-id <File_System_A_export_OCID> --export-options '
[{"source": "10.0.0.8", "require-privileged-source-port": "true", "access": "READ_
ONLY", "identitysquash": "NONE", "anonymousuid": "65534", "anonymousgid": "65534"}]'
```

## API Example

Set the source IP address 10.0.0.8 to READ\_ONLY in the export for file system A:

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.0.0.8",
 "requirePrivilegedSourcePort": true,
```

## CHAPTER 15 File Storage

```
 "access": "READ_ONLY",
 "identitySquash": "NONE",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
```

### SCENARIO C: IMPROVE FILE SYSTEM SECURITY

To increase security, you'd like to limit the root user's privileges when connecting to File System A. Use Identity Squash to remap root users to UID/GID 65534. In Unix-like systems, this UID/GID combination is reserved for *'nobody'*, a user with no system privileges.

### Console Example

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. ⓘ

Source	Ports	Access	Squash	UID	GID
0.0.0.0/0	Privileged ↕	Read/Write ↕	Root ↕	65534	65534
					+ Another Option

[Update](#)

### CLI Example

```
oci fs export update --export-id <File_System_A_export_OCID> --export-options '[{"source":"0.0.0.0/0","require-privileged-source-port":"true","access":"READ_WRITE","identitysquash":"ROOT","anonymousuid":"65534","anonymousgid":"65534"}]'
```

### API Example

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "0.0.0.0/0",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "ROOT",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```



#### Tip

If you don't want a file system to be visible to any clients, you can set all of the properties in the `exportOptions` array to empty values. For example,

```
{
 "exportOptions": [
 {
 "source": "",
 "requirePrivilegedSourcePort": "",
 "access": "",
 "identitySquash": ""
 }
]
}
```

### Using the Console

#### To set export options for a file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment. All of the file systems in the selected compartment are displayed.
3. Find the file system you want to set export options for, click the the Actions icon (three dots), and then click **View File System Details**.
4. In the **Exports** list, find the export you want to set export options in, click the the Actions icon (three dots), and then click **View Export Details**. If there is no export listed for the file system, you can create one. See [To create an export for a file system](#) for more information.



#### Tip

To be sure you be sure that you select export, check the following:

- **The export path:** This path uniquely identifies the file system within the mount target. [No two exports in a mount target can have the same export path, even if the exports are for the same file system.](#)
- **The mount target name:** File systems can be exported through more than one mount target. Be sure that you've selected the export for the correct mount target.

5. Click **Edit Export Options**.
6. Make one or more of these changes:

- Change an export option entry in the list.
  - Click **+Another Option** to create a new export option entry.
  - Click the Actions icon (three dots) for an entry and move it up or down in the list.
7. When you're done, click **Update**.

### Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To create an export

Open a command prompt and run `oci fs export create` to create an export for a specified file system within a specified export set. This example creates an export along with its NFS export options.

For example:

```
oci fs export create --export-set-id <export_set_OCID> --file-system-id <file_system_OCID> --path
"</pathname>" --export-options '[{"source":"10.0.0.0/16","requireprivilegedsourceport":"true","access":"READWRITE","identitysquash":"NO
NE","anonymousuid":"0","anonymousgid":"0"}]'
```



### Important

#### *Export Path Names*

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For any two export resources associated with the same export set, the path sequence for the first export resource can't contain the complete path element sequence of the second export sequence. Paths can't end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes.

Examples:

Acceptable:

`/example` and `/path`

`/example1` and `/example2`

Not Acceptable:

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../../example1`

### To update export options

Open a command prompt and run `oci fs export update`. To update export options for a specified file system, use `--export-options`.

## CHAPTER 15 File Storage

---

For example:

```
oci fs export update --export-id <export_OCID> --export-options '[{"source":"<0.0.0.0/0>","require-privileged-source-port":"true","access":"READ_ONLY","identity-squash":"ROOT","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

```
WARNING: Updates to export-options will replace any existing values. Are you sure you want to continue?
[y/N]: y
```



### Tip

If you don't want a file system to be visible to any clients, you can set all of the properties in Client Options to empty values. For example,

```
oci fs export update --export-id <export_OCID> --export-options '[{"source":"","require-privileged-source-port":"true","access":"READ_ONLY","identity-squash":"ROOT","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

### To list exports

Open a command prompt and run `oci fs export list` to list all exports in a specified compartment.

For example:

```
oci fs export list --compartment-id target_compartment_id
```

### To delete an export

Open a command prompt and run `oci fs export delete` to delete an export.

For example:

```
oci fs export delete --export-id export_OCID
```



### Warning

When you delete an export, you can no longer mount the file system using the file path specified in the deleted export.

### Using the API

- [CreateExport](#)
- [UpdateExport](#)
- [ListExports](#)
- [GetExport](#)
- [DeleteExport](#)

## Creating File Systems

You can create a shared file system in the cloud using the File Storage service. Network access to your file system is provided through a [mount target](#). Exports control how NFS clients access file systems when they connect to a mount target. File systems must have at least one [export](#) in one mount target for any instance to [mount](#) and use the file system. Typically, you create your first mount target when you create your first file system.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to create file systems. Since mount targets are network endpoints, users must also have "use" permissions for VNICs, private IPs, private DNS zones, and subnets to create or delete a mount target. See the [Policy Reference](#) for more information.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Prerequisites

Before you create a file system, you need:

- At least one Virtual Cloud Network (VCN) in a compartment. For more information, see [VCNs and Subnets](#).
- Correctly configured security list rules in the VCN subnet where you plan to create the file system's associated mount target. See [Security Lists](#) for information about how security lists work in Oracle Cloud Infrastructure. Use the instructions in [Configuring VCN Security List Rules for File Storage](#) to set up security lists for your file systems.

### Using the Console

#### To create a file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.

2. In the left-hand navigation, in the **List Scope** section, under **Compartment**, select a compartment.
3. Click **Create File System**.



### Note

File systems are encrypted by default. You cannot turn off encryption.

4. You can choose to accept the system defaults, or change them by clicking **Edit Details**.

- **File System Information:**

- **Name:** File Storage service creates a default name using "FileSystem-YYMMDD-HHMM". Optionally, change the default name for the file system. It doesn't have to be unique; an Oracle Cloud Identifier (OCID) uniquely identifies the file system.
- **Availability Domain:** The first availability domain selected in the left panel list is used as default.
- **Encryption:** File systems use Oracle-managed keys by default, which leaves all encryption-related matters to Oracle. Optionally, you can encrypt the data in this file system using your own Key Management encryption key. To use Key Management for your encryption needs, select **Encrypt using customer-managed keys** check box. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you

are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

- **Export Information**

Mount targets use exports to manage access to file systems. The path name uniquely identifies the file system within the mount target, and is used by an instance to mount the file system.

- **Export Path:** The File Storage service creates a default export path using the file system name. Optionally, replace the default export path name with a new path name, preceded by a forward slash (/). For example, /fss. This value specifies the mount path to the file system (relative to the mount target IP address or hostname). Avoid entering confidential information.



### Important

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For multiple file systems associated with a single mount target, the path sequence for the first file system cannot contain the complete path element sequence of the second file system path sequence. Paths cannot end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes. For example:

#### Acceptable:

`/example` and `/path`  
`/example` and `/example2`

#### Not Acceptable:

`/example` and `/example/path`  
`/` and `/example`  
`/example/`  
`/example/path/../../example1`



### Warning

If one file system associated to a mount target has '/' specified as an export path, you can't associate another file system with that mount target.



### Note

Export paths cannot be edited after the export is created. If you want to use a different export path, you must create a new export with the desired path. Optionally, you can then delete the export with the old path.

For more information, see [Paths in File Systems](#).

- **Use Secure Export Options:** Select to set the export options to require NFS clients to use a privileged port (1-1023) as its source port. This option enhances security because only a client with root privileges can use a privileged source port. After the export is created, you can edit the export options to adjust security. See [Working with NFS Export Options](#) for more information.



### Warning

Leaving the "Use Secure Export Options" setting disabled allows unprivileged users to read and modify any file or directory on the target file system.

- **Mount Target Information:**

File systems must be associated with a mount target to be mounted by an instance.

If you have one or more previously created mount targets in the availability domain, the File Storage service automatically chooses the most recently created mount target in the list. If you don't have a mount target in the selected availability domain, the File Storage service creates one using the following defaults.

- **Mount Target Name:** File Storage service creates a default mount target name using "Mount-YYYYMMDD-HHMM".
  - **Compartment:** The compartment you're currently working in.
  - **Virtual Cloud Network:** The first VCN listed in the current compartment is used as default.
  - **Subnet:** The most recently created subnet listed in the selected availability domain is used as default. Subnets can be either AD-specific or regional (regional ones have "*regional*" after the name). For more information, see [About Regional Subnets](#).
5. If you want to *accept the defaults* for the mount target, click **Create**. *The file system is created with the information displayed.* If you want to choose another mount target or change the default information, click the **Edit Details** link.
  6. In the **Mount Target Information** section, specify details for the mount target that is associated with the file system:

- **Select an Existing Mount Target:** Choose this option if you want to associate the file system with a mount target you already created. Choose the **Mount Target** from the list. Click the **click here** link in the dialog box if you want to enable compartment selection for the mount target.



### Tip

If there aren't any mount targets in the current combination of availability domain and compartment, this option is disabled. You can:

- Choose a different compartment.
  - Choose a different availability domain in the *File System Information* section.
  - Create a new mount target.
- **Create a New Mount Target:** Choose this option if you want to create a new mount target associated with this file system. By default, the mount target is created in your current compartment and you can use network resources in that compartment. Click the **click here** link in the dialog box if you want to enable compartment selection for the mount target, its VCN, or subnet resources.



### Important

The mount target is always in the same availability domain as the file system. While it is possible to access mount targets from any AD in a region, for optimal performance, your mount target and file system should be in the same availability domain as the Compute instances that access them. For more information, see [Regions and Availability Domains](#).

- **Create in Compartment:** Specify the compartment you want to create the mount target in.
- **New Mount Target Name:** Optionally, replace the default with a friendly name for the mount target. It doesn't have to be unique; an Oracle Cloud Identifier (OCID) uniquely identifies the mount target. Avoid entering confidential information.



### Note

The mount target name is different than the DNS hostname, which is specified in step 7.

- **Virtual Cloud Network Compartment:** The compartment containing the cloud network (VCN) in which to create the mount target.
- **Virtual Cloud Network:** Select the cloud network (VCN) where you want to create the new mount target.

- **Subnet Compartment:** Specify the compartment containing a subnet within the VCN to attach the mount target to.
- **Subnet:** Select a subnet to attach the mount target to. Subnets can be either AD-specific or regional (regional ones have "*regional*" after the name). For more information, see [About Regional Subnets](#).



### Warning

Each mount target requires three internal IP addresses in the subnet to function. Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Optionally, click **Show Advanced Options** to configure the mount target's advanced options.
    - **IP Address:** You can specify an unused IP address in the subnet you selected for the mount target.
    - **Hostname:** You can specify a hostname you want to assign to the mount target.



### Note

The File Storage service constructs a fully qualified domain name (FQDN) by combining the hostname with the FQDN of the subnet the mount target is located in.

For example,

```
myhostname.subnet123
```

```
.dnslabel.oraclevcn.com.
```

Once created, the hostname may be changed in the mount target's Details page. See [Managing Mount Targets](#) for more information.

### 8. Click **Create**.

The File Storage service typically creates the file system and mount target within seconds. Next, mount the file system from an instance so that you can read and write directories and files in your file system. See [Mounting File Systems](#) for instructions about obtaining mount commands for your operating system type and mounting your file system.

## Using the command line interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To create a file system

Open a command prompt and run `oci fs file-system create` to create a file system. For example:

```
oci fs file-system create --availability-domain <target_availability_domain> --display-name "<My File System>" --compartment-id <target_compartment_id>
```



### Warning

Avoid entering confidential information in the file system `display-name`.

The file system is created.

File systems use Oracle-managed keys by default, which leaves all encryption-related matters to Oracle. Optionally, you can encrypt the data in this file system using your own Key Management encryption key. For more information, see [Overview of Key Management](#).

For example:

```
oci fs file-system create --availability-domain AAbC:US-ASHBURN-AD-1 --display-name "My File System" --
compartment-id ocid1.compartment.oc1..<unique_id> --kms-key-id --kms-key-id ocid1.key.oc1.phx.<unique_
id>
```

## To create a mount target

You can create a mount target for file systems in a specified compartment and subnet. A file system can only be associated with a mount target in the same availability domain.



### Warning

Each mount target requires three internal IP addresses in the subnet to function. Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

Open a command prompt and run `oci fs mount-target create` to create a mount target.

For example:

```
oci fs mount-target create --availability-domain <target_availability_domain> --compartment-id <target_compartment_id> --subnet-id <subnet_OCID> --display-name "<My Mount Target>"
```



### Warning

Avoid entering confidential information in the mount target display-name.

## To create an export

An export is a file system together with the path that can be used to mount it. Each export resource belongs to one export set.

Open a command prompt and run `oci fs export create` to create an export for a specified file system within a specified export set.

For example:

```
oci fs export create --export-set-id <export_set_OCID> --file-system-id <file_system_OCID> --path "</pathname>"
```



### Important

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For multiple file systems associated with a single mount target, the path sequence for the first file system cannot contain the complete path element sequence of the second file system path sequence. Paths cannot end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes. For example:

Acceptable:

`/example` and `/path`

`/example` and `/example2`

Not Acceptable:

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../../example1`



### Warning

If one file system associated to a mount target has '/' specified as an export path, you can't associate another file system with that mount target.



### Note

Export paths cannot be edited after the export is created. If you want to use a different export path, you must create a new export with the desired path. Optionally, you can then delete the export with the old path.

For more information, see [Paths in File Systems](#).

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to create file systems:

- [CreateFileSystem](#)
- [CreateMountTarget](#)
- [CreateExport](#)

## Mounting File Systems

Users of Unix-style operating systems and Windows Server 2008 R2, 2012 R2, or 2016 can connect to a file system and write files. [Mount targets](#) serve as file system network access points for file systems. After your mount target is assigned an IP address, you can use it together with the file system [export path](#) to mount the file system. On the instance from which you want to mount the file system, you need to install an NFS client. For Unix-style operating systems, you create a mount point. When you mount the file system, the mount point effectively represents the root directory of the File Storage file system, allowing you to write

files to the file system from the instance. Windows operating systems use a drive letter assignment instead of a mount point to represent root access.

### Prerequisites

- The file system must have at least one export in at least one mount target. When you create a new file system, an export for the file system is created at the same time. See [Creating File Systems](#) for more information.
- Correctly configured security list rules in the VCN subnet where the file system's associated mount target resides. See [Security Lists](#) for information about how security lists work in Oracle Cloud Infrastructure. Use the instructions in [Configuring VCN Security List Rules for File Storage](#) to set up security lists for your file systems.

### Mounting File Systems From an Instance

[Mounting File Systems From Unix-Style Instances](#) (Including Oracle Linux DB instances)

[Mounting File Systems From Windows Instances](#)

### Obtaining Mount Command Samples

You can use the Console to get mount command samples that include all the information for a specific mount target and file system. Samples are available for the following operating system images:

- Oracle Linux
- CentOS
- Debian
- Red Hat Linux
- Ubuntu



### Warning

When mounting file systems, the following mount option combination is **not supported** by the File Storage service:

- `soft` when the file system is mounted with the read/write mount option (`-o rw`). **This combination can cause corruption of your data.**

The following mount options or mount option combinations are **not recommended** for use with the File Storage service:

- `soft` when the file system is mounted with the read-only mount option (`-o ro`) and the `timeo` has been specified as less than 300 seconds. **This combination can cause a profusion of I/O error responses.**
- `rsize`, `or wsize`. **These options cause issues with performance.**



### Note

When mounting file systems, Network Lock Manager (NLM) is enabled for file locking by default. The default requires no specified mount option. Typical NFS workloads function normally using the default.

Some applications might require you to specify the `nolock` mount option. Refer to your application documentation for best practices regarding this mount option.

## Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to obtain mount commands.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

## Using the Console

### To get mount command samples

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.

The Console displays a list of file systems that have already been created in the compartment, if any.

3. Find the file system you want to mount, click the Actions icon (three dots), and then click **View File System Details**.
4. In **Resources**, click **Exports**.
5. Find the export in the mount target you want to use to mount the file system, click the Actions icon (three dots), and then click **Mount Commands**.



### Tip

To be sure that you select the correct export, check the following:

- **The export path:** This path uniquely identifies the file system within the mount target. No two exports in a mount target can have the same export path, even if the exports are for the same file system.
- **The mount target name:** File systems can be exported through more than one mount target. Be sure that you've selected the export for the correct mount target.

6. In **Image**, choose the image of the Compute instance you want to mount the file system to.
7. Click the **Copy** link to copy the commands.

Next, mount the file system from a [Unix-style](#) or [Windows](#) instance.

### Mounting File Systems From Unix-Style Instances

Users of Ubuntu and Linux operating systems can use the command line to connect to a file system and write files. [Mount targets](#) serve as network access points for file systems. After your mount target is assigned an IP address, you can use it together with the [export path](#) to mount the file system. On the instance from which you want to mount the file system, you need to install an NFS client and create a mount point. When you mount the file system, the mount point effectively represents the root directory of the File Storage file system, allowing you to write files to the file system from the instance.

#### Prerequisites

- The file system must have at least one export in at least one mount target. When you create a new file system, an export for the file system is created at the same time. See [Creating File Systems](#) for more information.
- Correctly configured security list rules in the VCN subnet where the file system's associated mount target resides. See [Security Lists](#) for information about how security lists work in Oracle Cloud Infrastructure. Use the instructions in [Configuring VCN Security List Rules for File Storage](#) to set up security lists for your file systems.

#### Mounting File Systems

You can use the following instructions to construct your mount commands, or use the Console to get mount command samples that include all the information for a specific mount target and file system. For more information, see [To get mount command samples](#).



### Warning

When mounting file systems, the following mount option combination is **not supported** by the File Storage service:

- `soft` when the file system is mounted with the read/write mount option (`-o rw`). **This combination can cause corruption of your data.**

The following mount options or mount option combinations are **not recommended** for use with the File Storage service:

- `soft` when the file system is mounted with the read-only mount option (`-o ro`) and the `timeo` has been specified as less than 300 seconds. **This combination can cause a profusion of I/O error responses.**
- `rsize`, `or wsize`. **These options cause issues with performance.**



### Note

When mounting file systems, Network Lock Manager (NLM) is enabled for file locking by default. The default requires no specified mount option. Typical NFS workloads function normally using the default.

Some applications might require you to specify the `nolock` mount option. Refer to your application documentation for best practices regarding this mount option.

### To mount a file system from Ubuntu or Debian

1. Open a command window. Then, get the NFS client by copying and pasting the **Install Command** from the Console or type the following:

```
sudo apt-get install nfs-common
```

2. Create a mount point by copying and pasting the **Create Mount Point Command** from the Console or type the following, replacing `yourmountpoint` with the local directory from which you want to access your file system.

```
sudo mkdir -p /mnt/yourmountpoint
```

3. Mount the file system by copying and pasting the **Mount Command** from the Console or type the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `yourmountpoint` with the path to the local mount point. The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then `10.x.x.x:/` represents the full extent of the mount target.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
sudo mount -o nosuid,rsvport 10.x.x.x:/fs-export-path /mnt/yourmountpoint
```



### Warning

Omitting the `-o nosuid` option may allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.



### Note

The `-o rsvport` option is required when the "Require Privileged Source Port" export option is used and otherwise optional. It causes the mounting filesystem to connect from a privileged source port (1-1023). See [Working with NFS Export Options](#) for more information.

4. View the file system.

```
df -h
```

5. Write a file to the file system by typing the following. Replace `yourmountpoint` with the

path to the local mount point and `helloworld` with your file name.

```
sudo touch /mnt/yourmountpoint/helloworld
```

6. Verify that you can view the file by typing the following. Replace `yourmountpoint` with the path to the local mount point.

```
cd /mnt/yourmountpoint
```

```
ls
```

See [Mount Command Fails](#) in [Troubleshooting Your File System](#) for more information about common issues you may encounter.

### To mount a file system from Linux, Red Hat, or CentOS

1. Open a command window. Then, get the NFS client by copying and pasting the **Install Command** from the Console or typing the following:

```
sudo yum install nfs-utils
```

2. Create a mount point by copying and pasting the **Create Mount Point Command** from the Console or type the following, replacing `yourmountpoint` with the local directory from which you want to access your file system.

```
sudo mkdir -p /mnt/yourmountpoint
```

3. Mount the file system by copying and pasting the **Mount Command** from the Console or type the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `yourmountpoint` with the path to the local mount point. The export path is the path to the file system (relative to the mount target's IP address or hostname). If you did not specify a path when you associated the file system and mount target, then `10.x.x.x:/` represents the full extent of the mount target.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
sudo mount -o nosuid,rsvport 10.x.x.x:/fs-export-path /mnt/yourmountpoint
```



### Warning

Omitting the `-o nosuid` option may allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.



### Note

The `-o rsvport` option is required when the "Require Privileged Source Port" export option is used and otherwise optional. It causes the mounting filesystem to connect from a privileged source port (1-1023). See [Working with NFS Export Options](#) for more information.

4. View the file system.

```
df -h
```

5. Write a file to the file system by typing the following. Replace `yourmountpoint` with the

path to the local mount point and `helloworld` with your file name.

```
sudo touch /mnt/yourmountpoint/helloworld
```

6. Verify that you can view the file by typing the following. Replace `yourmountpoint` with the path to the local mount point.

```
cd /mnt/yourmountpoint
```

```
ls
```

See [Mount Command Fails](#) in [Troubleshooting Your File System](#) for more information about common issues you may encounter.

### To mount a file system from a Database VM instance

Database VM instances are built on Oracle Linux 6.8, unlike Oracle Linux Compute instances, which run on version 7.4. The NFS Utilities package is pre-installed on DB instances, but the Open Network Computing Remote Procedure Call (ONC RPC) `rpcbind` utility is disabled by default. Oracle Linux 6.8 does not have `systemd`, so DB instances are managed differently than OL compute instances. An Oracle DB instance comes with a set of `iptables` rules that excludes any non-database ports and need to be updated to allow mount target traffic.

1. SSH to the DB system.

```
ssh -i <private_key_path> opc@<db_system_ip_address>
```

2. Start the `rpcbind` service by typing the following:

```
sudo service rpcbind start
```

3. Use the `chkconfig` command to enable starting `rpcbind` service at system startup.

```
sudo chkconfig rpcbind on
```

4. Change the default configuration of `iptables` to include the mount target IP address and allow traffic by typing the following. Replace `10.x.x.x` with the local subnet address assigned to the mount target for the file system. Save the new `iptables` entries.

```
sudo iptables -A INPUT -p tcp -s 10.x.x.x -j ACCEPT
```

## CHAPTER 15 File Storage

---

```
sudo iptables -A OUTPUT -p tcp -s 10.x.x.x -j ACCEPT
```

```
sudo service iptables save
```

5. Create a mount point by typing the following, replacing `yourmountpoint` with the local directory from which you want to access your file system.

```
sudo mkdir -p /mnt/yourmountpoint
```

6. Mount the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `yourmountpoint` with an absolute path to a local mount point. The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then `10.x.x.x:/` represents the full extent of the mount target.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
sudo mount -t nfs -o nosuid,rsvport,tcp,vers=3 10.x.x.x:/fs-export-path /mnt/yourmountpoint
```



### Warning

Omitting the `-o nosuid` option may allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.



### Note

The `-o resvport` option is required when the “Require Privileged Source Port” export option is used and otherwise optional. It causes the mounting filesystem to connect from a privileged source port (1-1023). See [Working with NFS Export Options](#) for more information.

See [Mount Command Fails](#) in [Troubleshooting Your File System](#) for more information about common issues you may encounter.

### To auto-mount a shared file system

Auto-mount ensures that a file system is automatically re-mounted on an instance if it is rebooted.

1. Open a command window. Then, mount the file system using the steps described in the previous section.
2. Type the following command to get the file system entry point:

```
sudo cat /etc/mtab |grep -i nfs
```

3. Copy the file system entry point, and open the `/etc/fstab` file:

```
cd /etc
```

```
vi fstab
```

4. Add the following line to the `fstab` file:

```
<file_system_ip_address>:<file_system_path_name><your_local_mount_point> nfs
defaults,nofail,nosuid,resvport 0 0
```



### Warning

Omitting the `-o nosuid` option may allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.



### Important

Be sure to add the `nofail` option to each entry. This option ensures that an unavailable file system does not cause the instance reboot process to fail.



### Note

The `-o resvport` option is required when the "Require Privileged Source Port" export option is used and otherwise optional. It causes the mounting filesystem to connect from a privileged source port (1-1023). See [Working with NFS Export Options](#) for more information.

5. Save the fstab file.

See [Mount Command Fails](#) in [Troubleshooting Your File System](#) for more information about common issues you may encounter.

### Mounting File Systems From Windows Instances

Users of Windows Server 2008 R2, 2012 R2, or 2016 can mount a file system on any available drive letter using the mount target IP address and the file system [export path](#).

The Windows NFS client must be installed on the instance from which you want to mount the file system.



#### Warning

Installing the Windows NFS client may require a restart of your system.

Access to NFS file systems requires UNIX-style user and group identities, which are not the same as Windows user and group identities. To enable users to access NFS shared resources, Windows client for NFS accesses file systems anonymously, using `AnonymousGid` and `AnonymousUid`. On brand new file systems, write permissions are only granted to the root user. The `AnonymousGid` and `AnonymousUid` identity values must be configured to allow write access.



#### Warning

Updating the 'AnonymousGid' and 'AnonymousUid' values require registry changes to your system.

After you have installed the NFS client and correctly mapped user identities, you can mount the file system to any available drive letter using the command line or **Map network drive**. You can access your file system through the chosen drive letter to write files.

#### Prerequisites

- The file system must have at least one export in at least one mount target. When you create a new file system, an export for the file system is created at the same time. See

[Creating File Systems](#) for more information.

- Correctly configured security list rules in the VCN subnet where the file system's associated mount target resides. See [Security Lists](#) for information about how security lists work in Oracle Cloud Infrastructure. Use the instructions in [Configuring VCN Security List Rules for File Storage](#) to set up security lists for your file systems.



### Warning

When mounting file systems, the following mount option combination is **not supported** by the File Storage service:

- `soft` when the file system is mounted with the read/write mount option (`-o rw`). **This combination can cause corruption of your data.**

The following mount options or mount option combinations are **not recommended** for use with the File Storage service:

- `soft` when the file system is mounted with the read-only mount option (`-o ro`) and the `timeo` has been specified as less than 300 seconds. **This combination can cause a profusion of I/O error responses.**
- `rsize`, `orwsize`. **These options cause issues with performance.**



### Note

When mounting file systems, Network Lock Manager (NLM) is enabled for file locking by default. The default requires no specified mount option. Typical NFS workloads function normally using the default.

Some applications might require you to specify the `nolock` mount option. Refer to your application documentation for best practices regarding this mount option.

### Using Windows Command Prompt

#### To mount a file system from Windows Server 2008 R2 Command Prompt

If you are using Oracle-provided Windows images, the NFS client is already installed, and the correct user identities are mapped. Skip to step 7.

1. Install Services for NFS components.
  - a. Click **Start** go to **Administrative Tools**, then click **Server Manager**.
  - b. In the left pane, click **Roles**.
  - c. Under **Roles Summary** in the right pane, click **Add Roles**. The **Add Roles Wizard** appears. Click **Next**.
  - d. Select the **File Services** check box from the list and click **Next**. Review the information and click **Next**.
  - e. Select the **Services for Network File System** check box from the list, and click **Next**.

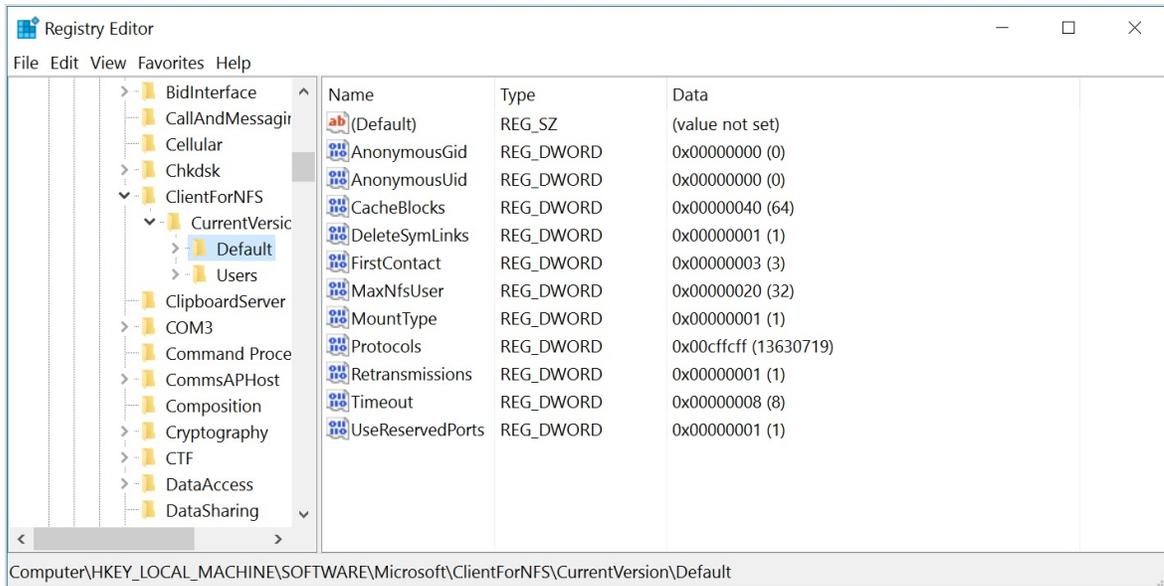
- f. Confirm your selections, then click **Install**. Click **Close** when the installation is complete.
2. Open the registry editor (regedit) and map the AnonymousGid and AnonymousUid to the root user.



### Warning

User identity mapping requires changes to your system registry.

- a. Click **Windows Search**.
  - b. Enter `regedit` in the **Search** field and press **Enter**.
  - c. Click **Yes** to allow changes to your device.
  - d. Click `HKEY_LOCAL_MACHINE`. Then, browse to:  
`Software\Microsoft\ClientForNFS\CurrentVersion\Default`.
3. Add a new DWORD32 registry entry for `AnonymousGid`:
  - a. Click **Edit**, and select **New DWORD (32 bit) Value**.
  - b. In the **Name** field, enter `AnonymousGid`. Leave the value at 0.
4. Repeat step 3 to add a second DWORD32 registry entry named `AnonymousUid` with a value of 0.



5. Open Windows Command Line (CMD) and run as **Administrator**:
  - a. Click **Start** .
  - b. Press **Ctrl+Shift** and click **Command Prompt**.
  - c. Click **Yes**.

 **Important**

If you've set export options for your file system to require clients to connect from a privileged source port (1-1023), then you must set the **UseReserverdPorts** registry key to **1**. For more information, see [Working with NFS Export Options](#).

6. In the Administrator: Windows Command Prompt (CMD) window, restart the NFS Client by typing the following:

```
nfsadmin client stop
```

```
nfsadmin client start
```

7. Close the Administrator: Windows Command Prompt (CMD) window. Open a **standard** Command Prompt Window:
  - a. Click **Start**, then click **Command Prompt**.



### Important

NFS file systems mounted as Administrator are not available to standard users.

8. In the standard Windows Command Line (CMD) window, mount the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
mount 10.x.x.x:/fs-export-path X:
```



### Important

The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then "/" represents the full extent of the mount target. In that case, you must use a "!" when mounting the file system. For example: `mount 10.0.0.0:!/ X:`

9. Write a file to the file system by typing the following. Replace `x` with the drive letter you used in step 8 and `helloworld` with your file name.

```
X:
```

```
echo > helloworld.txt
```

10. Verify that you can view the file by typing the following.

```
dir
```

See [Troubleshooting Windows NFS Client Connections](#) for more information on common issues you may encounter.

## To mount a file system from Windows Server 2012 R2 or 2016 Command Prompt

If you are using Oracle-provided Windows images, the NFS client is already installed, and the correct user identities are mapped. Skip to step 4.

1. Open **Windows PowerShell** and run as **Administrator**:
  - a. Go to **Start** and click the **Windows PowerShell** icon.
  - b. In Windows PowerShell, type the following to run as Administrator:

```
Start-Process powershell -Verb runAs
```

- c. In the **User Account Control** window, click **Yes**. A new Administrator: PowerShell window opens. You can close the standard PowerShell window to avoid confusing them.

2. In Administrator: PowerShell, get the NFS client by typing the following:

```
Install-WindowsFeature -Name NFS-Client
Set-ItemProperty HKLM:\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default -Name AnonymousUid -Value 0
Set-ItemProperty HKLM:\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default -Name AnonymousGid -Value 0
Stop-Service -Name NfsClnt
Restart-Service -Name NfsRdr
Start-Service -Name NfsClnt
```



### Important

If you've set export options for your file system to require clients to connect from a privileged source port (1-1023), then you must set the **UseReservedPorts** registry key to **1**.

For example:

```
Set-ItemProperty
HKLM:\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default -Name UseReservedPorts -Value 1
```

For more information, see [Working with NFS Export Options](#).

3. Close the Administrator: PowerShell window. Open a **standard** Command Prompt Window:
  - a. Click **Start**, then click **Command Prompt**.



### Important

NFS file systems mounted as Administrator are not available to standard users.

4. In the standard Windows Command Line (CMD) window, mount the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
mount 10.x.x.x:/fs-export-path X:
```



### Important

The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then "/" represents the full extent of the mount target. In that case, you must use a "!" when mounting the file system. For example: `mount 10.0.0.0:!/ X:`

5. Write a file to the file system by typing the following. Replace `x` with the drive letter you used in step 10 and `helloworld` with your file name.

```
X:
```

```
echo > helloworld.txt
```

6. Verify that you can view the file by typing the following.

```
dir
```

See [Troubleshooting Windows NFS Client Connections](#) for more information on common issues you may encounter.

### Using Windows File Explorer

#### To mount a file system from Windows Server 2008 R2 File Explorer

If you are using Oracle-provided Windows images, the NFS client is already installed, and the correct user identities are mapped. Skip to step 7.

1. Install Services for NFS components.
  - a. Click **Start** go to **Administrative Tools**, then click **Server Manager**.
  - b. In the left pane, click **Roles**.
  - c. Under **Roles Summary** in the right pane, click **Add Roles**. The **Add Roles Wizard** appears. Click **Next**.
  - d. Select the **File Services** check box from the list and click **Next**. Review the information and click **Next**.
  - e. Select the **Services for Network File System** check box from the list, and click **Next**.
  - f. Confirm your selections, then click **Install**. Click **Close** when the installation is complete.
2. Open the registry editor (regedit) to map the AnonymousGid and AnonymousUid to the root user.

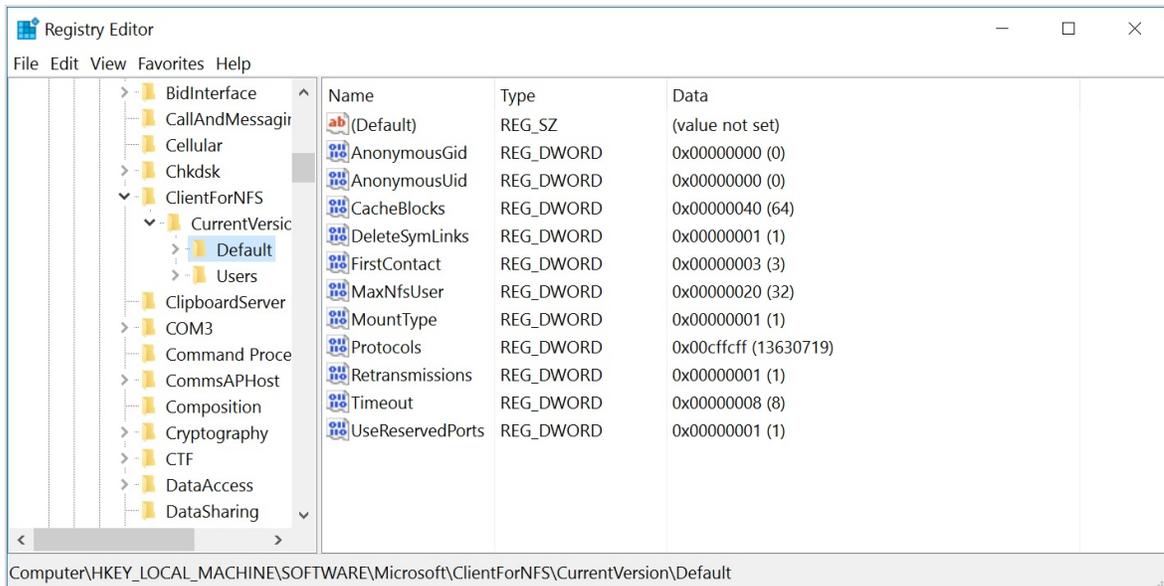


### Warning

User identity mapping requires changes to your system registry.

- a. Click **Windows Search**.
  - b. Enter `regedit` in the **Search** field and press **Enter**.
  - c. Click **Yes** to allow changes to your device.
  - d. Click `HKEY_LOCAL_MACHINE`. Then, browse to:  
`Software\Microsoft\ClientForNFS\CurrentVersion\Default`.
3. Add a new DWORD32 registry entry for `AnonymousGid`:

- a. Click **Edit**, and select **New DWORD (32 bit) Value**.
  - b. In the **Name** field, enter `AnonymousGid`. Leave the value at 0.
4. Repeat step 3 to add a second DWORD32 registry entry named `AnonymousUid` with a value of 0.



**Important**

If you've set export options for your file system to require clients to connect from a privileged source port (1-1023), then you must set the

**UseReservedPorts** registry key to **1**.

For more information, see [Working with NFS Export Options](#).

5. Open Windows Command Line (CMD) and run as **Administrator**:
  - a. Click **Start** .
  - b. Press **Ctrl+Shift** and click **Command Prompt**.
  - c. Click **Yes**.
6. In the Administrator: Windows Command Prompt (CMD) window, restart the NFS Client by typing the following:

```
nfsadmin client stop
```

```
nfsadmin client start
```

7. Open **File Explorer** and select **Computer**. Click the **Map network drive** tab.
8. Select the **Drive** letter that you want to assign to the file system.
9. In the **Folder** field, enter the following. Replace `10.x.x.x` with the local subnet IP address assigned to your mount target, and `fs-export-path` with the export path you specified when associating the file system with the mount target.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
\\10.x.x.x\fs-export-path
```



### Important

The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then "\" represents the full extent of the mount target. In that case, you must use a "!" when entering the file system folder path. For example: \\10.0.0.0\!

10. Click the **Finish** button when complete.

See [Troubleshooting Windows NFS Client Connections](#) for more information on common issues you may encounter.

### To mount a file system from Windows Server 2012 R2 or 2016 File Explorer

If you are using Oracle-provided Windows images, the NFS client is already installed, and the correct user identities are mapped. Skip to step 9.

1. Open **Windows PowerShell** and run as **Administrator**:
  - a. Go to **Start** and click the **Windows PowerShell** icon.
  - b. In Windows PowerShell, type the following to run as Administrator:

```
Start-Process powershell -Verb runAs
```
  - c. In the **User Account Control** window, click **Yes**. A new Administrator: PowerShell window opens. You can close the standard PowerShell window to avoid confusing them.
2. In Administrator: PowerShell, get the NFS client by typing the following:

```
Install-WindowsFeature -Name NFS-Client
```
3. If necessary, restart your system.

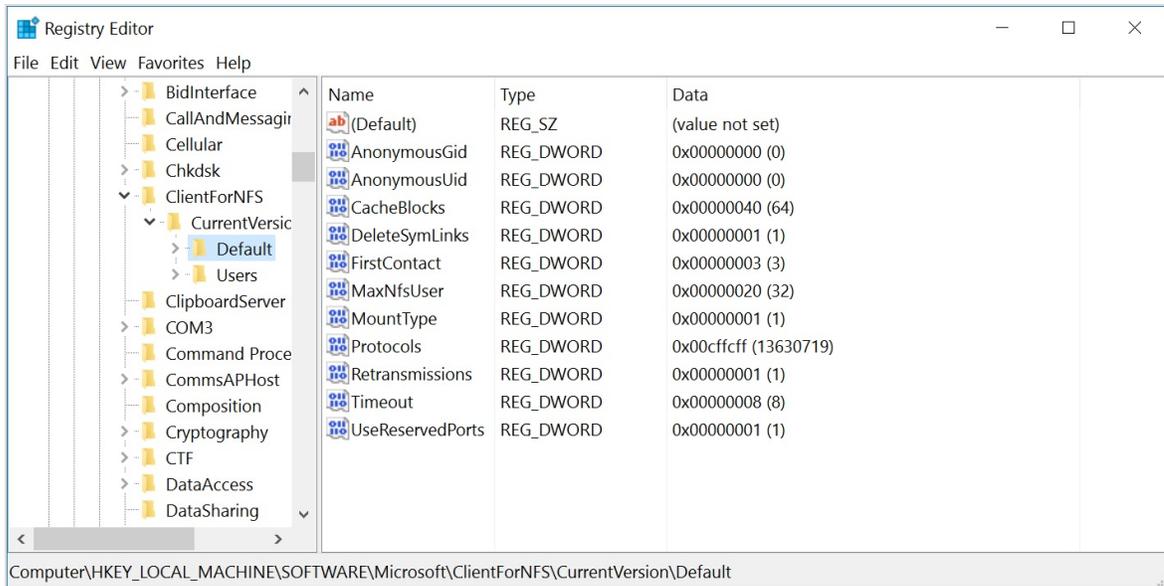
4. Open the registry editor (regedit) to map the AnonymousGid and AnonymousUid to the root user.



### Warning

User identity mapping requires changes to your system registry.

- a. Click **Windows Search**.
  - b. Enter `regedit` in the **Search** field and press **Enter**.
  - c. Click **Yes** to allow changes to your device.
  - d. Click `HKEY_LOCAL_MACHINE`. Then, browse to:  
`Software\Microsoft\ClientForNFS\CurrentVersion\Default`.
5. Add a new DWORD32 registry entry for `AnonymousGid`:
    - a. Click **Edit**, and select **New DWORD (32 bit) Value**.
    - b. In the **Name** field, enter `AnonymousGid`. Leave the value at 0.
  6. Repeat step 5 to add a second DWORD32 registry entry named `AnonymousUid` with a value of 0.



 **Important**

If you've set export options for your file system to require clients to connect from a privileged source port (1-1023), then you must set the **UseReserverdPorts** registry key to **1**. For more information, see [Working with NFS Export Options](#).

7. Open Windows Command Line (CMD) and run as Administrator:
  - a. Go to **Start** and scroll down to **Apps**.
  - b. In the **Windows System** section, press **Ctrl+Shift** and click **Command Prompt**.
8. In the Windows Command Line (CMD) window, restart the NFS Client by typing the

following:

```
nfsadmin client stop
```

```
nfsadmin client start
```

9. Open **File Explorer** and select **This PC**. In the **Computer** tab, select **Map network drive**.
10. Select the **Drive** letter that you want to assign to the file system.
11. In the **Folder** field, enter the following. Replace `10.x.x.x` with the local subnet IP address assigned to your mount target, and `fs-export-path` with the export path you specified when associating the file system with the mount target.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
\\10.x.x.x\fs-export-path
```



### Important

The export path is the path to the file system (relative to the mount target IP address or hostname). If you did not specify a path when you associated the file system and mount target, then "\" represents the full extent of the mount target. In that case, you must use a "!" when entering the file system folder path. For example: `\\10.0.0.0\!`

12. Click the **Finish** button when complete.

See [Troubleshooting Windows NFS Client Connections](#) for more information on common issues you may encounter.

## Managing File Systems

In the File Storage service, file systems are associated with a single compartment. When you select a compartment, the Console displays all file systems in the compartment. You can also see [exports](#) and [snapshots](#) associated with each file system. If there are no file systems in the compartment, see [Creating File Systems](#) for instructions about creating one.

The compartment has policies that indicate what actions a user can take to manage file system. UNIX permissions control what actions a user can take on the files stored in the file system. See [About Security](#) for more information.

Actions you can take to manage a file system include:

- Viewing file system details
- Editing file system settings
- Viewing associated file system resources
- Creating an export for the file system
- Deleting a file system

You can perform most administrative tasks for your file systems using the Console, Command Line Interface (CLI), or API. You can use the Console to list mount targets exporting a specific file system. Use the API or CLI if you want to list all mount targets in a compartment.

To access a file system, it must have at least one export in one mount target. Next, mount the file system from an instance, and then you can create directories and read and write files. For more information about creating an export for a file system, see [To create an export for a file system](#) in this topic. For more information about accessing your file system, see [Mounting File Systems](#).

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to manage file systems.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Moving File Systems to a Different Compartment

You can move file systems from one compartment to another. When you move a file system to a new compartment, its associated snapshots move with it. After you move the file system to the new compartment, inherent policies apply immediately and affect access to the file system and snapshots through the Console. Moving these resources doesn't affect access to file systems and snapshots from mounted instances. For more information, see [Managing Compartments](#).

### Details About Your File System

The file system details page provides the following information about your file system:

### FILE SYSTEM OCID

Every Oracle Cloud Infrastructure resource has an Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). You need your file system's OCID to use the Command Line Interface (CLI) or the API. You also need the OCID when contacting support.

### AVAILABILITY DOMAIN

When you create a file system, you specify the availability domain that it resides in. An availability domain is one or more data centers located within a region. You need your file system's availability domain to use the Command Line Interface (CLI) or the API. For more information, see [Regions and Availability Domains](#).

### CREATED

The date and time that the file system was created.

### COMPARTMENT

When you create a file system, you specify the compartment that it resides in. A compartment is a collection of related resources (such as cloud networks, compute instances, or file systems) that are only accessible to those groups that have been given permission by an administrator in your organization. You need your file system's compartment to use the Command Line Interface (CLI) or the API. For more information, see [Managing Compartments](#).

### UTILIZATION

Metered size of the file system that gets updated hourly.



#### Important

**There can be a delay of up to 1 hour** when reporting file system usage.

You can use `df` or `du` commands from your mounted instance command line application to view usage information about your file system.

## CHAPTER 15 File Storage

---

- `df` provides the amount of storage metered for your file system. Results are returned quickly, but can be up to 1 hour out of date.
- `du` provides the storage used by a directory hierarchy. The `du` command walks the directory tree, and if your hierarchy is large, it can take a long time to run and return results.

The results provided by `df` and `du` can differ for several reasons:

- `df` and `du` report snapshot utilization differently. Snapshots are copy-on-write, so each snapshot shares the blocks used by the unchanged data. The `df` command retrieves information provided by the File Storage service using the NFS FSSTAT call, which accounts correctly for the shared utilization. The `du` command can't detect this sharing, and reports the shared utilization once for each snapshot.
- `df` counts each file only once. `du` may count files with hard links more than once.
- File Storage needs 512 bytes for each directory entry and 8192 bytes for each symlink for metadata cost. `df` reports this cost usage, even on empty files. `du` reports empty files as using zero bytes.

### RESOURCES

Resources such as [exports](#) and [snapshots](#) that are associated with the file system are listed here. Click the resource type link to see a list of each individual resource. Each export in the list shows the file system's export path and mount target. [You need the export path to mount a file system.](#)



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Using the Console

#### To view file system details

The File Storage service displays a list of file systems in each compartment.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. To view information about a file system, find the file system, click the Actions icon (three dots), and then click **View File System Details**.

The Console displays metadata for the file system, exports and snapshots for the file system, and status for the file system and its exports in associated mount targets.

#### To change the file system name

You can change the display name of the file system.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. To view information about a file system, find the file system, click the Actions icon (three dots), and then click **View File System Details**.
4. Click **Rename**.
5. Enter the new file system name, and click **Rename**.

#### To create an export for a file system

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one

or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the left-hand navigation, in the **List Scope** section, under **Compartment**, select a compartment.
3. Click the name of the file system you want to create an export for, and click **Create Export**.



### Note

File systems are encrypted by default. You cannot turn off encryption.

4. You can choose to accept the system defaults, or change them by clicking **Edit Details**.
5. If you want to *accept the defaults* for the mount target, click **Create**. *The file system is created with the information displayed.* If you want to choose another mount target or change the default information, click the **Edit Details** link.
6. In the **Mount Target Information** section, specify details for the mount target that is associated with the file system:
  - **Select an Existing Mount Target:** Choose this option if you want to associate the file system with a mount target you already created. Choose the **Mount Target** from the list. Click the **click here** link in the dialog box if you want to enable compartment selection for the mount target.



### Tip

If there aren't any mount targets in the current combination of availability domain and compartment, this option is disabled. You can:

- Choose a different compartment.
- Create a new mount target.

- **Create a New Mount Target:** Choose this option if you want to create a new mount target associated with this file system. By default, the mount target is created in your current compartment and you can use network resources in that compartment. Click the **click here** link in the dialog box if you want to enable compartment selection for the mount target, its VCN, or subnet resources.



### Important

The mount target is always in the same availability domain as the file system. While it is possible to access mount targets from any AD in a region, for optimal performance, your mount target and file system should be in the same availability domain as the Compute instances that access them. For more information, see [Regions and Availability Domains](#).

- **Create in Compartment:** Specify the compartment you want to create the mount target in.
- **New Mount Target Name:** Optionally, replace the default with a friendly name

for the mount target. It doesn't have to be unique; an Oracle Cloud Identifier (OCID) uniquely identifies the mount target. Avoid entering confidential information.



### Note

The mount target name is different than the DNS hostname, which is specified in step 7.

- **Virtual Cloud Network Compartment:** The compartment containing the cloud network (VCN) in which to create the mount target.
- **Virtual Cloud Network:** Select the cloud network (VCN) where you want to create the new mount target.
- **Subnet Compartment:** Specify the compartment containing a subnet within the VCN to attach the mount target to.
- **Subnet:** Select a subnet to attach the mount target to. Subnets can be either AD-specific or regional (regional ones have "*regional*" after the name). For more information, see [About Regional Subnets](#).



### Warning

Each mount target requires three internal IP addresses in the subnet to function. Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Optionally, click **Show Advanced Options** to configure the mount target's advanced options.
    - **IP Address:** You can specify an unused IP address in the subnet you selected for the mount target.
    - **Hostname:** You can specify a hostname you want to assign to the mount target.



### Note

The File Storage service constructs a fully qualified domain name (FQDN) by combining the hostname with the FQDN of the subnet the mount target is located in.

For example,

```
myhostname.subnet123
```

```
.dnslabel.oraclevcn.com.
```

Once created, the hostname may be changed in the mount target's Details page. See [Managing Mount Targets](#) for more information.

### 8. Click **Create**.

Next, mount the file system from an instance so that you can read and write directories and files in your file system. See [Mounting File Systems](#) for instructions about obtaining mount commands for your operating system type and mounting your file system.

### To set the file system reported size

The File Storage service reports file system capacity as 8589934592 gibibytes (GiB) and 8589934592 gibiinodes (GiI) by default. Sometimes, application installers perform a space requirement check prior to running an installation process but cannot correctly interpret the reported size or reported inodes of the file system. When this occurs, you can define the file system size reported to the operating system by setting the **Reported Size** or **Reported Inodes** value in the file system's mount target. Typically, setting the size to 1024 GiB and the inodes to 1024 GiI permits successful installation.



### Important

Changing the **Reported Size** or **Reported Inodes** for a mount target affects all file systems exported by the mount target. **Changing these values does not limit the amount of data you can store.**

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. Find the mount target you're interested in, click the Actions icon (three dots), and then click **View File System Details**.
4. In **Exports**, click on the mount target name.
5. Click the **Reported Size (in GiB) Edit** or the **Reported Inodes (in Gil)** icon.
6. Enter the maximum free space in gibibytes or the maximum inodes in gibinodes you want the File Storage service to report.
7. Click the **Save** icon.



### Important

**There can be a delay of up to 1 hour** when reporting file system usage, either in the console or by using the `df` command.

### To manage tags for a file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.

2. In the **List Scope** section, select a compartment.
3. Find the file system you're interested in, click the Actions icon (three dots), and then click **View File System Details**.
4. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To move a file system to a different compartment

1. Open the Console,
2. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
3. In the **List Scope** section, select a compartment.
4. Find the file system in the list, click the the Actions icon (three dots), and then click **Change Compartment**.
5. Choose the destination compartment from the list.
6. Click **Change Compartment**.

The file system is moved immediately. Moving a file system doesn't affect mounted instances.

### To assign a key to a file system

File systems use Oracle-managed keys by default, which leaves all encryption-related matters to Oracle. Optionally, you can encrypt the data in this file system using your own Key Management encryption key. For more information, see [Overview of Key Management](#).

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the file system that you want to encrypt with a Key Management master encryption key.

3. From the list of file systems, click the file system name.
4. Next to **Encryption Key**, click **Edit**.
5. In **Encryption Type**, select **Encrypt using customer-managed keys**.
6. Choose the vault compartment, vault, key compartment, and key.
7. When you are finished, click **Save Changes**.

### To specify Oracle-managed keys for a file system

File systems use Oracle-managed keys by default, which leaves all encryption-related matters to Oracle. However, if you assign a Key Management key to a file system, you can later return the file system to using Oracle-managed keys for encryption. For more information, see [Overview of Key Management](#).

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the file system that you want to encrypt with a Key Management master encryption key.
3. From the list of file systems, click the file system name.
4. Next to **Encryption Key**, click **Edit**.
5. In **Encryption Type**, select **Encrypt using Oracle-managed keys**.
6. When you are finished, click **Save Changes**.

### To delete a file system

You can permanently delete a file system.



### Warning

You cannot undo this operation. Any data in a file system is permanently deleted with the file system. Snapshots of the file system are permanently deleted with the file system. You cannot recover a deleted file system or its snapshots.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. Find the file system you want to delete.
4. Click the Actions icon (three dots), and then click **View File System Details**.
5. Delete all of the file system's exports:
  - In **Exports**, select the check box for all exports listed, and then click **Delete**.
6. When all of the exports are deleted, click **Delete** to delete the file system.

The file system is deleted immediately, along with all of its snapshots.

## Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To list file systems

Open a command prompt and run `oci fs file-system list` to list all the file systems in a specified availability domain and compartment.

For example:

## CHAPTER 15 File Storage

---

```
oci fs file-system list --availability-domain <target_availability_domain> --compartment-id <target_compartment_id>
```

### To get a specific file system

Open a command prompt and run `oci fs file-system get` to retrieve information about a specific file system.

For example:

```
oci fs file-system get --file-system-id <file_system_OCID>
```

### To update a file system

Open a command prompt and run `oci fs file-system update` to update a specific file system's information.

For example:

```
oci fs file-system update --file-system-id <file_system_OCID> --display-name "<New File System Name>"
```



#### Warning

Avoid entering confidential information in the file system `display-name`.

### To create an export for a file system

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#).

## CHAPTER 15 File Storage

---

Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

Open a command prompt and run `oci fs export create` to create an export for a specified file system within a specified export set.

For example:

```
oci fs export create --export-set-id <export_set_OCID> --file-system-id <file_system_OCID> --path
"/<pathname>"
```



### Important

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For multiple file systems associated with a single mount target, the path sequence for the first file system cannot contain the complete path element sequence of the second file system path sequence. Paths cannot end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes. For example:

Acceptable:

`/example` and `/path`

`/example` and `/example2`

Not Acceptable:

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../../example1`



### Warning

If one file system associated to a mount target has '/' specified as an export path, you can't associate another file system with that mount target.



### Note

Export paths cannot be edited after the export is created. If you want to use a different export path, you must create a new export with the desired path. Optionally, you can then delete the export with the old path.

For more information, see [Paths in File Systems](#).

### To set the file system reported free space

Some existing application installers perform a capacity check before running an installation process. Sometimes an installation fails because of too much available capacity. The File Storage service currently reports 8 exabytes of available capacity by default for each file system.

Customers can define how much free capacity is reported as available to the operating system.

Open a command prompt and type in the following command:

```
oci fs export-set update --export-set-id <export_set_OCID> --max-fs-stat-bytes <number_of_bytes>
```



### Important

**The maximum free space setting affects each export in the export set. Setting the maximum free space does not limit the amount of data you can store.**

### To move a file system to a different compartment

```
oci fs file-system change-file-system-compartment --file-system-id <file_system_OCID> --compartment-id <destination_compartment_OCID>
```

### To update the key for a file system

Open a command prompt and run `oci fs file-system update` to update the file system with a new key.

```
oci fs file-system update --file-system-id <file_system_OCID> --kms-key-id <target_key_id>
```

For example:

```
oci fs file-system update --file-system-id ocid1.filesystem.oc1.phx.<unique_id> --kms-key-id ocid1.key.oc1.phx.<unique_id>
```

### To specify Oracle-managed keys for a file system

File systems use Oracle-managed keys by default, which leaves all encryption-related matters to Oracle. However, if you assign a Key Management key to a file system, you can later return the file system to using Oracle-managed keys for encryption.

Open a command prompt and run `oci fs file-system update`. Leave the `--kms-key-id` value **unspecified**.

```
oci fs file-system update --file-system-id <file_system_OCID> --kms-key-id ""
```

For example:

```
oci fs file-system update --file-system-id ocid1.filesystem.oc1.phx.<unique_id> --kms-key-id ""
```

### To delete a file system

You can delete a file system if no non-deleted export resources reference it. Deleting a file system also deletes all its snapshots.

Open a command prompt and run `oci fs file-system delete` to delete a file system.

For example:

```
oci fs file-system delete --file-system-id <file_system_OCID>
```



### Warning

You cannot undo this operation. Any data in a file system is permanently deleted with the file system. Snapshots of the file system are permanently deleted with the file system. You cannot recover a deleted file system or its snapshots.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage file systems:

- [ListFileSystems](#)
- [GetFileSystem](#)
- [UpdateFileSystem](#)
- [ChangeFileSystemCompartment](#)
- [DeleteFileSystem](#)

## Managing Mount Targets

This topic describes the basics of managing mount targets.

### Overview

Actions you can take to manage a mount target include:

- Viewing mount target details
- Obtaining mount command samples
- Creating a new export and file system
- Editing exports and export options
- Change the reported size of exported file systems
- Deleting a mount target

You can perform most administrative tasks for your mount targets using the Console, Command Line Interface (CLI), or API. You can use the Console to list mount targets exporting a specific file system. Use the API or CLI if you want to list all mount targets in a compartment.

### Mount Target

A mount target is an NFS endpoint that lives in a VCN subnet of your choice and provides network access for file systems. The mount target provides the IP address or DNS name that is used together with a unique export path to mount the file system. A single mount target can export many file systems. Typically, you [create your first mount target and export when you create your first file system](#). The mount target maintains an export set which contains all of the exports for its associated file systems.

### Exports

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

### NFS Export Options

NFS export options are a set of parameters within the export that specify the level of access granted to NFS clients when they connect to a mount target. An NFS export options entry within an export defines access for a single IP address or CIDR block range.

For more information, see [Working with NFS Export Options](#).

### Limitations and Considerations

- Each availability domain is limited to two mount targets by default. However, you can export up to 100 file systems through each mount target.

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

- Each mount target requires three internal IP addresses in the subnet to function. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.
- The File Storage service doesn't "reserve" the third IP address required for high availability failover. Use care when designing your subnets and file systems to ensure that sufficient IP addresses remain available for your mount targets.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to manage mount targets. Since mount targets are network endpoints, users must also have "use" permissions for VNICs, private IPs, private DNS zones, and subnets to create or delete a mount target. See the [Policy Reference](#) for more information.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Moving Mount Targets to a Different Compartment

You can move mount targets from one compartment to another. When you move a mount target to a new compartment, its associated export set and exports move with it. After you move the mount target to the new compartment, inherent policies apply immediately and affect access to the mount target, export set, and exports through the Console. Moving these resources doesn't affect access to file systems and snapshots from mounted instances. For more information, see [Managing Compartments](#).

### Details About Your Mount Target

The mount target details page provides the following information about your mount target:

#### **MOUNT TARGET OCID**

Every Oracle Cloud Infrastructure resource has an Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). You need your mount target's OCID to use the Command Line Interface (CLI) or the API. You also need the OCID when contacting support.

#### **CREATED**

The date and time that the mount target was created.

#### **AVAILABILITY DOMAIN**

When you create a mount target, you specify the availability domain that it resides in. An availability domain is one or more data centers located within a region. You need your mount target's availability domain to use the Command Line Interface (CLI) or the API. For more information, see [Regions and Availability Domains](#).

### **COMPARTMENT**

When you create a mount target, you specify the compartment that it resides in. A compartment is a collection of related resources (such as cloud networks, compute instances, or file systems) that are accessible only to those groups that have been given permission by an administrator in your organization. You need your mount target's compartment to use the Command Line Interface (CLI) or the API. For more information, see [Managing Compartments](#).

### **REPORTED SIZE (GiB)**

The maximum capacity in gibibytes reported by the file systems exported through this mount target. The File Storage service currently reports 8589934592 gibibytes (GiB) of available capacity by default. If you are installing an application that requires a specific reported size, you can change the reported size. Typically, setting the size to 1024 GiB is sufficient for most applications. This value is updated hourly. See [To set the file system reported size](#) for more information.

### **REPORTED INODES (GiI)**

The maximum capacity in gibinodes reported by the file systems exported through this mount target. The File Storage service currently reports gibinodes (GiI) of available inodes by default. If you are installing an application that requires specific reported inodes, you can change the reported inodes. Typically, setting the inodes to 1024 GiI is sufficient for most applications. This value is updated hourly. See [To set the file system reported size](#) for more information.

### **VIRTUAL CLOUD NETWORK**

The VCN that contains the subnet where the mount target VNIC resides.

### **SUBNET**

The subnet within the VCN where the mount target VNIC resides. Subnets can be either AD-specific or regional (regional ones have "*regional*" after the name). For more information, see [About Regional Subnets](#).

### **IP ADDRESS**

The IP address that was assigned to the mount target when it was created. You need your mount target's IP address to [mount associated file systems](#).

### **HOSTNAME**

The hostname that was assigned to the mount target, if any. For more information about hostnames, see [DNS in Your Virtual Cloud Network](#).

### **FULLY QUALIFIED DOMAIN NAME**

The hostname together with the subnet domain name. For more information, see [DNS in Your Virtual Cloud Network](#). If you specify a hostname, you can use the FDQN to mount the file system.

### **EXPORT SET OCID**

The OCID of the mount target's export set resource. Each mount target has one export set, which contains all of the exports for the mount target. You need your mount target's export set OCID when you perform export-related tasks in the Command Line Interface (CLI) or the API.

### **EXPORTS**

All of the mount target's exports are listed here. The export path and name of each file system is also listed. [You need the export path to mount a file system](#).



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Using the Console

#### To create a mount target



#### Important

While it is possible to access mount targets from any AD in a region, for optimal performance, your mount targets should be in the same availability domain as the Compute instances that access them. For more information, see [Regions and Availability Domains](#).

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.  
The Console displays a list of mount targets that have already been created in the compartment, if any.
3. Click **Create Mount Target**.
4. Enter the required mount target information. Click the **click here** link in the dialog box if you want to enable compartment selection for the mount target, its VCN, or subnet resources:
  - **New Mount Target Name:** Optionally, replace the default with a friendly name for the mount target. It doesn't have to be unique; an Oracle Cloud Identifier (OCID) uniquely identifies the mount target. Avoid entering confidential information.



### Note

The mount target name is different than the DNS hostname, which is specified in step 5.

- **Virtual Cloud Network Compartment:** The compartment containing the cloud network (VCN) in which to create the mount target.
- **Virtual Cloud Network:** Select the cloud network (VCN) where you want to create the new mount target.
- **Subnet Compartment:** Specify the compartment containing a subnet within the VCN to attach the mount target to.
- **Subnet:** Select a subnet to attach the mount target to. Subnets can be either AD-specific or regional (regional ones have "*regional*" after the name). For more information, see [About Regional Subnets](#).



### Warning

Each mount target requires three internal IP addresses in the subnet to function. Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

5. Optionally, click **Show Advanced Options** to configure the mount target's advanced options.

- **IP Address:** You can specify an unused IP address in the subnet you selected for the mount target.
- **Hostname:** You can specify a hostname you want to assign to the mount target.



### Note

The File Storage service constructs a fully qualified domain name (FQDN) by combining the hostname with the FQDN of the subnet the mount target is located in.

For example,

```
myhostname.subnet123.dnslabel.oraclevcn.com.
```

Once created, the hostname may be changed in the mount target's Details page. See [Managing Mount Targets](#) for more information.

6. Click **Create**.

### To view details of a mount target

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.  
The Console displays a list of mount targets that have already been created in the compartment, if any.
3. Find the mount target you're interested in, click the Actions icon (three dots), and then click **View Mount Target Details**.

### To change the mount target name

You can change the display name of the mount target.



#### Note

Changing the display name doesn't affect mounting file systems exported through the mount target.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.
3. To view information about a file system, find the file system, click the Actions icon (three dots), and then click **View Mount Target Details**.
4. Click **Rename**.
5. Enter the new mount target name, and click **Rename**.

### To create an export and a new file system

Exports control how NFS clients access file systems when they connect to a mount target. File systems must have at least one export in at least one mount target in order for [instances to mount the file system](#). The following steps create an export and a new file system. If you want to create an export for an *existing* file system, see [To create an export for a file system](#).

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the left-hand navigation, in the **List Scope** section, under **Compartment**, select a compartment.
3. Click the name of the mount target you want to create an export for, and click **Create Export**.



### Note

File systems are encrypted by default. You cannot turn off encryption.

4. You can choose to accept the system defaults, or change them by clicking **Edit Details**.
5. Click **Create**.

Next, mount the file system from an instance so that you can read and write directories and files in your file system. See [Mounting File Systems](#) for instructions about obtaining mount commands for your operating system type and mounting your file system.

### To set the file system reported size

The File Storage service reports file system capacity as 8589934592 gibibytes (GiB) and 8589934592 gibiinodes (GiI) by default. Sometimes, application installers perform a space requirement check prior to running an installation process but cannot correctly interpret the reported size or reported inodes of the file system. When this occurs, you can define the file system size reported to the operating system by setting the **Reported Size** or **Reported Inodes** value in the file system's mount target. Typically, setting the size to 1024 GiB and the inodes to 1024 GiI permits successful installation.



### Important

Changing the **Reported Size** or **Reported Inodes** for a mount target affects all file systems exported by the mount target. **Changing these values does not limit the amount of data you can store.**

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.
3. Find the mount target you're interested in, click the Actions icon (three dots), and then click **View Mount target Details**.
4. Click the **Reported Size (in GiB) Edit** or the **Reported Inodes (in GiI)** icon.
5. Enter the maximum size in gibibytes or the maximum inodes in gibinodes you want the File Storage service to report.
6. Click the **Save** icon.



### Important

**There can be a delay of up to 1 hour** when reporting file system usage, either in the console or by using the `df` command.

### To delete an export



### Note

Deleting an export does not impact the data stored in the associated file system. Deleting an export disconnects any instance that mounts the file system with the deleted export path. Mount targets that have no exports still count toward your service limit.

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.

2. In the **List Scope** section, select a compartment.
3. Find the mount target you're interested in, click the Actions icon (three dots), and then click **View Mount target Details**.
4. In **Exports**, find the export you want to delete.
5. Click the Actions icon (three dots), and then click **Delete**.

### To manage tags for a mount target

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.
3. Find the mount target you're interested in, click the Actions icon (three dots), and then click **View Mount Target Details**.
4. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To move a mount target to a different compartment

1. Open the Console,
2. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
3. In the **List Scope** section, select a compartment.
4. Find the mount target in the list, click the the Actions icon (three dots), and then click **Change Compartment**.
5. Choose the destination compartment from the list.
6. Click **Change Compartment**.

### To delete a mount target

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
2. In the **List Scope** section, select a compartment.
3. Find the mount target you want to delete.
4. Click the Actions icon (three dots), and then click **Delete**.



#### **Warning**

Deleting the mount target also deletes all of its exports of associated file systems. File systems are no longer available through the deleted mount target.

**Deleting a mount target has no effect on file system data or file system snapshots.**

### Using the Command Line

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To create a mount target

You can create a mount target for file systems in a specified compartment and subnet. A file system can only be associated with a mount target in the same availability domain.



### Warning

Each mount target requires three internal IP addresses in the subnet to function. Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

Open a command prompt and run `oci fs mount-target create` to create a mount target.

For example:

```
oci fs mount-target create --availability-domain <target_availability_domain> --compartment-id <target_compartment_id> --subnet-id <subnet_OCID> --display-name "<My Mount Target>"
```



### Warning

Avoid entering confidential information in the mount target `display-name`.

## To update a mount target

Open a command prompt and run `oci fs mount-target update` to update a specific mount target's information.

For example:

```
oci fs mount-target update --mount-target-id <mount_target_OCID> --display-name "<New Mount Target Name>"
```



### Warning

Avoid entering confidential information in the mount target `display-name`.

### To delete a mount target

Open a command prompt and run `oci fs mount-target delete` to delete a mount target. Deleting a mount target also deletes the mount target's VNICs.

For example:

```
oci fs mount-target delete --mount-target-id <mount_target_OCID>
```



### Warning

Deleting a mount target can cause any clients that have mounted an associated file system to hang. Be sure to have all clients unmount the file systems before deleting the mount target.

### To list mount targets

You cannot use the Console to list mount targets. Use the command line interface or the API from a host machine running a UNIX-style operating system.

Open a command prompt and run `oci fs mount-target list` to list all mount targets in a specified availability domain and compartment.

For example:

```
oci fs mount-target list --availability-domain <target_availability_domain> --compartment-id <target_compartment_OCID>
```

### To get a specific mount target

Open a command prompt and run `oci fs mount-target get` to retrieve information about a specific mount target.

For example:

```
oci fs mount-target get --mount-target-id <mount_target_OCID>
```

### To create an export

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

Open a command prompt and run `oci fs export create` to create an export for a specified file system within a specified export set.

For example:

```
oci fs export create --export-set-id <export_set_OCID> --file-system-id <file_system_OCID> --path
"</pathname>"
```



### Important

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For multiple file systems associated with a single mount target, the path sequence for the first file system cannot contain the complete path element sequence of the second file system path sequence. Paths cannot end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes. For example:

Acceptable:

`/example` and `/path`

`/example` and `/example2`

Not Acceptable:

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../../example1`



### Warning

If one file system associated to a mount target has '/' specified as an export path, you can't associate another file system with that mount target.



### Note

Export paths cannot be edited after the export is created. If you want to use a different export path, you must create a new export with the desired path. Optionally, you can then delete the export with the old path.

For more information, see [Paths in File Systems](#).

### To list exports

Open a command prompt and run `oci fs export list` to list all exports in a specified compartment.

For example:

```
oci fs export list --compartment-id <target_compartment_id>
```

### To get a specific export

Open a command prompt and run `oci fs export get` to retrieve information about a specific export.

For example:

```
oci fs export get --export-id <export_OCID>
```

### To delete an export

Open a command prompt and run `oci fs export delete` to delete an export.

For example:

```
oci fs export delete --export-id <export_OCID>
```



### Warning

When you delete an export, any file system referenced by the export is no longer accessible through the associated mount target.

### To list export sets

Open a command prompt and run `oci fs export-set list` to list all export sets in a specified availability domain and compartment.

For example:

```
oci fs export-set list --availability-domain <target_availability_domain> --compartment-id <target_compartment_OCID>
```

### To get a specific export set

Open a command prompt and run `oci fs export-set get` to retrieve information about a specific export set.

For example:

```
oci fs export-set get --export-set-id <export_set_OCID>
```

### To update an export set

Open a command prompt and run `oci fs export-set update` to update a specific export set's information.

For example:

```
oci fs export-set update --export-set-id <export_set_OCID> --display-name "<New Export Set Name>"
```

### To set the file system reported size

The File Storage service reports file system capacity as 8589934592 gibibytes (GiB) and 8589934592 gibiinodes (GiI) by default. Sometimes, application installers perform a space requirement check prior to running an installation process but cannot correctly interpret the reported size or reported inodes of the file system. When this occurs, you can define the file system size reported to the operating system by setting the **Reported Size** or **Reported Inodes** value in the export set of the file system's mount target. Typically, setting the size to 1024 GiB and the inodes to 1024 GiI permits successful installation.



#### Important

Changing the **Reported Size** or **Reported Inodes** for a mount target affects all file systems exported by the mount target. **Changing these values does not limit the amount of data you can store.**



#### Important

**There can be a delay of up to 1 hour** when reporting file system usage, either in the console or by using the `df` command.

Open a command prompt and type in the following command:

```
oci fs export-set update --export-set-id <export_set_OCID> --max-fs-stat-bytes <number_of_bytes>
```

### To move a mount target to a different compartment

```
oci fs mount-target change-mount-target-compartment --mount-target-id <mount_target_OCID> --compartment-id <destination_compartment_OCID>
```

### Using the API

- [CreateMountTarget](#)
- [UpdateMountTarget](#)
- [DeleteMountTarget](#)
- [GetMountTarget](#)
- [ListMountTargets](#)
- [ChangeMountTargetCompartment](#)
- [CreateExport](#)
- [DeleteExport](#)
- [GetExport](#)
- [ListExports](#)
- [UpdateExportSet](#)
- [GetExportSet](#)
- [ListExportSets](#)

### Working with NFS Export Options

This topic describes the basic features of NFS export options, and how to control client access to your file system.

#### **Overview**

NFS export options enable you to create more granular access control than is possible using just security list rules to limit VCN access. You can use NFS export options to specify access levels for IP addresses or CIDR blocks connecting to file systems through exports in a mount target. Access can be restricted so that each client's file system is inaccessible and invisible to the other, providing better security controls in multi-tenant environments.

Using NFS export option access controls, you can limit clients' ability to connect to the file system and view or write data. For example, if you want to allow clients to consume but not

update resources in your file system, you can set access to Read Only. You can also reduce client root access to your file systems and map specified User IDs (UIDs) and Group IDs (GIDs) to an anonymous UID/GID of your choice. For more information about how NFS export options work with other security layers, see [About Security](#).



### Tip

Watch a [video](#) about working with NFS export options in File Storage.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to manage NFS export options.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Exports

Exports control how NFS clients access file systems when they connect to a mount target. File systems are exported (made available) through mount targets. Each mount target maintains an export set which contains one or many exports. A file system may be exported through one or more mount targets. A file system must have at least one export in one mount target in order for instances to mount the file system. The information used by an export includes the file system OCID, mount target OCID, export set OCID, [export path](#), and client [export options](#). Typically, an export is created in a mount target when the file system is created. Thereafter, you can create additional exports for a file system in any mount target that resides in the same availability domain as the file system.

See [To create an export for a file system](#) for more information.

### **NFS Export Options**

NFS export options are a set of parameters within the export that specify the level of access granted to NFS clients when they connect to a mount target. An NFS export options entry within an export defines access for a single IP address or CIDR block range.

Each separate client IP address or CIDR block you want to define access for needs a separate export options entry in the export. For example, if you want to set options for NFS client IP addresses 10.0.0.6, 10.0.0.8, and 10.0.0.10, you need to create three separate entries, one for each IP address.

File Storage service considers the listed order of each export options entry for the export. During an NFS request by a client, File Storage service applies the first set of options that matches the client Source IP address. Only the first set is applied; the rest are ignored.

For example, consider the following two export options entries specifying access for an export:

Entry 1: Source: 10.0.0.0/16, Access: Read Only

Entry 2: Source: 10.0.0.8, Access: Read/Write

In this case, clients who connect to the export from IP address 10.0.0.8 have Read Only access. The request Source IP address is contained in the CIDR block specified in the first entry, and File Storage Service applies the options in the first match.



### Important

File systems can be associated with one or more exports, contained within one or more mount targets. If the client **source** IP address does not match any entry on the list for a single export, then that export is not visible to the client. However, the file system could be accessed through other exports on the same or other mount targets. **To completely deny client access to a file system, be sure that the client source IP address or CIDR block is not included in any export for any mount target associated with the file system.**

The following options can be set to control export access:

- **Source:** The IP address or CIDR block of a connecting NFS client.
- **Require Privileged Source Port (true/false):** This setting determines whether the NFS clients specified in **source** are required to connect from a privileged source port. Privileged ports are any port including 1-1023. On Unix-like systems, only the root user can open privileged ports. Setting this value to **true** disallows requests from unprivileged ports. The default for this setting is different depending on how the export is created. Creating an export without an explicit `ClientOption` array sets the `requirePrivilegedSourcePort` attribute of the client option to **false**. When you create a `ClientOption` array explicitly, `requirePrivilegedSourcePort` defaults to **true**. For example, creating an export in the Console using the default selections sets `requirePrivilegedSourcePort` to **false**. Creating an export in the API along with a `ClientOption` array sets `requirePrivilegedSourcePort` to **true**.



### Important

When **Require Privileged Source Port** is set to **true**, you also have to follow these additional configuration steps:

1. When mounting the file system from a Unix-like system, include the `resvport` option in your mount command when mounting. For example:

```
sudo mount -o resvport 10.x.x.x:/fs-export-path
/mnt/yourmountpoint
```

For more information, see [Mounting File Systems From Unix-Style Instances](#).

2. When mounting the file system from a Windows system, be sure the **UseReserverdPorts** registry key value is set to **1**.

For more information, see [Mounting File Systems From Windows Instances](#).

- **Access (Read\_Only, Read\_Write):** This setting specifies the **source** NFS client access. If unspecified, defaults to **Read\_Write**.
- **Identity Squash: (All, Root, None):** This setting determines whether the **source** clients accessing the file system have their User ID (UID) and Group ID (GID) remapped to **anonymousUid** and **anonymousGid**. If you choose **All**, all users and groups are remapped. If **Root**, only the root user UID/GID combination 0/0 is remapped. If **None**, no users are remapped. If unspecified, defaults to **None**.
- **anonymousUid:** This setting is used along with the **Identity Squash** option. When remapping users, you can use this setting to change the default anonymousUid of **65534** to any user ID of your choice.

- **anonymousGid:** This setting is used along with the **Identity Squash** option. When remapping groups, you can use this setting to change the default anonymousGid of **65534** to any group ID of your choice.

### Typical Access Control Scenarios

When you create file system and export, the NFS export options for that file system are set to the following defaults, which allow full access for all NFS client source connections. These defaults must be changed if you want to restrict access:

- **Source:** 0.0.0.0/0 (All)
- **Require Privileged Source Port:** False
- **Access:** Read\_Write
- **Identity Squash:** None

#### SCENARIO A: CONTROL HOST BASED ACCESS

Provide a managed hosted environment for two clients. The clients share a mount target, but each has their own file system, and cannot access each other's data. For example:

- Client A, who is assigned to CIDR block 10.0.0.0/24, requires Read/Write access to file system A, but not file system B.
- Client B, who is assigned to CIDR block 10.1.1.0/24, requires Read/Write access to file system B, but not file system A.
- Client C, who is assigned to CIDR block 10.2.2.0/24, has no access of any kind to file system A or file system B.
- Both file systems A and B are associated to a single mount target, MT1. Each file system has an export contained in the export set of MT1.

Since Client A and Client B access the mount target from different CIDR blocks, you can set the client options for both file system exports to allow access to only a single CIDR block. Client C is denied access by not including its IP address or CIDR block in the NFS export options for any export of either file system.

## Console Example

Set the export options for file system A to allow Read/Write access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. ⓘ

Source	Ports	Access	Squash	UID	GID
10.0.0.0/24	Privileged ↕	Read/Write ↕	None ↕	Not used	Not used

+ Another Option

Update

Set the export options for file system B to allow Read/Write access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. (i)

Source	Ports	Access	Squash	UID	GID
10.1.1.0/24	Privileged <span style="font-size: 0.8em;">↕</span>	Read/Write <span style="font-size: 0.8em;">↕</span>	None <span style="font-size: 0.8em;">↕</span>	Not used	Not used
					<span style="font-size: 0.8em;">⋮</span>
					<span style="font-size: 0.8em;">+ Another Option</span>

Update

## CLI Example

Set the export options for file system A to allow Read\_Write access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

```
oci fs export update --export-id <File_system_A_export_ID> --export-options '[{"source":"10.0.0.0/24","require-privileged-source-port":"true","access":"READ_WRITE","identity-squash":"NONE","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

Set the export options for file system B to allow Read\_Write access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

```
oci fs export update --export-id <File_system_B_export_ID> --export-options '[{"source":"10.1.1.0/24","require-privileged-source-port":"true","access":"READ_WRITE","identity-squash":"NONE","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

## API Example

Set the export options for file system A to allow READ\_WRITE access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system.

## CHAPTER 15 File Storage

---

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.0.0.0/24",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "NONE",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```

Set the export options for file system B to allow READ\_WRITE access only to Client B, who is assigned to CIDR block 10.1.1.0/24. Client A and Client C are not included in this CIDR block, and cannot access the file system.

```
PUT /<Current_API_Version>/exports/<File_System_B_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.1.1.0/24",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "NONE",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```

### SCENARIO B: LIMIT THE ABILITY TO WRITE DATA

Provide data to customers for consumption, but don't allow them to update the data.

For example, you'd like to publish a set of resources in file system A for an application to consume, but not change. The application connects from IP address 10.0.0.8.

## Console Example

Set the source IP address 10.0.0.8 to Read Only in the export for file system A:

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. (i)

Source	Ports	Access	Squash	UID	GID
10.0.0.8	Privileged ▾	Read Only ▾	None ▾	Not used	Not used
					⋮
					<a href="#">+ Another Option</a>

[Update](#)

## CLI Example

Set the source IP address 10.0.0.8 to READ\_ONLY in the export for file system A:

```
oci fs export update --export-id <File_System_A_export_OCID> --export-options '
[{"source":"10.0.0.8","require-privileged-source-port":"true","access":"READ_
ONLY","identitysquash":"NONE","anonymousuid":"65534","anonymousgid":"65534"}]'
```

## API Example

Set the source IP address 10.0.0.8 to READ\_ONLY in the export for file system A:

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "10.0.0.8",
 "requirePrivilegedSourcePort": true,
```

## CHAPTER 15 File Storage

```
"access": "READ_ONLY",
"identitySquash": "NONE",
"anonymousUid": 65534,
"anonymousGid": 65534
}
]
```

### SCENARIO C: IMPROVE FILE SYSTEM SECURITY

To increase security, you'd like to limit the root user's privileges when connecting to File System A. Use Identity Squash to remap root users to UID/GID 65534. In Unix-like systems, this UID/GID combination is reserved for *'nobody'*, a user with no system privileges.

### Console Example

Edit Export Options [help](#) [close](#)

Export options control how clients can access your file system. ⓘ

Source	Ports	Access	Squash	UID	GID
0.0.0.0/0	Privileged ↕	Read/Write ↕	Root ↕	65534	65534
					+ Another Option

[Update](#)

### CLI Example

```
oci fs export update --export-id <File_System_A_export_OCID> --export-options '[{"source":"0.0.0.0/0","require-privileged-source-port":"true","access":"READ_WRITE","identitysquash":"ROOT","anonymousuid":"65534","anonymousgid":"65534"}]'
```

### API Example

```
PUT /<Current_API_Version>/exports/<File_System_A_export_OCID>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "exportOptions": [
 {
 "source": "0.0.0.0/0",
 "requirePrivilegedSourcePort": true,
 "access": "READ_WRITE",
 "identitySquash": "ROOT",
 "anonymousUid": 65534,
 "anonymousGid": 65534
 }
]
}
```



#### Tip

If you don't want a file system to be visible to any clients, you can set all of the properties in the `exportOptions` array to empty values. For example,

```
{
 "exportOptions": [
 {
 "source": "",
 "requirePrivilegedSourcePort": "",
 "access": "",
 "identitySquash": ""
 }
]
}
```

### Using the Console

#### To set export options for a file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment. All of the file systems in the selected compartment are displayed.
3. Find the file system you want to set export options for, click the the Actions icon (three dots), and then click **View File System Details**.
4. In the **Exports** list, find the export you want to set export options in, click the the Actions icon (three dots), and then click **View Export Details**. If there is no export listed for the file system, you can create one. See [To create an export for a file system](#) for more information.



#### Tip

To be sure you be sure that you select export, check the following:

- **The export path:** This path uniquely identifies the file system within the mount target. [No two exports in a mount target can have the same export path, even if the exports are for the same file system.](#)
- **The mount target name:** File systems can be exported through more than one mount target. Be sure that you've selected the export for the correct mount target.

5. Click **Edit Export Options**.
6. Make one or more of these changes:

- Change an export option entry in the list.
  - Click **+Another Option** to create a new export option entry.
  - Click the Actions icon (three dots) for an entry and move it up or down in the list.
7. When you're done, click **Update**.

### Using the CLI

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To create an export

Open a command prompt and run `oci fs export create` to create an export for a specified file system within a specified export set. This example creates an export along with its NFS export options.

For example:

```
oci fs export create --export-set-id <export_set_OCID> --file-system-id <file_system_OCID> --path
"</pathname>" --export-options '[{"source":"10.0.0.0/16","requireprivilegedsourceport":"true","access":"READWRITE","identitysquash":"NO
NE","anonymousuid":"0","anonymousgid":"0"}]'
```



### Important

#### *Export Path Names*

The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For any two export resources associated with the same export set, the path sequence for the first export resource can't contain the complete path element sequence of the second export sequence. Paths can't end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes.

Examples:

Acceptable:

`/example` and `/path`

`/example1` and `/example2`

Not Acceptable:

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../example1`

### To update export options

Open a command prompt and run `oci fs export update`. To update export options for a specified file system, use `--export-options`.

## CHAPTER 15 File Storage

---

For example:

```
oci fs export update --export-id <export_OCID> --export-options '[{"source":"<0.0.0.0/0>","require-privileged-source-port":"true","access":"READ_ONLY","identity-squash":"ROOT","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

```
WARNING: Updates to export-options will replace any existing values. Are you sure you want to continue?
[y/N]: y
```



### Tip

If you don't want a file system to be visible to any clients, you can set all of the properties in Client Options to empty values. For example,

```
oci fs export update --export-id <export_OCID> --export-options '[{"source":"","require-privileged-source-port":"true","access":"READ_ONLY","identity-squash":"ROOT","anonymous-uid":"65534","anonymous-gid":"65534"}]'
```

### To list exports

Open a command prompt and run `oci fs export list` to list all exports in a specified compartment.

For example:

```
oci fs export list --compartment-id target_compartment_id
```

### To delete an export

Open a command prompt and run `oci fs export delete` to delete an export.

For example:

```
oci fs export delete --export-id export_OCID
```



### Warning

When you delete an export, you can no longer mount the file system using the file path specified in the deleted export.

### Using the API

- [CreateExport](#)
- [UpdateExport](#)
- [ListExports](#)
- [GetExport](#)
- [DeleteExport](#)

## Managing Snapshots

The File Storage service supports snapshots for data protection of your file system. Snapshots are a consistent, point-in-time view of your file systems. Snapshots are copy-on-write, and scoped to the entire file system. You can take as many snapshots as you need. Data usage is metered against differentiated snapshot data. If nothing has changed within the file system since the last snapshot was taken, the new snapshot does not consume more storage.

Snapshots are accessible under the root directory of the file system at `.snapshot/name`. For data protection, you can use a tool that supports NFSv3 to copy your data to a different [availability domain, region](#), file system, [object storage](#), or remote location.

For best performance, we recommend that you use the parallel tar (`partar`) and parallel copy (`parcp`) tools provided in the File Storage Parallel File Toolkit for this purpose. These tools work best with parallel workloads and requests. The Parallel File Toolkit is available for Oracle Linux, Red Hat Enterprise Linux, and CentOS. You can use `rsync` or regular `tar` for other operating system types. See [To install the Parallel File Tools suite](#) for more information.



### Tip

Watch a [video](#) about protecting data with snapshots in File Storage.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let users create, manage, and delete file systems](#) allows users to create and delete snapshots.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

## Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Using the Console

#### To create a snapshot

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. In the **File Systems** list, locate the file system you want to take a snapshot of. Click the Actions icon (three dots), and then click **View File System Details**.
4. In **Resources**, click **Snapshots**.
5. Click **Create Snapshot**.
6. Fill out the required information:
  - **Name:** Enter a name for the snapshot. It must be unique among all other snapshots for this file system. The name can't be changed. Avoid entering confidential information.
7. Click **Create Snapshot**.

The snapshot is accessible under the root directory of the file system at `.snapshot/name`.

#### To view details of a snapshot

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. In the **File Systems** list, locate the file system you took the snapshot of. Click the Actions icon (three dots), and then click **View File System Details**.
4. In **Resources**, click **Snapshots**.

5. In the **Snapshots** list, locate the snapshot you're interested in. Click the Actions icon (three dots), and then click **View Snapshot Details**.

### To create a snapshot from a Unix-style instance

You can create a snapshot from an instance that you've mounted the file system to. Snapshots are created under the root folder of your file system, in a hidden directory named `.snapshot`.

1. Connect to your instance and open a command window.
2. Navigate to your file system's hidden `.snapshot` directory. Type the following, replacing `yourmountpoint` with the name of the directory where you mounted the file system.

```
cd /mnt/yourmountpoint/.snapshot
```

3. Use the `mkdir` command to create a directory in the hidden `.snapshot` directory. The directory you create is the snapshot. Give the snapshot a name that will help you identify it. Avoid using confidential information in the snapshot name. For example:

```
mkdir snapshot-Jan1
```

4. Use the `ls` command to verify that your snapshot has been created in the `.snapshot` directory.

```
ls
```

### To install the Parallel File Tools suite

The Parallel File Tools suite provides parallel versions of `tar`, `rm`, and `cp`. These tools can run requests on large file systems in parallel, maximizing performance for data protection operations.

The tool suite is distributed as an RPM for Oracle Linux, Red Hat Enterprise Linux, and CentOS.

The toolkit includes:

- `partar`: Use this command to create and extract `tarballs` in parallel.
- `parrm`: You can use this command to recursively `remove` a directory in parallel.

## CHAPTER 15 File Storage

---

- `parcp`: Use this command to recursively `copy` a directory in parallel.

### To install Parallel File Tools on Oracle Linux:

```
sudo yum install -y fss-parallel-tools
```

### To install Parallel File Tools on CentOS and Red Hat 6.x:

```
sudo wget http://yum.oracle.com/public-yum-ol6.repo -O /etc/yum.repos.d/public-yum-ol6.repo
sudo wget http://yum.oracle.com/RPM-GPG-KEY-oracle-ol6 -O /etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
sudo yum --enablerepo=ol6_developer install fss-parallel-tools
```

### To install Parallel File Tools on CentOS and Red Hat 7.x:

```
sudo wget http://yum.oracle.com/public-yum-ol7.repo -O /etc/yum.repos.d/public-yum-ol7.repo
sudo wget http://yum.oracle.com/RPM-GPG-KEY-oracle-ol7 -O /etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
sudo yum --enablerepo=ol7_developer install fss-parallel-tools
```

### To display manual pages for each tool:

```
man partar
```

```
man parcp
```

```
man parrm
```

### To restore a snapshot

Snapshots are created under the root folder of your file system, in a hidden directory named `.snapshot`.

You can restore a file within the snapshot, or an entire snapshot using the `cp` command. Use the `-r` option when restoring a snapshot that contains subdirectories.

For example:

```
cp -r .snapshot/snapshot_name/* destination_directory_name
```

Optionally, you can use `rsync`, `tar`, or another tool that supports NFSv3 to copy your data to another remote location. For optimal performance, use the [Parallel File Tools](#).

For example:

```
parcp .snapshot/snapshot_name/* destination_directory_name
```

### To manage tags for a snapshot

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. In the **File Systems** list, locate the file system you took the snapshot of. Click the Actions icon (three dots), and then click **View File System Details**.
4. In **Resources**, click **Snapshots**.
5. In the **Snapshots** list, locate the snapshot you're interested in. Click the Actions icon (three dots), and then click **View Snapshot Details**.
6. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To delete a snapshot

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. In the **List Scope** section, select a compartment.
3. Find the file system with the snapshot you want to delete.
4. Click the Actions icon (three dots), and then click **View File System Details**.
5. In **Resources**, click **Snapshots**.
6. Find the snapshot you want to delete.
7. Click **Delete**.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#).

### To create a snapshot

You can create a snapshot of a file system. A snapshot is a point-in-time view of the file system. The snapshot is accessible at `./shapshot/name`.

Open a command prompt and run `oci fs snapshot create` to create a snapshot of a file system.

For example:

```
oci fs snapshot create --file-system-id <file_system_OCID> --name "<January1>"
```



#### Warning

Avoid entering confidential information in the snapshot name.

### To list snapshots

Open a command prompt and run `oci fs snapshot create` to list all snapshots associated with a specific file system.

For example:

```
oci fs snapshot list --file-system-id <file_system_OCID>
```

### To get a specific snapshot

Open a command prompt and run `oci fs snapshot get` to retrieve information about a specific snapshot.

For example:

```
oci fs snapshot get --snapshot-id <snapshot_OCID>
```

### To delete a snapshot

Open a command prompt and run `oci fs snapshot delete` to delete a snapshot.

For example:

```
oci fs snapshot delete --snapshot-id <snapshot_OCID>
```

### Using the API

- [CreateSnapshot](#)
- [ListSnapshots](#)
- [GetSnapshot](#)
- [DeleteSnapshot](#)

## Paths in File Systems

The File Storage service uses three kinds of paths :

1. **Export Paths** are part of the information contained in an export that makes a file system available through a mount target. The export path uniquely identifies the file system within the mount target, letting you associate up to 100 file systems behind a single mount target. The export path is used by an instance to mount (logically attach to) the file system. This path is unrelated to any path within the file system or the client instance. It exists solely as a way to distinguish one file system from another within a single mount target.

In this mount command example, `10.0.0.6` is the mount target IP address. `/example/path` is the unique export path that was specified when the file system was associated with a mount target during creation.

```
sudo mount 10.0.0.6:/example/path /mnt/mountpointA
```



#### Important



The path must start with a slash (/) followed by a sequence of zero or more slash-separated elements. For multiple file systems associated with a single mount target, the path sequence for the first file system cannot contain the complete path element sequence of the second file system path sequence. Paths cannot end in a slash. No path element can be a period (.) or two periods in sequence (..). Lastly, no path can exceed 255 bytes. For example:

**Acceptable:**

`/example` and `/path`

`/example` and `/example2`

**Not Acceptable:**

`/example` and `/example/path`

`/` and `/example`

`/example/`

`/example/path/../example1`



### Warning

If one file system associated to a mount target has '/' specified as an export path, you can't associate another file system with that mount target.



### Note

Export paths cannot be edited after the export is created. If you want to use a different export path, you must create a new export with the desired path. Optionally, you can then delete the export with the old path.

See [Managing Mount Targets](#) for more information about mount targets and exports.

2. **Mount Point Paths** are paths within a client instance to a locally accessible directory to which the remote file system is mounted.

In this mount command example, `/mnt/mountpointA` is the path to the directory on the client instance on which the external file system is mounted.

```
sudo mount 10.0.0.6:/example/path /mnt/mountpointA
```

See [Mounting File Systems](#) for more information.

3. **File System Paths** are paths to directories within the file system, and contain the contents of the file system. When the file system is mounted, you can create any directory structure within it you like. Snapshots of the file system can be accessed using the file system path, under the file system's root directory at `.snapshot/name`.

The following example shows the path to a snapshot called 'January 1' when navigating from the instance:

```
/mountpointA/.snapshot/January1
```

## Troubleshooting Your File System

These topics cover some common issues you may run into and how to address them.

### File System Setup

- [Mount Command Fails](#)
- [Mount Target Creation Fails](#)

- [Mount Target is in a Failed State](#)

### File System Management

- [Showmount Command Fails](#)
- [Symbolic Links \(Symlinks\) Produce Errors](#)
- [Removing File Locks from a Host that is No Longer Available](#)
- [Cannot Delete VCN - Mount Target VNIC Still Attached](#)

### Application Installation

- [Application Installation Fails Due to Too Much or Too Little Available Capacity](#)
- [Application Performance is Not as Expected](#)

### Windows NFS Connections

- [Create and Write to File System Fails using Windows NFS](#)
- [Mounted Drive is Not Visible in File Explorer](#)
- [Windows 2008 R2: UNC Access Delayed; "Network Error 53 Network path not found"](#)
- [Mounting from File Explorer Fails With "An Unexpected Error Occurred."](#)

### Application Installation Fails Due to Too Much or Too Little Available Capacity

The File Storage service reports file system capacity as 8589934592 gibibytes (GiB) and 8589934592 gibiinodes (GiI) by default. Sometimes, application installers perform a space requirement check prior to running an installation process but cannot correctly interpret the reported size or reported inodes of the file system. When this occurs, you can define the file system size reported to the operating system by setting the **Reported Size** or **Reported Inodes** value in the export set of the file system's mount target.



### Important

Changing the **Reported Size** or **Reported Inodes** for a mount target affects all file systems exported by the mount target. **Changing these values does not limit the amount of data you can store.**

If your application installation is failing because of too little available space, you can expand the reported available free space. If your application installation is failing because of too much reported available free space, you can reduce it. Typically, setting the size to 1024 GiB and the inodes to 1024 GiI permits successful installation.



### Important

**There can be a delay of up to 1 hour** when reporting file system usage, either in the console or by using the `df` command.

[To set the file system reported size](#) in the Console

[To set the file system reported size](#) in the CLI

### To set the reported free space in the API

You can use the [UpdateExportSet](#) operation to update the `MaxFsStatBytes`.

See [REST APIs](#) for more information.

### Application Performance is Not as Expected

Several factors can impact application performance:

- **Available bandwidth**

We recommend that you use bare metal Compute instances because instance bandwidth scales with the number of oCPU's. Bare metal Compute instances provide the greatest bandwidth. Virtual machines (VMs) are bandwidth limited based on the number of CPUs consumed. Single oCPU VM Compute instances provide the least bandwidth.

- **Latency**

Subnets can be either AD-specific or regional. The type of subnet you choose to create your File Storage resources in can affect latency. You can create File Storage resources in either type of subnet.

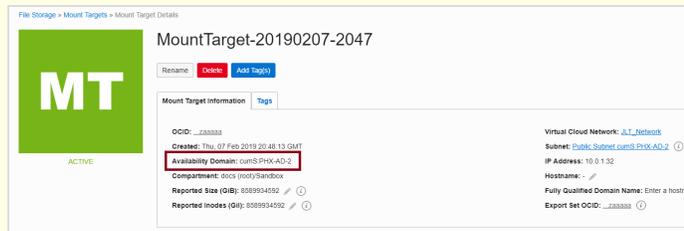
Regional subnets allow Compute instances to connect to any mount target in the subnet regardless of AD, with no additional routing configuration. However, to minimize latency, place mount targets in the same AD as Compute instances just as you would in an AD-specific subnet.

For more information, see [About Regional Subnets](#).



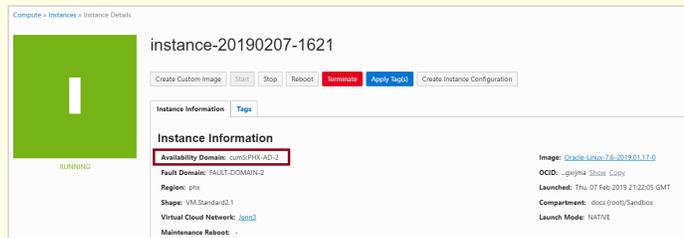
**Tip**

If you want to verify that your instance and mount target are in the same availability domain, you can view the availability domain for any mount target in its **Details** page, in the **Mount Target Information** tab:



A file system is always in the same subnet as its associated mount target.

You can also view the availability domain for any instance in its **Details** page, in the **Instance Information** tab:



- **Mount options**

By not providing explicit values for mount options such as `rsize` and `wsize`, the client and server can negotiate the window size for read and write operations that provide the best performance.

### Cannot Delete VCN- Mount Target VNIC Still Attached

A mount target is an NFS endpoint that lives in a VCN subnet of your choice and provides network access for the file systems that it exports. Each mount target has a VNIC to enable network access. Mount target VNICs that remain in a VCN must be deleted before you can delete the VCN.

Deleting a mount target also deletes all of the exports of associated file systems that exist in its export set. Data in the file systems is not affected, but the file systems are no longer available through the deleted mount target. You can create new exports for the file system in a different mount target and subnet.

For more information, see [Managing Mount Targets](#).

### To resolve this issue using the Console

1. Note the OCID in the error message you receive when you attempt to delete the VCN. Mount target OCIDs contain the identifier `mounttarget`. For example:

```
ocid1.mounttarget.oc1.phx.examplemounttargetid
```

2. Note the **Compartment** and **Subnet** information of the **VCN** you want to delete, and to assist navigation and choosing the correct mount target to delete.
3. Delete the mount target using the following steps:
  - a. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **Mount Targets**.
  - b. In the **List Scope** section, select a compartment.
  - c. Find the mount target you want to delete.
  - d. Click the Actions icon (three dots), and then click **Delete**.



### Warning

Deleting the mount target also deletes all of its exports of associated file systems. File systems are no longer available through the deleted mount target.



### Tip

In the Console, the mount target OCID can be seen in the mount target details page in the **Mount Target Information** tab. See [Managing Mount Targets](#) for more information about how to view the mount target details page. Be sure the mount target OCID seen on the details page matches the mount target OCID provided by the VCN delete process error message.

4. Delete the VCN.

### To resolve this issue using the API

1. Note the OCID in the error message you receive when you attempt to delete the VCN. Mount target OCIDs contain the identifier `mounttarget`. For example:

```
ocidl.mounttarget.oc1.phx.examplemounttargetid
```

2. Delete the mount target using the following steps:

- a. Use [DeleteMountTarget](#) to delete the mount target. For example:

```
DELETE /20171215/mountTargets/ocidl.mounttarget.oc1.phx.examplemounttargetid
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers
```

- b. You can use [GetMountTarget](#) to verify that the mount target has been deleted. For example:

```
GET
/20171215/mountTargets/ocidl.mounttarget.oc1.phx.examplemounttargetid?compartmentId=
<compartmentId>
Host: filestorage.us-phoenix-1.oraclecloud.com
<authorization and other headers>
```

The API should return Status 404 Not Found.

### To resolve this issue using the CLI

For general information about using the CLI, see [Command Line Interface \(CLI\)](#).

1. Note the OCID in the error message you receive when you attempt to delete the VCN. Mount target OCIDs contain the identifier `mounttarget`. For example:

```
ocidl.mounttarget.oc1.phx.examplemounttargetid
```

2. Delete the mount target using the following steps:

- a. Use `oci fs mount-target delete` to delete the mount target. For example:

```
oci fs mount-target delete --mount-target-id
ocidl.mounttarget.oc1.phx.examplemounttargetid
```

- b. You can use `oci fs export get` to verify that the mount target has been deleted. For example:

```
oci fs export get --export-id ocidl.mounttarget.oc1.phx.examplemounttargetid
```

The CLI should return a message indicating the mount target is not found. For example:

```
{
 "code": "NotAuthorizedOrNotFound",
 "message": "Authorization failed or requested resource not found.",
 "opc-request-id": "<requestID>",
 "status": 404
}
```

If you still can't delete the VCN, be sure there are no other resources remaining in the VCN that might prevent it. For more information, see [Subnet or VCN Deletion](#).

### Mount Command Fails

The export path of a file system must be correctly represented in your mount command, or the mount will fail.

The export path is specified when you create an export for the file system in a mount target. It uniquely identifies the file system within the mount target, letting you associate multiple file systems to a single mount target. The export path is appended to the mount target IP address, and used to mount the file system.

```
sudo mount 10.0.0.6:/example/path /mnt/mountpointA
```

In this example, `10.0.0.6:` is the mount target IP address, and `/example/path` is the export path. `/mnt/mountpointA` is the path to the directory on the client instance on which the external file system is mounted.



#### Tip

You can find all the export paths for a file system in the **Exports** list shown in its [Details page](#), together with associated mount target information.

- You can obtain the correct export path by copying mount commands directly from the file system export. These commands minimize the chance of a typing error. See [To get mount command samples](#) for more information.
- If one file system associated with a mount target uses an export path of `'/'`, it will prevent you from associating more file systems with that mount target. No two file systems associated with the same mount target can have an export path that contains a complete path of the other.

See [Paths in File Systems](#) for more information.

### Mount Target Creation Fails

Mount target creation can fail for various reasons:

- **You've exceeded your mount target limit.**

Each availability domain is limited to two mount targets by default. If you create both a file system and a mount target at the same time, it is possible for the file system to be successfully created but the mount target creation to fail because of this limitation.

You can create an export for the file system in a previously existing mount target. For more information, see [To create an export for a file system](#).

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

- **There aren't enough available IP addresses in your subnet.**

Each mount target requires three internal IP addresses in the subnet to function. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover.

Do not use /30 or smaller subnets for mount target creation because they do not have sufficient available IP addresses for mount target creation.

The File Storage service doesn't "reserve" the third IP address required for high availability failover, so use care when designing your subnets and file systems to ensure that sufficient IP addresses remain available for your mount targets in the future.

### Mount Target is in a Failed State

**Symptom:** A mount target reports a **Failed** state. File systems are not accessible using the mount target's IP address.

**Possible Cause:** There are insufficient unallocated IP addresses in the subnet. The mount target cannot fail over successfully.

Each mount target requires three internal IP addresses in the subnet to function. Two of the IP addresses are used during mount target creation. The third IP address must remain available for the mount target to use for high availability failover. The File Storage service doesn't "reserve" the third IP address required for high availability failover. Use care to ensure that

enough unallocated IP addresses remain available for your mount targets to use during failover.

### **Solution:**

1. Delete the failed mount target.

[To delete a mount target](#)

2. Export the file system through an active mount target. You can create a replacement mount target and then create an export for the file system, or create an export for the file system in a pre-existing mount target.
  - You can use the same export paths for the associated file systems as the previous mount target. However, the export path must be unique for each file system within the mount target.
  - If you create a replacement mount target, you can use the same IP address as the previous mount target, if available. Be sure to explicitly specify the desired IP address when you create the mount target.

[To create a mount target](#)

[To create an export for a file system](#)

3. If necessary, mount the file systems again.

[Mounting File Systems](#)



### **Note**

If a replacement mount target uses *exactly the same* IP address and export paths as previously existed in the deleted mount target, mounted instances reconnect automatically.

4. To prevent a recurrence of this issue, ensure that sufficient unallocated IP addresses remain available in the subnet.

### Removing File Locks from a Host that is No Longer Available

The File Storage service supports the removal of file locks from any file system. To request the removal of file locks on a file system:

1. Go to [My Oracle Support](#) and sign in.  
If you are not signed in directly to Oracle Cloud Support, click **Switch to Cloud Support** at the top of the page.
2. Click **Create Service Request**.
3. Select the following from the displayed menus:
  - **Service Type:** Select Oracle Cloud Infrastructure from the list.
  - **Service Name:** Select the appropriate option for your organization.
  - **Problem Type:** FSS File System Lock Removal Request.
4. Enter your contact information.
5. Enter a **Description**, and then enter the following required fields specific to your issue. For most Oracle Cloud Infrastructure issues you need to include the OCID (Oracle Cloud Identifier) for each resource you need help with. See [Locating IDs for Your Oracle Cloud Infrastructure Resources](#) for instructions on locating these.
  - Tenancy OCID
  - File System OCID
  - Mount Target OCID
  - Host IP Address

For help with any of the general fields in the service request or for information on managing your service requests, click **Help** at the top of the Oracle Cloud Support page.

### Showmount Command Fails

The File Storage service supports the `showmount -e` command. You can use the `showmount -e` command to show a list of NFS exports available from a specific mount target IP address.

For example, `$ showmount -e 10.0.0.0`

To enable the command, you must create a security list rule in the subnet containing the mount target. The rule must be a **stateful ingress** rule for **UDP** traffic with a **Destination Port Range** of **111**.

Click here for instructions about [Configuring VCN Security List Rules for File Storage](#).

For more information about this command, see the [Linux manual page about showmount](#).



### Important

Only `showmount -e` is supported. No other options are supported, and the `-e` option must be included.

## Symbolic Links (Symlinks) Produce Errors

The File Storage service fully supports the use of symbolic links. However, symbolic links are interpreted by the client and symlinks that point outside of the mounted File Storage system may be interpreted differently by each client and lead to unexpected results, such as broken links or pointing to the wrong file. Symbolic link targets that work on one client might be broken on another due to differences in file system layout or because clients mounted the file system using different mount targets.

Snapshots can also break symbolic links that point to a target outside the file system's root directory. This is because when you create a snapshot of a file system, it becomes available as a subdirectory of the `.snapshot` directory.

To minimize these potential issues, use a relative path as the target path when creating a symbolic link to a file in the network file system. Also, ensure that relative paths do not point to a target path outside the File Storage service root directory except when the target is on the local machine. If you must use a symbolic link that points to a target path outside the file system, use an absolute path starting with the client's root directory.

For example:

- Pointing to "/user/bin/example" **works**.
- Pointing to "/yourmountpoint/..." does **not** work.
- Pointing to "/home/user/yourmountpoint/..." does **not** work.

### Create and Write to File System Fails using Windows NFS



#### Important

To connect to file systems from Windows instances, the NFS Client must first be installed. Be sure to follow the installation procedure found in [Mounting File Systems From Windows Instances](#) before proceeding with troubleshooting.

**Symptom:** After installing Windows NFS client, you can successfully mount the file system from Windows, but any attempt to create or update a file in the file system fails.

**Cause 1:** Registry entries that map the `AnonymousGid` and `AnonymousUid` to the root user are missing or in the wrong place.

Access to NFS file systems requires UNIX-style user and group identities, which are not the same as Windows user and group identities. To enable users to access NFS shared resources, Windows client for NFS accesses file systems anonymously, using `AnonymousGid` and `AnonymousUid`. On brand new file systems, write permissions are only granted to the root user.

**Solution:** Verify that the correct registry entries are located in `HKEY_LOCAL_MACHINE\Software\Microsoft\ClientForNFS\CurrentVersion\Default`. If not, add the `AnonymousGid` and `AnonymousUid` registry entries to map them to the root user, and then remount the file system with the new user privileges.



**Tip**

You can verify the `AnonymousGid` and `AnonymousUid` are correctly set for a mounted file system by opening a Windows Command Line (CMD) window and typing the `mount` command without any arguments. A list of all mounted file systems and their properties is shown. The `AnonymousGid` (GID) and `AnonymousUid` (UID) values should appear as 0.

For example:

```
C:\>mount

Local Remote Properties

X: \\10.0.1.10\FileSystem UID=0, GID=0
 rsize=1048576,
 mount=soft,
 timeout=0.8
 retry=1,
 locking=yes
 fileaccess=755, lang=ANSI
 casesensitive=no
 sec=sys
```

If they appear as -2, they have not been correctly set. Proceed to the instructions below.

### To map the AnonymousGid and AnonymousUid to the root user

1. In the Windows Command Line (CMD) window, unmount the file system by typing the following. Replace `10.x.x.x` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.

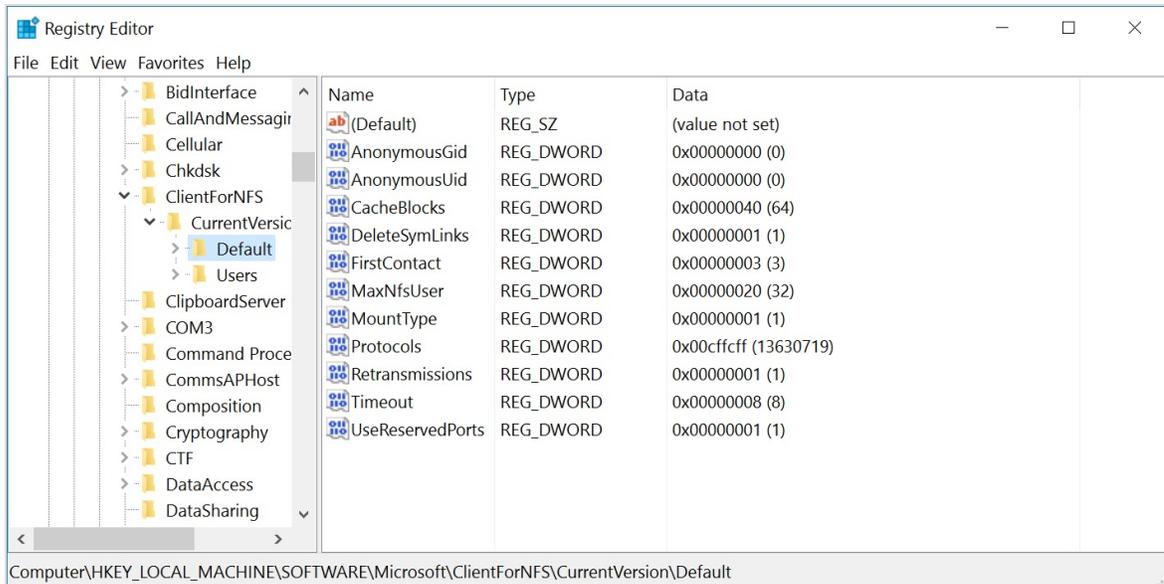


#### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
umount 10.x.x.x:/fs-export-path X:
```

2. Open the registry editor (regedit):
  - Click **Windows Search**.
  - Enter `regedit` in the **Search** field and press **Enter**.
  - Click **Yes** to allow changes to your device.
3. Click `HKEY_LOCAL_MACHINE`. Then, browse to:  
`Software\Microsoft\ClientForNFS\CurrentVersion\Default`.
4. Add a new DWORD32 registry entry for `AnonymousGid`:
  - Click **Edit**, and select **New DWORD (32 bit) Value**.
  - In the Name field, enter `AnonymousGid`. Leave the value at 0.
5. Repeat step 3 to add a second DWORD32 registry entry named `AnonymousUid` with a value of 0.



6. Open Windows Command Line (CMD) and run as Administrator:
  - Go to **Start** and scroll down to **Apps**.
  - In the **Windows System** section, press **Ctrl+Shift** and click **Command Prompt**.
7. In the Windows Command Line (CMD) window, restart the NFS Client by typing the following:

```
nfsadmin client stop
```

```
nfsadmin client start
```

8. Close the Administrator: Windows Command Prompt (CMD) window. Open a **standard** Command Prompt Window:
  - Click **Start**, then click **Command Prompt**.



### Important

NFS file systems mounted as Administrator are not available to standard users.

9. In the standard Windows Command Line (CMD) window, mount the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.

```
mount 10.x.x.x:/fs-export-path X:
```

**Cause 2:** A standard user is trying to access a file system that was mounted using the Administrator: Command Prompt (CMD). When mounting file systems, it isn't necessary to run the Command Prompt as Administrator.

**Solution:** Unmount the file system and then remount the file system using a standard Command Prompt. (CMD)

### To remount a file system with a standard Command Prompt (CMD)

1. Open Windows Command Line (CMD) and run as Administrator:
  - Go to **Start** and scroll down to **Apps**.
  - In the **Windows System** section, press **Ctrl+Shift** and click **Command Prompt**.
2. In the Administrator: Windows Command Line (CMD) window, **unmount** the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
umount 10.x.x.x:/fs-export-path X:
```

3. **Close** the Administrator: Windows Command Line (CMD) window.
4. Open a **standard** Command Prompt Window:
  - Click **Start**, then click **Command Prompt**.
5. In the **standard** Command Line (CMD) window, mount the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.

```
mount 10.x.x.x:/fs-export-path X:
```

## Mounted Drive is Not Visible in File Explorer



### Important

To connect to file systems from Windows instances, the NFS Client must first be installed. Be sure to follow the installation procedure found in [Mounting File Systems From Windows Instances](#) before proceeding with troubleshooting.

**Symptom:** After installing Windows NFS client, you can successfully mount the file system from Windows, but the file system drive is not visible in File Explorer.

**Cause:** A standard user is trying to access a file system that was mounted using the Administrator: Command Prompt (CMD). When mounting file systems, it isn't necessary to run the Command Prompt as Administrator.

**Solution:** Unmount the file system and then remount the file system using a standard Command Prompt. (CMD) See [To remount a file system with a standard Command Prompt \(CMD\)](#).

### Mounted Drive is Not Visible in PowerShell



#### Important

To connect to file systems from Windows instances, the NFS Client must first be installed. Be sure to follow the installation procedure found in [Mounting File Systems From Windows Instances](#) before proceeding with troubleshooting.

**Symptom:** After installing Windows NFS client, you can successfully mount the file system from either Windows File Explorer or the Command Prompt (CMD) using `mount` or `net use` commands. However, the file system drive is not visible in PowerShell.

**Cause:** A known issue exists where drives mapped from outside PowerShell aren't visible from within PowerShell.

**Solution:** Unmount the file system and remount the file system within PowerShell, using options to make it visible in File Explorer and in the CMD application.

### To unmount a file system using the CMD prompt

1. Open Windows Command Line (CMD) and run as Administrator:

- Go to **Start** and scroll down to **Apps**.
  - In the **Windows System** section, press **Ctrl+Shift** and click **Command Prompt**.
2. In the Administrator: Windows Command Line (CMD) window, **unmount** the file system by typing the following. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `x` with the drive letter of any available drive you want to map the file system to.



### Tip

IP address and export path information is available in the **Details** page of the mount target associated with your file system. See [To view details of a mount target](#) for more information.

```
umount 10.x.x.x:/fs-export-path X:
```

3. **Close** the Administrator: Windows Command Line (CMD) window.

## To map a drive in PowerShell and make it visible

You can map a drive in PowerShell and then use options to make it visible from File Explorer and the Windows Command Line (CMD).

1. Open **Windows PowerShell** and run as **Administrator**:
  - a. Go to **Start** and click the **Windows PowerShell** icon.
  - b. In Windows PowerShell, type the following to run as Administrator:

```
Start-Process powershell -Verb runAs
```

- c. In the **User Account Control** window, click **Yes**. A new Administrator: PowerShell window opens. You can close the standard PowerShell window to avoid confusing them.
2. Type the following cmdlet. Replace `10.x.x.x:` with the local subnet IP address assigned to your mount target, `fs-export-path` with the export path you specified when associating the file system with the mount target, and `X` with the drive letter of any available drive you want to map the file system to:

```
New-PSDrive X -PsProvider FileSystem -Root \\10.x.x.x:\fs-export-path -Persist
```

### Windows 2008 R2: UNC Access Delayed; "Network Error 53 Network path not found"



#### Important

To connect to file systems from Windows instances, the NFS Client must first be installed. Be sure to follow the installation procedure found in [Mounting File Systems From Windows Instances](#) before proceeding with troubleshooting.

**Symptom 1** : After installing NFS client on Windows 2008 R2 servers, mount fails with "Network Error 53 "Network path not found".

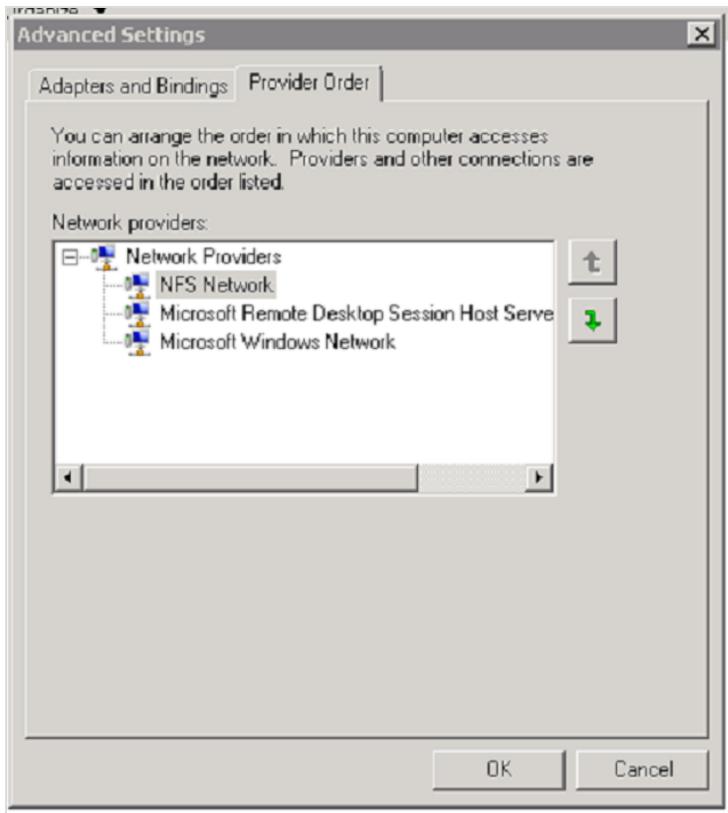
**Symptom 2:** After installing NFS client, connection to a file system using a Universal Naming Convention (UNC) path, is significantly delayed on Windows 2008 R2 servers.

**Cause 1:** Windows Network Provider has higher priority than Client for NFS Network Provider.

**Solution 1:** Change the Network Provider Order so that Client for NFS Network Provider is tried first.

### To change the Network Provider Order on Windows 2008 R2

1. Click **Start**, and then click **Network**.
2. **Right-click** or press **Shift+F10** and click **Properties**.
3. Click **Change Adapter Settings**. Press **Alt** and click **Advanced** in the menu.
4. In the **Advanced** menu, click **Advanced Settings**.
5. Click the **Provider Order** tab.
6. Select the **NFS Network** from the list of Network providers.
7. Click the **Up Arrow** to move the NFS Network provider to the top of the list.
8. Click **OK**.



**Cause 2:** The security list rules for the subnet where the mount target resides are not correctly set up.

Mounting a file system requires stateful **ingress** TCP ports 111, 2048, 2049, and 2050 as well as stateful **ingress** UDP ports 111 and 2048.

Stateful **egress** from TCP ports 111, 2048, 2049, and 2050 and stateful **egress** from UDP port 111 is also required.

**Solution 2:** Use the instructions in [Configuring VCN Security List Rules for File Storage](#) to set up the correct security list rules.

## Mounting from File Explorer Fails With "An Unexpected Error Occurred."



### Important

To connect to file systems from Windows instances, the NFS Client must first be installed. Be sure to follow the installation procedure found in [Mounting File Systems From Windows Instances](#) before proceeding with troubleshooting.

**Symptom:** The IP address and export path are correctly represented in the **Folder** field. When you click **Finish**, the system attempts to connect to the file system, but fails with an error: "The mapped network drive could not be created because the following error has occurred: An unexpected error occurred."

**Solution 1:** Reboot the instance, and [mount the file system again using File Explorer](#).

**Solution 2:** [Mount the file system using the Command Prompt](#).

# CHAPTER 16 Functions

This chapter explains how to create, deploy, and invoke functions using Oracle Functions.

## Overview of Functions

Oracle Functions is a fully managed, highly scalable, on-demand, Functions-as-a-Service platform, built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs. You don't have to worry about the underlying infrastructure because Oracle Functions will ensure your app is highly-available, scalable, secure, and monitored. With Oracle Functions, you can deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution.

Oracle Functions is based on Fn Project. Fn Project is an open source, container native, serverless platform that can be run anywhere - any cloud or on-premises. Fn Project is easy to use, supports every programming language, and is extensible and performant. You can download and install the open source distribution of Fn Project, develop and test a function locally, and then use the same tooling to deploy that function to Oracle Functions.

You can access Oracle Functions using the Console, a CLI, and a REST API. You can invoke the functions you deploy to Oracle Functions using the CLI or by making signed HTTP requests.

Oracle Functions is integrated with Oracle Cloud Infrastructure Identity and Access Management (IAM), which provides easy authentication with native Oracle Cloud Infrastructure identity functionality. See [Overview of Oracle Cloud Infrastructure Identity and Access Management](#).

To get set up and running quickly with Oracle Functions, see the [Quick Start Guide](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the REST API, see [REST APIs](#).

### Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

The following Oracle Functions resources emit events:

- applications
- functions

You can also have events in other services invoke functions in Oracle Functions. See [Invoking Oracle Functions from Other Oracle Cloud Infrastructure Services](#).

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Oracle Functions Capabilities and Limits

In your tenancy you can create a maximum of:

- 10 applications
- 20 functions

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

The maximum amount of data you can send to a function (the function's request payload) is 6MB. The maximum amount of data a function can return in response to a request (the function's response payload) is 6MB. These limits are fixed and cannot be changed.

Some other Oracle Functions capabilities and limits are also fixed. However, there are also a number that you can change. See [Changing Oracle Functions Default Behavior](#).

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For more information about policies for Oracle Functions, see:

- [Create Policies to Control Access to Network and Function-Related Resources](#)
- [Details for Functions](#)

### Oracle Functions Concepts

This topic describes key concepts you need to understand when using Oracle Functions.

#### Functions Developers

Oracle Cloud Infrastructure users who use Oracle Functions to create and deploy functions are referred to as 'functions developers'. To use Oracle Functions, functions developers must have Oracle Cloud Infrastructure user accounts. Their user accounts must belong to groups to which appropriate policies grant access to function-related resources.

See [Create Groups and Users to use with Oracle Functions, if they don't exist already](#).

#### Applications

In Oracle Functions, an application is:

- a logical grouping of functions
- a common context to store configuration variables that are available to all functions in the application

When you define an application in Oracle Functions, you specify the subnets in which to run the functions in the application.

Oracle Functions shows applications and their functions in the Console.

See [Creating Applications](#).

### Functions

In Oracle Functions, functions are:

- small but powerful blocks of code that generally do one simple thing
- grouped into applications
- stored as Docker images in a specified Docker registry
- invoked in response to a CLI command or signed HTTP request

When you deploy a function to Oracle Functions using the Fn Project CLI, the function is built as a Docker image and pushed to a specified Docker registry.

A definition of the function is stored as metadata in the Oracle Functions server. The definition describes how the function is to be executed and includes:

- the Docker image to pull when the function is invoked
- the maximum length of time the function is allowed to execute for
- the maximum amount of memory the function is allowed to consume

Oracle Functions shows functions, and the applications into which they are grouped, in the Console.

See [Creating, Deploying, and Invoking a Helloworld Function](#).

### Invocations

In Oracle Functions, a function's code is run (or executed) when the function is called (or invoked). You can invoke a function that you've deployed to Oracle Functions from:

- The Fn Project CLI.
- The Oracle Cloud Infrastructure SDKs.
- Signed HTTP requests to the function's invoke endpoint. Every function has an invoke endpoint.
- Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services.

When a function is invoked for the first time, Oracle Functions pulls the function's Docker image from the specified Docker registry, runs it as a Docker container, and executes the function. If there are subsequent requests to the same function, Oracle Functions directs those requests to the same container. After a period being idle, the Docker container is removed.

Oracle Functions shows information about function invocations in metric charts.

See [Invoking Functions](#).

### Triggers

A trigger is the result of an action elsewhere in the system, that sends a request to invoke a function in Oracle Functions. For example, an event in the Events service might cause a trigger to send a request to Oracle Functions to invoke a function. Alternatively, a trigger might send regular requests to invoke a function on a defined, time-based schedule.

A function might not be associated with any triggers, or it can be associated with one or multiple triggers.

### How Oracle Functions Works

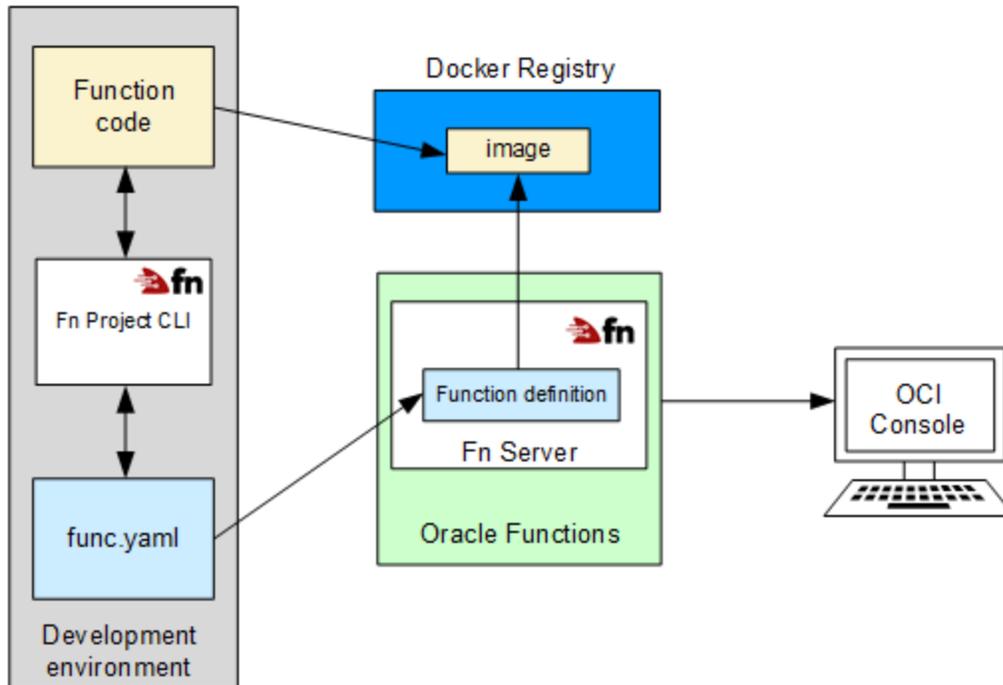
This topic describes how Oracle Functions works when you deploy a function, and when you invoke a function.

### What Happens When You Deploy a Function to Oracle Functions?

When you have written the code for a function and it's ready to deploy, you can use a single Fn Project CLI command to perform all the deploy operations in sequence:

- building a Docker image from the function
- providing a definition of the function in a `func.yaml` file that includes:
  - the maximum length of time the function is allowed to execute for
  - the maximum amount of memory the function is allowed to consume
- pushing the image to the specified Docker registry
- uploading function metadata (including the memory and time restrictions, and a link to the image in the Docker registry) to the Fn Server
- adding the function to the list of functions shown in the Console

The above process of deploying a function to Oracle Functions is shown in the diagram.



## What Happens When You Invoke a Function?

You can invoke a function that you've deployed to Oracle Functions from:

- The Fn Project CLI.
- The Oracle Cloud Infrastructure SDKs.
- Signed HTTP requests to the function's invoke endpoint. Every function has an invoke endpoint.
- Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services.

When a function is invoked for the first time, Oracle Functions first verifies the request with the IAM service. Assuming the request passes authentication and authorization checks, Oracle Functions then passes the request to the Fn Server, which uses the function definition to:

## CHAPTER 16 Functions

---

- identify the Docker image of the function to pull from the Docker registry
- execute the function by running the function's image as a container on an instance in a subnet associated with the application to which the function belongs

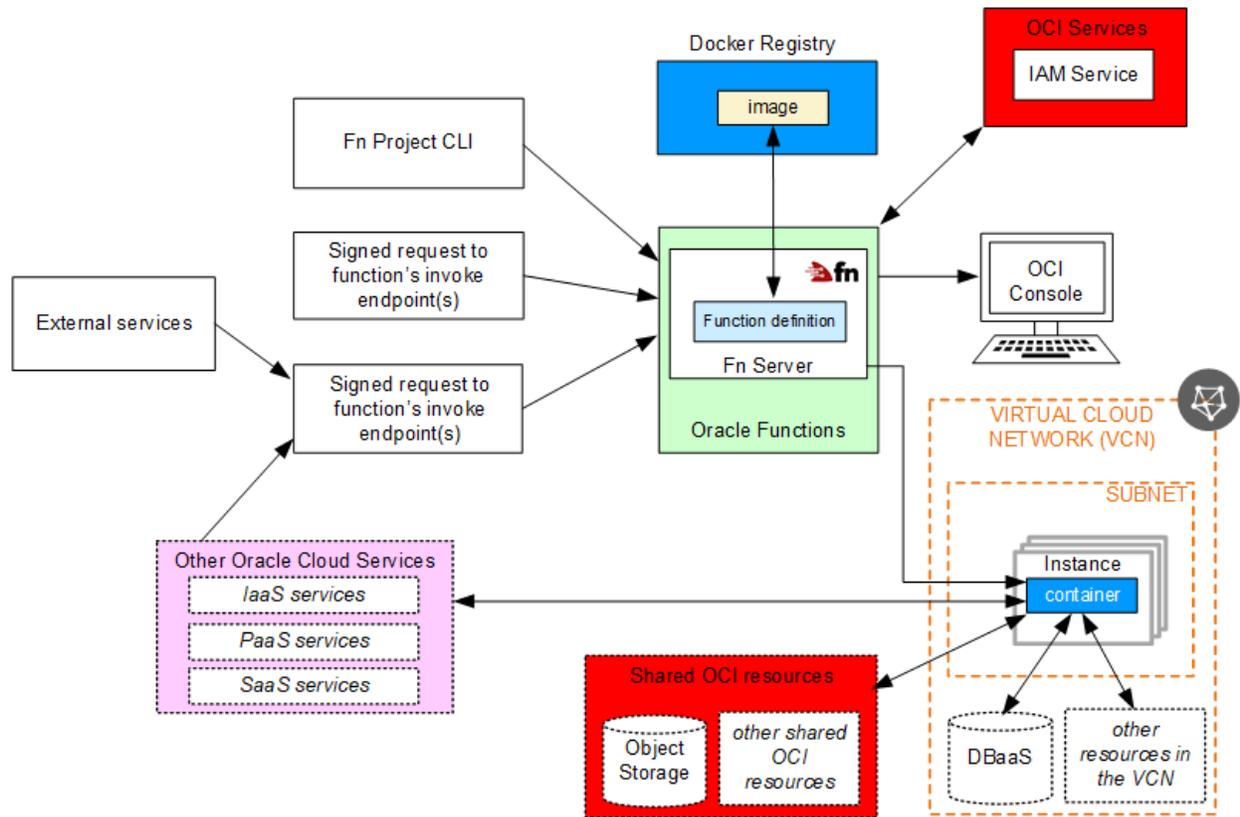
When the function is executing inside the container, the function can read from and write to other resources and services running in the same subnet (for example, Database as a Service). The function can also read from and write to other shared resources (for example, Object Storage), and other Oracle Cloud Services. You can specify the maximum length of time the function is allowed to execute by setting a timeout in the `func.yaml` file or in the Console.

Oracle Functions stores the function's logs in Oracle Cloud Infrastructure or in an external logging destination.

When the function has finished executing and after a period being idle, the Docker container is removed. If Oracle Functions receives another call to the same function before the container is removed, the second request is routed to the same running container. If Oracle Functions receives a call to a function that is currently executing inside a running container, Oracle Functions scales horizontally to serve both incoming requests and a second Docker container is started.

Oracle Functions shows information about function invocations in metric charts.

The above process of invoking a function is shown in the diagram.



## Preparing for Oracle Functions

Before you can deploy functions to Oracle Functions:

- A tenancy administrator must have configured your Oracle Cloud Infrastructure tenancy for function development. There are a number of different tenancy configuration tasks to complete. For more information, see [Configuring Your Tenancy for Function Development](#).

When your tenancy is configured, you will have access, via a suitable policy and user account, to a compartment that has a VCN with at least one public subnet (and an

internet gateway) or at least one private subnet (and a service gateway). For more information about these network components, see [Overview of Networking](#).

You will also have access to a Docker registry in which to store images. This documentation assumes you will be using Oracle Cloud Infrastructure Registry as your Docker registry and provides instructions accordingly. For more information, see [Overview of Registry](#).

- You must have configured your client environment for functions development. There are a number of different client environment configuration tasks to complete. For more information, see [Configuring Your Client Environment for Function Development](#).

To get set up and running quickly with Oracle Functions, see the [Quick Start Guide](#).

## Availability by Region Name and Region Code

Oracle Functions is available in the following regions. Note that you have to use region name and region code in some commands. In some cases, you might have to use shortened versions of availability domain names.

Region	Region Identifier	Region Code	Shortened Availability Domain Names
US East (Ashburn)	us-ashburn-1	iad	<ul style="list-style-type: none"> <li>• US-ASHBURN-AD-1</li> <li>• US-ASHBURN-AD-2</li> <li>• US-ASHBURN-AD-3</li> </ul>
Germany Central (Frankfurt)	eu-frankfurt-1	fra	<ul style="list-style-type: none"> <li>• EU-FRANKFURT-1-AD-1</li> <li>• EU-FRANKFURT-1-AD-2</li> <li>• EU-FRANKFURT-1-AD-3</li> </ul>
UK South (London)	uk-london-1	lhr	<ul style="list-style-type: none"> <li>• UK-LONDON-1-AD-1</li> <li>• UK-LONDON-1-AD-2</li> <li>• UK-LONDON-1-AD-3</li> </ul>

Region	Region Identifier	Region Code	Shortened Availability Domain Names
India West (Mumbai)	ap-mumbai-1	bom	<ul style="list-style-type: none"> <li>• AP-MUMBAI-1-AD-1</li> </ul>
US West (Phoenix)	us-phoenix-1	phx	<ul style="list-style-type: none"> <li>• PHX-AD-1</li> <li>• PHX-AD-2</li> <li>• PHX-AD-3</li> </ul>
Brazil East (Sao Paulo)	sa-saopaulo-1	gru	<ul style="list-style-type: none"> <li>• SA-SAOPAULO-1-AD-1</li> </ul>
South Korea Central (Seoul)	ap-seoul-1	icn	<ul style="list-style-type: none"> <li>• AP-SEOUL-1-AD-1</li> </ul>
Australia East (Sydney)	ap-sydney-1	syd	<ul style="list-style-type: none"> <li>• AP-SYDNEY-1-AD-1</li> </ul>
Japan East (Tokyo)	ap-tokyo-1	nrt	<ul style="list-style-type: none"> <li>• AP-TOKYO-1-AD-1</li> </ul>
Canada Southeast (Toronto)	ca-toronto-1	yyz	<ul style="list-style-type: none"> <li>• CA-TORONTO-1-AD-1</li> </ul>
Switzerland North (Zurich)	eu-zurich-1	zrh	<ul style="list-style-type: none"> <li>• EU-ZURICH-1-AD-1</li> </ul>

## Configuring Your Tenancy for Function Development

Before you can start using Oracle Functions to create and deploy functions, you have to set up your tenancy for function development.

When a tenancy is created, an Administrators group is automatically created for the tenancy. Users that are members of the Administrators group can perform any operation on resources in the tenancy. Oracle Functions users are typically not members of the Administrators group, and do not have to be. However, a member of the Administrators group does need to perform a number of administrative tasks to enable users to use Oracle Functions.

## CHAPTER 16 Functions

---

To set up your tenancy for function development, you have to complete the following tasks in the order shown in this checklist (the instructions in the topics below assume that you are a tenancy administrator):

Task #	Tenancy Configuration Task	Done?
1	<a href="#">Create Groups and Users to use with Oracle Functions, if they don't exist already</a>	
2	<a href="#">Create Compartments to Own Network Resources and Oracle Functions Resources in the Tenancy, if they don't exist already</a>	
3	<a href="#">Create the VCN and Subnets to Use with Oracle Functions, if they don't exist already</a>	
4	<a href="#">Create Policies to Control Access to Network and Function-Related Resources</a> , and more specifically: <ul style="list-style-type: none"><li>• <a href="#">Create a Policy to Give Oracle Functions Users Access to Oracle Cloud Infrastructure Registry Repositories</a></li><li>• <a href="#">Create a Policy to Give Oracle Functions Users Access to Function-Related Resources</a></li><li>• <a href="#">Create a Policy to Give Oracle Functions Users Access to Network Resources</a></li><li>• <a href="#">Create a Policy to Give the Oracle Functions Service Access to Network Resources</a></li><li>• <a href="#">Create a Policy to Give the Oracle Functions Service Access to Repositories in Oracle Cloud Infrastructure Registry</a></li></ul>	

Click each of the links in turn, and follow the instructions.

When you have set up your tenancy for function development, the next step is to set up your client development environment (see [Configuring Your Client Environment for Function Development](#)).

### Create Groups and Users to use with Oracle Functions, if they don't exist already

Before users can start using Oracle Functions to create and deploy functions, as a tenancy administrator you have to create Oracle Cloud Infrastructure user accounts, along with a group to which the user accounts belong. Later on, you'll define policies to give the group (and the user accounts that belong to it) access to function-related resources. If a suitable group and user accounts already exist, there's no need to create new ones.

To create groups and users to use with Oracle Functions:

1. Log in to the Console as a tenancy administrator.
2. If a suitable group for Oracle Functions users doesn't exist already, create such a group as follows:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**. A list of the groups in your tenancy is displayed.
  - b. Click **Create Group** and create a new group (see [To create a group](#)). Give the group a meaningful name (for example, `acme-functions-developers`) and description. Avoid entering confidential information.
3. If suitable user accounts for Oracle Functions users don't exist already, create users as follows:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. A list of the users in your tenancy is displayed.
  - b. Click **Create User** and create one or more new users (see [To create a user](#)).
4. If they haven't been added already, add users to the group to use Oracle Functions as follows:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. A list of the users in your tenancy is displayed.
  - b. Select one or more users and add them to the group authorized to use Oracle Functions (see [To add a user to a group](#)).

### **Create Compartments to Own Network Resources and Oracle Functions Resources in the Tenancy, if they don't exist already**

Before users can start using Oracle Functions to create and deploy functions, as a tenancy administrator you have to create:

- a compartment to own network resources (a VCN, a public or private subnet, and other resources such as an internet gateway or service gateway, a route table, security lists)
- a compartment to own function-related resources (functions, applications)

Note that the same compartment can own both network resources and function-related resources. Alternatively, you can create two separate compartments for network resources and function-related resources.

If suitable compartments already exist, there's no need to create new ones.

To create a compartment to own network resources and/or function-related resources in the tenancy:

1. Log in to the Console as a tenancy administrator.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**. A list of the compartments in your tenancy is displayed.
3. Click **Create Compartment** and create a new compartment (see [To create a compartment](#)). Give the compartment a meaningful name (for example, `acme-network`, `acme-functions-compartment`) and description. Avoid entering confidential information.

### **Create the VCN and Subnets to Use with Oracle Functions, if they don't exist already**

Before users can start using Oracle Functions to create and deploy functions, a VCN containing the subnets in which to create functions and applications must already exist. The VCN can be, but need not be, owned by the same compartment to which other function-related resources will belong.

Each subnet in the VCN must have a CIDR block that provides at least a certain minimum number of free IP addresses, as follows:

## CHAPTER 16 Functions

---

- AD-specific subnets must have a minimum of 12 free IP addresses
- regional subnets must have a minimum of 32 free IP addresses

Note that Oracle strongly recommends each subnet has a CIDR block that provides more than the minimum number of free IP addresses.

To support the largest possible number of concurrent connections, Oracle also strongly recommends that the security lists used by subnets in the VCN only have stateless rules.

If a suitable VCN already exists, there's no need to create a new one.

If you do decide to create a new VCN, you have several options, including the following:

- You can create the new VCN and have related resources created automatically at the same time (as described in this topic). In this case, three public subnets and an internet gateway are created, and a new route rule is added to the default route table.
- You can create just the VCN initially, and then create the related resources yourself later (see [VCNs and Subnets](#)). In this case, you can choose whether to create public subnets and an internet gateway (see [Internet Gateway](#)), or private subnets and a service gateway (see [Access to Oracle Services: Service Gateway](#)). For example, if you don't want to expose traffic over the public internet, create private subnets and a service gateway.

Note that to use an external logging destination like Papertrail, you have to create a VCN with public subnets (see [Storing and Viewing Function Logs](#)).

To create a VCN to use with Oracle Functions (with related resources created automatically):

1. Log in to the Console as a tenancy administrator.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
3. Choose the compartment that will own the network resources (on the left side of the page). For example, `acme-network`.

The VCN can be, but need not be, owned by the same compartment to which other function-related resources will belong. The page updates to display only the resources in that compartment.

4. Click **Create Virtual Cloud Network** to create a new VCN.
5. In the **Create Virtual Cloud Network** dialog box, enter the following:
  - a. **Name:** A meaningful name for the cloud network, such as `acme-functions-vcn`. The name doesn't have to be unique, but it cannot be changed later in the Console. Avoid entering confidential information.
  - b. **Create Virtual Cloud Network Plus Related Resources:** For convenience, select this option to create the VCN and associated resources (all with default properties) in a single operation.
  - c. **Use DNS Hostnames in this VCN:** Select this option.
6. Click **Create Virtual Cloud Network** to create the VCN, along with the related resources (three public subnets and an internet gateway).

### Create Policies to Control Access to Network and Function-Related Resources

Before users can start using Oracle Functions to create and deploy functions, as a tenancy administrator you have to create a number of Oracle Cloud Infrastructure policies to grant access to function-related and network resources. You have to:

- [Create a Policy to Give Oracle Functions Users Access to Oracle Cloud Infrastructure Registry Repositories](#)
- [Create a Policy to Give Oracle Functions Users Access to Function-Related Resources](#)
- [Create a Policy to Give Oracle Functions Users Access to Network Resources](#)
- [Create a Policy to Give the Oracle Functions Service Access to Network Resources](#)
- [Create a Policy to Give the Oracle Functions Service Access to Repositories in Oracle Cloud Infrastructure Registry](#)

See [Details for Functions](#) for more information about policies.

## CHAPTER 16 Functions

### SUMMARY OF POLICIES TO CREATE FOR ORACLE FUNCTIONS

<b>Policy to give:</b>	<b>Where to create the policy:</b>	<b>Statement:</b>	<b>More information and examples:</b>
Users access to repositories in Oracle Cloud Infrastructure Registry	Root compartment	Allow group <group-name> to manage repos in tenancy	<a href="#">Create a Policy to Give Oracle Functions Users Access to Oracle Cloud Infrastructure Registry Repositories</a>
Users access to function-related resources	Compartment that owns function-related resources	Allow group <group-name> to manage functions-family in compartment <compartment-name>  Allow group <group-name> to read metrics in compartment <compartment-name>	<a href="#">Create a Policy to Give Oracle Functions Users Access to Function-Related Resources</a>
Users access to network resources	Compartment that owns network resources	Allow group <group-name> to use virtual-network-family in compartment <compartment-name>	<a href="#">Create a Policy to Give Oracle Functions Users Access to Network Resources</a>

Policy to give:	Where to create the policy:	Statement:	More information and examples:
Oracle Functions service access to network resources	Root compartment	Allow service FaaS to use virtual-network-family in compartment <compartment-name>	<a href="#">Create a Policy to Give the Oracle Functions Service Access to Network Resources</a>
Oracle Functions service access to repositories in Oracle Cloud Infrastructure Registry	Root compartment	Allow service FaaS to read repos in tenancy	<a href="#">Create a Policy to Give the Oracle Functions Service Access to Repositories in Oracle Cloud Infrastructure Registry</a>

#### CREATE A POLICY TO GIVE ORACLE FUNCTIONS USERS ACCESS TO ORACLE CLOUD INFRASTRUCTURE REGISTRY REPOSITORIES

When Oracle Functions users work with functions, they have to access repositories in Oracle Cloud Infrastructure Registry. Users can only access repositories that the groups to which they belong have been granted access. To enable users to access a repository, you must create an identity policy to grant the groups access to that repository.

To create a policy to give Oracle Functions users access to repositories in Oracle Cloud Infrastructure Registry:

1. Log in to the Console as a tenancy administrator and create a new policy in the root compartment:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
  - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `acme-functions-developers-ocir-access`).

2. Specify a policy statement to give the group access to repositories in Oracle Cloud Infrastructure Registry:

```
Allow group <group-name> to manage repos in tenancy
```

where `<group-name>` is the name of the group to which users using Oracle Functions belong.

For example:

```
Allow group acme-functions-developers to manage repos in tenancy
```

The above policy statement gives the group permission to manage all repositories in the tenancy. If you consider this to be too permissive, then you can restrict the repositories to which the group has access by including a `where` clause in the `manage repos` statement. Note that if you do include a `where` clause, you must also include a second statement in the policy to enable the group to inspect all repositories in the tenancy (when using the Console).

For example, the following policy statements restrict the group to accessing only repositories with names that start 'acme-web-app', but also enables the group to inspect all repositories in the tenancy:

```
Allow group acme-functions-developers to inspect repos in tenancy
```

```
Allow group acme-functions-developers to manage repos in tenancy where all
{target.repo.name=/acme-web-app*/ }
```

3. Click **Create**.

### CREATE A POLICY TO GIVE ORACLE FUNCTIONS USERS ACCESS TO FUNCTION-RELATED RESOURCES

When Oracle Functions users create functions and applications, they have to specify a compartment for those function-related resources (including for metrics emitted by Oracle Functions). Users can only specify a compartment that the groups to which they belong have been granted access. To enable users to specify a compartment, you must create an identity policy to grant the groups access to that compartment.

To create a policy to give Oracle Functions users access to function-related resources in the compartment that will own those resources:

1. Log in to the Console as a tenancy administrator and create a new policy in the compartment that will own Oracle Functions resources:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
  - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `acme-functions-developers-manage-access`).
2. Specify a policy statement to give the group access to all function-related resources in the compartment:

```
Allow group <group-name> to manage functions-family in compartment <compartment-name>
```

For example:

```
Allow group acme-functions-developers to manage functions-family in compartment acme-functions-compartment
```

3. Specify a second policy statement to give the group access to metrics emitted by Oracle Functions:

```
Allow group <group-name> to read metrics in compartment <compartment-name>
```

For example:

```
Allow group acme-functions-developers to read metrics in compartment acme-functions-compartment
```

4. Click **Create**.

### CREATE A POLICY TO GIVE ORACLE FUNCTIONS USERS ACCESS TO NETWORK RESOURCES

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. Users can only specify VCNs and subnets in compartments that the groups to which they belong have been granted access. To enable users to specify a VCN and subnet, you must create an identity policy to grant the groups access to the compartment.

To create a policy to give Oracle Functions users access to network resources:

1. Log in to the Console as a tenancy administrator and create a new policy in the compartment that will own network resources:

- a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
  - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `acme-functions-developers-manage-network-access`).
2. Specify a policy statement to give the group access to the network resources in the compartment:

```
Allow group <group-name> to use virtual-network-family in compartment <compartment-name>
```

For example:

```
Allow group acme-functions-developers to use virtual-network-family in compartment acme-network
```

3. Click **Create**.

### CREATE A POLICY TO GIVE THE ORACLE FUNCTIONS SERVICE ACCESS TO NETWORK RESOURCES

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. To enable the Oracle Functions service to create the function or application in the specified VCN and subnet, you must create an identity policy to grant the Oracle Functions service access to the compartment to which the network resources belong.

To create a policy to give the Oracle Functions service access to network resources:

1. Log in to the Console as a tenancy administrator.
2. Create a new policy in the root compartment:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
  - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `functions-service-network-access`).
  - c. Specify a policy statement to give the Oracle Functions service access to the network resources in the compartment:

```
Allow service FaaS to use virtual-network-family in compartment <compartment-name>
```

For example:

```
Allow service FaaS to use virtual-network-family in compartment acme-network
```

### 3. Click **Create**.

#### **CREATE A POLICY TO GIVE THE ORACLE FUNCTIONS SERVICE ACCESS TO REPOSITORIES IN ORACLE CLOUD INFRASTRUCTURE REGISTRY**

The Oracle Functions service must have read access to images stored for functions in repositories in Oracle Cloud Infrastructure Registry. To enable the Oracle Functions service to access repositories in Oracle Cloud Infrastructure Registry, you must create an identity policy.

To create a policy to give the Oracle Functions service access to repositories in Oracle Cloud Infrastructure Registry:

1. Log in to the Console as a tenancy administrator.
2. Create a new policy in the root compartment:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
  - b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `functions-service-repos-access`).
  - c. Specify a policy statement to give the Oracle Functions service access to all repositories in the tenancy:

```
Allow service FaaS to read repos in tenancy
```

The above policy statement gives the Oracle Functions service access to all repositories in the tenancy. If you consider this to be too permissive, then you can restrict the repositories to which Oracle Functions has access by including a `where` clause in the `read repos` statement.

For example, the following policy statement restricts Oracle Functions to accessing only repositories with names that start 'acme-web-app':

```
Allow service FaaS to read repos in tenancy where all {target.repo.name=/acme-web-app*/ }
```

### 3. Click **Create**.

## Configuring Your Client Environment for Function Development

Before you can start using Oracle Functions to create and deploy functions, you have to set up your client environment for function development. Note that prior to setting up your client environment, you must already have set up your tenancy (see [Configuring Your Tenancy for Function Development](#)).

To set up your client environment for function development, you have to complete the following tasks in the order shown in this checklist:

Task #	Development Environment Configuration Task	Done?
1	<a href="#">1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions</a>	
2	<a href="#">2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File</a>	
3	<a href="#">3. Create and Configure a Copy of oci-curl</a>	
4	<a href="#">4. Install Docker for Use with Oracle Functions</a>	
5	<a href="#">5. Install the Fn Project CLI</a>	
6	<a href="#">6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure</a>	
7	<a href="#">7. Set the Context for the Fn Project CLI Using the oracle.profile Parameter</a>	
8	<a href="#">8. Generate an Auth Token to Enable Login to Oracle Cloud Infrastructure Registry</a>	
9	<a href="#">9. Start Docker</a>	
10	<a href="#">10. Log in to Oracle Cloud Infrastructure Registry</a>	

Click each of the links in the checklist in turn, and follow the instructions.

When you have completed all of the development environment configuration tasks, confirm your configuration is correct (see [Verifying Your Configuration for Function Development](#)).

### **1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions**

Before using Oracle Functions, you have to set up an Oracle Cloud Infrastructure API signing key.

The instructions in this topic assume:

- you are using Linux
- you are following Oracle's recommendation to provide a passphrase to encrypt the private key

For more information and other options, see [Required Keys and OCIDs](#).

The instructions below describe how to create a new `~/.oci` directory, how to generate a new private key file and public key file in that `~/.oci` directory, how to upload the public key to Oracle Cloud Infrastructure to create a new API signing key, and how to obtain a fingerprint for the public API key. Be aware that instructions and examples elsewhere in this documentation assume the `~/.oci` directory exists and contains the private and public key files.

If your user account already has an API signing key, create the `~/.oci` directory if it doesn't exist, and then go straight to [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#).

If your user account doesn't already have an API signing key, follow the steps below, but note the following:

- If the `~/.oci` directory doesn't exist, create it.
- If the `~/.oci` directory already exists, go straight to the step below that instructs you to generate a private key file.
- If the `~/.oci` directory already exists and already contains a private key file and public key file, and you know the passphrase that was used to encrypt the existing private key

file, there's no need to create new private and public key files. Instead, go straight to the step below that instructs you to create a new API signing key and upload the public key value to Oracle Cloud Infrastructure to obtain a fingerprint.

- If you already have a private key file and public key file but they are not in the `~/oci` directory, and you know the passphrase that was used to encrypt the existing private key file, there's no need to create new private and public key files. Having created the `~/oci` directory if it doesn't exist, go straight to the step below that instructs you to create a new API signing key and upload the public key value to Oracle Cloud Infrastructure to obtain a fingerprint.

To set up an API signing key:

1. Log in to your development environment as a functions developer.
2. In a terminal window, confirm that the `~/oci` directory does not already exist. For example, by entering:

```
ls ~/oci
```

3. Assuming the `~/oci` directory does not already exist, create it. For example, by entering:

```
mkdir ~/oci
```

4. Generate a private key encrypted with a passphrase that you provide by entering:

```
$ openssl genrsa -out ~/oci/<private-key-file-name>.pem -aes128 2048
```

where `<private-key-file-name>` is a name of your choice for the private key file (for example, `john_api_key_private.pem`).

For example:

```
$ openssl genrsa -out ~/oci/john_api_key_private.pem -aes128 2048
```

```
Generating RSA private key, 2048 bit long modulus
```

```
....+++
```

```
.....+++
```

```
e is 65537 (0x10001)
```

```
Enter pass phrase for /Users/johndoe/.oci/john_api_key_private.pem:
```

5. When prompted, enter a passphrase to encrypt the private key file. Be sure to make a note of the passphrase you enter, as you will need it later.
6. When prompted, re-enter the passphrase to confirm it.
7. Confirm that the private key file has been created in the directory you specified. For example, by entering:

```
$ ls -l ~/.oci/john_api_key_private.pem
-rw-r--r-- 1 johndoe staff 1766 Jul 14 00:24 /Users/johndoe/.oci/john_api_key_private.pem
```

8. Change permissions on the file to ensure that only you can read it. For example, by entering:

```
$ chmod go-rwx ~/.oci/john_api_key_private.pem
```

9. Generate a public key (in the same location as the private key file) by entering:

```
$ openssl rsa -pubout -in ~/.oci/<private-key-file-name>.pem -out ~/.oci/<public-key-file-name>.pem
```

where:

- `<private-key-file-name>` is what you specified earlier as the name of the private key file (for example, `john_api_key_private.pem`)
- `<public-key-file-name>` is a name of your choice for the public key file (for example, `john_api_key_public.pem`)

For example:

```
$ openssl rsa -pubout -in ~/.oci/john_api_key_private.pem -out ~/.oci/john_api_key_public.pem
Enter pass phrase for /Users/johndoe/.oci/john_api_key_private.pem:
```

10. When prompted, enter the same passphrase you previously entered to encrypt the private key file.
11. Confirm that the public key file has been created in the directory you specified. For example, by entering:

```
$ ls -l ~/.oci/
```

## CHAPTER 16 Functions

---

```
-rw----- 1 johndoe staff 1766 Jul 14 00:24 john_api_key_private.pem
-rw-r--r-- 1 johndoe staff 451 Jul 14 00:55 john_api_key_public.pem
```

12. Copy the contents of the public key file you just created. For example, by entering:

```
$ cat ~/.oci/john_api_key_public.pem | pbcopy
```

13. Having created the API key pair, upload the public key value to Oracle Cloud Infrastructure:
  - a. Log in to the Console as the Oracle Cloud Infrastructure user who will be using Oracle Functions to create and deploy functions.
  - b. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
  - c. On the **API Keys** page, click **Add Public Key**.
  - d. Paste the public key's value into the window and click **Add**.  
The key is uploaded and its fingerprint is displayed (for example, d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:13).
  - e. (Optional) Note the fingerprint value. You'll use the fingerprint in a subsequent configuration task, so you might want to copy it to a convenient and secure location.

When you have completed the steps in this topic, go on to [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#).

### 2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File

Before using Oracle Functions, you must have an Oracle Cloud Infrastructure CLI configuration file that contains the credentials of the user account that you will be using to create and deploy functions. These user account credentials are referred to as a 'profile'.

By default, the Oracle Cloud Infrastructure CLI configuration file is located at `~/.oci/config`. You might already have a configuration file as a result of installing the Oracle Cloud Infrastructure CLI. However, you don't need to have installed the Oracle Cloud Infrastructure CLI in order to use Oracle Functions.

## CHAPTER 16 Functions

---

The Oracle Cloud Infrastructure CLI configuration file can contain several profiles. If you already have a configuration file containing one or more profiles, you have to add a new profile to the existing file for the Oracle Cloud Infrastructure user who will be using Oracle Functions to create and deploy functions.

The instructions in this topic assume:

- you are using Linux
- you have already completed the steps in [1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions](#)

To create a profile in the Oracle Cloud Infrastructure CLI configuration file for the user account that you will be using to create and deploy functions:

1. Log in to your development environment as a functions developer.
2. In a terminal window, confirm the contents of the `~/.oci` directory. For example, by entering:

```
$ ls -l ~/.oci/

-rw----- 1 johndoe staff 1766 Jul 14 00:24 john_api_key_private.pem
-rw-r--r-- 1 johndoe staff 451 Jul 14 00:55 john_api_key_public.pem
```

3. Do one of the following, depending on whether the `~/.oci` directory already contains a file called `config`:
  - If the `~/.oci` directory already contains a file called `config`, open the file in a text editor.
  - If the `~/.oci` directory doesn't yet contain a file called `config`, create the file and open it in a text editor. For example, by entering:

```
$ vim ~/.oci/config
```

4. Add a new profile to the `~/.oci/config` file as follows:

```
[<profile-name>]
user=<user-ocid>
fingerprint=<public-key-fingerprint>
key_file=<full-path-to-private-key-pem-file>
tenancy=<tenancy-ocid>
```

## CHAPTER 16 Functions

---

```
region=<region-identifier>
pass_phrase=<passphrase>
```

where:

- <profile-name> is a name of your choosing for the profile.
- <user-ocid> is the OCID of the Oracle Cloud Infrastructure user account you will be using to create and deploy functions. See [Where to Get the Tenancy's OCID and User's OCID](#).
- <public-key-fingerprint> is the fingerprint of the public API key value that you uploaded in the Console in [1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions](#).
- <full-path-to-private-key-pem-file> is the full path to the private key file that you created in [1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions](#).
- <tenancy-ocid> is the OCID of the tenancy in which you will be creating and deploying functions. See [Where to Get the Tenancy's OCID and User's OCID](#).
- <region-identifier> is the identifier of the Oracle Cloud Infrastructure region in which you will be creating and deploying functions. For example, us-phoenix-1. See [Availability by Region Name and Region Code](#) for the list of region identifiers.
- <passphrase> is the passphrase you entered in [1. Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions](#)).

For example:

```
[john-oci-profile]
user=ocidl.user.oc1..aaaaaaaas...7ap
fingerprint=d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:13
key_file=~/.oci/john_api_key_private.pem
tenancy=ocidl.tenancy.oc1..aaaaaaaap...keq
region=us-phoenix-1
pass_phrase=<your-passphrase>
```

When you have completed the steps in this topic, go on to [3. Create and Configure a Copy of oci-curl](#).

### 3. Create and Configure a Copy of oci-curl

You can use a bash script provided by Oracle (commonly referred to as oci-curl) to invoke a function. The oci-curl script creates a signed request, based on credentials you provide in the body of the script. For more information about oci-curl, see [Request Signatures - Bash](#).

To use oci-curl to invoke a function, you must provide the credentials of an Oracle Cloud Infrastructure user that has been granted access to resources in the same tenancy and belonging to the same compartment as the function (see [Create Policies to Control Access to Network and Function-Related Resources](#)).

Typically you'll want to invoke a function as the functions developer that's configured for your development environment. The instructions below assume that is the case.

The instructions in this topic assume you have already completed the steps in [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#).

To create and configure a copy of oci-curl:

1. Log in to your development environment as a functions developer.
2. Create a copy of the oci-curl script file in your development environment and add your credentials to the file as follows:
  - a. In a browser navigate to [https://docs.cloud.oracle.com/iaas/Content/Resources/Assets/signing\\_sample\\_bash.txt](https://docs.cloud.oracle.com/iaas/Content/Resources/Assets/signing_sample_bash.txt) to see the oci-curl code as raw text.
  - b. Select all the text and copy it.
  - c. In a text editor, open a new file in a convenient location. For example, in a terminal window, you might create a new subdirectory in your home directory and open a new file in that directory by entering:

```
$ cd ~

$ mkdir oci-curl

$ vim ~/oci-curl/oci-curl.sh
```

The name and location of the new file is up to you, but the following instructions assume `~/oci-curl/oci-curl.sh`.

- d. Paste the `oci-curl` script code that you copied earlier into the new file.
  - e. Save the file but leave it open so you can add your credentials.
3. Replace the sample credentials in the `oci-curl.sh` file with those of the user account that you want to invoke functions, as follows:
- a. Locate the following lines in the `oci-curl.sh` file that contain sample credential values:

```
TODO: update these values to your own
local tenancyId="ocidl.tenancy.oc1..aaaaaaaab_____dsq";
local authUserId="ocidl.user.oc1..aaaaaaaas_____o3r";
local keyFingerprint="20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34";
local privateKeyPath="/Users/someuser/.oci/oci_api_key.pem";
```



### Tip

Typically you'll want to invoke a function as the functions developer that's configured for your development environment. If that's the case, open the `~/oci/config` file so you can easily copy values from there to replace the sample values in the `oci-curl.sh` file.

- b. Change the value of the `tenancyId` parameter in the `oci-curl.sh` file by replacing the sample value in quotes with the OCID of the tenancy in which the function has been deployed. For example:

```
local tenancyID="ocidl.tenancy.oc1..aaaaaaaap_____keq";
```

- c. Change the value of the `authUserID` parameter in the `oci-curl.sh` file by replacing the sample value in quotes with the OCID of the user account that you want to run the function. The user account must have access to resources in the same tenancy and belonging to the same compartment as the function. For example:

```
local authUserId="ocid1.user.oc1..aaaaaaaas_____7ap";
```

- d. Change the value of the `keyFingerprint` parameter in the `oci-curl.sh` file by replacing the sample value in quotes with the fingerprint of the user's public key uploaded to Oracle Cloud Infrastructure. For example:

```
local keyFingerprint="d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:";
```

- e. Change the value of the `privateKeyPath` parameter in the `oci-curl.sh` file by replacing the sample value in quotes with the full path to the private key file that is paired with the public key for which you provided the fingerprint. For example:

```
local privateKeyPath="/Users/johndoe/.oci/john_api_key_private.pem";
```

4. Save and close the `oci-curl.sh` file.



### Note

When you've deployed functions to Oracle Functions and you want to use `oci-curl` to invoke them, you will first have to run the `source` command to set up the current shell environment for `oci-curl`. See [Invoking Functions](#).

When you have completed the steps in this topic, go on to [4. Install Docker for Use with Oracle Functions](#).

## 4. Install Docker for Use with Oracle Functions

Before using Oracle Functions, a version of Docker supported by Fn Project must be installed in your development environment. If Docker is not already installed, or the installed version of Docker is not supported, you'll have to install or upgrade Docker.

The instructions in this topic assume you have already completed the steps in [3. Create and Configure a Copy of `oci-curl`](#).

## CHAPTER 16 Functions

---

To confirm that a supported version of Docker is installed in your development environment:

1. Log in to your development environment as a functions developer.
2. In a terminal window, confirm that Docker is installed by entering:

```
$ docker version
```

3. Do one of the following, depending on the message you see:
  - If you see an error message indicating that Docker is not installed, you have to install Docker before proceeding to the next step. See the [Docker documentation](#) for information about installing Docker on your platform. If your platform is Oracle Linux, see [Oracle Container Runtime for Docker User's Guide](#).
  - If you see a message indicating the version of Docker that's installed, go to the next step.
4. Assuming Docker is installed, go to the [Fn Project home page on GitHub](#) to confirm that the installed version of Docker is at least the minimum version specified in the [Pre-requisites section](#).

If the installed version of Docker is not supported by Fn Project, you have to upgrade the version of Docker before proceeding. See the [Docker documentation](#) for information about upgrading Docker on your platform. If your platform is Oracle Linux, see [Oracle Container Runtime for Docker User's Guide](#).

When you have completed the steps in this topic, go on to [5. Install the Fn Project CLI](#).

### 5. Install the Fn Project CLI

Before using Oracle Functions, the Fn Project CLI must be installed in your development environment.

You can install the Fn Project CLI in a number of different ways according to your environment.

The instructions in this topic assume you have already completed the steps in [4. Install Docker for Use with Oracle Functions](#).

To install the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. Open the [README.md](#) file in the fnproject/cli repository on GitHub and follow the appropriate instructions for installing the Fn Project CLI in your development environment. As a convenient overview, the instructions are summarized below:
  - In a MacOS environment using Homebrew, install the Fn Project CLI by entering:

```
$ brew install fn
```

- In a Linux or MacOS environment, install the Fn Project CLI by entering:

```
$ curl -LSs https://raw.githubusercontent.com/fnproject/cli/master/install | sh
```

If prompted for a password, enter the superuser's password.

- In a Linux, MacOS, or Windows environment, install the Fn Project CLI by downloading the binary from the [Releases](#) page and running it.
3. In a terminal window, confirm that the CLI has been installed by entering:

```
$ fn version
```

Assuming the Fn Project CLI has been installed correctly, you'll see a message indicating the version of the CLI that has been installed.

When you have completed the steps in this topic, go on to [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#).

### 6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure

Before using Oracle Functions, you have to configure the Fn Project CLI to connect to your Oracle Cloud Infrastructure tenancy.

When the Fn Project CLI is initially installed, it's configured for a local development 'context'. To configure Fn Project CLI to connect to your Oracle Cloud Infrastructure tenancy instead, you have to create a new context. The context specifies Oracle Functions endpoints, the OCID of the compartment to which deployed functions will belong, and the address of the Docker registry to and from which to push and pull images.

## CHAPTER 16 Functions

---

You can define multiple contexts, each stored in a different context file in .yaml format. By default, the individual context files are stored in the `~/.fn/contexts` directory. The `~/.fn/config.yaml` file specifies which context file Fn Project uses.

To create a new context, you can create a new context file manually and edit the `~/.fn/config.yaml` file by hand to point to that file. Alternatively, you can use the Fn Project CLI to interactively create the new context file and instruct the Fn Project CLI to start using that file, as described below.

The instructions in this topic assume you have already completed the steps in [5. Install the Fn Project CLI](#).

To create a new context file using the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. In a terminal window, create the new Fn Project CLI context for Oracle Cloud Infrastructure by entering:

```
$ fn create context <my-context> --provider oracle
```

where `<my-context>` is a name of your choosing. For example:

```
$ fn create context johns-oci-context --provider oracle
```

3. Specify that the Fn Project CLI is to use the new context by entering:

```
$ fn use context <my-context>
```

where `<my-context>` is the name you specified in the previous step. For example:

```
$ fn use context johns-oci-context
```

4. Configure the new context with the OCID of the compartment that you want to own the deployed functions (you might have created a new compartment specifically for this purpose, see [Create Compartments to Own Network Resources and Oracle Functions Resources in the Tenancy, if they don't exist already](#)) by entering:

```
$ fn update context oracle.compartment-id <compartment-ocid>
```

For example:

```
$ fn update context oracle.compartment-id ocid1.compartment.oc1..aaaaaaaarvdfa72n...
```

5. Configure the new context with the api-url endpoint to use when calling the API by entering:

```
$ fn update context api-url <api-endpoint>
```

where `<api-endpoint>` is one of the endpoints in the list of Functions endpoints in [Functions API](#), in the format `https://functions.<region-identifier>.oci.oraclecloud.com`. Note that this is the preferred format. An older format (`https://functions.<region-identifier>.oraclecloud.com`) is still supported, but is not preferred. The `<region-identifier>` in `<api-endpoint>` is the identifier of the Oracle Cloud Infrastructure region in which you'll be creating and deploying functions. For example, `us-phoenix-1`.

For example:

```
$ fn update context api-url https://functions.us-phoenix-1.oci.oraclecloud.com
```

6. Configure the new context with the address of the Docker registry that you want to use with Oracle Functions by entering:

```
$ fn update context registry <region-code>.ocir.io/<tenancy-namespace>/<repo-name>
```

where:

- `<region-code>` is the code of the Oracle Cloud Infrastructure Registry region. For example, `phx` for Phoenix. See [Availability by Region Name and Region Code](#) for the list of region codes.  
Oracle recommends that the Docker registry you specify is in the same region as the subnet on which you intend functions to run.
- `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy in which to create repositories (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`).
- `<repo-name>` is a repository name to pre-pend to the names of functions that you deploy.

For example:

```
$ fn update context registry phx.ocir.io/ansh81vrulzp/acme-repo
```

7. (Optional) Verify the Fn Project CLI context you've created by viewing the context file. For example, by entering:

```
$ more ~/.fn/contexts/johns-oci-context.yaml

api-url: https://functions.us-phoenix-1.oci.oraclecloud.com
provider: oracle
registry: phx.ocir.io/ansh81vrulzp/acme-repo
```

When you have completed the steps in this topic, go on to [7. Set the Context for the Fn Project CLI Using the oracle.profile Parameter](#).

### 7. Set the Context for the Fn Project CLI Using the oracle.profile Parameter

Before using Oracle Functions, you have to configure the Fn Project CLI to use the new profile you added to the Oracle Cloud Infrastructure CLI configuration file `~/.oci/config` (see [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#)). The profile you added contains the credentials of the user account you'll be using to create and deploy functions.

Note that unless you specify otherwise, the Fn Project CLI will attempt to use a profile in the `~/.oci/config` file named `default`.

The instructions in this topic assume you have already completed the steps in [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#).

To configure the Fn Project CLI to use the profile you've created for use with Oracle Functions:

1. Log in to your development environment as a functions developer.
2. In a terminal window, configure the Fn Project CLI context with the name of the profile you've created for use with Oracle Functions by entering:

```
$ fn update context oracle.profile <profile-name>
```

For example:

```
$ fn update context oracle.profile john-oci-profile
```

When you have completed the steps in this topic, go on to [8. Generate an Auth Token to Enable Login to Oracle Cloud Infrastructure Registry](#).

### 8. Generate an Auth Token to Enable Login to Oracle Cloud Infrastructure Registry

Before using Oracle Functions, the user account you'll be using to create and deploy functions must have an Oracle Cloud Infrastructure auth token. You use the auth token as the password when logging Docker in to Oracle Cloud Infrastructure Registry.

The instructions in this topic assume you have already completed the steps in [7. Set the Context for the Fn Project CLI Using the `oracle.profile` Parameter](#).

If the user account already has an auth token, go straight on to [9. Start Docker](#). Otherwise, if the user account does not have an auth token, generate an auth token now.

To generate an auth token for the user account you'll be using to create and deploy functions:

1. Log in to the Console as a functions developer.
2. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
3. On the **Auth Tokens** page, click **Generate Token**.
4. In the **Generate Token** dialog:
  - a. Enter a meaningful description for the auth token. For example, `John's auth token for use with Oracle Functions`. Avoid entering confidential information.
  - b. Click **Generate Token**. The new auth token is displayed. For example, `6<! )N_____6MqX`.
5. Copy the auth token immediately to a secure location from where you can retrieve it later, because you won't see the auth token again in the Console.
6. Close the **Generate Token** dialog.

When you have completed the steps in this topic, go on to [9. Start Docker](#).

### 9. Start Docker

Before using Oracle Functions, Docker must be running in your development environment. If it is not running, you must start Docker before proceeding.

The instructions in this topic assume you have already completed the steps in [8. Generate an Auth Token to Enable Login to Oracle Cloud Infrastructure Registry](#).

To verify that Docker is running:

1. Log in to your development environment as a functions developer.
2. In a terminal window, launch the standard hello-world Docker image as a container to confirm that Docker is running by entering:

```
$ docker run hello-world
```

3. Do one of the following, depending on the message you see:
  - If you see an error message indicating that Docker is not running, you have to start the Docker daemon before proceeding. See the [Docker documentation](#) for information about starting Docker on your platform.
  - If you see an error message indicating that the network timed out while trying to connect and advising you to check your internet connection or whether you are behind a proxy, your development environment might be behind a corporate proxy server or firewall. In which case, you will probably need to set the `http_proxy`, `https_proxy`, and `no_proxy` environment variables. Ask your network administrator for advice.
  - If you see a message like the one shown below, Docker is already running and you can proceed:

```
Hello from Docker.
```

```
This message shows that your installation appears to be working correctly.
```

When you have completed the steps in this topic, go on to [10. Log in to Oracle Cloud Infrastructure Registry](#).

### 10. Log in to Oracle Cloud Infrastructure Registry

Before using Oracle Functions, you have to log Docker in to the Docker registry in which you are going to store your functions as Docker images. This is the Docker registry specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

You can store functions in public and private repositories in Oracle Cloud Infrastructure Registry, an Oracle-managed registry built on top of Oracle Cloud Infrastructure.

When you log Docker into a Docker registry, you have to provide the appropriate authentication details. For example, in the case of Oracle Cloud Infrastructure Registry, you have to provide the tenancy Object Storage namespace, the user name, and the user's auth token.

The instructions in this topic assume you have already completed the steps in [9. Start Docker](#).

To log Docker into Oracle Cloud Infrastructure Registry:

1. Log in to your development environment as a functions developer.
2. In a terminal window, log in to Oracle Cloud Infrastructure Registry by entering:

```
$ docker login <region-code>.ocir.io
```

where `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)). For example, `phx` for Phoenix. See [Availability by Region Name and Region Code](#) for the list of region codes.

For example:

```
$ docker login phx.ocir.io
```

3. When prompted for **Username**, enter the name of the user you will be using with Oracle Functions to create and deploy functions, in the format:

```
<tenancy-namespace>/<username>
```

where `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy in which to create repositories (as shown on the **Tenancy Information** page). For example, `ansh81vrulzp/jdoe@acme.com`.

Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`).

If your tenancy is federated with Oracle Identity Cloud Service, use the format `<tenancy-namespace>/oracleidentitycloudservice/<username>`.

You must have already generated an Oracle Cloud Infrastructure auth token for the user you specify (see [8. Generate an Auth Token to Enable Login to Oracle Cloud Infrastructure Registry](#)).

4. When prompted for **Password**, enter the user's Oracle Cloud Infrastructure auth token. Having entered the password, Docker might warn you that the password is stored unencrypted in the Docker configuration file. The warning includes a [link to the Docker documentation](#) where you can find out how to configure a credential helper. Oracle recommends you review the information in the [Docker documentation](#) and consider using an external credentials store for increased security.

When you have completed the steps in this topic, you have completed the configuration tasks for your client environment. Go on to [Verifying Your Configuration for Function Development](#) to confirm that the Fn Project CLI can communicate with the API endpoint.

### Verifying Your Configuration for Function Development

Before using Oracle Functions, it's a good idea to confirm that you have successfully completed:

- the tasks for [Configuring Your Tenancy for Function Development](#)
- the tasks for [Configuring Your Client Environment for Function Development](#)

If you have successfully completed the configuration tasks, the Fn Project CLI will be able to communicate with the API endpoint.

To confirm that the Fn Project CLI can communicate with the API endpoint:

1. Log in to your development environment as a functions developer.
2. In a terminal window, try and view a list of applications that have been defined in Oracle Functions by entering:

```
$ fn list apps
```

3. If you see either of the following, you can proceed to create and deploy functions because your system is configured correctly:
  - A message indicating that no applications have been found, which is expected if this is the first time the tenancy has been configured for Oracle Functions.
  - A list of applications that have already been created, which is expected if other users are already using the tenancy for functions development.
4. If you see an error message, it's likely that the Fn Project CLI cannot communicate with the API endpoint due to some incorrect configuration. Do the following:
  - Review the configuration tasks to confirm you completed them as instructed (see [Configuring Your Tenancy for Function Development](#) and [Configuring Your Client Environment for Function Development](#)).
  - Review the solutions for common problems (see [Troubleshooting Oracle Functions](#)).

### Using the Fn Project CLI with Oracle Functions

Oracle Functions is powered by the Fn Project open source engine. As a result, you can use the Fn Project CLI to perform create, read, update, and delete operations on Oracle Functions.

To enable you to use the Fn Project CLI with Oracle Functions, you perform a number of preparatory tasks. See [Configuring Your Client Environment for Function Development](#).

Most Fn Project CLI commands have a similar syntax:

```
fn [global options] <command> [command options] [subcommands] [arguments]
```

For example, to:

- list all the available applications, use the command `fn list apps`
- create an application, use a command like `fn create app acmeapp --annotation oracle.com/oci/subnetIds='["ocid1.subnet.oc1.phx.aaaaaaaacnh..."]'`
- invoke a function, use a command like `fn invoke helloworld-app helloworld-func`

## CHAPTER 16 Functions

---

- change the profile that the Fn Project CLI uses for its context, use a command like `fn update context oracle.profile john-oci-profile`

To see a complete list of Fn Project CLI commands, you can:

- Log in to your development environment as a functions developer and enter `fn --help` or `fn -h` in a terminal window.
- In a web browser, go to the [Fn Project CLI documentation](#).

To see detailed information about individual Fn Project CLI commands, you can:

- Log in to your development environment as a functions developer and enter `fn <command> [subcommand] --help` or `fn <command> [subcommand] -h` in a terminal window. For example:
  - `fn create --help`
  - `fn update app -h`
- In a web browser, go to the [Fn Project CLI documentation](#) and select the command from the list.

From time to time, new versions of the Fn Project CLI are released. To:

- See which version of the Fn Project CLI is currently installed and whether it is the most recent version, log in to your development environment as a functions developer and enter `fn version` in a terminal window. The Fn Project CLI version number is displayed. If a more recent version of the Fn Project CLI is available, the number of the latest available version is also displayed.
- Upgrade the Fn Project CLI to the most recent version, reinstall the Fn Project CLI by following the instructions in [5. Install the Fn Project CLI](#).

## Creating, Deploying, and Invoking a Helloworld Function

You can start off using Oracle Functions by using Fn Project CLI commands to:

- create a simple helloworld function written in java
- push the image to the Docker registry that's configured for Oracle Functions

- deploy the function to an application in Oracle Functions
- invoke the function



### Tip

If you aren't able to successfully complete one of the steps in this topic, review the solutions for common problems (see [Troubleshooting Oracle Functions](#)).

To get started with Oracle Functions:

1. Confirm that you have completed the prerequisite steps for using Oracle Functions, as described in [Preparing for Oracle Functions](#). Specifically, that you have:
  - set up your tenancy (see [Configuring Your Tenancy for Function Development](#))
  - set up your development environment (see [Configuring Your Client Environment for Function Development](#))
2. Log in to the Console as a functions developer.
3. Use the Console to create a new application in Oracle Functions:
  - a. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
  - b. Select the region you intend to use for Oracle Functions. Oracle recommends that you use the same region as the Docker registry specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
  - c. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).  
The **Applications** page shows the applications already defined in the compartment.
  - d. Click **Create Application** and specify:

- The name for the new application as helloworld-app.
  - The VCN and subnet (or subnets, up to a maximum of three) in which to run the function. For example, a VCN called acme-vcn-01 and a public subnet called Public Subnet IHsY:US-PHOENIX-AD-1). If a regional subnet has been defined, best practice is to select that subnet to make failover across availability domains simpler to implement. If a regional subnet has not been defined and you need to meet high availability requirements, select multiple subnets. Oracle recommends that subnets are in the same region as the Docker registry specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
- e. Click **Create**.
4. Log in to your development environment as a functions developer.
  5. In a terminal window, create a helloworld java function by entering:

```
$ fn init --runtime java helloworld-func
```

A directory called helloworld-func is created, containing:

- A function definition file called func.yaml, containing the minimum amount of information required to build and run the function. See the [Fn Project documentation](#) to find out about the additional parameters you can include in a func.yaml file.
  - A /src directory containing source files and directories for the helloworld function (including /src/main/java/com/example/fn/HelloFunction.java).
  - A Maven configuration file called pom.xml that specifies the project artifacts and dependencies required to compile the function from the source files.
6. Change directory to the newly created helloworld-func directory.
  7. Enter the following single Fn Project command to build the function and its dependencies as a Docker image called helloworld-func, push the image to the specified Docker registry, and deploy the function to Oracle Functions in the helloworld-app:

```
$ fn -v deploy --app helloworld-app
```

The `-v` option simply shows more detail about what Fn Project commands are doing (see [Using the Fn Project CLI with Oracle Functions](#)).

8. (Optional) Assuming the specified Docker registry is Oracle Cloud Infrastructure Registry, use the Console to confirm that the `helloworld-func` image has been pushed to Oracle Cloud Infrastructure Registry successfully:
  - a. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
  - b. Choose the registry's region.

You see all the repositories in the registry to which you have access. The image you pushed is in a new repository with a name constructed from:

    - the repository name in the address of the Docker registry in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#))
    - the name of the `helloworld-func` image

For example, the new repository might be called `acme-repo/helloworld-func`.
  - c. Click the name of the new repository. You see details of the `helloworld-func` image that's been pushed to Oracle Cloud Infrastructure Registry.
9. (Optional) Use the Console to confirm that the function has been deployed to Oracle Functions successfully:
  - a. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
  - b. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

The **Applications** page shows that an application called `helloworld-app` has been created.
  - c. Click the `helloworld-app` application to see the functions within it.

The **Functions** page shows that the `helloworld-func` function has been deployed to Oracle Functions.
10. In a terminal window, invoke the `helloworld-func` function by entering:

```
$ fn invoke helloworld-app helloworld-func
```

The 'Hello World !' message is displayed.

Congratulations! You've successfully created and deployed your first function to Oracle Functions!

## Creating Applications

You can create applications in Oracle Functions in readiness for deploying functions. An application need not contain any functions.

You can create applications using the Console, the Fn Project CLI, and the API.

### Using the Console

To create a new application in Oracle Functions using the Console:

1. Confirm that you have completed the prerequisite steps for using Oracle Functions, as described in [Preparing for Oracle Functions](#). Specifically, that you have:
  - set up your tenancy (see [Configuring Your Tenancy for Function Development](#))
  - set up your development environment (see [Configuring Your Client Environment for Function Development](#))
2. Log in to the Console as a functions developer.
3. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
4. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
5. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).  
The **Applications** page shows the applications already defined in the compartment.
6. Click **Create Application** and specify:

- A name for the new application (for example, acmeapp). Avoid entering confidential information.
  - The VCN and subnet (or subnets, up to a maximum of three) in which to run functions. For example, a VCN called acme-vcn-01 and a public subnet called Public Subnet IHSY:US-PHOENIX-AD-1). If a regional subnet has been defined, best practice is to select that subnet to make failover across availability domains simpler to implement. If a regional subnet has not been defined and you need to meet high availability requirements, select multiple subnets. Oracle recommends that the subnets are in the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
7. Click **Create**.
- The new application appears in the list of applications.

### Using Fn Project CLI Commands

To create a new application in Oracle Functions using the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. In a terminal window, create a new application by entering:

```
$ fn create app <app-name> --annotation oracle.com/oci/subnetIds='["<subnet-ocid>"]'
```

where:

- `<app-name>` is the name of the new application. Avoid entering confidential information.
- `<subnet-ocid>` is the OCID of the subnet (or subnets, up to a maximum of three) in which to run functions. If a regional subnet has been defined, best practice is to select that subnet to make failover across availability domains simpler to implement. If a regional subnet has not been defined and you need to meet high availability requirements, specify multiple subnets (enclose each OCID in double quotes separated by commas, in the format `'["<subnet-ocid>","<subnet-ocid>"]'`). Oracle recommends that the subnets are in the

same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

For example:

```
$ fn create app acmeapp --annotation oracle.com/oci/subnetIds='
["ocid1.subnet.oc1.phx.aaaaaaacnh..."]'
```

An application is created in Oracle Functions, in the tenancy and region implied by the subnet OCID and belonging to the compartment specified in the Fn Project CLI context file.

3. Verify that the new application has been created by entering:

```
$ fn list apps
```

For example:

```
$ fn list apps
acmeapp
```

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage applications:

- [CreateApplication](#)
- [DeleteApplication](#)
- [GetApplication](#)
- [UpdateApplication](#)

## Creating and Deploying Functions

You use Fn Project CLI commands to create and deploy functions to Oracle Functions.



### Tip

If you aren't able to successfully complete one of the steps in this topic, review the solutions for common problems (see [Troubleshooting Oracle Functions](#)).

## Using Fn Project CLI Commands

To create and deploy a function to Oracle Functions using Fn Project CLI commands:

1. Confirm that you have completed the prerequisite steps for using Oracle Functions, as described in [Preparing for Oracle Functions](#). Specifically, that you have:
  - set up your tenancy (see [Configuring Your Tenancy for Function Development](#))
  - set up your development environment (see [Configuring Your Client Environment for Function Development](#))
2. If the application to which you want to add the function doesn't yet exist in Oracle Functions, create it now using the Fn Project CLI or the Console. For example, you might create a new application called acmeapp. See [Creating Applications](#).
3. Log in to your development environment as a functions developer.
4. In a terminal window, change directory to the directory containing the function code.
5. Initialize the function by entering:

```
$ fn init --runtime <runtime-language> <function-name>
```

where:

- `<runtime-language>` is one of the supported runtime languages (currently go, java, node, and python are supported)
- `<function-name>` is the name to use as the function name. If you don't specify a function name, the name of the current directory (in lower case) is used. Avoid entering confidential information.

For example:

```
$ fn init --runtime java acme-func
```

A directory is created with the function name you specified, containing:

- A function definition file called `func.yaml`, containing the minimum amount of information required to build and run the function. See the [Fn Project documentation](#) to find out about the additional parameters you can include in a `func.yaml` file.
- A `/src` directory containing source files and directories.
- A Maven configuration file called `pom.xml` that specifies the project artifacts and dependencies required to compile the function from the source files.

Note that depending on the runtime language you specify, the `fn init` command might create an `/example` directory containing code for a `helloworld` application. As a matter of good practice, you'll probably want to delete the `/example` directory.

6. Change directory to the newly created directory.
7. Enter the following single Fn Project command to build the function and its dependencies as a Docker image, push the image to the specified Docker registry, and deploy the function to Oracle Functions:

```
$ fn -v deploy --app <app-name>
```

where `<app-name>` is the name of the application in Oracle Functions to which you want to add the function. For example:

```
$ fn -v deploy --app acmeapp
```

The `-v` option simply shows more detail about what Fn Project commands are doing (see [Using the Fn Project CLI with Oracle Functions](#)).

Note that you can build, push, and deploy the function using separate Fn Project commands, instead of the single `fn deploy` command.

8. (Optional) Assuming the specified Docker registry is Oracle Cloud Infrastructure Registry, use the Console to confirm that the image has been pushed to Oracle Cloud Infrastructure Registry successfully:
  - a. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.

- b. Choose the registry's region.

You see all the repositories in the registry to which you have access. The image you pushed is in a new private repository with a name constructed from:

- the repository name in the address of the Docker registry in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#))
- the name of the image you pushed

For example, the new repository might be called `acme-repo/acme-func`.

- c. Click the name of the new repository. You see details of the image that's been pushed to Oracle Cloud Infrastructure Registry
9. (Optional) Use the Console to confirm that the function has been deployed to Oracle Functions successfully:
    - a. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
    - b. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).  
The **Applications** page shows the applications in the compartment, including the one you specified in the `fn deploy` command.
    - c. Click the name of the application you specified in the `fn deploy` command to see the functions within it.  
The **Functions** page shows that the function has been deployed to Oracle Functions.

## Creating Functions from Existing Docker Images

You can create a new function definition in the Oracle Functions server in different ways:

- Using the Console or the Fn Project CLI command `fn create function` to create a new function based on an existing Docker image that has already been pushed to the Docker registry (as described in this topic).

## CHAPTER 16 Functions

---

- Using the single Fn Project CLI command `fn deploy` to build a new Docker image, push the image to the Docker registry, and create a new function based on the image in one step (as described in [Creating and Deploying Functions](#)).
- Using the API (see [CreateFunction](#)).

When using the Console or the `fn create function` command to create a new function based on an existing Docker image, you specify function metadata to store in the Oracle Functions server. For example, the maximum length of time the function is allowed to execute for.

The existing image on which you base a new function must be suitable for use with Oracle Functions. Typically, to build and push a suitable image, you or somebody else will use Fn Project CLI commands and/or Docker CLI commands. For example, having written your function code and a `func.yaml` file containing function metadata (perhaps based on the template `helloworld` function and `func.yaml` created using `fn init`), you can:

- Use `fn build` to build a new Docker image from the function.
- Use `docker push` to push the image to the Docker registry.

With the image in the Docker registry, you can then use the Console to create a function based on the image, as described in this topic.

### Using the Console

To use the Console to create a new function in the Oracle Functions server from an existing Docker image that has already been pushed to the Docker registry:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

The **Applications** page shows the applications defined in the compartment.

5. Click the name of the application in which you want to create the new function.
6. Click **Create Function** and specify:
  - **Name:** A name for the new function.
  - **Image:** The existing image in the Oracle Cloud Infrastructure Registry in your currently selected region. You first select the image repository, and then the image version.
  - **Memory:** The maximum amount of memory the function can use during execution.
  - **Timeout:** The maximum amount of time the function will be allowed to run for.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create** to create the new function in the Oracle Functions server.

The new function is shown in the Console, in the list of functions in the application you selected.

### Using Fn Project CLI Commands

To use the Fn Project CLI to create a new function in the Oracle Functions server from an existing Docker image that has already been pushed to the Docker registry:

1. Log in to your development environment as a functions developer.
2. In a terminal window, create a new function by entering:

```
$ fn create function <app-name> <function-name> <image-name>
```

where:

## CHAPTER 16 Functions

---

- `<app-name>` is the name of an existing application in which to create the new function.
- `<function-name>` is the name of the new function you want to create. Avoid entering confidential information.
- `<image-name>` is the name of the existing image in the Docker registry on which to base the new function.

For example:

```
$ fn create function acmeapp acme-func phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.3
```

A new function is created in Oracle Functions, based on the existing image and with the name you specified

3. Verify that the new function has been created by entering:

```
$ fn list functions <app-name>
```

For example:

```
$ fn list functions acme-app
```

NAME	IMAGE
acme-func	phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.3

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage functions:

- [CreateFunction](#)
- [DeleteFunction](#)
- [GetFunction](#)
- [UpdateFunction](#)

# Viewing Functions and Applications

Having deployed functions to Oracle Functions, you'll typically want to view the functions you've deployed, along with other functions in the same application and different applications. For example, you might want to see:

- all the applications in a compartment
- details of the image for a given function

You can view applications and functions using the Console, the Fn Project CLI, and the API.

## Using the Console

To view details of applications and functions deployed to Oracle Functions using the Console:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment containing the applications and functions that you want to see information about.

The **Applications** page shows all the applications in the compartment you selected.

5. Click the name of an application to see the functions within it.

The **Functions** page shows details for all the functions within the application you selected, including:

- the Docker image created for each function
  - when the function was last updated
6. Click the name of a function on the **Functions** page to see additional information about that function, including the values of timeout and memory configuration parameters.

### Using Fn Project CLI Commands

To view details of applications and functions deployed to Oracle Functions using the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. If you want to see details about applications, in a terminal window:
  - Enter the following command to see a simple list of applications:

```
$ fn list apps
```

For example:

```
$ fn list apps
```

```
acme-app
```

- Enter the following command to see more detail about a particular application:

```
$ fn inspect app <app-name>
```

For example:

```
$ fn inspect app acme-app
```

```
{
 "annotations": {
 "oracle.com/oci/appCode": "fht7ns4mn2q",
 "oracle.com/oci/compartmentId": "ocid1.compartment.oc1..aaaaaaaaw_____nyq",
 "oracle.com/oci/subnetIds": [
 "ocid1.subnet.oc1.phx.aaaaaaaao..."
],
 "oracle.com/oci/tenantId": "ocid1.tenancy.oc1..aaaaaaaap...keq"
 },
 "created_at": "2018-07-13T17:54:34.000Z",
 "id": "ocid1.fnapp.oc1.phx.aaaaaaaaf_____r3ca",
 "name": "acme-app",
 "updated_at": "2018-07-13T17:54:34.000Z"
}
```

3. If you want to see details about functions, in a terminal window:

- Enter the following command to see a simple list of functions in a particular application:

```
$ fn list functions <app-name>
```

For example:

```
$ fn list functions acme-app
```

NAME	IMAGE
acme-func	phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.3
acme-func-dev	phx.ocir.io/ansh81vrulzp/acme-repo/acme-func-dev:0.0.7
acme-func-test	phx.ocir.io/ansh81vrulzp/acme-repo/acme-func-test:0.0.6

- Enter the following command to see more detail about a particular function:

```
$ fn inspect function <app-name> <function-name>
```

For example:

```
$ fn inspect function acme-app acme-func
```

```
{
 "annotations": {
 "fnproject.io/fn/invokeEndpoint": "https://fht7ns4mn2q.us-phoenix-1.functions.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke",
 "oracle.com/oci/compartmentId": "ocid1.compartment.oc1..aaaaaaaaw____nyq"
 },
 "app_id": "ocid1.fnapp.oc1.phx.aaaaaaaaf____r3ca",
 "created_at": "2018-07-26T12:50:53.000Z",
 "format": "default",
 "id": "ocid1.fnfunc.oc1.phx.aaaa____uxoa",
 "image": "phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.3",
 "memory": 128,
 "name": "acme-func",
 "timeout": 30,
 "updated_at": "2018-07-26T13:59:18.000Z"
}
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to see details about applications and functions:

- [ListApplications](#)
- [ListFunctions](#)

### Invoking Functions

You can invoke a function that you've deployed to Oracle Functions in different ways:

- Using the Fn Project CLI.
- Using the Oracle Cloud Infrastructure CLI.
- Using the Oracle Cloud Infrastructure SDKs.
- Making a signed HTTP request to the function's invoke endpoint. Every function has an invoke endpoint.

Each of the above invokes the function via requests to the API. Any request to the API must be authenticated by including a signature and the OCID of the compartment to which the function belongs in the request header. Such a request is referred to as a 'signed' request. The signature includes Oracle Cloud Infrastructure credentials in an encrypted form.

If you use the Fn Project CLI or the Oracle Cloud Infrastructure CLI to invoke a function, authentication is handled for you. See [Using the Fn Project CLI to Invoke Functions](#) and [Using the Oracle Cloud Infrastructure CLI to Invoke Functions](#).

If you use an Oracle Cloud Infrastructure SDK to invoke a function, you can use the SDK to handle authentication. See [Using SDKs to Invoke Functions](#).

If you make a signed HTTP request to a function's invoke endpoint, you'll have to handle authentication yourself by including a signature and the OCID of the compartment to which the function belongs in the request header. You can do this in different ways:

- Using the Oracle Cloud Infrastructure CLI `raw-request` command. See [Sending a Signed Request to a Function's Invoke Endpoint \(using the Oracle Cloud Infrastructure CLI `raw-request` command\)](#).
- Using a bash script provided by Oracle (commonly referred to as `oci-curl`). See [Sending a Signed Request to a Function's Invoke Endpoint \(using `oci-curl`\)](#).
- Writing code to programmatically sign requests. For information about the required credentials and how to sign the requests, see [Request Signatures](#).



### Tip

If you aren't able to successfully complete one of the steps in this topic, review the solutions for common problems (see [Troubleshooting Oracle Functions](#)).

## Using the Fn Project CLI to Invoke Functions

To invoke a function deployed to Oracle Functions using the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. In a terminal window, enter:

```
$ fn invoke <app-name> <function-name>
```

where:

- `<app-name>` is the name of the application containing the function you want to invoke
- `<function-name>` is the name of the function you want to invoke

For example:

```
$ fn invoke helloworld-app helloworld-func
Hello World !
```



### Tip

If you want to pass arguments and values to a function, prefix the `fn invoke` command with `echo -n '<argument>=<value>' |`

If the function is expecting the argument and value as JSON, use a valid JSON format. For example:

```
$ echo -n '{"name":"John"}' | fn invoke helloworld-app
helloworld-func
Hello John !
```

## Using the Oracle Cloud Infrastructure CLI to Invoke Functions

If you have installed the Oracle Cloud Infrastructure CLI, you can use it to send API requests to invoke functions. Among other things, the Oracle Cloud Infrastructure CLI will facilitate Oracle Cloud Infrastructure authentication. For information about using the Oracle Cloud Infrastructure CLI, see [Command Line Interface \(CLI\)](#).

These instructions assume:

- you have already installed and configured the Oracle Cloud Infrastructure CLI
- you want to invoke a function as the functions developer that's configured for your development environment

To invoke a function using the Oracle Cloud Infrastructure CLI:

1. Log in to your development environment as a functions developer.
2. In a terminal window, enter:

```
$ oci fn function invoke <function-ocid> --file "<output-filepath>" --body "<request-parameters>"
```

where:

## CHAPTER 16 Functions

---

- `<function-ocid>` is the OCID of the function you want to invoke. To find out a function's OCID, use the `fn inspect` command to see the value of the function's `id` property (see [Viewing Functions and Applications](#)).
- `<output-filepath>` is the path and name of a file to write the response to. To write the response to stdout, specify `--file "-"`
- `<request-parameters>` are optionally arguments and values to pass to the function. If the function is expecting arguments and values as JSON, use a valid JSON format. For example, `--body '{"name": "John"}'`. Note that you must include `--body ""` in the request, even if there are no request parameters to pass.

For example:

- ```
$ oci fn function invoke --function-id ocid1.fnfunc.oc1.phx.aaaa____uxoa --file "-" --body ""  
Hello World !
```
- ```
$ oci fn function invoke --function-id ocid1.fnfunc.oc1.phx.aaaa____uxoa --file "-" --body '{"name": "John"}'
Hello John !
```

## Using SDKs to Invoke Functions

If you're writing a program to invoke a function in a language for which an Oracle Cloud Infrastructure SDK exists, Oracle recommends you use that SDK to send API requests to invoke the function. Among other things, the SDK will facilitate Oracle Cloud Infrastructure authentication.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [InvokeFunction](#) API operation to invoke functions.

### Obtaining a Function's Invoke Endpoint

When invoking a function using `oci-curl` or the Oracle Cloud Infrastructure CLI `raw-request` command, you have to specify the function's invoke endpoint.

To obtain a function's invoke endpoint:

1. Log in to your development environment as a functions developer.
2. In a terminal window, enter:

```
$ fn inspect function <app-name> <function-name>
```

where:

- `<app-name>` is the name of the application containing the function for which you want to obtain the invoke endpoint
- `<function-name>` is the name of the function for which you want to obtain the invoke endpoint

For example:

```
$ fn inspect function helloworld-app helloworld-func
{
 "annotations": {
 "fnproject.io/fn/invokeEndpoint": "https://fht7ns4mn2q.us-phoenix-
1.functions.oci.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa____
uxoa/actions/invoke",
 ...
 }
}
```

The function's invoke endpoint is the value of `"fnproject.io/fn/invokeEndpoint"`. For example, `"https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke"` (abbreviated for readability).

### Sending a Signed Request to a Function's Invoke Endpoint (using the Oracle Cloud Infrastructure CLI `raw-request` command)

If you have installed the Oracle Cloud Infrastructure CLI, you can use it to send API requests to invoke functions. Among other things, the CLI will facilitate Oracle Cloud Infrastructure authentication. For more information about using the Oracle Cloud Infrastructure CLI, see [Command Line Interface \(CLI\)](#).

These instructions assume:

- you have already installed and configured the Oracle Cloud Infrastructure CLI
- you want to invoke a function as the functions developer that's configured for your development environment

To invoke a function deployed to Oracle Functions by sending a signed request to the function's invoke endpoint using the Oracle Cloud Infrastructure CLI `raw-request` command:

1. Log in to your development environment as a functions developer.
2. Obtain the function's invoke endpoint (see [Obtaining a Function's Invoke Endpoint](#)).  
For example, "fnproject.io/fn/invokeEndpoint": "https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa\_\_\_\_uxoa/actions/invoke" (abbreviated for readability).
3. Use the Oracle Cloud Infrastructure CLI `raw-request` command to invoke the function by sending a signed POST request to the function's invoke endpoint by entering:

```
$ oci raw-request --http-method POST --target-uri <invoke-endpoint> --request-body "<request-parameters>"
```

where:

- `<invoke-endpoint>` is the endpoint you obtained in the earlier step.
- `<request-parameters>` are optionally arguments and values to pass to the function. If the function is expecting arguments and values as JSON, use a valid JSON format. Note that you must include `--request-body ""` in the request, even if there are no request parameters to pass.

For example:

- ```
$ oci raw-request --http-method POST --target-uri https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocidl.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke --request-body ""  
Hello World !
```
- ```
$ oci raw-request --http-method POST --target-uri https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocidl.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke --request-body '{"name":"John"}'
Hello John !
```

4. Assuming a passphrase was provided to encrypt the API signing key (as recommended by Oracle), enter the passphrase when prompted.

### Sending a Signed Request to a Function's Invoke Endpoint (using oci-curl)

When you followed the instructions to prepare your client environment for Oracle Functions, you installed and configured `oci-curl` in readiness for using it to invoke functions. Among other things, `oci-curl` will facilitate Oracle Cloud Infrastructure authentication.

These instructions assume:

- you have already configured `oci-curl` appropriately (see [7. Set the Context for the Fn Project CLI Using the `oracle.profile` Parameter](#))
- you want to invoke a function as the functions developer that's configured for your development environment

To invoke a function deployed to Oracle Functions by sending a signed request to the function's invoke endpoint using `oci-curl`:

1. Log in to your development environment as a functions developer.
2. Obtain the function's invoke endpoint (see [Obtaining a Function's Invoke Endpoint](#)).

For example, `"fnproject.io/fn/invokeEndpoint": "https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocidl.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke"` (abbreviated for readability).

3. In a terminal window, use the `source` command to set up the current shell environment for `oci-curl` by entering:

```
$ source <path-to-script>/oci-curl.sh
```

where `<path-to-script>` is the path to the location of the `oci-curl.sh` script (see [3. Create and Configure a Copy of oci-curl](#)). For example:

```
$ source ~/oci-curl/oci-curl.sh
```

4. In the same terminal window in which you entered the `source` command, use `oci-curl` to invoke the function by sending a signed POST request to the function's invoke endpoint by entering:

```
oci-curl "<invoke-endpoint-host>" post <filename> "<invoke-endpoint-path>"
```

where:

- `<invoke-endpoint-host>` is the first half of the endpoint you obtained in the earlier step, excluding `https://` and up to (and including) `.oraclecloud.com`. For example, `fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com`
- `<filename>` is the name of a file containing data to pass to the function (you must specify a file, even if it's empty) . For example, `payload.json`
- `<invoke-endpoint-path>` is the second half of the endpoint you obtained in the earlier step, from (but not including) `.oraclecloud.com` onwards. For example, `/20181201/functions/ocidl.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke`

For example, combining the previous examples, you might enter:

```
$ oci-curl "fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com" post payload.json
"/20181201/functions/ocidl.fnfunc.oc1.phx.aaaa____uxoa/actions/invoke"
```

5. Assuming a passphrase was provided to encrypt the API signing key (as recommended by Oracle), enter the passphrase when prompted.

## Storing and Viewing Function Logs

When a function you've deployed to Oracle Functions is invoked, you'll typically want to store the function's logs so that you can review them later. You specify where Oracle Functions

stores a function's logs by setting a logging policy for the application containing the function. You can specify that Oracle Functions:

- Stores logs in Oracle Cloud Infrastructure. Note that until an Oracle Cloud Infrastructure logging service is released, Oracle Functions stores logs as files in a storage bucket in Oracle Cloud Infrastructure Object Storage.
- Stores logs by exporting them to an external logging destination like Papertrail. Note that to use an external logging destination, you must have set up a VCN with public subnets and an internet gateway (see [Create the VCN and Subnets to Use with Oracle Functions, if they don't exist already](#)).

You set application logging policies in the Console.

### Using the Console

To store logs for the functions in an application:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment containing the application with functions for which you want to store the logs.

The **Applications** page shows all the applications in the compartment you selected.

5. To set the logging policy:
  - for a new application, click **Create Application** and set properties for the new application
  - for an existing application, click the name of the application, and then click **Edit Application**

By default, **Logging Policy** is set to **None**, so logs are not stored.

6. To store logs, select one of the following **Logging Policy** options:
  - Select the **Log to Object Storage** option to store logs as files in a storage bucket in Oracle Cloud Infrastructure Object Storage.
  - Select the **Syslog URL** option to store logs in an external logging destination, and enter the syslog URL to which to export the logs. For example,  
`tcp://my.papertrail.com:4242`
7. Click **Save**.

Whenever a function is invoked in this application, its logs are stored according to the logging policy that you specified.

To view the logs for a function that have been stored in a storage bucket in Oracle Cloud Infrastructure Object Storage:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment containing the application with the function for which you want to see the logs.  
The **Applications** page shows all the applications in the compartment you selected.
5. Click the name of the application containing the function for which you want to see the logs.
6. On the **Application Information** tab, click **Logs: Object Storage** to see the **Bucket Details** page in a separate browser tab. The **Bucket Details** page shows a list of files containing the logs for functions in the compartment you selected. Each log file name includes the OCID of the associated function.

If the storage bucket doesn't contain the log files you're expecting, or you see an error message saying that the storage bucket doesn't exist, note the following:

- A storage bucket is only created after the first log file has been written.
  - After a function has been invoked, it takes around 15 minutes for the function's log files to be written to Object Storage and appear in a storage bucket. So you won't see a storage bucket for a function until around 15 minutes after the function has been invoked for the first time.
  - If you have waited for longer than 15 minutes and you still don't see a storage bucket for a function, double-check that the function includes log statements.
7. (Optional) If a large number of log files are shown in the list, enter `log/<function-OCID>` in the **Search** field to reduce the list to a manageable size.
  8. Click the name of a log file to download it.

## Updating Functions

Having previously created a function definition in the Oracle Functions server, you can change some, but not all, of the function's properties. For example, you can change the maximum length of time a function is allowed to execute for, but you cannot change the function's name.

You can change the Docker image on which a function is based. If you do want to change the image, the replacement image must be suitable for use with Oracle Functions, and must have already been pushed to the Docker registry. With the replacement image in the Docker registry, you can then update a function's definition so that it is based on the replacement image, as described in this topic. If the replacement image has the same name and tag as the image on which the function was originally based, see [Notes About Image Digests](#).

You can update functions using the Console, the Fn Project CLI, and the API.

### Using the Console to update an existing function

To use the Console to update an existing function in the Oracle Functions server:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.

3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

The **Applications** page shows the applications defined in the compartment.

5. Click the name of the application containing the existing function that you want to update.
6. Click the name of the function that you want to update.
7. Click **Edit** and update some or all of the following properties:
  - **Image:** The existing image in the Oracle Cloud Infrastructure Registry in your currently selected region. You first select the image repository, and then the image version. If the image has the same name and tag as the image on which the function was originally based, see [Notes About Image Digests](#).
  - **Memory:** The maximum amount of memory the function can use during execution.
  - **Timeout:** The maximum amount of time the function will be allowed to run for.
8. Click **Save** to update the function in the Oracle Functions server.

The function's updated properties are shown in the Console.

### Using Fn Project CLI Commands

To use the Fn Project CLI to update an existing function in the Oracle Functions server:

1. Log in to your development environment as a functions developer.
2. In a terminal window, update properties of an existing function by entering:

```
$ fn update function <app-name> <function-name> <image-name> --<property> <value>
```

where:

## CHAPTER 16 Functions

- `<app-name>` is the name of an existing application containing the existing function.
- `<function-name>` is the name of the existing function you want to update.
- `<image-name>` (optionally) is the name of an existing image in the Docker registry that you now want to base the function on, instead of the previously specified image. If the image has the same name and tag as the image on which the function was originally based, see [Notes About Image Digests](#).
- `<property> <value>` (optionally) is the property you want to update, and the new value you want it to have. Enter `fn update function --help` to see a list of properties and valid values.

For example:

```
$ fn update function acmeapp acme-func phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.4 --
timeout 60
```

```
$ fn update function acmeapp acme-func --memory 256
```

The properties of the existing function are updated with the values you specified.

### 3. Verify that the function has been updated by entering:

```
$ fn inspect function <app-name> <function-name>
```

For example:

```
$ fn inspect function acme-app acme-func

{
 "annotations": {
 "fnproject.io/fn/invokeEndpoint": "https://fht7ns4mn2q.us-phoenix-
1.functions.oc1.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa____
uxoa/actions/invoke",
 "oracle.com/oci/compartmentId": "ocid1.compartment.oc1..aaaaaaaaw____nyq"
 },
 "app_id": "ocid1.fnapp.oc1.phx.aaaaaaaaaf____r3ca",
 "created_at": "2018-07-26T12:50:53.000Z",
 "format": "default",
 "id": "ocid1.fnfunc.oc1.phx.aaaa____uxoa",
 "image": "phx.ocir.io/ansh81vrulzp/acme-repo/acme-func:0.0.4",
 "memory": 256,
```

```
"name": "acme-func",
"timeout": 60,
"updated_at": "2018-07-26T13:59:18.000Z"
}
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [UpdateFunction](#) API operation to update functions.

### Notes About Image Digests

Images in a Docker registry are identified by repository, name, and a tag. In addition, Docker gives each version of an image a unique alphanumeric digest. When pushing an updated Docker image, it's recommended best practice to give the updated image a new tag to identify it, rather than reusing an existing tag. However, even if you push an updated image and give it the same name and tag as an earlier version, the newly pushed version will have a different digest to the earlier version.

When you create a function with Oracle Functions, you specify the name and tag of a particular version of an image on which to base the function. To avoid later inconsistencies, Oracle Functions also records the unique digest of that particular version of the image.

By default, if you push an updated version of an image to the Docker registry with the same name and tag as the original version of the image on which a function is based, Oracle Functions continues to use the original digest to pull the original version of the image. This might be the behavior you require. However, if you want Oracle Functions to pull the later version of the image, you can explicitly change the digest that Oracle Functions uses to identify which version of the image to pull in one of the following ways:

- Use the `fn update function` command and specify the original name and tag of the version of the image on which you want the function to be based. For example:

```
fn update function acmeapp acme-func phx.ocir.io/ansh81vrulzp/acme-
repo/acme-func:0.0.4
```

Oracle Functions will update the digest recorded for the image on which the function is based to be the digest of the image in the Docker registry that has the name and tag you specify.

- Use the `fn update function` command and specify the digest of the version of the image on which you want the function to be based. For example:

```
fn update function acmeapp acme-func --annotation
oracle.com/oci/imageDigest='"sha256:8af7cb8d7_____c498c0"'
```

Oracle Functions will update the digest recorded for the image on which the function is based to be the digest you specify.

- Use the Console and click **Edit Function** on the **Function Information** tab, re-select the original name and tag of the version of the image on which the function is currently based, and click **Save Changes**. Oracle Functions will update the digest recorded for the image on which the function is based.
- Use the Oracle Cloud Infrastructure API or an Oracle Cloud Infrastructure SDK (for more information, see [REST APIs](#) and [Software Development Kits and Command Line Interface](#)).

## Deleting Applications and Functions

You can delete applications and functions in Oracle Functions that you or other functions developers have created, provided you have been granted the necessary permission (FN\_APP\_DELETE or FN\_FUNCTION\_DELETE as appropriate).

Note the following:

- Deleting a function does not delete the Docker image on which the function is based. To delete the image, you have to delete it explicitly (see [Deleting an Image](#)).
- Deleting applications and functions is permanent. You cannot undelete an application or function that you've deleted.

- Deleting a function does not necessarily enable you to immediately delete the subnet and VCN in which the function runs. Expect to wait up to 30 minutes after the function was last invoked before you can delete the associated network resources.

You can delete applications and functions using the Console, the Fn Project CLI, and the API.

### Using the Console

When using the Console to delete applications and functions, note that:

- when you delete an application, all of its functions are also deleted
- you're always prompted to confirm deletion because you cannot undelete an application or function later

To delete applications and functions in Oracle Functions using the Console:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).  
The **Applications** page shows the applications defined in the compartment.
5. To delete an application, and all of its functions:
  - a. Click the name of the application you want to delete.
  - b. On the **Application Detail** page, click **Delete** and confirm you want to delete the application as follows:
    - If the application does not have functions within it, click **Delete** to confirm that you want to delete the application.

- If the application does have functions within it, you are shown a list of the functions in the application. To delete the application, enter `DELETE <APPLICATION-NAME>` in the text box, and click **Delete**.

Note that deleting an application and all of its functions does not delete the Docker images on which the functions are based. To delete the images, you have to delete them explicitly (see [Deleting an Image](#)).

6. To delete a function:
  - a. Click the name of the application containing the function you want to delete.
  - b. On the **Application Detail** page, click the name of the function you want to delete.
  - c. On the **Function Detail** page, click **Delete** and confirm you want to delete the function.

Note that deleting a function does not delete the Docker image on which the function is based. To delete the image, you have to delete it explicitly (see [Deleting an Image](#)).

## Using Fn Project CLI Commands

When using the Fn Project CLI to delete applications and functions, note that you cannot delete an application if it contains functions (you must delete the functions first).

To delete applications and functions in Oracle Functions using the Fn Project CLI:

1. Log in to your development environment as a functions developer.
2. To delete an application:
  - a. In a terminal window, enter:

```
$ fn delete app <app-name>
```

where `<app-name>` is the name of the application to delete.

For example:

```
$ fn delete app acmeapp
```

- b. Verify that the application has been deleted by entering:

```
$ fn list apps
```

3. To delete a function:

- a. In a terminal window, enter:

```
$ fn delete function <app-name> <function-name>
```

where:

- `<app-name>` is the name of the application containing the function you want to delete.
- `<function-name>` is the name of the function you want to delete.

For example:

```
$ fn delete function acmeapp acme-func
```

- b. Verify that the function has been deleted by entering:

```
$ fn list functions <app-name>
```

For example:

```
$ fn list functions acmeapp
```

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to delete applications and functions:

- [DeleteApplication](#)
- [DeleteFunction](#)

# Passing Custom Configuration Parameters to Functions

The code in functions you deploy to Oracle Functions will typically require values for different parameters. Some pre-defined parameters are available to your functions as environment variables. But you'll often want your functions to use parameters that you've defined yourself. For example, you might create a function that reads from and writes to a database. The function will require a database connect string, comprising a username, password, and hostname. You'll probably want to define username, password, and hostname as parameters that are passed to the function when it's invoked.

To pass user-defined parameters to a function deployed in Oracle Functions, you create key-value pairs known as custom configuration parameters. You can create custom configuration parameters that are:

- application-wide, meaning they are passed to every function in an application
- function-specific, meaning they are passed to the particular function for which they are defined (function-specific parameters override application-wide parameters with the same name)

To create custom configuration parameters, you can use:

- the `config:` section of a function's `func.yaml` file, to define function-specific custom configuration parameters
- the Console and the Fn Project CLI, to define both application-wide and function-specific custom configuration parameters

Oracle Functions combines all the custom configuration parameters (both application-wide and function-specific) in the application into a single, serially-encoded configuration object with a maximum allowable size of 4Kb.

### Using the Console

To specify custom configuration parameters to pass to functions using the Console:

1. Log in to the Console as a functions developer.
2. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
3. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
4. Select the compartment specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).  
The **Applications** page shows the applications defined in the compartment.
5. Click the name of the application containing functions to which you want to pass custom configuration parameters:
  - To pass one or more custom configuration parameters to every function in the application, click **Configuration** to see the **Configuration** section for the application.
  - To pass one or more custom configuration parameters to a particular function, click the function's name to see the **Configuration** section for the function.
6. In the **Configuration** section, specify details for the first custom configuration parameter:
  - **Key:** The name of the custom configuration parameter. The name must only contain alphanumeric characters and underscores, and must not start with a number. For example, `username`
  - **Value:** A value for the custom configuration parameter. The value must only contain printable unicode characters. For example, `jdoe`
7. Click the plus button to save the new custom configuration parameter.  
Oracle Functions combines the key-value pairs for all the custom configuration parameters (both application-wide and function-specific) in the application into a single, serially-encoded configuration object with a maximum allowable size of 4Kb. You

cannot save the new custom configuration parameter if the size of the serially-encoded configuration object would be greater than 4Kb.

8. (Optional) Enter additional custom configuration parameters as required.

### Using Fn Project CLI Commands

To specify custom configuration parameters to pass to functions using the Fn Project CLI:

1. Log in to your development environment as a functions developer and open a terminal window.
2. To specify one or more custom configuration parameters to pass to every function in an existing application, enter:

```
$ fn config app <app-name> <key> <value>
```

where:

- `<app-name>` is the name of the application containing the functions to which you want to pass the custom configuration parameter.
- `<key>` is the name of the custom configuration parameter. The name must only contain alphanumeric characters and underscores, and must not start with a number.
- `<value>` is the value to give to the custom configuration parameter. The value must only contain printable unicode characters.

For example:

```
$ fn config app acmeapp username jdoe
```

Note the following:

- You can also define application-wide custom configuration parameters when you create a new application using the `fn create app` command.
- Oracle Functions combines the key-value pairs for all the custom configuration parameters (both application-wide and function-specific) in the application into a single, serially-encoded configuration object with a maximum allowable size of 4Kb.

3. To specify one or more custom configuration parameters to pass to a particular function, enter:

```
$ fn config function <app-name> <function-name> <key> <value>
```

where:

- `<app-name>` is the name of the application containing the function to which you want to pass the custom configuration parameter.
- `<function-name>` is the name of the function to which to pass the custom configuration parameter.
- `<key>` is the name of the custom configuration parameter. The name must only contain alphanumeric characters and underscores, and must not start with a number.
- `<value>` is the value to give to the custom configuration parameter. The value must only contain printable unicode characters.

For example:

```
$ fn config function acmeapp acme-func username jdoe
```

Note the following:

- You can also define function-specific custom configuration parameters when you create a new function using the `fn create function` command.
- Oracle Functions combines the key-value pairs for all the custom configuration parameters (both application-wide and function-specific) in the application into a single, serially-encoded configuration object with a maximum allowable size of 4Kb.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to define custom configuration parameters:

- [CreateFunction](#)
- [UpdateFunction](#)
- [CreateApplication](#)
- [UpdateApplication](#)

### Accessing File Systems from Running Functions

A function you've deployed to Oracle Functions can access the file system of the container in which it's running as follows:

- the function can read files from all directories
- the function can write files to the /tmp directory

For example, you might want a function to download an Excel file and then read its contents. To meet this requirement, you might create a function that writes the file to the /tmp directory in the container's filesystem, and then subsequently reads the file.

When writing files to the /tmp directory, the /tmp directory is generally always writable. However, the maximum allowable size of the /tmp directory depends on the maximum memory threshold specified for the function:

Maximum memory threshold for the function (MB)	Maximum allowed size of /tmp (MB)	Maximum allowed number of files (inodes) in /tmp
<b>128MB</b>	32MB	1024
<b>256MB</b>	64MB	2048
<b>512MB</b>	128MB	4096
<b>1024MB</b>	256MB	8192

Note that the /tmp directory might be shared by multiple invocations of the function. A file written by an earlier invocation of a function could still exist when the function is invoked a second time. It is your responsibility to delete any files to avoid unexpected behavior.

## Accessing Other Oracle Cloud Infrastructure Resources from Running Functions

When a function you've deployed to Oracle Functions is running, it can access other Oracle Cloud Infrastructure resources. For example:

- You might want a function to get a list of VCNs from the Networking service.
- You might want a function to read data from an Object Storage bucket, perform some operation on the data, and then write the modified data back to the Object Storage bucket.

To enable a function to access another Oracle Cloud Infrastructure resource, you have to include the function in a dynamic group, and then create a policy to grant the dynamic group access to that resource. For more information about dynamic groups, including the permissions required to create them, see [Managing Dynamic Groups](#).

Having set up the policy and the dynamic group, you can then include a call to a 'resource principal provider' in your function code. The resource principal provider uses a resource provider session token (RPST) that enables the function to authenticate itself with other Oracle Cloud Infrastructure services. The token is only valid for the resources to which the dynamic group has been granted access.

Note also that the token is cached for 15 minutes. So if you change the policy or the dynamic group, you will have to wait for 15 minutes to see the effect of your changes.

Oracle recommends that you use the resource principal provider included in the Oracle Cloud Infrastructure SDK. However, you might be writing a function in a language that the Oracle Cloud Infrastructure SDK does not support. Or you might simply not want to use the Oracle Cloud Infrastructure SDK. In either case, you can write your own custom resource principal provider to enable a function to authenticate itself with other Oracle Cloud Infrastructure services, using files and environment variables in the container in which the function is executing.

## Using the Console

To enable a running function to access other Oracle Cloud Infrastructure resources:

1. Log in to the Console and create a new dynamic group:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Dynamic Groups**.
  - b. Follow the instructions in [To create a dynamic group](#), and give the dynamic group a name (for example, `acme-func-dyn-grp`).
  - c. When specifying a rule for the dynamic group, consider the following examples:

- If you want all functions in a compartment to be able to access a resource, enter a rule similar to the following that adds all functions in the compartment with the specified compartment OCID to the dynamic group:

```
ALL {resource.type = 'fnfunc', resource.compartment.id =
'ocidl.compartment.oc1..aaaaaaa23_____smwa'}
```

- If you want a specific function to be able to access a resource, enter a rule similar to the following that adds the function with the specified OCID to the dynamic group:

```
resource.id = 'ocidl.fnfunc.oc1.iad.aaaaaaaacq_____dnya'
```

- If you want all functions with a specific defined tag to be able to access a resource, enter a rule similar to the following that adds all functions with the defined tag to the dynamic group :

```
ALL {resource.type = 'fnfunc', tag.department.operations.value = '45'}
```

Note that free-form tags are not supported. For more information about tagging, see [Resource Tags](#).

- d. Click **Create Dynamic Group**.

Having created a dynamic group that includes the function, you can now create a policy to give the dynamic group access to the required Oracle Cloud Infrastructure resource.

2. Create a new policy:

- a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
- b. Follow the instructions in [To create a policy](#), and give the policy a name (for example, `acme-func-dyn-grp-policy`).
- c. When specifying a policy statement, consider the following examples:
  - If you want functions in the `acme-func-dyn-grp` to be able to get a list of all the VCNs in the tenancy, enter a rule similar to the following:

```
allow dynamic-group acme-func-dyn-grp to inspect vcns in tenancy
```

- If you want functions in the `acme-func-dyn-grp` to be able to read and write to a particular Object Storage bucket, enter a rule similar to the following:

```
allow dynamic-group acme-func-dyn-grp to manage objects in compartment acme-storage-compartment where all {target.bucket.name='acme-functions-bucket'}
```

- If you want functions in the `acme-func-dyn-grp` to be able to read and write to all resources in a compartment, enter a rule similar to the following:

```
allow dynamic-group acme-func-dyn-grp to manage all-resources in compartment acme-storage-compartment
```

- d. Click **Create** to create the new policy.
3. Include a resource principal provider in the function code to enable the function to authenticate with other Oracle Cloud Infrastructure services. See:
    - [Example: Adding the Oracle Resource Principal Provider to a Python Function to Get a List of VCNs from the Networking Service](#)
    - [Example: Adding a Custom Resource Principal Provider to a Function](#)

### Example: Adding the Oracle Resource Principal Provider to a Python Function to Get a List of VCNs from the Networking Service

Having added a function to a dynamic group, and created a policy that allows the dynamic group to list the VCNs in the tenancy, you could include code similar to the following example

to get a list of VCNs from the Networking service. This example uses the Oracle resource principal provider to extract credentials from the RPST token.

```
import io
import json

from fdk import response
import oci

def handler(ctx, data: io.BytesIO=None):
 signer = oci.auth.signers.get_resource_principals_signer()
 resp = do(signer)
 return response.Response(ctx,
 response_data=json.dumps(resp),
 headers={"Content-Type": "application/json"})

def do(signer):
 # List VCNs -----
 client = oci.core.VirtualNetworkClient({}, signer=signer)
 try:
 vcns = client.list_vcns(signer.compartment_id)
 vcns = [[v.id, v.display_name] for v in vcns.data]
 except Exception as e:
 vcns = str(e)
 return {"vcns": vcns, }
```

### Example: Adding a Custom Resource Principal Provider to a Function

Oracle recommends that you use the resource principal provider included in the Oracle Cloud Infrastructure SDK. However, you might be writing a function in a language that the Oracle Cloud Infrastructure SDK does not support. Or you might simply not want to use the Oracle Cloud Infrastructure SDK. In either case, you can write your own custom resource principal provider to enable a function to authenticate itself with other Oracle Cloud Infrastructure services, using files and environment variables in the container in which the function is executing.

The container in which a function executes includes a directory tree that holds Oracle Cloud Infrastructure compatible credentials, specifically:

- A resource principal session token (RPST) in a file named **rpst**. The RPST token is formatted as a [JWT token](#), and includes claims that identify the function's host tenancy and compartment.
- A private key for use in making requests to Oracle Cloud Infrastructure services on behalf of the function, in a file named **private.pem**.

The following environment variables are set inside the container in which the function executes:

- `OCI_RESOURCE_PRINCIPAL_VERSION`, containing the value `2.2`.
- `OCI_RESOURCE_PRINCIPAL_RPST`, containing the absolute path to the **rpst** file (including the filename).
- `OCI_RESOURCE_PRINCIPAL_PRIVATE_PEM`, containing the absolute path to the **private.pem** file (including the filename).
- `OCI_RESOURCE_PRINCIPAL_REGION`, containing the region identifier in which the function is deployed (for example, `us-phoenix-1`).

To enable a function to access another Oracle Cloud Infrastructure service, add code to the function so that it can authenticate itself with the other resource:

1. Add code that loads the RPST token from the path in the `OCI_RESOURCE_PRINCIPAL_RPST` environment variable.
2. Add code that loads the private key from the path in the `OCI_RESOURCE_PRINCIPAL_PRIVATE_PEM` environment variable.
3. Add code that uses the RPST token and the private key to create an Oracle Cloud Infrastructure request signature (see [Request Signatures](#)).
4. Add code that constructs the request to the other Oracle Cloud Infrastructure resource.

If necessary, you can identify:

- The endpoints of other Oracle Cloud Infrastructure services in the same (local) region as the function, using the region identifier in the `OCI_RESOURCE_PRINCIPAL_REGION` environment variable.

## CHAPTER 16 Functions

---

- The function's host tenancy and compartment, using the `res_tenant` and `res_compartment` claims in the RPST token.

For example, the sample Python function below includes a custom resource principal provider that extracts credentials from the RPST token. It then submits a GET request to the IAM API's `getTenancy` operation to return the OCID of the function's tenancy.

```
#!/usr/bin/env python3

import base64
import email.utils
import hashlib
import httpsig_cffi.sign
import json
import logging
import os.path
import re
import requests.auth
import urllib.parse

LOG = logging.getLogger(__name__)

The following class is derived from
https://docs.cloud.oracle.com/iaas/Content/API/Concepts/signingrequests.htm#Python

class SignedRequestAuth(requests.auth.AuthBase):
 """A requests auth instance that can be reused across requests"""
 generic_headers = [
 "date",
 "(request-target)",
 "host"
]
 body_headers = [
 "content-length",
 "content-type",
 "x-content-sha256",
]
 required_headers = {
 "get": generic_headers,
 "head": generic_headers,
```

```
"delete": generic_headers,
"put": generic_headers + body_headers,
"post": generic_headers + body_headers,
}

def __init__(self, key_id, private_key):
 # Build a httpsig_cffi.requests_auth.HTTPSignatureAuth for each
 # HTTP method's required headers
 self.signers = {}
 for method, headers in self.required_headers.items():
 signer = httpsig_cffi.sign.HeaderSigner(
 key_id=key_id, secret=private_key,
 algorithm="rsa-sha256", headers=headers[:])
 use_host = "host" in headers
 self.signers[method] = (signer, use_host)

def inject_missing_headers(self, request, sign_body):
 # Inject date, content-type, and host if missing
 request.headers.setdefault(
 "date", email.utils.formatdate(usegmt=True))
 request.headers.setdefault("content-type", "application/json")
 request.headers.setdefault(
 "host", urllib.parse.urlparse(request.url).netloc)

 # Requests with a body need to send content-type,
 # content-length, and x-content-sha256
 if sign_body:
 body = request.body or ""
 if "x-content-sha256" not in request.headers:
 m = hashlib.sha256(body.encode("utf-8"))
 base64digest = base64.b64encode(m.digest())
 base64string = base64digest.decode("utf-8")
 request.headers["x-content-sha256"] = base64string
 request.headers.setdefault("content-length", len(body))

def __call__(self, request):
 verb = request.method.lower()
 # nothing to sign for options
 if verb == "options":
 return request
 signer, use_host = self.signers.get(verb, (None, None))
 if signer is None:
```

## CHAPTER 16 Functions

---

```
 raise ValueError(
 "Don't know how to sign request verb {}".format(verb))

 # Inject body headers for put/post requests, date for all requests
 sign_body = verb in ["put", "post"]
 self.inject_missing_headers(request, sign_body=sign_body)

 if use_host:
 host = urllib.parse.urlparse(request.url).netloc
 else:
 host = None

 signed_headers = signer.sign(
 request.headers, host=host,
 method=request.method, path=request.path_url)
 request.headers.update(signed_headers)
 return request

def rp_auther():
 if os.environ['OCI_RESOURCE_PRINCIPAL_VERSION'] != "2.2":
 raise EnvironmentError('{} must be set to the value "2.2"'.format('OCI_RESOURCE_PRINCIPAL_
VERSION'))
 rpst = os.environ['OCI_RESOURCE_PRINCIPAL_RPST']
 if os.path.isabs(rpst):
 with open(rpst) as f:
 rpst = f.read()
 private_key = os.environ['OCI_RESOURCE_PRINCIPAL_PRIVATE_PEM']
 if os.path.isabs(private_key):
 with open(private_key) as f:
 private_key = f.read()
 return get_claims(rpst), SignedRequestAuth('ST${}'.format(rpst), private_key)

def get_claims(rpst):
 """Parse an RPST as a JWT; return a dictionary of claims

 The claims that are important are: sub, res_compartment, and res_tenant.
 These carry the resource OCID together with its location.
 """
 s = rpst.split('.')[1]
 s += "=" * ((4 - len(s) % 4) % 4) # Pad to a multiple of 4 characters
```

```
return json.loads(base64.b64decode(s).decode('utf-8'))

Use RP credentials to make a request
region = os.environ['OCI_RESOURCE_PRINCIPAL_REGION']
claims, rp_auth = rp_auther()

response = requests.get("https://identity.{}.oraclecloud.com/20160918/tenancies/{}".format(region,
claims['res_tenant']), auth=rp_auth)
print(response.json())
```

## Permissions Granted to Containers Running Functions

When a function you've deployed to Oracle Functions is invoked, it runs inside a container. The operations that a container can perform are determined by the user ID (UID) and group ID (GID) specified when the container is started. If a UID or GID is not specified, the container runs processes as the root user, with all the default capabilities enabled.

When starting a container to run a function, Oracle Functions always specifies a user named 'fn' with a UID of 1000, and a group name 'fn' with a GID of 1000. No privileges are granted to UID 1000 and GID 1000, so the container (and the function running inside it) does not acquire the default capabilities listed in the [Docker documentation](#). In addition, the container is prevented from gaining privileges.

As a result, do not create and deploy functions that:

- depend on capabilities that are unavailable
- depend on privilege elevation (for example, `su`, `sudo` or `setuid`)

If you are using your own Dockerfile, include the following lines:

```
groupadd --gid 1000 fn && \
adduser --uid 1000 --gid fn fn
```

For example:

```
FROM oraclelinux:7-slim

RUN yum -y install oracle-release-el7 oracle-nodejs-release-el7 && \
 yum-config-manager --disable ol7_developer_EPEL && \
```

```
yum -y install oracle-instantclient19.3-basiclite nodejs && \
rm -rf /var/cache/yum && \
groupadd --gid 1000 fn && \
adduser --uid 1000 --gid fn fn

WORKDIR /function
ADD . /function/
RUN npm install

CMD exec node func.js
```

Note that if you do not include the `groupadd` and `adduser` lines in the above example Dockerfile, you will see the following error message:

```
cx_Oracle.DatabaseError: ORA-12560: TNS:protocol adapter error
```

## Invoking Oracle Functions from Other Oracle Cloud Infrastructure Services

You can invoke functions in Oracle Functions from other Oracle Cloud Infrastructure services. Typically, you'll want an event in another service to trigger a request to invoke a function defined in Oracle Functions.

This functionality is currently available in the Events service. For more information, see [Overview of Events](#).

## Changing Oracle Functions Default Behavior

You can change several aspects of Oracle Functions default behavior using configuration parameters and environment variables.

Depending on the parameter, you can override a default value by specifying an alternative value in the following ways (note the order of precedence):

- by adding an entry to the `func.yaml` file (which overrides default values)
- by explicitly setting an environment variable (which overrides values set in the `func.yaml` file)

## CHAPTER 16 Functions

- by including a command option when you invoke the function using the Fn Project CLI (which overrides values set in environment variables or in the func.yaml file)

The following table indicates the parameters you can set, the default value, and where the default value can be overridden.

Parameter Description	Default Value	Units	func.yaml Parameter	Environment Variable	Fn CLI option	Notes
<b>Maximum time a function will be allowed to run</b>	30	Seconds	timeout	n/a	-- timeout	Maximum value: 120
<b>Maximum memory threshold for a function</b>	128	MB	memory	FN_MEMORY	-- memory	One of: <ul style="list-style-type: none"><li>• 128</li><li>• 256</li><li>• 512</li><li>• 1024</li></ul> If this limit is exceeded during execution, the function is stopped and an error message is logged.

For more information about the above parameters, and other configuration parameters, see [Func files](#) in the [Fn Project documentation](#).

## Differences between Oracle Functions and Fn Project

In general, Oracle Functions and Fn Project are very similar. However there are some differences, as detailed below.

### Differences in Authentication When Making API Calls

When you use the Oracle Cloud Infrastructure API with Oracle Functions, in the request header you have to provide:

- the OCID of the compartment to which the function belongs
- Oracle Cloud Infrastructure authentication details

### Differences When Invoking Functions

To invoke a function deployed to Oracle Functions, you have to explicitly specify an Oracle Cloud Infrastructure endpoint (unless you're using the Fn Project CLI).

For example, when you use `oci-curl` to invoke a function, you have to send a request to the function's invoke endpoint (for example `https://fht7ns4mn2q.us-phoenix-1.functions.oci.oraclecloud.com/20181201/functions/ocid1.fnfunc.oc1.phx.aaaa__uxoa/actions/invoke`).

You can obtain the appropriate endpoint by making a call to the API, either directly or by using the Fn Project CLI command:

```
$ fn inspect function <app-name> <function-name>
```

### Additional Context Configuration Parameters in Oracle Functions

As well as supporting Fn Project context configuration parameters, Oracle Functions also has some additional parameters, as shown in the following table.

## CHAPTER 16 Functions

Additional Parameter	Set in	Value	Notes
provider	A context configuration .yaml file in ~/.fn/contexts	oracle	<p>Enables Oracle Functions rather than Fn Project functionality. When provider is set to oracle, the following parameters are valid:</p> <ul style="list-style-type: none"> <li>• oracle.compartment-id</li> <li>• oracle.profile</li> </ul> <p>See <a href="#">6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure</a>.</p>
oracle.compartment-id	A context configuration .yaml file in ~/.fn/contexts	<compartment-ocid>	<p>Specifies the OCID of the Oracle Cloud Infrastructure compartment that owns function-related resources.</p> <p>See <a href="#">6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure</a>.</p>
oracle.profile	A context configuration .yaml file in ~/.fn/contexts	<profile-name>	<p>Specifies which profile to use from the ~/.oci/config file. If not set, the profile named default is used.</p> <p>See <a href="#">7. Set the Context for the Fn Project CLI Using the oracle.profile Parameter</a></p>

### Use of Annotations

When you're creating and viewing Oracle Functions resources using the Fn Project CLI, annotations enable you to identify and specify associated Oracle Cloud Infrastructure resources.

For example:

- When you're using the Fn Project CLI to create a new application, you use the `--annotation` parameter to specify the OCID of the subnet in which to run the function.
- When you're using the Fn Project CLI to view the properties of a function, the `annotations` element shows the OCID of the compartment that owns the function.

Note that unlike other configuration parameters and environment variables, annotation values cannot be passed as arguments to running Docker containers.

### Troubleshooting Oracle Functions

This topic covers common issues related to Oracle Functions and how you can address them.

#### Using `DEBUG=1` to see more details about an error

If you encounter an unexpected error when using an Fn Project CLI command, you can find out more about the problem by starting the command with the string `DEBUG=1` and running the command again. For example:

```
$ DEBUG=1 fn invoke helloworld-app helloworld-func
```

Note that `DEBUG=1` must appear before the command, and that `DEBUG` must be in upper case.

#### Using `--display-call-id` when invoking functions to aid issue resolution

If you encounter an issue when invoking a function, you can engage with Oracle Support. Oracle Support can investigate the issue more efficiently if you provide the call id of the function invocation. You can obtain the call id using the `--display-call-id` command option. For example:

## CHAPTER 16 Functions

---

```
$ fn invoke helloworld-app helloworld-func --display-call-id
```

```
Call ID: 01CS23SDG71BT2N9GZJ002DQM5
```

```
Hello World !
```

### Creating a new application displays an error message in the New Application dialog

If you've already reached the limit for the number of applications in your tenancy, you might see a message similar to the following in the **New Application** dialog when trying to create a new application:

#### **Unable to create your app, please try again.**

Double-check how many applications already exist in your tenancy. Compare that with the number of applications you're allowed to create. See [Oracle Functions Capabilities and Limits](#).

If you've exceeded the number of applications allowed in your tenancy, consider:

- Deleting unwanted applications (see [Deleting Applications and Functions](#)).
- Requesting an increase to the application limit (see [Service Limits](#) for instructions).

### Oracle Functions attempts to interact with docker.io

If you see a message similar to the following when deploying a function, double-check that your development environment doesn't have the FN\_REGISTRY environment variable set to your Docker username:

```
The push refers to repository [docker.io. ...
. . .
denied: requested access to the resource is denied
Fn: error running docker push, are you logged into docker?: exit status 1
See fn <command> --help' for more information.
```

## CHAPTER 16 Functions

---

If you have used the open source Fn Project platform, you might have followed instructions in the [Fn Project documentation](#) to set the `FN_REGISTRY` environment variable to your Docker username to enable interaction with the official Docker registry.

The `FN_REGISTRY` environment variable overrides the value of the `registry` option in your Fn Project CLI context.

To use the Fn Project CLI with Oracle Functions, do one of the following:

- Unset the `FN_REGISTRY` environment variable.
- Override the `FN_REGISTRY` environment variable using the `--registry` global option whenever you enter an Fn Project CLI command that interacts with Oracle Cloud Infrastructure Registry.

### Running `fn version` shows that a more recent version of the Fn Project CLI is available

If you see a message similar to the following when you enter the `fn version` command, a more recent version of the Fn Project CLI is available:

```
$ fn version

Client version: 0.5.33 is not latest: 0.5.34
Server version: ?
```

To upgrade the Fn Project CLI to the most recent version, reinstall the Fn Project CLI by following the instructions in [5. Install the Fn Project CLI](#).

### Deploying a function to Oracle Functions returns "Fn: Missing subnets annotation" message

When you deploy a function to Oracle Functions, you might see the following message:

```
$ fn deploy --app joes-helloworld-app
Deploying helloworld-func to app: joes-helloworld-app
.
.
.
Fn: Missing subnets annotation
```

## CHAPTER 16 Functions

---

If you see the `Fn: Missing subnets` annotation message, confirm that you entered the correct application name. For example:

- the application might not be in the compartment currently specified by the Fn Project CLI context
- the application might have existed previously, but has subsequently been deleted

### Running Fn Project CLI commands returns a 401 error

If you see a message similar to the following when running an Fn Project CLI command, double-check that the credentials specified for your current profile in the `~/.oci/config` file are authenticating you correctly:

```
$ fn list apps
Fn: [GET /apps][401] ListApps default &{Fields: Message:Not authenticated}
```

For example:

- Does `user` specify the OCID of your Oracle Cloud Infrastructure user account?
- Does `fingerprint` specify the fingerprint of the public API key value uploaded to the Console?
- Does `key_file` specify the full path to the private key file?

See [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#). Also see [API Errors](#).

### Running Fn Project CLI commands returns a 404 error

If you see a message similar to the following when running an Fn Project CLI command, double-check that you are authorized to access function-related and network resources:

```
$ fn list apps
Fn: [GET /apps][404] ListApps default &{Fields: Message:Resource is not authorized or not found}
```

For example:

- Does `oracle.compartment-id` in your current context correctly specify the OCID of the compartment that owns deployed functions?
- Have policies been set up correctly to give group access to function-related and network resources?
- Does your user account belong to the group to which access to function-related and network resources has been granted?
- Has a policy been set up to give Oracle Functions access to network resources?

See [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#) and [Create Policies to Control Access to Network and Function-Related Resources](#). Also see [API Errors](#).

### Invoking a function returns a `FunctionInvokeSyslogUnavailable` message and a 502 error

Oracle Functions enables you to export a function's logs to an external logging destination (like Papertrail) by setting a syslog URL for the application. See [Storing and Viewing Function Logs](#).

If the syslog URL is invalid or unreachable, you will see the following error when you invoke the function:

```
{"code":"FunctionInvokeSyslogUnavailable","message":"Syslog endpoint unavailable"}
```

```
Fn: Error invoking function. status: 502 message: Syslog endpoint unavailable
```

To confirm that the external logging destination's URL is the cause of the error:

1. Update the application to unset the syslog URL. For example:
  - Using the Console, click **Edit Application** on the **Application Details** page and select **None** from the **Logging Policy** options.
  - Using the Fn Project CLI, enter `fn update app helloworld-app --syslog-url ""`
2. Deploy the function you want to run. See [Creating and Deploying Functions](#).
3. Invoke the function. See [Invoking Functions](#).

If the function runs successfully, the external logging destination's URL is not reachable from the subnet in which the function is running. Double-check that:

- The external logging destination's URL is valid.
- The external logging destination's URL is publicly accessible.
- The subnet in which the function is running has outbound access to the public internet.

### Invoking a function returns a `FunctionInvokeImageNotAvailable` message and a 502 error

When you invoke a function, Oracle Functions pulls the corresponding image from Oracle Cloud Infrastructure Registry using the VCN and subnets specified for the application.

If Oracle Functions is unable to pull the image, the following message is returned when you invoke a function:

```
{"code":"FunctionInvokeImageNotAvailable","message":"Failed to pull function image"}
```

```
Fn: Error invoking function. status: 502 message: Failed to pull function image
```

Possible solutions:

- Double-check that the image specified for the function still exists in the specified location in Oracle Cloud Infrastructure Registry.
- Double-check that Oracle Cloud Infrastructure is available (this message is returned if Oracle Cloud Infrastructure is unexpectedly unavailable).
- Double-check that the VCN includes an internet gateway or service gateway. For Oracle Functions to be able to access Oracle Cloud Infrastructure Registry to pull an image, the VCN must include an internet gateway or a service gateway, as follows:
  - If public subnets were specified for the application, the VCN must also include an internet gateway.
  - If private subnets were specified for the application, the VCN must also include a service gateway.

If an internet gateway or service gateway has not been defined for the VCN already, define one now.

### Invoking a function returns a FunctionInvokeSubnetOutOfIPs message and a 502 error

When you invoke a function that you've deployed to Oracle Functions, you might see the following error message:

```
{"code":"FunctionInvokeSubnetOutOfIPs","message":"subnet ocid1.subnet.oc1.phx.aaaaaaaac... is out of IPs"}
Fn: Error invoking function. status: 502 message: subnet ocid1.subnet.oc1.phx.aaaaaaaac... is out of IPs
```

If you see this error, double-check that each subnet in the VCN has at least the required minimum number of free IP addresses specified in [Create the VCN and Subnets to Use with Oracle Functions, if they don't exist already](#).

### Invoking a function returns a FunctionInvokeSubnetNotAvailable message and a 502 error

When you invoke a function that you've deployed to Oracle Functions, you might see the following error message:

```
{"code":"FunctionInvokeSubnetNotAvailable","message":"subnet ocid1.subnet.oc1.phx.aaaaaaaac... does not exist or Oracle Functions is not authorized to use it"}
Fn: Error invoking function. status: 502 message: subnet ocid1.subnet.oc1.phx.aaaaaaaac... does not exist or Oracle Functions is not authorized to use it
```

If you see this error:

- Double-check that a policy has been created to give Oracle Functions access to network resources. See [Create a Policy to Give the Oracle Functions Service Access to Network Resources](#).
- Double-check that the subnet specified for the application still exists.

### Invoking a function returns a Gateway Time-out message and a 504 error

When you invoke a function that you've deployed to Oracle Functions, you might see the following error message:

```
<html>
<head><title>504 Gateway Time-out</title></head>
<body bgcolor="white">
<center><h1>504 Gateway Time-out</h1></center>
<hr><center></center>
</body>
</html>
```

If you see this error, it's likely that a policy has not been created to give Oracle Functions access to network resources. See [Create a Policy to Give the Oracle Functions Service Access to Network Resources](#).

### Invoking a function returns `FunctionInvokeContainerInitFail` and 'Container initialization timed out' messages, and a 504 error

When you invoke a function that you've deployed to Oracle Functions, the function execution is subject to a maximum memory threshold. If this limit is exceeded, function execution stops and the following error message is returned:

```
{"code":"FunctionInvokeContainerInitFail","message":"Container failed to initialize, please ensure you are using the latest fdk and check the logs"}
```

```
Fn: Error invoking function. status: 504 message: Container failed to initialize, please ensure you are using the latest fdk and check the logs
```

If you see this error, increase the maximum memory threshold when you invoke the function. Valid values for the maximum memory threshold are 128MB, 256MB, 512MB, and 1024MB (see [Changing Oracle Functions Default Behavior](#)).

For example, to set a function's maximum memory threshold to 256MB, do one of the following:

- Click **Edit Function** on the **Function Details** page in the Console, and select **256** from the **Memory (in MBs)** drop-down list.
- Use the following syntax when invoking the function using the Fn Project CLI. This will set the maximum memory threshold to 256MB for the current function invocation:

```
$ fn invoke <app-name> <function-name> --memory 256
```

- Add the following line to the function's `func.yaml` file. This will set the maximum

memory threshold to 256MB whenever the function is invoked:

```
memory: 256
```

Note that if you edit the `func.yaml` file, you must re-deploy the function to Oracle Functions before invoking it again.

It's a good idea to use the latest version of the Fn Project CLI when creating a helloworld Python function. When you enter the `fn init --runtime python <function-name>` command to create the helloworld function, the line `memory: 256` is added to the `func.yaml` file automatically.

### Invoking a function returns a `FunctionInvokeTimeout` message and a 504 error

When you invoke a function that you've deployed to Oracle Functions, the function is only allowed to run for a certain amount of time. If this time limit is exceeded, function execution stops and the following error message is returned:

```
{"code":"FunctionInvokeTimeout","message":"Timed out"}
```

```
Fn: Error invoking function. status: 504 message: Timed out
```

If you see this error, increase the maximum time a function is allowed to run for. For example, to set the maximum time to 120 seconds, do one of the following:

- Click **Edit Function** on the **Function Details** page in the Console, and enter **120** in the **Timeout** field.
- Use the following syntax when invoking the function using the Fn Project CLI. This will set the time limit to 120 seconds for the current function invocation:

```
$ fn invoke <app-name> <function-name> --timeout 120
```

- Add the following line to the function's `func.yaml` file. This will set the maximum time limit to 120 seconds whenever the function is invoked:

```
timeout: 120
```

Note that if you edit the `func.yaml` file, you must re-deploy the function to Oracle Functions before invoking it again.

### Running Fn Project CLI commands returns an X509: decryption password incorrect error

If you see a message similar to the following when running an Fn Project CLI command, double-check that the `pass_phrase` specified for your current profile in the `~/.oci/config` file is correct:

```
$ fn list apps
Fn: x509: decryption password incorrect
```

See [2. Create a Profile in the Oracle Cloud Infrastructure CLI Configuration File](#).

### Deploying an application returns an "unauthorized: incorrect username or password" message

When deploying an application, you might see a message similar to the following:

```
$ fn -v deploy --app acme-app

Deploying go-app to app: acme-app
Bumped to version 0.0.2
Building image phx.ocir.io/ansh81vrulzp/acme-repo/go-app:0.0.2
FN_REGISTRY: phx.ocir.io/ansh81vrulzp/acme-repo
Current Context: acme-functions-compartment
Sending build context to Docker daemon 5.12kB
Step 1/10 : FROM fnproject/go:dev as build-stage
Get https://registry-1.docker.io/v2/fnproject/go/manifests/dev: unauthorized: incorrect username or password
```

The message indicates an unnecessary and unsuccessful attempt to log in to Docker Hub. To resolve this situation, log out from Docker using the following command:

```
$ docker logout
```

Having logged out from Docker, re-run the command to deploy the application.

### When running Oracle Functions on Ubuntu, Docker login returns an "error getting credentials - err: exit status 1..." message

When you configure your development environment for Oracle Functions, you have to install Docker (see [4. Install Docker for Use with Oracle Functions](#)). If your development environment is running Ubuntu, when you follow the subsequent instructions to log in to Oracle Cloud Infrastructure Registry using Docker (see [10. Log in to Oracle Cloud Infrastructure Registry](#)), you might see a message similar to the following:

```
error getting credentials - err: exit status 1, out: Error spawning command line 'dbus-launch --autolaunch=d7159335070ef1c0854c75de55c8f588 --binary-syntax --close-stderr': Child process exited with code 1
```

For more information about this Docker issue, including likely causes and possible resolutions, see <https://github.com/docker/docker-credential-helpers/issues/60>.

### Deploying a function returns a ListTriggers message and a 500 error

When deploying a function that you've previously created using an earlier version of the Fn Project CLI, you might see a message similar to the following:

```
Fn: [GET /triggers][500] ListTriggers default &{Fields: Message:Internal server error}
```

This message indicates that the function's `func.yaml` file contains one or more HTTP trigger definitions. Oracle Functions does not currently support HTTP triggers. To deploy the function, remove the `triggers:` section from the `func.yaml` file.

To avoid creating new `func.yaml` files containing trigger definitions, follow the instructions in [5. Install the Fn Project CLI](#) to upgrade the Fn Project CLI to the most recent version.

### Performing Docker-related operations with the Fn Project CLI displays an "Error response from daemon... unknown: Unauthorized" message

To enable the Fn Project CLI to access the Docker registry specified in the Fn Project CLI context, the local Docker client (the Docker daemon on Linux) in your development environment must be logged in to that Docker registry. If the Docker client is not logged in to the Docker registry, you see a message similar to the following:

```
Error response from daemon: Get https://phx.ocir.io/v2/: unknown: Unauthorized
```

Follow the instructions in [10. Log in to Oracle Cloud Infrastructure Registry](#) to log the Docker client in to the appropriate Oracle Cloud Infrastructure Registry, an Oracle-managed Docker registry available in a number of different regions.

## Function Metrics

You can monitor the health, capacity, and performance of functions you've deployed to Oracle Functions by using metrics, alarms, and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_faas` (the Oracle Functions service).

Resources: functions

## Overview of the Oracle Functions Service Metrics

Oracle Functions monitors function execution, and collects and reports metrics such as:

- The number of times a function is invoked.
- The length of time a function runs for.
- The number of times a function failed.
- The number of requests to invoke a function that returned a '429 Too Many Requests' error in the response (known as 'throttled function invocations').

## Prerequisites

IAM policies: To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't

have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

For more information about the policy statement required to access metrics emitted by Oracle Functions, see [Create a Policy to Give Oracle Functions Users Access to Function-Related Resources](#).

### Available Metrics: oci\_faas

The metrics listed in the following tables are automatically available for any functions you create. You do not need to enable monitoring on the resource to get these metrics.

Oracle Functions metrics include the following dimensions:

**APPLICATIONID**

The OCID of the application containing functions.

**RESOURCEID**

The OCID of the function.

**RESPONSETYPE**

The response when a function is invoked (one of Success, Error, or Throttled).

## CHAPTER 16 Functions

Metric	Metric Display Name	Unit	Description	Dimensions
FunctionExecutionDuration	<b>Function Duration</b>	ms	Total function execution duration. Expressed in milliseconds.	applicationId resourceId
FunctionInvocationCount	<b>Function Invocations</b>	count	Total number of function invocations.	applicationId resourceId
FunctionResponseCount	This metric is used in the following default metric charts:  <b>Errors</b> (with <code>responseType = "Error"</code> )  <b>Throttles</b> (with <code>responseType = "Throttled"</code> )	count	Total number of function responses.	applicationId resourceId responseType

### Using the Console

#### To view default metric charts for a single function

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.
2. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).
3. Select the compartment containing the application with functions for which you want to view metrics.

The **Applications** page shows all the applications in the compartment you selected.

4. Click the name of the application containing the function for which you want to view metrics.
5. Click the name of the function for which you want to view metrics.
6. Under **Resources**, click **Metrics**.

The Metrics page displays a chart for each metric that is emitted by the metric namespace for Oracle Functions. For more information about the emitted metrics, see [Available Metrics: oci\\_faas](#).

#### Not seeing the function metrics data you expect?

If you don't see the metrics data for a function that you expect, see the following possible causes and resolutions.

## CHAPTER 16 Functions

---

Problem	Possible Cause	Resolution
Missing functions: A function I invoked is missing from the <b>Invocations</b> chart.	The chart range (time period or x-axis window) does not cover the time of invocation.	Adjust the chart range or time period as necessary.
Gaps in metrics data: The chart line is discontinuous. I want to see data in the charts as a continuous line over time, but the line has gaps in it.	No metrics data exist in the times indicated by the gaps.	Smooth out the display by increasing the chart interval to see if gaps are removed.

Problem	Possible Cause	Resolution
Empty charts: The <b>Errors</b> and <b>Throttles</b> charts never show data.	No metrics data exists for these charts in the specified chart range. No errors have occurred, and no requests have been throttled. Empty <b>Errors</b> and <b>Throttles</b> charts are expected.	Not applicable.
Throttles data: The <b>Throttles</b> chart shows data. What should I do?	Data in the <b>Throttles</b> chart indicates at least one request to invoke a function returned a '429 Too Many Requests' error in the response.	Resubmit the throttled invocation requests. Submit future invocation requests less frequently.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
 For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for all functions in an application

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Functions**.

2. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

3. Select the compartment containing the application for which you want to view function metrics.

The **Applications** page shows all the applications in the compartment you selected.

4. Click the name of the application for which you want to view function metrics.

5. Under **Resources**, click **Metrics**.

The Metrics page displays a chart for each metric that is emitted by the metric namespace for Oracle Functions. For more information about the emitted metrics, see [Available Metrics: oci\\_faas](#).

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for all the functions in all the applications in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see [6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure](#)).

3. Select the compartment containing the applications for which you want to view function metrics.

4. For **Metric Namespace**, select **oci\_faas**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace. For more information about the emitted metrics, see [Available Metrics: oci\\_faas](#).

## CHAPTER 16 Functions

---

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 17 Health Checks

This chapter explains how to monitor the health of your endpoints.

## Overview of the Health Checks Service

The Oracle Cloud Infrastructure Health Checks service provides users with high frequency external monitoring to determine the availability and performance of any publicly facing service, including hosted websites, API endpoints, or externally facing load balancers. By using Health Checks, users can ensure that they are immediately aware of any availability issue affecting their customers.

## Health Checks Service Components

The following list describes the key components used in creating a health check.

### **MONITORS**

Monitors allow you to continuously monitor the health of public-facing endpoints. You can configure monitors to use either HTTP and ping protocols.

### **ON-DEMAND PROBES**

On-demand probes allow you to execute a one-time probe to assess the health of a public-facing endpoint. You can configure on-demand probes to use either or both HTTP and ping protocols. This feature is currently only available via the [REST API](#).

### **VANTAGE POINTS**

Vantage points are geographic locations from which monitors and probes can be executed to your specified target. Oracle Cloud Infrastructure maintains dozens of vantage points around the world.

### **PROTOCOLS**

The Health Checks service allows you to configure both HTTP and ping type monitors. Each type has respective protocols.

### Ways to Access the Health Checks Service

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide.

To access the Console, you must use a supported browser. You can use the Console link at the top of this page to go to the sign-in page. Enter your tenancy, user name, and your password.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Health Checks Service Capabilities and Limits

The Oracle Cloud Infrastructure Health Checks service is limited to 1000 endpoint tests per account.

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for Health Checks, see [Details for the Health Checks Service](#).

#### Policy examples:

- To enable all operations on Health Checks for all users in a tenant :

```
Allow any-user to manage health-check-family in tenancy
```

- To enable all operations on Health Checks for all users in a compartment:

```
Allow any-user to manage health-check-family in compartment <Compartment Name>
```

- To enable all operations on Health Checks for a specific user group:

```
Allow group <Your Group Name> to manage health-check-family in compartment <Compartment Name>
```

### Moving Health Checks to a Different Compartment

You can move health checks from one compartment to another. When you move a health check to a new compartment, its associated monitor and test results moves with it. After the move, health checks are accessible through the SDK, CLI, and Console. For more information, see [Managing Compartments](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about available Health Checks service metrics and how to view them, see [Health Checks Metrics](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

## Getting Started With the Health Checks API

The Health Checks service allows you to configure and deploy monitors and on-demand probes using the Health Checks API. Use the following guide to learn how to set up monitors and probes then retrieve their results using the [REST API](#).



#### Note

Monitors, metrics, and probes created with the API, SDK and CLI are associated with the region where they were configured. While using the API, you must perform monitor updates (including compartment changes), metrics retrieval, and probe results retrieval in the region where they were configured. However, you can get a list of currently configured monitors and monitor details in every region, no matter where the monitors were configured.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Endpoints

The Health Checks API can be accessed via the following endpoints:

- <https://healthchecks.ap-mumbai-1.oraclecloud.com/20180501>
- <https://healthchecks.ap-seoul-1.oraclecloud.com/20180501>
- <https://healthchecks.ap-sydney-1.oraclecloud.com/20180501>
- <https://healthchecks.ap-tokyo-1.oraclecloud.com/20180501>
- <https://healthchecks.ca-toronto-1.oraclecloud.com/20180501>
- <https://healthchecks.eu-frankfurt-1.oraclecloud.com/20180501>
- <https://healthchecks.eu-zurich-1.oraclecloud.com/20180501>
- <https://healthchecks.sa-saopaulo-1.oraclecloud.com/20180501>
- <https://healthchecks.uk-london-1.oraclecloud.com/20180501>
- <https://healthchecks.us-ashburn-1.oraclecloud.com/20180501>
- <https://healthchecks.us-phoenix-1.oraclecloud.com/20180501>

### Available Protocols For Probes and Monitors

You can configure monitors and probes to use HTTP or ping requests. You will need to ensure that the endpoint being monitored is configured to accept the specified protocol.

## CHAPTER 17 Health Checks

---

**HTTP** - Configure a GET or HEAD request using HTTP/1.1 to test the target for availability. The probe results are returned in JSON and include the HTTP Status Code and DNS lookup, connection and response timings.

**HTTPS** - Configure an encrypted HTTPS GET or HEAD request to test the availability of any secure hosted target. Defaults to port 443. The probe results are returned in JSON and include the HTTP Status Code and DNS lookup, connection and response timings.

**ICMP** - Configure an ICMP echo request ping. The results include the round trip time (RTT) latency.

**TCP** - Configure a TCP handshake to the specified end point. You should be sure to own this endpoint as testing this connection can be costly to the recipient. The results include the round trip time (RTT) latency.

### Create A Monitor

Monitors allow you to monitor the health of endpoints over time. The following example shows how to create an HTTPS monitor that checks the health of `www.example.com` at an interval of every 30 seconds using a GET request.

```
POST /20180501/httpMonitors
{
 "compartmentId":"ocid1.compartment.oc1..<unique_ID>",
 "protocol":"HTTPS",
 "port":443,
 "targets":[
 "www.example.com"
],
 "timeoutInSeconds":30,
 "method":"GET",
 "displayName":"Example HTTP monitor",
 "intervalInSeconds":30
}
```

Targets can be either hostnames or IP addresses and the `path` field can be used to specify an optional path, such as `www.example.com/project/help.htm`. Optionally, you can specify which geographic locations you would like the monitor to launch from by using the

## CHAPTER 17 Health Checks

---

vantagePointNames field. At least one vantage point must be listed when using this field. For a list of available vantage points, see [Vantage Points](#).

A 200 response will be returned with the successful creation of a probe and the results of the probe can be retrieved from the URL in the resultsUrl field of the response.

```
{
 "id":"ocidl.httpmonitor.OC2...<unique_ID>",
 "compartmentId":"ocidl.compartment.oc1...<unique_ID>",
 "resultsUrl":"https://healthchecks.us-ashburn-
1.oraclecloud.com/20180501/httpProbeResults/ocidl.httpmonitor.OC2...<unique_ID>",
 "targets":[
 "www.example.com",
 "www.oracle.com"
],
 "vantagePointNames":[
 "ibm-sjc",
 "aws-dub",
 "dgo-nyc"
],
 "protocol":"HTTPS",
 "timeoutInSeconds":30,
 "displayName":"Example Monitor",
 "intervalInSeconds":30,
 "isEnabled":true
}
```

For more information about creating an HTTP monitor, see [CreateHttpMonitor](#).



### Note

You can configure a similar style monitors using TCP or ICMP protocols. For more information, see [CreatePingMonitor](#).

### Create An On-Demand Probe

Probes are one-off health assessments of an endpoint that can be deployed at anytime. The following example shows how to create an on-demand HTTP probe that checks the health of `www.example.com` with a GET request.

```
POST /20180501/httpProbeResults

{
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "protocol": "HTTP",
 "targets": [
 "www.example.com"
],
 "timeoutInSeconds": 30,
 "method": "GET"
}
```

Targets can be either hostnames or IP addresses and the `path` field can be used to specify an optional path, such as `www.example.com/project/help.htm`. Additionally, you can specify which geographic locations you would like the probe to launch from by using the `vantagePointNames` field. For a list of available vantage points, see [Vantage Points](#).

A 200 response will be returned with the successful creation of a probe and the results of the probe can be retrieved from the URL in the `resultsUrl` field of the response. It will take a few moments for results to display once the tests have been configured.

```
{
 "id": "ocidl.pingprobe.OC2..<unique_ID>",
 "compartmentId": "ocidl.compartment.oc1..<unique_ID>",
 "resultsUrl": "https://healthchecks.us-ashburn-1.oraclecloud.com/20180501/pingProbeResults/ocidl.pingprobe.OC2..<unique_ID>",
 "targets": [
 "www.example.com"
],
 "vantagePointNames": [
 "ibm-sjc",
 "aws-dub",
 "dgo-nyc"
],
 "protocol": "ICMP",
}
```

```
"timeoutInSeconds":30
}
```

For more information about creating a probe, see [CreateOnDemandHttpProbe](#).



### Note

You can configure similar style probes using TCP or ICMP protocols. For more information, see [CreateOnDemandPingProbe](#).

## Retrieving Probe And Monitor Results

Probe and monitor results can be retrieved from URL in the `resultsUrl` field of a monitor or probe creation response. It will take a few moments for results to display once the tests have been configured. Results can also be retrieved at anytime using the following methods:

- [ListPingProbeResults](#) - For monitors or on-demand probes using TCP or ICMP protocols.
- [ListHttpProbeResults](#) - For monitors or on-demand probes using HTTP protocols.

Retrieving results for an on-demand probe or monitor requires the probe or monitor's configuration ID as a parameter. On-demand probe and configuration IDs are assigned upon their creation and are returned in the `id` field of the POST response. You can also use the [ListHttpMonitor](#) method to retrieve a list of currently configured monitors and probes using HTTP protocols. Use the [ListPingMonitors](#) method to retrieve a list of currently configured monitors and probes using TCP and ICMP protocols.

The following is an example of results retrieved using `GET /httpProbeResults/{probeConfigurationId}`.

```
{
 "key": "651b9f3a46041cace0530204060ae27e",
 "probeConfigurationId": "ocidl.httpmonitor.OC2..<unique_ID>",
 "startTime": 1517323711505,
 "target": "www.example.com",
 "vantagePointName": "dgo-nyc",
 "protocol": "HTTPS",
```

## CHAPTER 17 Health Checks

---

```
"connection": {
 "connectDuration": 114,
 "secureConnectDuration": 99,
 "address": "93.184.216.34",
 "port": 443
},
"dns":{
 "domainLookupDuration": 29,
 "addresses": [
 "93.184.216.34",
 "2606:2800:220:1:248:1893:25c8:1946"
]
},
"statusCode": 200,
"fetchStart": 1517323711505,
"domainLookupStart": 1517323711505,
"domainLookupEnd": 1517323711534,
"connectStart": 1517323711535,
"secureConnectionStart": 1517323711550,
"connectEnd": 1517323711649,
"requestStart": 1517323711649,
"responseStart": 1517323711673,
"responseEnd": 1517323711676,
"duration": 171,
"encodedBodySize": 1270,
"isTimedOut": false,
"isHealthy": true
}
```

### Vantage Points

Vantage points are geographic locations from which monitors and probes can be launched. Oracle Cloud Infrastructure maintains vantage points on the infrastructure of cloud providers around the world, including AWS, IBM, and Azure. The list below is a sampling of the vantage points available. The list of vantage points is dynamic and changes frequently. Use the [ListHealthChecksVantagePoints](#) method to return a list of available vantage points.

## CHAPTER 17 Health Checks

---

<b>Provider</b>	<b>Location</b>	<b>Name</b>
Amazon	Singapore	aws-sin
Amazon	Sao Paulo	aws-sao
Amazon	Dublin	aws-dub
Amazon	San Francisco	aws-sfo
Azure	Dublin	azr-dub
Azure	Amsterdam	azr-ams
Azure	Singapore	azr-sin
Azure	Sydney	azr-syd
Digital Ocean	Toronto	dgo-yyz
Digital Ocean	Frankfurt	dgo-fra
Digital Ocean	New York City	dgo-nyc
Digital Ocean	Biratnagar, Nepal	dgo-blr
Google	Taiwan	goo-tpe
Google	Brussels	goo-bru
Google	Council Bluffs, IA	goo-cbf
Google	Charleston, SC	goo-chs
IBM	San Jose, CA	ibm-sjc
IBM	Tokyo	ibm-hnd
IBM	Dallas, TX	ibm-dfw

Provider	Location	Name
IBM	Hong Kong	ibm-hkg
Rack Space	Ashburn, VA	rck-iad
Rack Space	Dallas, TX	rck-dfw
Rack Space	London	rck-lhr
Rack Space	Sydney	rck-syd

## Managing Health Checks

The Health Checks service allows you to monitor the health of IP addresses and hostnames, as measured from geographic vantage points of your choosing, using HTTP and ping probes. After configuring a health check, you can view the monitor's results. The results include the location from which the host was monitored, the availability of the endpoint, and the date and time the test was performed.

### Using the Console

#### To add a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Click **Create Health Check**.
3. In the **Create Health Check** dialog box, enter the following:
  - **Health Check Name:** The name used for the health check. Avoid entering confidential information.
  - **Compartment:** Select the compartment the health check runs in.

- **Target(s):** The IP address of the host being monitored. (Optional) Click **+ Add Target** to add multiple targets in succession.
- **Vantage Points:** Select the location from which the health of the target is monitored. No more than ten vantage points can be added.
- **Request Type:** Select the type of request sent to monitor the target.
- **Protocol:** The network protocol used to interact with your endpoint, such as HTTP protocol, which initializes an HTTP handshake with your endpoint.
- **Port:** The port for the monitor to look for a connection. The default is port 80 for HTTP. For HTTPS, use port 443.
- **Path:** The specific path on the target to be monitored.
- **Header Name:** (Optional) The name displayed in the request header as part of the health check. Avoid entering confidential information.
- **Header Value:** (Optional) Specifies the data requested by the header. Click **+ Add Header** to add multiple headers in succession.
- **Method:** Select the HTTP method used for the health check.
- **Timeout:** Select the maximum time to wait for a reply before marking the health check as failed.
- **Interval:** Select the period of time between health checks of the target.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

#### 4. Click **Create Health Check**.

The health check is added to the health check list. To view more details, click the health check name. It will take a few moments for results to display once the tests have been configured.

### To edit a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Select the check box for the health check you want to edit.



#### Tip

To help find a health check, you can enter the name of the health check in the **Search** field.

3. Select **Edit** from the **Actions** drop-down menu.
4. In the Edit Health Check dialog box, make the needed changes, and then click **Edit Health Check**.

### To disable a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Select the check box for the health check you want to disable.



#### Tip

To help find a health check, you can enter the name of the health check in the **Search** field.

3. Select **Disable** from the **Actions** drop-down menu.

The status of the health check changes to **Disabled** in the health check list.

### To duplicate a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Select the check box for the health check you want to duplicate.
3. Select **Duplicate** from the **Actions** drop-down menu.
4. In the Create Health Check dialog box, make any updates to the duplicated health check, and then click **Create Health Check**.

### To delete a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Select the check box for the health check you want to delete.



#### Tip

To help find a health check, you can enter the name of the health check in the **Search** field.

3. Select **Delete** from the **Actions** drop-down menu.
4. In the confirmation dialog box, click **Delete**.

### To view the history of a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Click the name of the health check you want to view.  
The Health Check history displays a list of results for the past 90 days.



### Tip

To help find a result, you can use the **Start Date**, **Start Time**, **End Date**, **End Time**, and **Targets** filter options.

3. Click the drop-down arrow beside the **Timestamp** to view the monitor result details. You can use the API to download the data.

### To manage tags for a health check

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Click the name of the health check you want to view.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To move a health check to a different compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. In the **Scope** section, select a compartment.
3. Find the health check in the list, click the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

For more information, see [Managing Compartments](#).



### Tip

If your health checks are continually failing, please ensure that you have permission to monitor the host and that the ports on the host have been configured to receive traffic from Health Checks.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- Use the [CreateHTTPMonitor](#) operation to create a Health Check monitor that uses the HTTP protocol.
- Use the [CreatePingMonitor](#) operation to create a Health Check monitor that uses the ping protocol.
- Use the [ListHealthChecksVantagePoints](#) to retrieve a list of available vantage points from which to execute monitors.
- Use the [UpdateHttpMonitor](#) operation to update the configuration of an HTTP health check monitor. You can also use this operation to disable an HTTP monitor by setting the `isEnabled` field to `false`.
- Use the [UpdatePingMonitor](#) operation to update the configuration of ping health check monitor. You can also use this operation to disable a ping monitor by setting the `isEnabled` field to `false`.
- Use the [DeleteHttpMonitor](#) operation to remove an HTTP health check monitor from your setup.
- Use the [DeletePingMonitor](#) operation to remove a ping health check monitor from your setup.

- Use the [ListHttpProbeResults](#) operation to retrieve the results of an HTTP health check monitor.
- Use the [ListPingProbeResults](#) operation to retrieve results of a ping health check monitor.

## Health Checks Metrics

You can monitor the health, capacity, and performance of your health checks by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_healthchecks` (the Health Checks service).

### Overview of the Health Checks Service Metrics

Oracle Cloud Infrastructure Health Checks provides users with high frequency external monitoring to determine the availability and performance of any publicly facing service, including hosted websites, API endpoints, or externally facing load balancers. The Health Checks service metrics help you monitor the performance of your endpoints over a 24 hour period.

### Prerequisites

- IAM policies: To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_healthchecks

The metrics listed in the following table are automatically available for each health check that you create. You do not need to enable monitoring on the health check to get these metrics.

Each metric includes the following dimensions:

**RESOURCEID**

The OCID of the policy to which the metric applies.

Metric	Metric Display Name	Unit	Description	Dimensions
BasicCount	<b>Basic Health Checks</b>	count	The total number of active basic health checks.	resourceID resourceDisplayName
PremiumCount	<b>Premium Health Checks</b>	count	The total number of active premium health checks.	

## CHAPTER 17 Health Checks

Metric	Metric Display Name	Unit	Description	Dimensions
HTTP.StatusCode	<b>HTTP(S) Response Status Code</b>	count	The HTTP response code.	target vantagePoint resourceId resourceDisplayName statusCode2xx statusCode3xx statusCode4xx statusCode5xx protocol errorMessage
PING.isHealthy	<b>Success Rate of Ping Test</b>	percent	Displays availability of end point being monitored.	target vantagePoint resourceId resourceDisplayName protocol errorMessage icmpCode

## CHAPTER 17 Health Checks

Metric	Metric Display Name	Unit	Description	Dimensions
HTTP.DNSLookupTime	<b>HTTP(S) DNS Lookup Time</b>	ms	The time taken for domain name lookup in milliseconds.	target vantagePoint resourceId resourceDisplayName protocol errorMessage
HTTP.TCPConnectTime.Full	<b>HTTP(S) Connection Duration</b>	ms	The total duration in milliseconds from start of the request until response is fully consumed or the connection is closed.	
HTTP.TCPConnectTime.SSL	<b>HTTP(S) Secure Connection Duration</b>	ms	The total duration in milliseconds from start of secure connection to end of connection.	

## CHAPTER 17 Health Checks

---

<b>Metric</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Description</b>	<b>Dimensions</b>
HTTP.RequestTime	<b>HTTP(S) Request Duration</b>	ms	The total duration of the request in milliseconds.	
HTTP.ResponseTime	<b>HTTP(S) Response Duration</b>	ms	The total duration of response in milliseconds.	
HTTP.TotalDuration	<b>HTTP(S) Total Duration</b>	ms	The total duration of the test run in milliseconds.	

Metric	Metric Display Name	Unit	Description	Dimensions
HTTP.isHealthy	<b>Success Rate of HTTP(S) Test</b>	percent	Displays if the end point being monitored is up or down.	
PING.Latency	<b>Ping Latency Measurement</b>	ms	Latency measurement for ping test in milliseconds.	

## Using the Console

To view metric charts for resources related to a health check monitor

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Health Checks**.
2. Click the name of the health check you want to view metrics for.
3. Click **Metrics**.

To view health check metric charts using monitoring

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.  
For **Metric Namespace**, select **oci\_healthchecks**.

2. Select a metric to view from the **Metric Name** field.
3. Select a qualifier specified in the Dimension Name field. For example, the dimension `resourceId` is specified in the metric definition for `BasicCount`.
4. Select the value you want to use for the specified dimension in the **Dimension Value** field. For example, the resource identifier for your instance of interest.
5. Click **Update Chart**.  
The chart will be updated with the metrics that have been requested.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 18 IAM

This chapter explains how to set up administrators, users, and groups and specify their permissions.

## Overview of Oracle Cloud Infrastructure Identity and Access Management

Oracle Cloud Infrastructure Identity and Access Management (IAM) lets you control who has access to your cloud resources. You can control what type of access a group of users have and to which specific resources. This section gives you an overview of IAM components and an example scenario to help you understand how they work together.



### Note

This document uses the term "you" broadly to mean any administrator in your company who has access to work with IAM.

## Components of IAM

IAM uses the components described in this section. To better understand how the components fit together, see [Example Scenario](#).

### RESOURCE

The cloud objects that your company's employees create and use when interacting with Oracle Cloud Infrastructure. For example: compute instances, block storage volumes, virtual cloud networks (VCNs), subnets, route tables, etc.

### **USER**

An individual employee or system that needs to manage or use your company's Oracle Cloud Infrastructure resources. Users might need to launch instances, manage remote disks, work with your virtual cloud network, etc. End users of your application are not typically IAM users. Users have one or more IAM credentials (see [User Credentials](#)).

### **GROUP**

A collection of users who all need the same type of access to a particular set of resources or compartment.

### **DYNAMIC GROUP**

A special type of group that contains resources (such as compute instances) that match rules that you define (thus the membership can change dynamically as matching resources are created or deleted). These instances act as "principal" actors and can make API calls to services according to policies that you write for the dynamic group.

### **COMPARTMENT**

A collection of related resources. Compartments are a fundamental component of Oracle Cloud Infrastructure for organizing and isolating your cloud resources. You use them to clearly separate resources for the purposes of measuring usage and billing, access (through the use of policies), and isolation (separating the resources for one project or business unit from another). A common approach is to create a compartment for each major part of your organization. For more information, see "Setting Up Your Tenancy" in the *Oracle Cloud Infrastructure Getting Started Guide*.

### **TENANCY**

The root compartment that contains *all* of your organization's Oracle Cloud Infrastructure resources. Oracle automatically creates your company's tenancy for you. Directly within the tenancy are your IAM entities (users, groups, compartments, and some policies; you can also put policies into compartments inside the tenancy). You place the other types of cloud resources (e.g., instances, virtual networks, block storage volumes, etc.) inside the compartments that you create.

### **POLICY**

A document that specifies who can access which resources, and how. Access is granted at the group and compartment level, which means you can write a policy that gives a group a specific type of access within a specific compartment, or to the tenancy itself. If you give a group access to the tenancy, the group automatically gets the same type of access to all the compartments inside the tenancy. For more information, see [Example Scenario](#) and [How Policies Work](#). The word "policy" is used by people in different ways: to mean an individual statement written in the policy language; to mean a collection of statements in a single, named "policy" document (which has an Oracle Cloud ID (OCID) assigned to it); and to mean the overall body of policies your organization uses to control access to resources.

### **HOME REGION**

The region where your IAM resources reside. All IAM resources are global and available across all regions, but the master set of definitions reside in a single region, the home region. You must make changes to your IAM resources in your home region. The changes will be automatically propagated to all regions. For more information, see [Managing Regions](#).

### **FEDERATION**

A relationship that an administrator configures between an identity provider and a service provider. When you federate Oracle Cloud Infrastructure with an identity provider, you manage users and groups in the identity provider. You manage authorization in Oracle Cloud Infrastructure's IAM service. Oracle Cloud Infrastructure tenancies are federated with Oracle Identity Cloud Service by default.

## Services You Can Control Access To

You can write policies to control access to all of the [services](#) within Oracle Cloud Infrastructure.

### The Administrators Group and Policy

When your company signs up for an Oracle account and Identity Domain, Oracle sets up a *default administrator* for the account. This person will be the first IAM user for your company and will be responsible for initially setting up additional administrators. Your tenancy comes with a group called *Administrators*, and the default administrator automatically belongs in this group. You can't delete this group, and there must always be at least one user in it.

Your tenancy also automatically has a policy that gives the Administrators group access to all of the Oracle Cloud Infrastructure API operations and all of the cloud resources in your tenancy. You can neither change nor delete this policy. Any other users you put into the Administrators group will have full access to all of the services. This means they can create and manage IAM resources such as, groups, policies, and compartments. And they can create and manage the cloud resources such as virtual cloud networks (VCNs), instances, block storage volumes, and any other new types of Oracle Cloud Infrastructure resources that become available in the future.

### Example Scenario

The goal of this scenario is to show how the different IAM components work together, and basic features of policies.

In this scenario, Acme Company has two teams that will be using Oracle Cloud Infrastructure resources for infrastructure: Project A and Project B. In reality, your company may have many more.

Acme Company plans to use a single virtual cloud network (VCN) for both teams, and wants a network administrator to manage the VCN.

Acme Company also wants the Project A team and Project B team to *each* have their own set of instances and block storage volumes. The Project A team and Project B teams shouldn't be able to use each other's instances. These two teams also shouldn't be allowed to change anything about the VCN set up by the network administrator. Acme Company wants each team to have administrators for that team's resources. The administrators for the Project A team can decide who can use the Project A cloud resources, and how. Same for the Project B team.

### Acme Company Gets Started with Oracle Cloud Infrastructure

Acme Company signs up to use Oracle Cloud Infrastructure and tells Oracle that an employee named Wenpei will be the default administrator. In response, Oracle:

- Creates a tenancy for Acme Company (see the following diagram).
- Creates an IAM user account for Wenpei in the tenancy.
- Creates the Administrators group in the tenancy and places Wenpei in that group.
- Creates a policy in Acme Company's tenancy that gives the Administrators group access to manage all of the resources in the tenancy. Here's that policy:

```
Allow group Administrators to manage all-resources in tenancy
```

#### The default setup for a new tenancy:

The screenshot displays the configuration for a new tenancy. At the top, it is labeled 'CompanyA Tenancy'. Below this, there are three sections:

- Policies attached to the tenancy:** A red dot indicates one policy is attached: *Allow group Administrators to manage all-resources in tenancy*.
- Users:** A table lists the user 'Wenpei'.
- Groups:** A table lists a group named 'Administrators', which contains the user 'Wenpei'.

#### The Default Administrator Creates Some Groups and Another Administrator

Wenpei next creates several groups and users (see the following diagram). She:

## CHAPTER 18 IAM

- Creates groups called *NetworkAdmins*, *A-Admins*, and *B-Admins* (these last two are for Project A and Project B within the company)
- Creates a user called Alex and puts him in the Administrators group.
- Leaves the new groups empty.

To learn how to create groups, see [Working with Groups](#). To learn how to create users and put them in groups, see [Working with Users](#).

The screenshot displays the 'CompanyA Tenancy' configuration page. It is divided into three main sections: Policies, Users, and Groups. The Policies section shows a single policy: 'Allow group Administrators to manage all-resources in tenancy'. The Users section lists two users: Alex and Wenpei. The Groups section shows four groups: Administrators (containing Wenpei and Alex), NetworkAdmins, A-Admins, and B-Admins.

CompanyA Tenancy			
<b>Policies attached to the tenancy:</b>			
● Allow group <b>Administrators</b> to manage all-resources in tenancy			
<b>Users</b>	Alex Wenpei		
<b>Groups</b>			
<b>Administrators</b>	<b>NetworkAdmins</b>	<b>A-Admins</b>	<b>B-Admins</b>
Wenpei Alex			

### The Default Administrator Creates Some Compartments and Policies

Wenpei next creates compartments to group resources together (see the following diagram). She:

- Creates a compartment called *Networks* to control access to the Acme Company's VCN, subnets, IPsec VPN, and other components from Networking.
- Creates a compartment called *Project-A* to organize Project A team's cloud resources and control access to them.

- Creates a compartment called *Project-B* to organize Project B team's cloud resources and control access to them.

To learn how to manage compartments, see [Working with Compartments](#).

Wenpei then creates a policy to give the administrators for each compartment their required level of access. She attaches the policy to the tenancy, which means that only users with access to manage policies in the tenancy can later update or delete the policy. In this scenario, that is only the Administrators group. The policy includes multiple statements that:

- Give the NetworkAdmins group access to manage networks and instances (for the purposes of easily testing the network) in the Networks compartment
- Give both the A-Admins and B-Admins groups access to use the networks in the Networks compartment (so they can create instances into the network).
- Give the A-Admins group access to manage all resources in the Project-A compartment.
- Give the B-Admins group access to manage all resources in the Project-B compartment.

Here's what that policy looks like (notice it has multiple statements in it):

```
Allow group NetworkAdmins to manage virtual-network-family in compartment Networks
Allow group NetworkAdmins to manage instance-family in compartment Networks

Allow group A-Admins,B-Admins to use virtual-network-family in compartment Networks

Allow group A-Admins to manage all-resources in compartment Project-A

Allow group B-Admins to manage all-resources in compartment Project-B
```

Notice the difference in the verbs (*manage*, *use*), as well as the resources (*virtual-network-family*, *instance-family*, *all-resources*). For more information about them, see [Verbs](#) and [Resource-Types](#). To learn how to create policies, see [To create a policy](#).

**Important**

A-Admins and B-Admins can *use* the virtual-network-family in the compartment Networks. However, they can't create instances in that compartment. They can only create instances in the Project-A or Project-B compartment. Remember, a compartment is a logical grouping, not a physical one, so resources that make up or reside on the same VCN can belong to different compartments.

Acme Company wants to let the administrators of the Project-A and Project-B compartments decide which users can use the resources in those compartments. So Wenpei creates two more groups: A-Users and B-Users. She then adds six more statements that give the compartment admins the required access they need in order to add and remove users from those groups:

```
Allow group A-Admins to use users in tenancy where target.group.name='A-Users'
Allow group A-Admins to use groups in tenancy where target.group.name='A-Users'

Allow group B-Admins to use users in tenancy where target.group.name='B-Users'
Allow group B-Admins to use groups in tenancy where target.group.name='B-Users'

Allow group A-Admins,B-Admins to inspect users in tenancy
Allow group A-Admins,B-Admins to inspect groups in tenancy
```

Notice that this policy doesn't let the project admins *create* new users or manage credentials for the users. It lets them decide which existing users can be in the A-Users and B-Users groups. The last two statements are necessary for A-Admins and B-Admins to list all the users and groups, and confirm which users are in which groups.

### CompanyA Tenancy

**Policies attached to the tenancy:**

- Allow group **Administrators** to manage all-resources in tenancy
- Allow group **NetworkAdmins** to manage virtual-network-family in compartment Networks
- Allow group **NetworkAdmins** to manage instance-family in compartment Networks
- Allow group **A-Admins,B-Admins** to use virtual-network-family in compartment Networks
- Allow group **A-Admins** to manage all-resources in compartment Project-A
- Allow group **B-Admins** to manage all-resources in compartment Project-B
- Allow group **A-Admins** to use users in tenancy where target.group.name='A-Users'
- Allow group **A-Admins** to use groups in tenancy where target.group.name='A-Users'
- Allow group **B-Admins** to use users in tenancy where target.group.name='B-Users'
- Allow group **B-Admins** to use groups in tenancy where target.group.name='B-Users'
- Allow group **A-Admins,B-Admins** to inspect users in tenancy
- Allow group **A-Admins,B-Admins** to inspect groups in tenancy

---

**Users**      Alex  
                  Wenpei

---

**Groups**

<b>Administrators</b>	<b>NetworkAdmins</b>	<b>A-Admins</b>	<b>B-Admins</b>
Wenpei Alex			
		<b>A-Users</b>	<b>B-Users</b>

---

**Compartments**

<b>Networks</b>	<b>Project-A</b>	<b>Project-B</b>
-----------------	------------------	------------------

### **An Administrator Creates New Users**

At this point, Alex is in the Administrators group and now has access to create new users. So he provisions users named Leslie, Jorge, and Cheri and places them in the NetworkAdmins, A-Admins, and B-Admins groups, respectively. Alex also creates other users who will eventually be put in the A-Users and B-Users groups by the admins for Project A and Project B.

**CompanyA Tenancy**

Policies attached to the tenancy:

- Allow group **Administrators** to manage all-resources in tenancy
- Allow group **NetworkAdmins** to manage virtual-network-family in compartment Networks
- Allow group **NetworkAdmins** to manage instance-family in compartment Networks
- Allow group **A-Admins,B-Admins** to use virtual-network-family in compartment Networks
- Allow group **A-Admins** to manage all-resources in compartment Project-A
- Allow group **B-Admins** to manage all-resources in compartment Project-B
- Allow group **A-Admins** to use users in tenancy where target.group.name='A-Users'
- Allow group **A-Admins** to use groups in tenancy where target.group.name='A-Users'
- Allow group **B-Admins** to use users in tenancy where target.group.name='B-Users'
- Allow group **B-Admins** to use groups in tenancy where target.group.name='B-Users'
- Allow group **A-Admins,B-Admins** to inspect users in tenancy
- Allow group **A-Admins,B-Admins** to inspect groups in tenancy

---

**Users**

	Alex	Fred	Jorge	Tarik
	Cheri	Helali	Laura	Wenpei
	Dylan	Jenna	Leslie	

---

**Groups**

<b>Administrators</b>	<b>NetworkAdmins</b>	<b>A-Admins</b>	<b>B-Admins</b>
Wenpei Alex	Leslie	Jorge	Cheri
		<b>A-Users</b>	<b>B-Users</b>

---

**Compartments**

<b>Networks</b>	<b>Project-A</b>	<b>Project-B</b>
-----------------	------------------	------------------

### The Network Admin Sets Up the Network

Leslie (in the NetworkAdmins group) has access to manage `virtual-network-family` and `instance-family` in the Networks compartment. She creates a virtual cloud network (VCN) with a single subnet in that compartment. She also sets up an Internet gateway for the VCN, and updates the VCN's route table to allow traffic via that gateway. To test the VCN's connectivity to the on-premises network, she launches an instance in the subnet in the VCN. As part of the launch request, she must specify which compartment the instance should reside in. She specifies the Networks compartment, which is the only one she has access to. She then confirms connectivity from the on-premises network to the VCN by logging in to the instance via SSH from the on-premises network.

Leslie terminates her test instance and lets Jorge and Cheri know that the VCN is up and running and ready to try out. She lets them know that their compartments are named Project-A and Project-B respectively. For more information about setting up a cloud network, see [Overview of Networking](#). For information about launching instances into the network, see [Overview of the Compute Service](#).

### Compartment Admins Set Up Their Compartments

Jorge and Cheri now need to set up their respective compartments. Each admin needs to do the following:

- Launch instances in their own compartment
- Put users in their "users" group (e.g., A-Users)
- Decide the type of access to give those users, and accordingly attach a policy to their compartment

Jorge and Cheri both launch instances into the subnet in the VCN, into their respective team's compartments. They create and attach block volumes to the instances. Only the compartment admins can launch/terminate instances or attach/detach block volumes in their respective team's compartments.



### **Important**

#### *Network Topology and Compartment Access Are Different Concepts*

It's important to understand the difference between the network topology of the VCN and the access control that the compartments provide. The instances Jorge launched reside in the VCN from a network topology standpoint. But from an access standpoint, they're in the Project-A compartment, not the Networks compartment where the VCN is. Leslie (the Networks admin) can't terminate or reboot Jorge's instances, or launch new ones into the Project-A compartment. But Leslie controls the instances' network, so she controls what traffic will be routed to them. If Jorge had specified the Networks compartment instead of the Project-A compartment when launching his instances, his request would have been denied. The story is similar for Cheri and the Project-B compartment.

But it's also important to note that Wenpei and Alex in the Administrators group do have access to the resources inside the compartments, because they have access to manage all kinds of resources in the tenancy. Compartments inherit any policies attached to their parent compartment (the tenancy), so the Administrators access also applies to all compartments *within* the tenancy.

Next, Jorge puts several of the users that Alex created into the A-Users group. Cheri does the same for B-Users.

Then Jorge writes a policy that gives users the level of access they need in the Project-A compartment.

```
Allow group A-Users to use instance-family in compartment Project-A
Allow group A-Users to use volume-family in compartment Project-A
Allow group A-Users to inspect virtual-network-family in compartment Networks
```

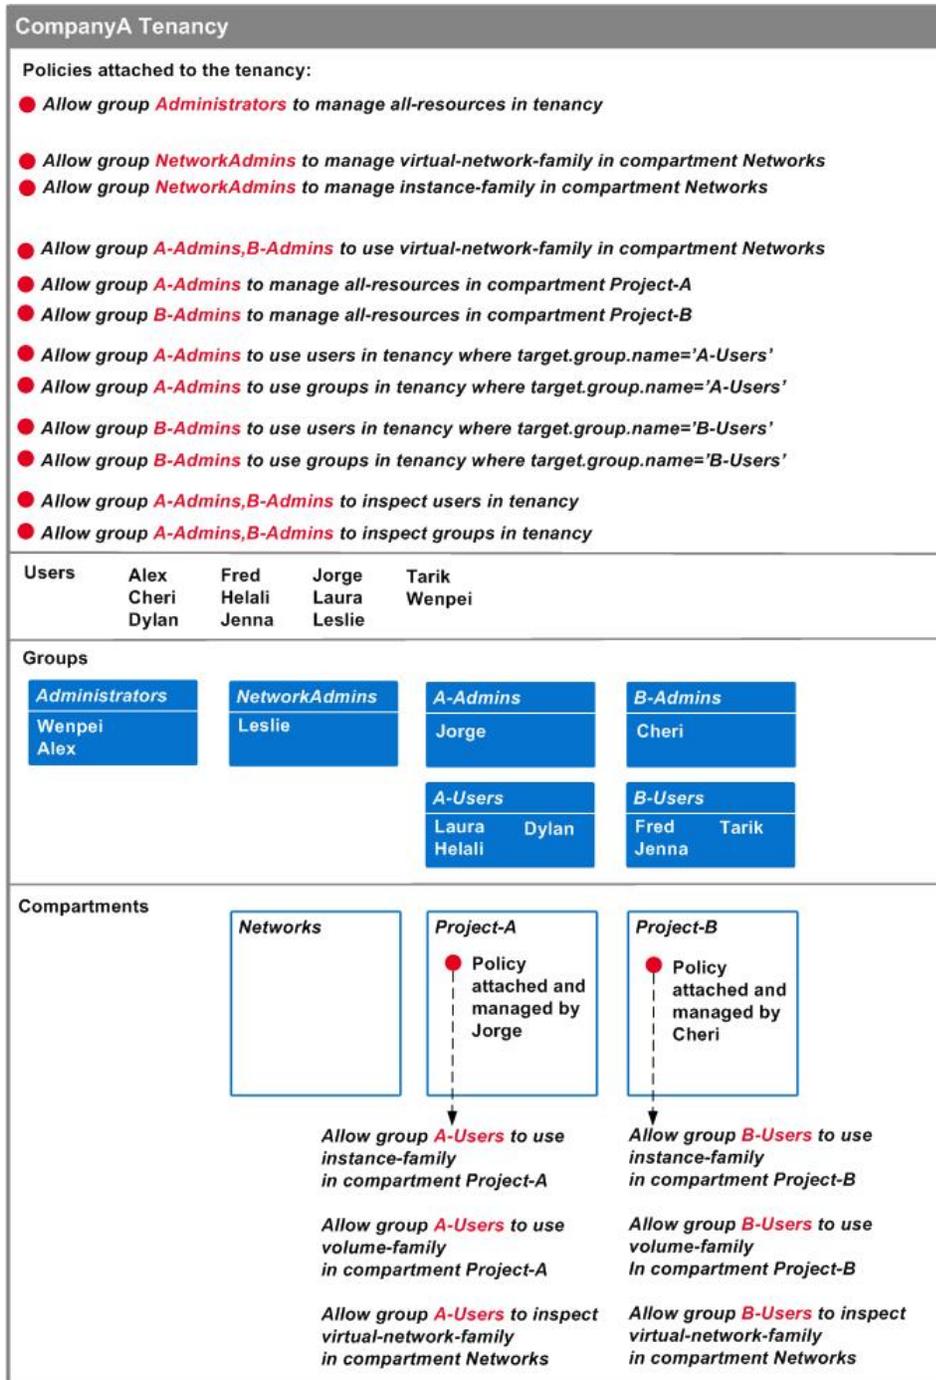
This lets them use existing instances (with attached block volumes) that the compartment admins already launched in the compartment, and stop/start/reboot them. It does not let A-Users create/delete or attach/detach any volumes. To give that ability, the policy would need to include `manage volume-family`.

Jorge attaches this policy to the Project-A compartment. Anyone with the ability to manage policies *in the compartment* can now modify or delete this policy. Right now, that is only the A-Admins group (and the Administrators group, which can do anything throughout the tenancy).

Cheri creates and attaches her own policy to the Project-B compartment, similar to Jorge's policy:

```
Allow group B-Users to use instance-family in compartment Project-B
Allow group B-Users to use volume-family in compartment Project-B
Allow group B-Users to inspect virtual-network-family in compartment Networks
```

Now the A-Users and B-Users can work with the existing instances and attached volumes in the Project-A and Project-B compartments, respectively. Here's what the layout looks like:



For more information about basic and advanced features of policies, see [How Policies Work](#). For examples of other typical policies your organization might use, see [Common Policies](#).

### Viewing Resources by Compartment in the Console

In the Console, you view your cloud resources *by compartment*. This means that after you sign in to the Console, you'll choose which compartment to work in (there's a list of the compartments you have access to on the left side of the page). Notice that compartments can be nested inside other compartments. The page will update to show that compartment's resources that are within the current region. If there are none, or if you don't have access to the resource in that compartment, you'll see a message.

This experience is different when you're viewing the lists of users, groups, dynamic groups, and federation providers. Those reside in the tenancy itself (the root compartment), not in an individual compartment.

As for policies, they can reside in either the tenancy or a compartment, depending on where the policy *is attached*. Where it's attached controls who has access to modify or delete it. For more information, see [Policy Attachment](#).

### The Scope of IAM Resources

Oracle Cloud Infrastructure uses the concepts of regions and availability domains (see [Regions and Availability Domains](#)). Some resources are available regionally, whereas others are available only within a certain availability domain. IAM resources (users, groups, dynamic groups, compartments, tag namespaces, federation providers, and policies) are global and available across all regions. See [Managing Regions](#).

### Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

The following IAM resources emit events:

- Authentication policies
- Credentials
- Dynamic groups
- Groups
- Identity Providers
- Multi-factor Authentication TOTP Devices
- Policies
- Users

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

### Limits on IAM Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

### Getting Started with Policies

If you're new to Oracle Cloud Infrastructure Identity and Access Management (IAM) policies, this topic gives guidance on how to proceed.

#### If You're Doing a Proof-of-Concept

If you're just trying out Oracle Cloud Infrastructure or doing a proof-of-concept project with infrastructure resources, you may not need more than a few administrators with full access to everything. In that case, you can simply create any new users you need and add them to the Administrators group. The users will be able to do anything with any kind of resource. And you can create all your resources directly in the tenancy (the root compartment). You don't need to create any compartments yet, or any other policies beyond the Tenant Admin Policy, which automatically comes with your tenancy and can't be changed.



#### Note

Don't forget to add your new users to the Administrators group; it's easy to forget to do that after creating them.

#### If You're Past the Proof-of-Concept Phase

If you're past the proof-of-concept phase and want to restrict access to your resources, first:

- Make sure you're familiar with the basic IAM components, and read through the example scenario: [Overview of Oracle Cloud Infrastructure Identity and Access](#)

### [Management](#)

- Think about how to organize your resources into compartments: See "Setting Up Your Tenancy" in the *Oracle Cloud Infrastructure Getting Started Guide*
- Learn the basics of how policies work: [How Policies Work](#)
- Check out some typical policies: [Common Policies](#)
- Read the FAQs below

### Policy FAQs

**Which of the services within Oracle Cloud Infrastructure can I control access to through policies?**

All of them, including IAM itself. You can find specific details for writing policies for each service in the [Policy Reference](#).

**Can users do anything without an administrator writing a policy for them?**

Yes. All users can automatically do these things without an explicit policy:

- Change or reset their own Console password.
- Manage their own API signing keys and other credentials.

**Why should I separate resources by compartment? Couldn't I just put all the resources into one compartment and then use policies to control who has access to what?**

You could put all your resources into a single compartment and use policies to control access, but then you would lose the benefits of measuring usage and billing by compartment, simple policy administration at the compartment level, and clear separation of resources between projects or business units.

### Can I control or deny access to an individual user?

Yes. However, there are a couple things to know first:

- Enterprise companies typically have multiple users that need similar permissions, so policies are designed to give access to *groups* of users, not individual users. A user gains access by being in a group.
- Policies are designed to *allow* access; there's no explicit "deny" when you write a policy.

If you need to restrict a particular user's access, you can:

- Remove the user from the particular group of interest
- Delete the user entirely from IAM (you have to remove the user from all groups first)

### How do I delete a user?

First ensure the user isn't in any groups. Only then can you delete the user.

### How do I delete a compartment?

See [Deleting Compartments](#).

### How can I tell which policies apply to a particular group or user?

You need to look at the individual statements in all your policies to see which statements apply to which group. There's not currently an easy way to get this information.

### How can I tell which policies apply to a particular compartment?

You need to look at the individual statements in all the policies in the tenancy to see if any apply to the particular compartment. You also need to look at any policies in the compartment itself. Policies in any of the sibling compartments *cannot* refer to the compartment of interest, so you don't need to check those policies.

## How Policies Work

This topic describes how policies work and the basic features.

### Overview of Policies

A *policy* is a document that specifies who can access which Oracle Cloud Infrastructure resources that your company has, and how. A policy simply allows a group to work in certain ways with specific types of resources in a particular compartment. If you're not familiar with users, groups, or compartments, see [Overview of Oracle Cloud Infrastructure Identity and Access Management](#).

In general, here's the process an IAM administrator in your organization needs to follow:

1. Define users, groups, and one or more compartments to hold the cloud resources for your organization.
2. Create one or more policies, each written in the policy language. See [Common Policies](#).
3. Place users into the appropriate groups depending on the compartments and resources they need to work with.
4. Provide the users with the one-time passwords that they need in order to access the Console and work with the compartments. For more information, see [User Credentials](#).

After the administrator completes these steps, the users can access the Console, change their one-time passwords, and work with specific cloud resources as stated in the policies.

### Policy Basics

To govern control of your resources, your company will have at least one policy. Each policy consists of one or more policy *statements* that follow this basic syntax:

```
Allow group <group_name> to <verb> <resource-type> in compartment <compartment_name>
```

Notice that the statements always begin with the word `Allow`. Policies only *allow* access; they cannot deny it. Instead there's an implicit deny, which means by default, users can do nothing

and have to be granted access through policies. (There's one exception to this rule; see [Can users do anything without an administrator writing a policy for them?](#))

An administrator in your organization defines the groups and compartments in your tenancy. Oracle defines the possible verbs and resource-types you can use in policies (see [Verbs](#) and [Resource-Types](#)).

In some cases you'll want the policy to apply to the tenancy and not a compartment inside the tenancy. In that case, change the end of the policy statement like so:

```
Allow group <group_name> to <verb> <resource-type> in tenancy
```

For more details about the syntax, see [Policy Syntax](#).

For information about how many policies you can have, see [Service Limits](#).

### A Few Examples

Let's say your administrator creates a group called *HelpDesk* whose job is to manage users and their credentials. Here is a policy that enables that:

```
Allow group HelpDesk to manage users in tenancy
```

Notice that because users reside in the tenancy (the root compartment), the policy simply states the word `tenancy`, without the word `compartment` in front of it.

Next, let's say you have a compartment called *Project-A*, and a group called *A-Admins* whose job is to manage all of the Oracle Cloud Infrastructure resources in the compartment. Here's an example policy that enables that:

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

Be aware that the policy directly above includes the ability to write policies *for that compartment*, which means *A-Admins* can control access to the compartment's resources. For more information, see [Policy Attachment](#).

If you wanted to limit *A-Admins'* access to only launching and managing compute instances and block storage volumes (both the volumes and their backups) in the *Project-A* compartment, but the network itself lives in the *Networks* compartment, then the policy could instead be:

```
Allow group A-Admins to manage instance-family in compartment Project-A
```

## CHAPTER 18 IAM

---

```
Allow group A-Admins to manage volume-family in compartment Project-A
```

```
Allow group A-Admins to use virtual-network-family in compartment Networks
```

The third statement with the `virtual-network-family` resource-type enables the instance launch process, because the cloud network is involved. Specifically, the launch process creates a new VNIC and attaches it to the subnet where the instance resides.

For additional examples, see [Common Policies](#).

### Details about Specifying Groups and Compartments

Typically you'll specify a group or compartment by name in the policy. However, you can use the OCID instead. Just make sure to add "id" before the OCID. For example:

```
Allow group
 id ocid1.group.oc1..aaaaaaaaqjihfhvxmumrl3isyxjw3n6c4rzwskaawuc7i5xwe6s7qmnsbc6a
to manage instance-family in compartment Project-A
```

You can specify multiple groups separated by commas:

```
Allow group A-Admins, B-Admins to manage instance-family in compartment Projects-A-and-B
```

### Verbs

Oracle defines the possible verbs you can use in your policies. Here's a summary of the verbs, from least amount of access to the most:

Verb	Types of Access Covered	Target User
inspect	Ability to list resources, without access to any confidential information or user-specified metadata that may be part of that resource. <b>Important:</b> The operation to list policies includes the contents of the policies themselves, and the list operations for the Networking resource-types return all the information (e.g., the contents of security lists and route tables).	Third-party auditors
read	Includes <code>inspect</code> plus the ability to get user-specified metadata and the actual resource itself.	Internal auditors
use	Includes <code>read</code> plus the ability to work with existing resources (the actions vary by resource type). Includes the ability to update the resource, except for resource-types where the "update" operation has the same effective impact as the "create" operation (e.g., <code>UpdatePolicy</code> , <code>UpdateSecurityList</code> , etc.), in which case the "update" ability is available only with the <code>manage</code> verb. In general, this verb does not include the ability to create or delete that type of resource.	Day-to-day end users of resources
manage	Includes all permissions for the resource.	Administrators

The verb gives a certain general type of access (e.g., `inspect` lets you list and get resources). When you then join that type of access with a particular resource-type in a policy (e.g., `Allow group XYZ to inspect compartments in the tenancy`), then you give that group access to a specific set of permissions and API operations (e.g., `ListCompartments`, `GetCompartment`). For more examples, see [Details for Verbs + Resource-Type Combinations](#). The [Policy Reference](#) includes a similar table for each service, giving you a list of exactly which API operations are covered for each combination of verb and resource-type.

There are some special exceptions or nuances for certain resource-types.

**Users:** Access to both `manage users` and `manage groups` lets you do anything with users and groups, including creating and deleting users and groups, and adding/removing users from groups. To add/remove users from groups without access to creating and deleting users and groups, only both `use users` and `use groups` are required. See [Common Policies](#).

**Policies:** The ability to update a policy is available only with `manage policies`, not `use policies`, because updating a policy is similar in effect to creating a new policy (you can overwrite the existing policy statements). In addition, `inspect policies` lets you get the full contents of the policies.

**Object Storage objects:** `inspect objects` lets you list all the objects in a bucket and do a HEAD operation for a particular object. In comparison, `read objects` lets you download the object itself.

**Load Balancing resources:** Be aware that `inspect load-balancers` lets you get *all* information about your load balancers and related components (backend sets, etc.).

### Networking resources:

Be aware that the `inspect` verb not only returns general information about the cloud network's components (for example, the name and OCID of a security list, or of a route table). It also includes the contents of the component (for example, the actual rules in the security list, the routes in the route table, and so on).

Also, the following types of abilities are available only with the `manage` verb, not the `use` verb:

- Update (enable/disable) `internet-gateways`
- Update `security-lists`
- Update `route-tables`
- Update `dhcp-options`
- Attach a dynamic routing gateway (DRG) to a virtual cloud network (VCN)
- Create an IPSec connection between a DRG and customer-premises equipment (CPE)
- Peer VCNs

**Important**

Each VCN has various components that directly affect the behavior of the network (route tables, security lists, DHCP options, Internet Gateway, and so on). When you create one of these components, you establish a relationship between that component and the VCN, which means you must be allowed in a policy to both create the component and manage the VCN itself. However, the ability to *update* that component (to change the route rules, security list rules, and so on) does NOT require permission to manage the VCN itself, even though changing that component can directly affect the behavior of the network. This discrepancy is designed to give you flexibility in granting least privilege to users, and not require you to grant excessive access to the VCN just so the user can manage other components of the network. Be aware that by giving someone the ability to update a particular type of component, you're implicitly trusting them with controlling the network's behavior.

**Resource-Types**

Oracle also defines the resource-types you can use in your policies. First, there are *individual* types. Each individual type represents a specific type of resource. For example, the `vcns` resource-type is specifically for virtual cloud networks (VCNs).

To make policy writing easier, there are *family* types that include multiple individual resource-types that are often managed together. For example, the `virtual-network-family` type brings together a variety of types related to the management of VCNs (e.g., `vcns`, `subnets`, `route-tables`, `security-lists`, etc.). If you need to write a more granular policy

that gives access to only an individual resource-type, you can. But you can also easily write a policy to give access to a broader range of resources.

In another example: Block Volume has `volumes`, `volume-attachments`, and `volume-backups`. If you need to give access to only making backups of volumes, you can specify the `volume-backups` resource-type in your policy. But if you need to give broad access to all of the Block Volume resources, you can specify the family type called `volume-family`. For a full list of the resource-types, see [Resource-Types](#).



### Important

If a service introduces new individual resource-types, they will typically be included in the family type for that service. For example, if Networking introduces a new individual resource-type, it will be automatically included in the definition of the `virtual-network-family` resource type. For more information about future changes to the definitions of resource-types, see [Policies and Service Updates](#).

Note that there are other ways to make policies more granular, such as the ability to specify conditions under which the access is granted. For more information, see [Advanced Policy Features](#).

### Access that Requires Multiple Resource-Types

Some API operations require access to multiple resource-types. For example, `LaunchInstance` requires the ability to create instances and work with a cloud network. The `CreateVolumeBackup` operation requires access to both the volume and the volume backup. That means you'll have separate statements to give access to each resource-type (for an example, see [Let volume backup admins manage only backups](#)). These individual statements do not have to be in the same policy. And a user can gain the required access from being in different groups. For example, George could be in one group that gives the required level of access to the `volumes` resource-type, and in another group that gives the required access to

the `volume-backups` resource-type. The sum of the individual statements, regardless of their location in the overall set of policies, gives George access to `CreateVolumeBackup`.

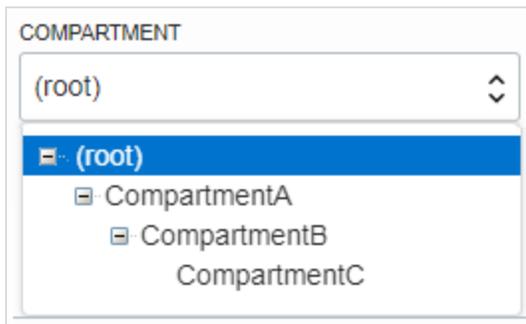
### Policy Inheritance

A basic feature of policies is the concept of *inheritance*: Compartments inherit any policies from their parent compartment. The simplest example is the Administrators group, which automatically comes with your tenancy (see [The Administrators Group and Policy](#)). There's a built-in policy that enables the Administrators group to do anything in the tenancy:

```
Allow group Administrators to manage all-resources in tenancy
```

Because of policy inheritance, the Administrators group can also do anything in *any* of the compartments in the tenancy.

To illustrate further, consider a tenancy with three levels of compartments: CompartmentA, CompartmentB, and CompartmentC, shown here:



Policies that apply to resources in CompartmentA also apply to resources in CompartmentB and CompartmentC. So this policy:

```
Allow group NetworkAdmins to manage virtual-network-family in compartment CompartmentA
```

allows the group NetworkAdmins to manage VCNs in CompartmentA, CompartmentB, and CompartmentC.

### Policy Attachment

Another basic feature of policies is the concept of *attachment*. When you create a policy you must attach it to a compartment (or the tenancy, which is the root compartment). **Where you attach it controls who can then modify it or delete it.** If you attach it to the tenancy (in other words, if the policy *is in* the root compartment), then anyone with access to manage policies in the tenancy can then change or delete it. Typically that's the Administrators group or any similar group you create and give broad access to. Anyone with access only to a child compartment cannot modify or delete that policy.

If you instead attach the policy to a child compartment, then anyone with access to manage the policies *in that compartment* can change or delete it. In practical terms, this means it's easy to give compartment administrators (i.e., a group with access to `manage all-resources` in the compartment) access to manage their own compartment's policies, without giving them broader access to manage policies that reside in the tenancy. For an example that uses this kind of compartment administrator design, see [Example Scenario](#). (Recall that because of policy inheritance, users with access to manage policies in the tenancy automatically have the ability to manage policies in compartments inside the tenancy.)

The process of attaching the policy is easy (whether attaching to a compartment or the tenancy): If you're using the Console, when you add the policy to IAM, simply make sure you're in the desired compartment when you create the policy. If you're using the API, you specify the OCID of the desired compartment (either the tenancy or other compartment) as part of the request to create the policy.

When you attach a policy to a compartment, you must be in that compartment *and* you must indicate directly in the statement which compartment it applies to. If you are not in the compartment, you'll get an error if you try to attach the policy to a different compartment. Notice that attachment occurs during policy creation, which means a policy can be attached to only one compartment. To learn how to attach a policy to a compartment, see [To create a policy](#).

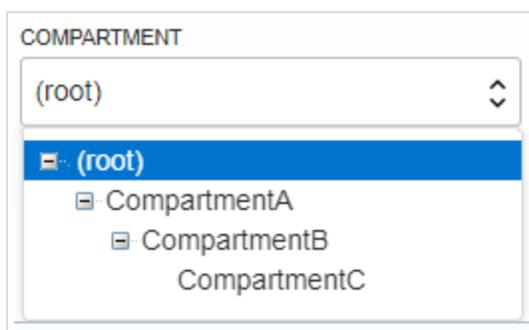
### Policies and Compartment Hierarchies

As described in the previous section, a policy statement must specify the compartment for which access is being granted (or the tenancy). Where you create the policy determines who

can update the policy. If you attach the policy to the compartment or its parent, you can simply specify the compartment name. If you attach the policy further up the hierarchy, you must specify the path. The format of the path is each compartment name (or OCID) in the path, separated by a colon:

*<compartment\_level\_1>:<compartment\_level\_2>: . . . <compartment\_level\_n>*

For example, assume you have a three-level compartment hierarchy, shown here:



You want to create a policy to allow NetworkAdmins to manage VCNs in CompartmentC. If you want to attach this policy to CompartmentC or to its parent, CompartmentB, write this policy statement:

```
Allow group NetworkAdmins to manage virtual-network-family in compartment CompartmentC
```

However, if you want to attach this policy to CompartmentA (so that only administrators of CompartmentA can modify it), write this policy statement that specifies the path:

```
Allow group NetworkAdmins to manage virtual-network-family in compartment CompartmentB:CompartmentC
```

To attach this policy to the tenancy, write this policy statement that specifies the path from CompartmentA to CompartmentC:

```
Allow group NetworkAdmins to manage virtual-network-family in compartment
CompartmentA:CompartmentB:CompartmentC
```

### Policies and Service Updates

It's possible that the definition of a verb or resource-type could change in the future. For example, let's say that the `virtual-network-family` resource-type changes to include a new kind of resource that's been added to Networking. By default, your policies automatically stay current with any changes in service definition, so any policy you have that gives access to `virtual-network-family` would automatically include access to the newly added resource. If you'd prefer a different behavior, see [Policy Language Version](#).

### Writing Policies for Each Service

The [Policy Reference](#) includes details of the specific resource-types for each service, and which verb + resource-type combination gives access to which API operations.

## Common Policies

This section includes some common policies you might want to use in your organization.



#### Note

These policies use example group and compartment names. Make sure to replace them with your own names.

### Let the Help Desk manage users

**Type of access:** Ability to create, update, and delete users and their credentials. It does not include the ability to put users in groups (see [Let group admins manage group membership](#)).

**Where to create the policy:** In the tenancy, because users reside in the tenancy.

```
Allow group HelpDesk to manage users in tenancy
```

### Let auditors inspect your resources

**Type of access:** Ability to list the resources in all compartments. Be aware that:

- The operation to list IAM policies includes the contents of the policies themselves
- The list operations for Networking resource-types return all the information (for example, the contents of security lists and route tables)
- The operation to list instances requires the `read` verb instead of `inspect`, and the contents include the user-provided metadata.
- The operation to view Audit service events requires the `read` verb instead of `inspect`.

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), auditors can then inspect both the tenancy and all compartments beneath it. Or you could choose to give auditors access to only specific compartments if they don't need access to the entire tenancy.

```
Allow group Auditors to inspect all-resources in tenancy
```

```
Allow group Auditors to read instances in tenancy
```

```
Allow group Auditors to read audit-events in tenancy
```

### Let network admins manage a cloud network

**Type of access:** Ability to manage all components in Networking. This includes cloud networks, subnets, gateways, virtual circuits, security lists, route tables, and so on. If the network admins need to launch instances to test network connectivity, see [Let users launch Compute instances](#) on this page.

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), NetworkAdmins can then manage a cloud network in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy.

```
Allow group NetworkAdmins to manage virtual-network-family in tenancy
```

### Let network admins manage load balancers

**Type of access:** Ability to manage all components in Load Balancing. If the group needs to launch instances, see [Let users launch Compute instances](#) on this page.

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), NetworkAdmins can then manage load balancers in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy.

```
Allow group NetworkAdmins to manage load-balancers in tenancy
```

If the group uses the Console to manage load balancers, an additional policy to use the associated networking resources is required:

```
Allow group NetworkAdmins to manage load-balancers in tenancy
```

```
Allow group NetworkAdmins to use virtual-network-family in tenancy
```

If a particular group needs to update existing load balancers (for example, modify the backend set) but not create or delete them, use this statement:

```
Allow group LBUUsers to use load-balancers in tenancy
```

### Let users launch Compute instances

**Type of access:** Ability to do everything with instances launched into the cloud network and subnets in compartment XYZ, and attach/detach any existing volumes that already exist in compartment ABC. The first statement also lets the group create and manage instance images in compartment ABC. If the group doesn't need to attach/detach volumes, you can delete the third statement.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of the individual compartments (ABC and XYZ) to have control over the individual policy statements for their compartments, see [Policy Attachment](#).

```
Allow group InstanceLaunchers to manage instance-family in compartment ABC
```

```
Allow group InstanceLaunchers to read app-catalog-listing in tenancy
```

```
Allow group InstanceLaunchers to use volume-family in compartment ABC
```

```
Allow group InstanceLaunchers to use virtual-network-family in compartment XYZ
```

### Let users manage Compute dedicated virtual machine hosts

**Type of access:** Ability to create, update, and delete dedicated virtual machine hosts as well as launch instances on dedicated virtual machine hosts.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the dedicated virtual machine hosts and instances in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group DedicatedVMHostAdmins to manage dedicated-vm-hosts in tenancy
```

```
Allow group DedicatedVMHostAdmins to manage instances in tenancy
```

### Let users launch Compute instances on dedicated virtual machine hosts

**Type of access:** Ability to launch instances on dedicated virtual machine hosts.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the dedicated virtual machine hosts and instances in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group DedicatedVMHostAdmins to use dedicated-vm-hosts in tenancy
```

```
Allow group DedicatedVMHostAdmins to manage instances in tenancy
```

### Let users manage Compute instance configurations, instance pools, and cluster networks

**Type of access:** Ability to do all things with instance configurations, instance pools, and cluster networks in all compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the instance configurations, instance pools, and cluster networks in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group InstancePoolAdmins to manage compute-management-family in tenancy
```

If a group needs to create instance configurations using existing instances as a template, and uses the API, SDKs, or command line interface (CLI) to do this, add the following statements to the policy:

```
Allow group InstancePoolAdmins to read instance-family in tenancy
```

```
Allow group InstancePoolAdmins to inspect volumes in tenancy
```

If a particular group needs to start, stop, or reset the instances in existing instance pools, but not create or delete instance pools, use this statement:

```
Allow group InstancePoolUsers to use instance-pools in tenancy
```

### Let users manage Compute autoscaling configurations

**Type of access:** Ability to create, update, and delete autoscaling configurations.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the autoscaling configurations in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group AutoscalingAdmins to manage auto-scaling-configurations in tenancy
```

```
Allow group AutoscalingAdmins to manage instance-pools in tenancy
```

### Let users list and subscribe to images from the Partner Image catalog

**Type of access:** Ability to list and create subscriptions to images in the Partner Image catalog. It does not include the ability to create instances using images from the Partner Image catalog (see [Let users launch Compute instances](#)).

**Where to create the policy:** In the tenancy. To reduce the scope of access to just creating subscriptions in a particular compartment, specify that compartment instead of the tenancy in the third statement.

```
Allow group CatalogSubscribers to inspect app-catalog-listing in tenancy
```

```
Allow group CatalogSubscribers to read app-catalog-listing in tenancy
```

```
Allow group CatalogSubscribers to manage app-catalog-listing in tenancy
```

### Let volume admins manage block volumes, backups, and volume groups

**Type of access:** Ability to do all things with block storage volumes, volume backups, and volume groups in all compartments with the exception of copying volume backups across regions. This makes sense if you want to have a single set of volume admins manage all the volumes, volume backups, and volume groups in all the compartments. The second statement is required in order to attach/detach the volumes from instances.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes/backups and instances in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeAdmins to manage volume-family in tenancy
```

```
Allow group VolumeAdmins to use instance-family in tenancy
```

If the group needs to also copy volume backups across regions, add the following statement to the policy:

```
Allow group VolumeAdmins to copy volume-backups in tenancy
```

### Let volume backup admins manage only backups

**Type of access:** Ability to do all things with volume backups, but not create and manage volumes themselves. This makes sense if you want to have a single set of volume backup admins manage all the volume backups in all the compartments. The first statement gives the

required access to the volume that is being backed up; the second statement enables creation of the backup (and the ability to delete backups). **The third statement enables the creation and management of user defined backup policies; the fourth statement enables assignment and removal of assignment of backup policies.**

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes and backups in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeBackupAdmins to use volumes in tenancy

Allow group VolumeBackupAdmins to manage volume-backups in tenancy

Allow group VolumeBackupAdmins to manage backup-policies in tenancy

Allow group VolumeBackupAdmins to manage backup-policy-assignments in tenancy
```

If the group will be using the Console, the following policy gives a better user experience:

```
Allow group VolumeBackupAdmins to use volumes in tenancy

Allow group VolumeBackupAdmins to manage volume-backups in tenancy

Allow group VolumeBackupAdmins to inspect volume-attachments in tenancy

Allow group VolumeBackupAdmins to inspect instances in tenancy

Allow group VolumeBackupAdmins to manage backup-policies in tenancy

Allow group VolumeBackupAdmins to manage backup-policy-assignments in tenancy
```

The last two statements are not necessary in order to manage volume backups. However, they enable the Console to display all the information about a particular volume and the available backup policies.

### Let users create a volume group

**Type of access:** Ability to create a volume group from a set of volumes.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes and volume groups in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeGroupCreators to inspect volumes in tenancy
Allow group VolumeGroupCreators to manage volume-groups in tenancy
```

### Let users clone a volume group

**Type of access:** Ability to clone a volume group from an existing volume group.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes and volume groups in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeGroupCloners to inspect volumes in tenancy
Allow group VolumeGroupCloners to manage volume-groups in tenancy
Allow group VolumeGroupCloners to manage volumes in tenancy
```

### Let users create a volume group backup

**Type of access:** Ability to create a volume group backup.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes/backups and volume groups/volume group backups in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeGroupBackupAdmins to inspect volume-groups in tenancy
Allow group VolumeGroupBackupAdmins to manage volumes in tenancy
Allow group VolumeGroupBackupAdmins to manage volume-group-backups in tenancy
Allow group VolumeGroupBackupAdmins to manage volume-backups in tenancy
```

### Let users restore a volume group backup

**Type of access:** Ability to create a volume group by restoring a volume group backup.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the volumes/backups and volume groups/volume group backups in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group VolumeGroupBackupAdmins to inspect volume-group-backups in tenancy
Allow group VolumeGroupBackupAdmins to read volume-backups in tenancy
Allow group VolumeGroupBackupAdmins to manage volume-groups in tenancy
Allow group VolumeGroupBackupAdmins to manage volumes in tenancy
```

### Let users create, manage, and delete file systems

**Type of access:** Ability to create, manage, or delete a file system. Administrative functions for a file system include the ability to rename or delete it or disconnect from it.

**Where to create the policy:** In the tenancy, so that the ability to create, manage, or delete a file system is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of these administrative functions to file systems in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group StorageAdmins to manage file-family in tenancy

Allow group StorageAdmins to manage file-family in compartment ABC
```

### Let users create file systems

**Type of access:** Ability to create a file system.

**Where to create the policy:** In the tenancy, so that the ability to create a file system is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of these administrative functions to file systems in a particular compartment, specify that compartment instead of the tenancy.

## CHAPTER 18 IAM

---

```
Allow group Managers to manage file-systems in tenancy
```

```
Allow group Managers to read mount-targets in tenancy
```

The second statement is required when users create a file system using the Console. It enables the Console to display a list of mount targets that the new file system can be associated with.

### Let Object Storage admins manage buckets and objects

**Type of access:** Ability to do all things with Object Storage buckets and objects in all compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the buckets and objects in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group ObjectAdmins to manage buckets in tenancy
```

```
Allow group ObjectAdmins to manage objects in tenancy
```

### Let users write objects to Object Storage buckets

**Type of access:** Ability to write objects to any Object Storage bucket in compartment ABC (imagine a situation where a client needs to regularly write log files to a bucket). This consists of the ability to list the buckets in the compartment, list the objects in a bucket, and create a new object in a bucket. Although the second statement gives broad access with the `manage` verb, that access is then scoped down to only the `OBJECT_INSPECT` and `OBJECT_CREATE` permissions with the condition at the end of the statement.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of compartment ABC to have control over the policy, see [Policy Attachment](#).

```
Allow group ObjectWriters to read buckets in compartment ABC
```

```
Allow group ObjectWriters to manage objects in compartment ABC where any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}
```

**Access limited to a specific bucket:** To limit access to a specific bucket in a particular compartment, add the condition `where target.bucket.name='<bucket_name>'`. The following policy allows the user to list all the buckets in a particular compartment, but they can only list the objects in and upload objects to BucketA:

```
Allow group ObjectWriters to read buckets in compartment ABC

Allow group ObjectWriters to manage objects in compartment ABC where all {target.bucket.name='BucketA',
any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}}
```

For more information about using conditions, see [Advanced Policy Features](#).

### Let users download objects from Object Storage buckets

**Type of access:** Ability to download objects from any Object Storage bucket in compartment ABC. This consists of the ability to list the buckets in the compartment, list the objects in a bucket, and read existing objects in a bucket.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of compartment ABC to have control over the policy, see [Policy Attachment](#).

```
Allow group ObjectReaders to read buckets in compartment ABC

Allow group ObjectReaders to read objects in compartment ABC
```

**Access limited to a specific bucket:** To limit access to a specific bucket in a particular compartment, add the condition `where target.bucket.name='<bucket_name>'`. The following policy allows the user to list all buckets in a particular compartment, but they can only read the objects in and download from BucketA:

```
Allow group ObjectReaders to read buckets in compartment ABC

Allow group ObjectReaders to read objects in compartment ABC where target.bucket.name='BucketA'
```

For more information about using conditions, see [Advanced Policy Features](#).

### Let database admins manage DB systems

**Type of access:** Ability to do all things with the DB system resources in all compartments.

This makes sense if you want to have a single set of database admins manage all the bare metal, virtual machine, and Exadata DB systems in all the compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the database systems in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group DatabaseAdmins to manage database-family in tenancy
```

### Let database admins manage Exadata Cloud at Customer instances

**Type of access:** Ability to do all things with the Exadata Cloud at Customer resources in all compartments. This makes sense if you want to have a single set of database admins manage all the Exadata Cloud at Customer systems in all the compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the Exadata Cloud at Customer systems in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group ExaCCAdmins to manage database-family in tenancy
```

### Let database and fleet administrators manage Autonomous Databases

**Type of access:** Ability to do all things with Autonomous Database instances in all compartments. This makes sense if you want to have a single set of database admins manage all the Autonomous Database databases in all the compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the Autonomous Databases databases in a particular compartment, specify that compartment instead of the tenancy.

**Example 1:** For [fleet admins](#). Allows Autonomous Database fleet administrator access to **both workload types** (Autonomous Data Warehouse and Autonomous Transaction

Processing), and to [dedicated deployment infrastructure resources](#) (Autonomous Container Databases and Autonomous Exadata Infrastructure resources).

```
Allow group DatabaseAdmins to manage autonomous-database-family in tenancy
```

If you need to restrict access to the Autonomous Exadata Infrastructure and Autonomous Container Database resource types (applicable only to [dedicated deployment](#) instances), you can do so by creating separate policy statements for database administrators that allow access to only Autonomous Databases and their backups. Because a policy statement can only specify one resource type, you will need to create separate statements for the database and backup resources.

**Example 2:** For [database admins](#). Allows Autonomous Database database administrator access to databases and backups of **both workload types**, but denies access to dedicated deployment infrastructure resources (Autonomous Container Databases and Autonomous Exadata Infrastructure resources).

```
Allow group ADB-Admins to manage autonomous-database in tenancy
```

```
Allow group ADB-Admins to manage autonomous-backup in tenancy
```

To reduce the scope of access to either the Autonomous Data Warehouse or Autonomous Transaction Processing workload types, use a `where` clause, as in the following examples:

**Example 3:** For Autonomous Transaction Processing database admins. Limits Autonomous Database access to **Autonomous Transaction Processing** databases and backups.

```
Allow group ADB-Admins to manage autonomous-database in tenancy where target.workloadType = 'OLTP'
```

```
Allow group ADB-Admins to manage autonomous-backup in tenancy where target.workloadType = 'OLTP'
```

**Example 4:** For Autonomous Data Warehouse admins. Limits Autonomous Database access to **Autonomous Data Warehouse** databases and backups.

```
Allow group ADB-Admins to manage autonomous-database in tenancy where target.workloadType = 'DW'
```

```
Allow group ADB-Admins to manage autonomous-backup in tenancy where target.workloadType = 'DW'
```

### Let database admins manage Autonomous Data Warehouse databases

Deprecated. See [Let database and fleet administrators manage Autonomous Databases](#) for

sample Autonomous Database policies covering the Autonomous Data Warehouse workload type.

### Let security admins manage vaults and keys

**Type of access:** Ability to do all things with the Key Management service in all compartments. This makes sense if you want to have a single set of security admins manage all the vaults and keys in all compartments.

**Where to create the policy:** In the tenancy, so that access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the keys and vaults in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group SecurityAdmins to manage vaults in tenancy
```

```
Allow group SecurityAdmins to manage keys in tenancy
```

### Let security admins manage all keys in a specific vault in a compartment

**Type of access:** Ability to do all things with keys in a specific vault in compartment ABC.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of the individual compartment (ABC) to have control over the individual policy statements for their compartment, see [Policy Attachment](#).

```
Allow group SecurityAdmins to manage keys in compartment ABC where target.vault.id='<vault_OCID>'
```

### Let security admins use a specific key in a compartment

**Type of access:** Ability to list, view, and perform cryptographic operations with a specific key in a compartment.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of the individual compartment (ABC) to have control over the individual policy statements for their compartment, see [Policy Attachment](#).

```
Allow group SecurityAdmins to use keys in compartment ABC where target.key.id='<key_OCID>'
```

## Let a user group delegate key usage in a compartment

**Type of access:** Ability to associate an Object Storage bucket, Block Volume volume, or a File Storage file system with a specific key authorized for use in a specific compartment. With this policy, a user in the specified group does not have permission to use the key itself. Rather, by association, the key can be used by Object Storage, Block Volume, or File Storage on behalf of the user to create or update an encrypted bucket, volume, or file system and to encrypt or decrypt data in the bucket, volume, or file system. This policy requires that you also have a companion policy that lets Object Storage, Block Volume, or File Storage use the key to perform cryptographic operations.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of the individual compartment (ABC) to have control over the individual policy statements for their compartment, see [Policy Attachment](#).

```
Allow group ObjectWriters, VolumeWriters, FileWriters to use key-delegate in compartment ABC where
target.key.id = '<key_OCID>'
```

## Let Block Volume, Object Storage, File Storage services encrypt and decrypt volumes, volume backups, buckets, and file systems

**Type of access:** Ability to list, view, and perform cryptographic operations with all keys in compartment ABC. Because Object Storage is a regional service, it has regional endpoints. As such, you must specify the regional service name for each region where you're using Object Storage with Key Management encryption. This policy also requires that you have a companion policy that allows a user group to use the delegated key that Object Storage, Block Volume, or File Storage will use.

**Where to create the policy:** The easiest approach is to put this policy in the tenancy. If you want the admins of the individual compartment (ABC) to have control over the individual policy statements for their compartment, see [Policy Attachment](#).

```
Allow service blockstorage, objectstorage-<region_name>, FssOclProd to use keys in compartment ABC where
target.key.id = '<key_OCID>'
```

For Object Storage, refer to the service in policies in the following ways, depending on your region:

- objectstorage-us-phoenix-1
- objectstorage-us-ashburn-1
- objectstorage-eu-frankfurt-1
- objectstorage-uk-london-1
- objectstorage-Japan East (Tokyo)

For File Storage, the service name used in the policy is `FssOclProd`.

To determine the region name value of an Oracle Cloud Infrastructure region, see [Regions and Availability Domains](#).

### Let group admins manage group membership

**Type of access:** Ability to manage the membership of a group. Does not include the ability to create or delete users, or manage their credentials (see [Let the Help Desk manage users](#)).

The first two statements let GroupAdmins list all the users and groups in the tenancy, list which users are in a particular group, and list what groups a particular user is in.

The last two statements together let GroupAdmins change a group's membership. The condition at the end of the last two statements lets GroupAdmins manage membership to all groups except the Administrators group (see [The Administrators Group and Policy](#)). You should protect membership to that group because it has power to do anything throughout the tenancy.

It might seem that the last two statements should also cover the basic listing functionality that the first two statements enable. To better understand how conditions work and why you also need the first two statements, see [Variables that Aren't Applicable to a Request Result in a Declined Request](#).

**Where to create the policy:** In the tenancy, because users and groups reside in the tenancy.

```
Allow group GroupAdmins to inspect users in tenancy
```

```
Allow group GroupAdmins to inspect groups in tenancy
```

## CHAPTER 18 IAM

---

```
Allow group GroupAdmins to use users in tenancy where target.group.name != 'Administrators'
```

```
Allow group GroupAdmins to use groups in tenancy where target.group.name != 'Administrators'
```

### Let users manage their own passwords and credentials

No policy is required to let users manage *their own* credentials. All users have the ability to change and reset their own passwords, manage their own API keys, and manage their own auth tokens. For more information, see [User Credentials](#).

### Let a compartment admin manage the compartment

**Type of access:** Ability to manage all aspects of a particular compartment. For example, a group called A-Admins could manage all aspects of a compartment called Project-A, including writing additional policies that affect the compartment. For more information, see [Policy Attachment](#). For an example of this kind of setup and additional policies that are useful, see [Example Scenario](#).

**Where to create the policy:** In the tenancy.

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

### Restrict admin access to a specific region

**Type of access:** Ability to manage resources in a specific region. Remember that IAM resources must be managed in the home region. If the specified region is not the home region, then the Admin will not be able to manage IAM resources. For more information about the home region, see [Managing Regions](#).

**Where to create the policy:** In the tenancy.

```
Allow group PHX-Admins to manage all-resources in tenancy where request.region='phx'
```

The preceding policy allows PHX-Admins to manage all aspects of all resources in US West (Phoenix).

Members of the PHX-Admins group can only manage IAM resources if the tenancy's home region is US West (Phoenix).

### Restrict user access to view only summary announcements

**Type of access:** Ability to view the summary versions of announcements about the operational status of Oracle Cloud Infrastructure services.

**Where to create the policy:** In the tenancy.

```
Allow group AnnouncementListeners to inspect announcements in tenancy
```

The preceding policy allows AnnouncementListeners to view a list of summary announcements.

### Let users view details of announcements

**Type of access:** Ability to view the details of announcements about the operational status of Oracle Cloud Infrastructure services.

**Where to create the policy:** In the tenancy.

```
Allow group AnnouncementReaders to read announcements in tenancy
```

The preceding policy allows AnnouncementReaders to view a list of summary announcements and the details of specific announcements.

### Let streaming users manage streams

**Type of access:** Ability to do all things with the Streaming service in all compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the streams in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group StreamAdmins to manage streams in tenancy
```

### Let streaming users publish messages to streams

**Type of access:** Ability to produce messages to streams with the Streaming service in all compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the streams in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group StreamUsers to use stream-push in tenancy
```

### Let streaming users publish messages to a specific stream

**Type of access:** Ability to produce messages to a stream with the Streaming service.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the streams in a particular compartment, specify that compartment instead of the tenancy.

```
allow group StreamUsers to use stream-pull in tenancy where target.stream.id = '<stream_OCID>'
```

### Let streaming users consume messages from streams

**Type of access:** Ability to consume messages from streams with the Streaming service in all compartments.

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the streams in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group StreamUsers to use stream-pull in tenancy
```

### Let users view metric definitions in a compartment

**Type of access:** Ability to view metric definitions in a specific compartment. For more information about metrics, see [Metrics Feature Overview](#).

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the metric definitions in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group MetricReaders to inspect metrics in compartment ABC
```

### Let users access monitoring metrics in a compartment

**Type of access:** Ability to view and retrieve monitoring metrics for supported resources in a specific compartment. For more information about metrics, see [Metrics Feature Overview](#).

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just the metrics in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group MetricReaders to read metrics in compartment ABC
```

### Restrict user access to a specific metric namespace

**Type of access:** Ability to view and retrieve monitoring metrics for resources under a specific metric namespace. For more information about metrics, see [Metrics Feature Overview](#).

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to the specified metric namespace to just within a particular compartment, specify that compartment instead of the tenancy.

```
Allow group MetricReaders to read metrics in compartment ABC where target.metrics.namespace='oci_computeagent'
```

The preceding policy allows `MetricReaders` to view and retrieve metric data points from all [monitoring-enabled](#) Compute instances in the `ABC` compartment.

### Let users publish custom metrics

**Type of access:** Ability to publish custom metrics under a specific metric namespace to the Monitoring service. For instructions, see [Publishing Custom Metrics](#).

**Where to create the policy:** In the tenancy, so that the access is easily granted to all compartments by way of [policy inheritance](#). To reduce the scope of access to just metrics in a particular compartment, specify that compartment instead of the tenancy.

```
Allow group MetricPublishers to use metrics in tenancy where
target.metrics.namespace='mycustomnamespace'
```

The preceding policy allows `MetricPublishers` to publish data points for the custom metric namespace `mycustomnamespace` in the tenancy.

### Let instances make API calls to access monitoring metrics in the tenancy

**Type of access:** Ability to call the Monitoring API for access to monitoring metrics. The instances on which API requests originate must be members of the dynamic group indicated in the policy. For more information, see [Calling Services from an Instance](#) and [Metrics Feature Overview](#).

**Where to create the policy:** In the tenancy.

```
Allow dynamic-group MetricInstances to read metrics in tenancy
```

The preceding policy allows applications that are running on Compute instances in the dynamic group `MetricInstances` to send API requests to the Monitoring service in the tenancy.

### Let users view alarms

**Type of access:** Ability to view alarms for supported resources in tenancy. Does not include the ability to create alarms or to create or delete topics. For more information about alarms, see [Alarms Feature Overview](#).

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), AlarmUsers can then view alarms in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy.

```
Allow group AlarmUsers to read alarms in tenancy
```

```
Allow group AlarmUsers to read metrics in tenancy
```

### Let users manage alarms

**Type of access:** Ability to view and create alarms with existing notification topics for supported resources in the tenancy. Does not include the ability to create or delete topics. For more information about alarms, see [Alarms Feature Overview](#).

All statements are required to let AlarmUsers create alarms.

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), AlarmUsers can then view and create alarms in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy.

```
Allow group AlarmUsers to manage alarms in tenancy
```

```
Allow group AlarmUsers to read metrics in tenancy
```

```
Allow group AlarmUsers to use ons-topics in tenancy
```

### Let users manage alarms and create topics

**Type of access:** Ability to view and create alarms (with new or existing topics) for supported resources in tenancy. Also includes the ability to create subscriptions in the tenancy, to publish messages (broadcast notification messages) to all subscriptions in the tenancy, and to move alarms to different compartments in the tenancy. For more information about alarms, see [Alarms Feature Overview](#).

**Where to create the policy:** In the tenancy. Because of the concept of [policy inheritance](#), AlarmUsers can then view and create alarms in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy.

```
Allow group AlarmUsers to manage alarms in tenancy
```

```
Allow group AlarmUsers to read metrics in tenancy
```

```
Allow group AlarmUsers to manage ons-topics in tenancy
```

### Allow a group to manage topics

**Type of access:** Ability to get, create, update, and delete topics in the tenancy, as well as move topics to different compartments in the tenancy. Also includes the ability to create subscriptions in the tenancy and to publish messages (broadcast notification messages) to all subscriptions in the tenancy.

**Where to create the policy:** In the tenancy.

```
Allow group A-Admins to manage ons-topics in tenancy
```

### Allow a group to manage topic subscriptions

**Type of access:** Ability to list, create, update, and delete subscriptions for topics in the tenancy. Ability to move subscriptions to different compartments in the tenancy.

**Where to create the policy:** In the tenancy.

```
Allow group A-Admins to manage ons-subscriptions in tenancy
```

### Allow a group to publish messages to topics

**Type of access:** Ability to broadcast notification messages to all subscriptions in the tenancy, as well as list, create, update, and delete subscriptions in the tenancy.

**Where to create the policy:** In the tenancy.

## CHAPTER 18 IAM

---

```
Allow group A-Admins to use ons-topics in tenancy
```

### Let users list Events rules in a compartment

**Type of access:** Ability to list Events rules.

**Where to create the policy:** In the tenancy.

```
Allow group RuleReaders to read cloudevents-rules in tenancy
```

The preceding policy allows RuleReaders to list rules in the tenancy.

### Let admins manage Events rules in a compartment

**Type of access:** Ability to manage Events rules, including creating, deleting and updating rules.

**Where to create the policy:** In the tenancy.

This line gives the user inspect access to resources in compartments to select actions.

```
allow group <RuleAdmins> to inspect compartments in tenancy
```

This line gives the user access to defined tags to apply filter tags to rules.

```
allow group <RuleAdmins> to use tag-namespaces in tenancy
```

These lines give the user access to Streaming resources for actions

```
allow group <RuleAdmins> to inspect streams in tenancy
allow group <RuleAdmins> to use stream-push in tenancy
allow group <RuleAdmins> to use stream-pull in tenancy
```

These lines give the user access to Functions resources for actions.

```
allow group <RuleAdmins> to use virtual-network-family in tenancy
allow group <RuleAdmins> to manage function-family in tenancy
```

This line give the user access to Notifications topics for actions.

```
allow group <RuleAdmins> to use ons-topic in tenancy
```

This line gives the user manage access to rules for Events.

```
allow group <RuleAdmins> to manage cloudevents-rules in tenancy
```

## Advanced Policy Features

This section describes policy language features that let you grant more granular access.

### Conditions

As part of a policy statement, you can specify one or more *conditions* that must be met in order for access to be granted. For a simple example, see [Let group admins manage group membership](#).

Each condition consists of one or more predefined variables that you specify values for in the policy statement. Later, when someone requests access to the resource in question, if the condition in the policy is met, it evaluates to *true* and the request is allowed. If the condition is not met, it evaluates to *false* and the request is not allowed.

There are two types of variables: those that are relevant to the request itself, and those relevant to the resource being acted upon in the request, also known as the *target*. The name of the variable is prefixed accordingly with either `request` or `target` followed by a period. For example, there's a request variable called `request.operation` to represent the API operation being requested. This variable lets you write a broad policy statement, but add a condition based on the specific API operation. For an example, see [Let users write objects to Object Storage buckets](#).

**Important**

Condition matching is case insensitive. This is important to remember when writing conditions for resource types that allow case-sensitive naming. For example, the Object Storage service allows you to create both a bucket named "BucketA" and a bucket named "bucketA" in the same compartment. If you write a condition that specifies "BucketA", it will apply also to "bucketA", because the condition matching is case insensitive.

**Variables that Aren't Applicable to a Request Result in a Declined Request**

If the variable is *not applicable* to the incoming request, the condition evaluates to *false* and the request is declined. For example, here are the basic policy statements that together let someone add or remove users from any group except Administrators:

```
Allow group GroupAdmins to use users in tenancy
 where target.group.name != 'Administrators'
```

```
Allow group GroupAdmins to use groups in tenancy
 where target.group.name != 'Administrators'
```

Given the above policy, if GroupAdmins tried to call a general API operation for users such as `ListUsers` or `UpdateUser` (which lets you change the user's description), the request would be declined, even though those API operations are covered by `use users`. This is because the above policy statement for `use users` also includes the `target.group.name` variable, but the `ListUsers` or `UpdateUser` request doesn't involve specifying a group. There is no `target.group.name` for those requests, so the request is declined.

If you want to also grant access to general user API operations that don't involve a particular group, you would need an additional statement that gives the level of access you want to grant, *but without the condition*. For example, if you want to grant access to `ListUsers`, you need this additional statement:

```
Allow group GroupAdmins to inspect users in tenancy
```

Or if you want to grant access to `UpdateUser`, you need this additional statement (which also covers `ListUsers` because the `use` verb includes the capabilities of the `inspect` verb):

```
Allow group GroupAdmins to use users in tenancy
```

This general concept also applies to groups (e.g., `ListGroups` and `UpdateGroup`), and any other resource type with target variables.

For more information about the syntax of conditions, see [Conditions](#). For a list of all the variables you can use in policies, see the tables in the [Policy Reference](#).

## Permissions

*Permissions* are the atomic units of authorization that control a user's ability to perform operations on resources. Oracle defines all the permissions in the policy language. When you write a policy giving a group access to a particular [verb](#) and resource-type, you're actually giving that group access to one or more predefined permissions. The purposes of verbs is to simplify the process of granting multiple related permissions that cover a broad set of access or a particular operational scenario. The next sections give more details and examples.

### Relation to Verbs

To understand the relationship between permissions and verbs, let's look at an example. A policy statement that allows a group to `inspect volumes` actually gives the group access to a permission called `VOLUME_INSPECT` (permissions are always written with all capital letters and underscores). In general, that permission enables the user to get information about block volumes.

As you go from `inspect > read > use > manage`, the level of access generally increases, and the permissions granted are cumulative. The following table shows the permissions included with each verb for the `volumes` resource-type. Notice that no additional permissions are granted going from `inspect` to `read`.

Inspect Volumes	Read Volumes	Use Volumes	Manage Volumes
VOLUME_INSPECT	VOLUME_INSPECT	VOLUME_INSPECT VOLUME_UPDATE VOLUME_WRITE	VOLUME_INSPECT VOLUME_UPDATE VOLUME_WRITE VOLUME_CREATE VOLUME_DELETE

The [policy reference](#) lists the permissions covered by each verb for each given resource-type. For example, for block volumes and other resources covered by the Core Services, see the tables in [Details for Verb + Resource-Type Combinations](#). The left column of each of those tables lists the permissions covered by each verb. The other sections of the policy reference include the same kind of information for the other services.

### Relation to API Operations

Each API operation requires the caller to have access to one or more permissions. For example, to use either `ListVolumes` or `GetVolume`, you must have access to a single permission: `VOLUME_INSPECT`. To attach a volume to an instance, you must have access to multiple permissions, some of which are related to the `volumes` resource-type, some to the `volume-attachments` resource-type, and some related to the `instances` resource-type:

- `VOLUME_WRITE`
- `VOLUME_ATTACHMENT_CREATE`
- `INSTANCE_ATTACH_VOLUME`

The policy reference lists which permissions are required for each API operation. For example, for the Core Services API operations, see the table in [Permissions Required for Each API Operation](#).

### Understanding a User's Access

The policy language is designed to let you write simple statements involving only verbs and resource-types, without having to state the desired permissions in the statement. However, there may be situations where a security team member or auditor wants to understand the specific permissions a particular user has. The tables in the [policy reference](#) show each verb and the associated permissions. You can look at the groups the user is in and the policies applicable to those groups, and from there compile a list of the permissions granted. However, having a list of the permissions isn't the complete picture. Conditions in a policy statement can scope a user's access beyond individual permissions (see the next section). Also, each policy statement specifies a particular compartment and can have conditions that further scope the access to only certain resources in that compartment.

### Scoping Access with Permissions or API Operations

In a policy statement, you can use [conditions](#) combined with permissions or API operations to reduce the scope of access granted by a particular verb.

For example, let's say you want group XYZ to be able to list, get, create, or update groups (i.e., change their description), but not delete them. To list, get, create, and update groups, you need a policy with `manage groups` as the verb and resource-type. According to the table in [Details for Verbs + Resource-Type Combinations](#), the permissions covered are:

- GROUP\_INSPECT
- GROUP\_UPDATE
- GROUP\_CREATE
- GROUP\_DELETE

To restrict access to only the desired permissions, you could add a condition *that explicitly states the permissions you want to allow*:

```
Allow group XYZ to manage groups in tenancy

where any {request.permission='GROUP_INSPECT',
 request.permission='GROUP_CREATE',
 request.permission='GROUP_UPDATE'}
```

An alternative would be a policy that *allows all permissions except* GROUP\_DELETE:

## CHAPTER 18 IAM

---

```
Allow group XYZ to manage groups in tenancy where request.permission != 'GROUP_DELETE'
```

However, with this approach, be aware that any new permissions the service might add in the future would automatically be granted to group XYZ. Only GROUP\_DELETE would be omitted.

Another alternative would be to write a condition *based on the specific API operations*. Notice that according to the table in [Permissions Required for Each API Operation](#), both ListGroups and GetGroup require only the GROUP\_INSPECT permission. Here's the policy:

```
Allow group XYZ to manage groups in tenancy
```

```
where any {request.operation='ListGroups',
 request.operation='GetGroup',
 request.operation='CreateGroup',
 request.operation='UpdateGroup'}
```

It can be beneficial to use permissions instead of API operations in conditions. In the future, if a new API operation is added that requires one of the permissions listed in the permissions-based policy above, that policy will already control XYZ group's access to that new API operation.

But notice that you can further scope a user's access to a permission by *also* specifying a condition based on API operation. For example, you could give a user access to GROUP\_INSPECT, but then only to ListGroups.

```
Allow group XYZ to manage groups in tenancy
```

```
where all {request.permission='GROUP_INSPECT',
 request.operation='ListGroups'}
```

### Policy Language Version

As mentioned in [Policies and Service Updates](#), by default your policies stay current with any changes to the definitions of verbs and resources as the services change. If you'd prefer to limit access according to the definitions that were current on a specific date, you can do that. When creating a policy, you can specify a date, and that is considered its "version". You can also update the version for an existing policy. For more information, see [To create a policy](#) and also [To update the version date for an existing policy](#).

## Policy Syntax

The overall syntax of a policy statement is as follows:

```
Allow <subject> to <verb> <resource-type> in <location> where <conditions>
```

Spare spaces or line breaks in the statement have no effect.

For limits on the number of policies and statements, see [Service Limits](#).

### Subject

Specify one or more comma-separated groups by name or OCID. Or specify `any-user` to cover all users in the tenancy.

**Syntax:** `group <group_name> | group id <group_ocid> | dynamic-group <dynamic-group_name> | dynamic-group id<dynamic-group_ocid> | any-user`

#### Examples:

- To specify a single group by name:

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

- To specify multiple groups by name (a space after the comma is optional):

```
Allow group A-Admins, B-Admins to manage all-resources in compartment Projects-A-and-B
```

- To specify a single group by OCID (the OCID is shortened for brevity):

```
Allow group
 id ocid1.group.oc1..aaaaaaaaqjihfvxmum...awuc7i5xwe6s7qmnsbc6a
to manage all-resources in compartment Project-A
```

- To specify multiple groups by OCID (the OCIDs are shortened for brevity):

```
Allow group
 id ocid1.group.oc1..aaaaaaaaqjihfvxmumrl...wuc7i5xwe6s7qmnsbc6a,
 id ocid1.group.oc1..aaaaaaaavhea5mellwzb...66yfxv1462tdgx2oecyq
```

```
to manage all-resources in compartment Projects-A-and-B
```

- To specify any user in the tenancy:

```
Allow any-user to inspect users in tenancy
```

### Verb

Specify a single verb. For a list of verbs, see [Verbs](#). Example:

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

### Resource-Type

Specify a single resource-type, which can be one of the following:

- An individual resource-type (e.g., *vcns*, *subnets*, *instances*, *volumes*, *etc.*)
- A family resource-type (e.g., *virtual-network-family*, *instance-family*, *volume-family*, *etc.*)
- *all-resources*: Covers all resources in the compartment (or tenancy).

A family resource-type covers a variety of components that are typically used together. This makes it easier to write a policy that gives someone access to work with various aspects of your cloud network.

For a list of the available resource-types, see [Resource-Types](#).

**Syntax:** *<resource\_type>* | *all-resources*

#### Examples:

- To specify a single resource-type:

```
Allow group HelpDesk to manage users in tenancy
```

- To specify multiple resource-types, use separate statements:

```
Allow group A-Users to manage instance-family in compartment Project-A
```

```
Allow group A-Users to manage volume-family in compartment Project-A
```

- To specify all resources in the compartment (or tenancy):

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

### Location

Specify a single compartment or compartment path by name or OCID. Or simply specify `tenancy` to cover the entire tenancy. Remember that users, groups, and compartments reside in the tenancy. Policies can reside in (i.e., be attached to) either the tenancy or a child compartment.



#### Note

*Granting Access to Specific Regions or Availability Domains*

To create a policy that gives access to a specific region or availability domain, use the `request.region` or `request.ad` variable with a condition. See [Conditions](#).

The location is required in the statement. If you want to attach a policy to a compartment, you must be in that compartment when you create the policy. For more information, see [Policy Attachment](#).

To specify a compartment that is not a direct child of the compartment you are attaching the policy to, specify the path to the compartment, using the colon (:) as a separator. For more information, see [Policies and Compartment Hierarchies](#).

**Syntax:** [ `tenancy` | `compartment` *<compartment\_name>* | `compartment id` *<compartment\_ocid>* ]

#### Examples:

- To specify a compartment by name:

```
Allow group A-Admins to manage all-resources in compartment Project-A
```

- To specify a compartment by OCID:

```
Allow group
 id ocid1.group.oc1..aaaaaaaavhea5mellwzbmplwrpum46xfc73sb4rm66yfxvl462tdgx2oecyq
to manage all-resources in compartment
 id ocid1.compartment.oc1..aaaaaaaayzfq...4fmameqh7lclihrvur7xq
```

- To specify multiple compartments, use separate statements:

```
Allow group InstanceAdmins to manage instance-family in compartment Project-A

Allow group InstanceAdmins to manage instance-family in compartment Project-B
```

- To specify multiple compartments by OCID, use separate statements:

```
Allow group id
 ocd1.group.oc1..aaaaaaaavhea5mell...b4rm66yfxvl462tdgx2oecyq
to manage all-resources in compartment id
 ocid1.compartment.oc1..aaaaaaaayzfqi...ameq4h7lclihrvur7xq

Allow group id
 ocd1.group.oc1..aaaaaaaavhea5mell...b4rm66yfxvl462tdgx2oecyq
to manage all-resources in compartment id
 ocid1.compartment.oc1..aaaaaaaaphfjutov5s...vyypl1btctehngq756a
```

- To specify a compartment that is not a direct child of the compartment where you are attaching the policy, specify the path:

```
Allow group InstanceAdmins to manage instance-family in compartment Project-A:Project-A2
```

## Conditions

Specify one or more conditions. Use `any` or `all` with multiple conditions for a logical OR or AND, respectively.

**Syntax for a single condition:** `variable =|!= value`

**Syntax for multiple conditions:** `any|all {<condition>,<condition>,...}`

**Important**

Condition matching is case insensitive. This is important to remember when writing conditions for resource types that allow case-sensitive naming. For example, the Object Storage service allows you to create both a bucket named "BucketA" and a bucket named "bucketA" in the same compartment. If you write a condition that specifies "BucketA", it will apply also to "bucketA", because the condition matching is case insensitive.

For a list of variables supported by all the services, see [General Variables for All Requests](#). Also see the details for each service in the [Policy Reference](#). Here are the types of values you can use in conditions:

Type	Examples
String	'johnsmith@example.com' 'ocidl.compartment.oc1..aaaaaaaaph...ctehnqg756a' (single quotation marks are required around the value)
Pattern	/HR*/ (matches strings that start with "HR") /*HR/ (matches strings that end with "HR") /*HR*/ (matches strings that contain "HR")

Examples:

**Note**

In the following examples, the statements that specify the condition do not let GroupAdmins actually list all the users and groups, therefore statements including the `inspect` verb are added for completeness. To understand why this is required, see [Variables that Aren't Applicable to a Request Result in a Declined Request](#).

- A single condition.

The following policy enables the GroupAdmins group to create, update, or delete any groups with names that start with "A-Users-":

```
Allow group GroupAdmins to manage groups in tenancy where target.group.name = /A-Users-*/
Allow group GroupAdmins to inspect groups in tenancy
```

The following policy enables the GroupAdmins group to manage the membership of any group besides the Administrators group. (Note that you must include separate statements for `inspectaccess` because the `target.group.name` variable is not used by the ListUsers and ListGroups operations):

```
Allow group GroupAdmins to inspect users in tenancy
Allow group GroupAdmins to use users in tenancy where target.group.name != 'Administrators'
Allow group GroupAdmins to inspect groups in tenancy
Allow group GroupAdmins to use groups in tenancy where target.group.name != 'Administrators'
```

The following policy enables the NetworkAdmins group to manage cloud networks in any compartment except the one specified:

```
Allow group NetworkAdmins to manage virtual-network-family in tenancy where target.compartment.id
!= 'ocid1.compartment.oc1..aaaaaaaayzfqeibduyox6icmdol6zyar3ugly4fameq4h7lcdlihrvur7xq'
```

- Multiple conditions.

The following policy lets GroupAdmins create, update, or delete any groups whose names start with "A-", except for the A-Admins group itself:

```
Allow group GroupAdmins to manage groups in tenancy where all {target.group.name=/A-
*/ ,target.group.name!='A-Admins'}

Allow group GroupAdmins to inspect groups in tenancy
```

## Policy Reference

This reference includes:

- [Verbs](#): A list of the available actions to pair with a resource-type
- [Resource-Types](#): A list of the main resource-types
- [General Variables for All Requests](#): Variables you can use when writing policies for any resource-type
- [Details for Analytics Cloud](#)
- [Details for the Announcements Service](#)
- [Details for the Audit Service](#)
- [Details for Container Engine for Kubernetes](#)
- [Details for the Core Services](#) (this includes Networking, Compute, and Block Volume)
- [Details for the Database Service](#)
- [Details for the DNS Service](#)
- [Details for Digital Assistant](#)
- [Details for the Email Service](#)
- [Details for the Events Service](#)
- [Details for the File Storage Service](#)
- [Details for Functions](#)
- [Details for the Health Checks Service](#)
- [Details for IAM](#)
- [Details for Integration](#)
- [Details for the Key Management Service](#)

- [Details for Load Balancing](#)
- [Details for Monitoring](#)
- [Details for the Notifications Service](#)
- [Details for Object Storage, Archive Storage, and Data Transfer](#)
- [Details for Registry](#)
- [Details for Resource Manager](#)
- [Details for the Search Service](#)
- [Details for the Streaming Service](#)
- [Details for the WAF Service](#)

For instructions on how to create and manage policies using the Console or API, see [Managing Policies](#).

## Verbs

The verbs are listed in order of least amount of ability to most. The exact meaning of a each verb depends on which resource-type it's paired with. The tables later in this section show the API operations covered by each combination of verb and resource-type.

Verb	Types of Access Covered	Target User
<code>inspect</code>	Ability to list resources, without access to any confidential information or user-specified metadata that may be part of that resource. <b>Important:</b> The operation to list policies includes the contents of the policies themselves, and the list operations for the Networking resource-types return all the information (e.g., the contents of security lists and route tables).	Third-party auditors
<code>read</code>	Includes <code>inspect</code> plus the ability to get user-specified metadata and the actual resource itself.	Internal auditors

Verb	Types of Access Covered	Target User
use	Includes <code>read</code> plus the ability to work with existing resources (the actions vary by resource type). Includes the ability to update the resource, except for resource-types where the "update" operation has the same effective impact as the "create" operation (e.g., <code>UpdatePolicy</code> , <code>UpdateSecurityList</code> , etc.), in which case the "update" ability is available only with the <code>manage</code> verb. In general, this verb does not include the ability to create or delete that type of resource.	Day-to-day end users of resources
manage	Includes all permissions for the resource.	Administrators

## Resource-Types

The family resource-types are listed below. For the individual resource-types that make up each family, follow the links.

- `all-resources`: All Oracle Cloud Infrastructure resource-types
- `cluster-family`: See [Details for Container Engine for Kubernetes](#)
- `database-family`: See [Details for the Database Service](#)
- `dns`: See [Details for the DNS Service](#)
- `file-family`: See [Details for the File Storage Service](#)
- `instance-family`: See [Details for the Core Services](#)
- `object-family`: See [Details for Object Storage, Archive Storage, and Data Transfer](#)
- `virtual-network-family`: See [Details for the Core Services](#)
- `volume-family`: See [Details for the Core Services](#)

IAM has no family resource-type, only individual ones. See [Details for IAM](#).

## General Variables for All Requests

You use variables when adding conditions to a policy. For more information, see [Conditions](#). Here are the general variables applicable to all requests.

Name	Type	Description
<code>request.user.id</code>	Entity (OCID)	The OCID of the requesting user.
<code>request.user.mfaTotpVerified</code>	Boolean	Whether the user has been verified by multi-factor authentication (MFA). To restrict access to only MFA-verified users, add the condition  where <code>request.user.mfaTotpVerified = 'true'</code>  See <a href="#">Managing Multi-Factor Authentication</a> for information on setting up MFA.
<code>request.groups.id</code>	List of entities (OCIDs)	The OCIDs of the groups the requesting user is in.
<code>target.compartment.id</code>	Entity (OCID)	The OCID of the compartment containing the primary resource.  <b>Note:</b> <code>target.compartment.id</code> and <code>target.compartment.name</code> cannot be used with a "List" API operation to filter the list based on the requesting user's access to the compartment.

## CHAPTER 18 IAM

---

Name	Type	Description
<code>target.compartment.name</code>	String	The name of the compartment specified in <code>target.compartment.id</code> .
<code>request.operation</code>	String	The API operation name being requested (for example, <a href="#">ListUsers</a> ).
<code>request.permission</code>	String	The underlying permission being requested (see <a href="#">Permissions</a> ).

Name	Type	Description
request.region	String	<p>The 3-letter key for the region the request is made in. Allowed values are:</p> <ul style="list-style-type: none"><li>• BOM - use for India West (Mumbai)</li><li>• FRA - use for Germany Central (Frankfurt)</li><li>• GRU - use for Brazil East (Sao Paulo)</li><li>• IAD - use for US East (Ashburn)</li><li>• ICN - use for South Korea Central (Seoul)</li><li>• LHR - use for UK South (London)</li><li>• NRT - use for Japan East (Tokyo)</li><li>• PHX - use for US West (Phoenix)</li><li>• SYD - use for Australia East (Sydney)</li><li>• YYZ - use for Canada Southeast (Toronto)</li><li>• ZRH - use for Switzerland North (Zurich)</li></ul>
request.ad	String	<p>The name of the availability domain the request is made in. To get a list of availability domain names, use the <a href="#">ListAvailabilityDomains</a> operation.</p>

## Details for Analytics Cloud

For the details for writing policies to control access to Analytics Cloud, see [Details for Analytics Cloud](#).

## Details for the Announcements Service

This topic covers details for writing policies to control access to the Announcements service.

### Resource-Types

- `announcements`

### Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `announcements` resource-type includes the same permissions and API operations as the `inspect` verb, plus the `ANNOUNCEMENT_READ` permission and an additional API operation, `GetAnnouncement`. However, the `use` verb and `manage` verbs cover no extra permissions or API operations compared to `read`.

## announcements

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
ANNOUNCEMENT_LIST	ListAnnouncements	none

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i>	<i>INSPECT +</i>	<i>none</i>
ANNOUNCEMENT_READ	GetAnnouncement	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListAnnouncements	ANNOUNCEMENT_LIST
GetAnnouncement	ANNOUNCEMENT_READ

### Details for the Audit Service

This topic covers details for writing policies to control access to the Audit service.

### Resource-Types

`audit-events`

## Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `use` and `manage` verbs for the `audit-events` resource-type cover no extra permissions or API operations compared to the `read` verb.

### AUDIT-EVENTS

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<i>none</i>	<i>none</i>	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
AUDIT_EVENT_READ	ListEvents	<i>none</i>

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListEvents	AUDIT_EVENT_READ

### Details for the Core Services

This topic covers details for writing policies to control access to the Core Services (Networking, Compute, and Block Volume).

#### Resource-Types

##### Networking

###### *AGGREGATE RESOURCE-TYPE*

virtual-network-family

###### *INDIVIDUAL RESOURCE-TYPES*

vcns

subnets

route-tables

network-security-groups

security-lists

dhcp-options

private-ips

public-ips

ipv6s

internet-gateways

nat-gateways

service-gateways

local-peering-gateways (which includes local-peering-from, and local-peering-to)

remote-peering-connections (which includes remote-peering-from, and remote-peering-to)

drgs

drg-attachments

cpes

ipsec-connections

cross-connects

cross-connect-groups

virtual-circuits

vnics

vnic-attachments

### COMMENTS

A policy that uses `<verb> virtual-network-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `virtual-network-family`.

## Compute

*INSTANCE-FAMILY AGGREGATE RESOURCE-TYPE*

## CHAPTER 18 IAM

---

The `instance-family` aggregate resource-type covers these individual resource-types:

`app-catalog-listing`

`console-histories`

`instances`

`instance-console-connection`

`instance-images`

`volume-attachments` (includes only the permissions required for attaching volumes to instances)

### *COMPUTE-MANAGEMENT-FAMILY AGGREGATE RESOURCE-TYPE*

The `compute-management-family` aggregate resource-type covers these individual resource-types:

`instance-configurations`

`instance-pools`

`cluster-networks`

### *ADDITIONAL INDIVIDUAL RESOURCE-TYPES*

`auto-scaling-configurations`

`work-requests`

`dedicated-vm-hosts`

### *COMMENTS*

A policy that uses `<verb> instance-family` or `<verb> compute-management-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types in the family.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type.

### Block Volume

#### *AGGREGATE RESOURCE-TYPE*

volume-family

#### *INDIVIDUAL RESOURCE-TYPES*

volumes

volume-attachments

volume-backups

boot-volume-backups

backup-policies

backup-policy-assignments

volume-groups

volume-group-backups

#### *COMMENTS*

A policy that uses `<verb> volume-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `volume-family`.

### Supported Variables

The Core Services support all the general variables, plus the ones listed here. For more information about general variables supported by Oracle Cloud Infrastructure services, see [General Variables for All Requests](#).

Variable	Variable Type	Comments
<code>target.boot-volume.kms-key.id</code>	String	Use this variable to control whether Compute instances can be launched with boot volumes that were created without a Key Management master encryption key.

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` and `use` verbs for the `vcns` resource-type cover no extra permissions or API operations compared to the `inspect` verb. However, the `manage` verb includes several extra permissions and API operations.

#### FOR VIRTUAL-NETWORK-FAMILY RESOURCE TYPES

##### VCNS

###### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>VCN_READ</code>	<code>ListVcns</code> <code>GetVcn</code>	<code>CreateNatGateway</code> , <code>DeleteNatGateway</code>  (both also need <code>manage nat-gateways</code> and <code>manage vcns</code> )  <b>Note:</b> The above operations in this cell are totally covered with just <code>manage virtual-network-family</code> .

###### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

## CHAPTER 18 IAM

---

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*



## CHAPTER 18 IAM

---

### **Permissions**

*USE +*

VCN\_ATTACH

VCN\_DETACH

VCN\_UPDATE

VCN\_CREATE

VCN\_DELETE

VCN\_MOVE

### **APIs Fully Covered**

*USE +*

CreateVcn

UpdateVcn

DeleteVcn

ChangeVcnCompartment



## subnets

## INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
SUBNET_READ	ListSubnets GetSubnet	none

## READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

## USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + SUBNET_ATTACH SUBNET_DETACH	no extra	LaunchInstance (also need use vnics, use network-security-groups, and manage instance-family) TerminateInstance (also need manage instance-family, and use volumes if a volume is attached) AttachVnic (also need manage instances, use network-security-groups, and either use vnics or use instance-family) DetachVnic (also need manage instances and either use vnics or use instance-family) CreatePrivateIp, DeletePrivateIp (both also need use private-ips and use vnics)

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>USE +</i></p> <p>SUBNET_CREATE</p> <p>SUBNET_UPDATE</p> <p>SUBNET_DELETE</p> <p>SUBNET_MOVE</p>	<p><i>no extra</i></p> <p>ChangeSubnetCompartment</p>	<p><i>USE +</i></p> <p>CreateSubnet, DeleteSubnet (both also need manage vcn's, manage route-tables, manage security-lists, manage dhcp-options)</p> <p>UpdateSubnet (also need manage route-tables if changing which route table is associated with the subnet, manage security-lists if changing which security lists are associated with the subnet, and manage dhcp-options if changing which set of DHCP options is associated with the subnet)</p> <p><b>Note:</b> The above operations in this cell are covered with just <code>manage virtual-network-family</code>.</p>

route-tables

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
<p>ROUTE_TABLE_READ</p>	<p>ListRouteTables</p> <p>GetRouteTable</p>	<p><i>none</i></p>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>no extra</i></p>	<p><i>none</i></p>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>no extra</i></p>	<p><i>none</i></p>

## MANAGE

**Permissions***USE +*

ROUTE\_TABLE\_ATTACH

ROUTE\_TABLE\_DETACH

ROUTE\_TABLE\_UPDATE

ROUTE\_TABLE\_CREATE

ROUTE\_TABLE\_DELETE

ROUTE\_TABLE\_MOVE

**APIs Fully Covered***no extra*

ChangeRouteTableCompartment

**APIs Partially Covered**

CreateRouteTable, DeleteRouteTable

(both also need `manage vcns`, `manage internet-gateways`, `manage drgs`, `manage private-ips`, `manage local-peering-gateways`, `use nat-gateways`, `use service-gateways`)

UpdateRouteTable (also need `manage internet-gateways`, `manage drgs`, `manage private-ips`, `manage local-peering-gateways`, `use nat-gateways`, `use service-gateways`)

CreateSubnet, DeleteSubnet (both also need `manage vcns`, `manage subnets`, `manage security-lists`, `manage dhcp-options`)

UpdateSubnet (if changing which route table is associated with the subnet, also need `manage subnets`)

**Note:** All of the above operations in this cell are totally covered with just `manage virtual-network-family`.

## network-security-groups

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
NETWORK_SECURITY_GROUP_INSPECT	<i>none</i>	AddNetworkSecurityGroupSecurityRules and UpdateNetworkSecurityGroupSecurityRules (both also need <code>manage network-security-groups</code> )

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i> NETWORK_SECURITY_GROUP_READ	<i>INSPECT +</i> GetNetworkSecurityGroup ListNetworkSecurityGroups	<i>no extra</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i> NETWORK_SECURITY_GROUP_LIST_SECURITY_RULES NETWORK_SECURITY_GROUP_LIST_MEMBERS NETWORK_SECURITY_GROUP_UPDATE_MEMBERS	<i>READ +</i> ListNetworkSecurityGroupSecurityRules ListNetworkSecurityGroupVnics	<i>READ +</i> LaunchInstance (also need <code>manage instances, read instance-images, use vnics, use subnets, and read app-catalog-listing</code> ) AttachVnic (also need <code>manage instances, and use subnets</code> ) UpdateVnic (also need <code>use vnics</code> )

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>USE +</i>
NETWORK_SECURITY_GROUP_UPDATE	UpdateNetworkSecurityGroup	CreateNetworkSecurityGroup,
NETWORK_SECURITY_GROUP_CREATE	ChangeNetworkSecurityGroupCompartment	DeleteNetworkSecurityGroup (both also
NETWORK_SECURITY_GROUP_DELETE	AddNetworkSecurityGroupSecurityRules	need manage vcns)
NETWORK_SECURITY_GROUP_MOVE	UpdateNetworkSecurityGroupSecurityRules	<b>Note:</b> Both of the above operations in this cell
NETWORK_SECURITY_GROUP_UPDATE_SECURITY_RULES	RemoveNetworkSecurityGroupSecurityRules	are totally covered with just manage
		virtual-network-family.

### security-lists

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
SECURITY_LIST_READ	ListSecurityLists GetSecurityList	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**MANAGE**

**Permissions**

USE +  
 SECURITY\_LIST\_ATTACH  
 SECURITY\_LIST\_DETACH  
 SECURITY\_LIST\_UPDATE  
 SECURITY\_LIST\_CREATE  
 SECURITY\_LIST\_DELETE  
 SECURITY\_LIST\_MOVE

**APIs Fully Covered**

USE +  
 UpdateSecurityList  
**Note:** Ability to update a security list is available only with the `manage` verb, not the `use` verb.  
 ChangeSecurityListCompartment

**APIs Partially Covered**

CreateSecurityList,  
 DeleteSecurityList (both also need `manage vcn`s)  
 CreateSubnet, DeleteSubnet (both also need `manage vcn`s, `manage subnets`, `manage route-tables`, `manage dhcp-options`)  
 UpdateSubnet (if changing which security lists are associated with the subnet, also need `manage subnets`)  
**Note:** All of the above operations in this cell are totally covered with just `manage virtual-network-family`.

dhcp-options

**INSPECT**

**Permissions**

DHCP\_READ

**APIs Fully Covered**

ListDhcpOptions  
 GetDhcpOptions

**APIs Partially Covered**

none

**READ**

**Permissions**

no extra

**APIs Fully Covered**

no extra

**APIs Partially Covered**

none

**USE**

**Permissions**

no extra

**APIs Fully Covered**

no extra

**APIs Partially Covered**

none

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>USE +</i>
DHCP_ATTACH	UpdateDhcpOptions	CreateDhcpOptions, DeleteDhcpOptions
DHCP_DETACH	<b>Note:</b> Ability to update a set of DHCP options is	(both also need <code>manage vcns</code> )
DHCP_UPDATE	available only with the <code>manage</code> verb, not the	CreateSubnet, DeleteSubnet (also need
DHCP_CREATE	<code>use</code> verb.	<code>manage vcns</code> , <code>manage subnets</code> , <code>manage</code>
DHCP_DELETE	ChangeDhcpOptionsCompartment	<code>route-tables</code> , <code>manage security-lists</code> )
DHCP_MOVE		UpdateSubnet (if changing which set of DHCP
		options is associated with the subnet, also
		need <code>manage subnets</code> )
		<b>Note:</b> All of the above operations in this cell are
		totally covered with just <code>manage virtual-</code>
		<code>network-family</code> .

### private-ips

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
PRIVATE_IP_READ	ListPrivateIps	<i>none</i>
	GetPrivateIp	
	For ephemeral public IPs only:	
	ListPublicIps,	
	GetPublicIpByPrivateIpId,	
	GetPublicIpByIpAddress	

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## CHAPTER 18 IAM

### USE

#### Permissions

*READ* +  
PRIVATE\_IP\_UPDATE  
PRIVATE\_IP\_ASSIGN  
PRIVATE\_IP\_UNASSIGN  
PRIVATE\_IP\_CREATE  
PRIVATE\_IP\_DELETE  
PRIVATE\_IP\_ASSIGN\_PUBLIC\_IP  
PRIVATE\_IP\_UNASSIGN\_PUBLIC\_IP

#### APIs Fully Covered

*READ* +  
UpdatePrivateIp  
For ephemeral public IPs: UpdatePublicIp,  
CreatePublicIp, DeletePublicIp

#### APIs Partially Covered

CreatePrivateIp, DeletePrivateIp (both also need use subnets and use vnics)  
For reserved public IPs: UpdatePublicIp, CreatePublicIp, DeletePublicIp (all also need manage public-ips)  
**Note:** The above operations in this cell are totally covered with just use virtual-network-family.

### MANAGE

#### Permissions

*USE* +  
PRIVATE\_IP\_ROUTE\_TABLE\_ATTACH  
PRIVATE\_IP\_ROUTE\_TABLE\_DETACH

#### APIs Fully Covered

*no extra*

#### APIs Partially Covered

*USE* +  
CreateRouteTable, DeleteRouteTable  
(both also need manage vcns, manage internet-gateways, manage drgs, and manage route-tables, manage local-peering-gateways, use nat-gateways, use service-gateways)  
UpdateRouteTable (also need manage internet-gateways, manage drgs, manage route-tables, manage local-peering-gateways, use nat-gateways, use service-gateways)  
**Note:** The above operations in this cell are totally covered with just manage virtual-network-family.

## public-ips

## INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<i>none</i>	<i>none</i>	<i>none</i>

## READ

Permissions	APIs Fully Covered	APIs Partially Covered
PUBLIC_IP_READ	For reserved public IPs only: ListPublicIps, GetPublicIpByPrivateIpId, GetPublicIpByIpAddress Permissions for listing/getting ephemeral public IPs are part of the <a href="#">private-ip</a> permissions.	<i>none</i>

## USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + PUBLIC_IP_ASSIGN_PRIVATE_IP PUBLIC_IP_UNASSIGN_PRIVATE_IP	<i>no extra</i>	For reserved public IPs: UpdatePublicIp, CreatePublicIp, DeletePublicIp (all of these also need use private-ips and manage public-ips). <b>Note:</b> The above operations in this cell are totally covered with just manage virtual-network-family.

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i> PUBLIC_IP_UPDATE PUBLIC_IP_CREATE PUBLIC_IP_DELETE	<i>no extra</i>	<i>USE +</i> For reserved public IPs: UpdatePublicIp, CreatePublicIp, DeletePublicIp (all of these also need use private-ips). <b>Note:</b> The above operations in this cell are totally covered with just <code>manage virtual-network-family</code> .

### ipv6s

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<i>none</i>	<i>none</i>	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
IPV6_READ	GetIpv6	ListIpv6s (also need inspect vnics and inspect subnets to list IPv6s by VNIC and subnets) <b>Note:</b> The above operation in this cell is totally covered with just <code>use virtual-network-family</code> .

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>USE +</i></p> <p>IPV6_UPDATE</p> <p>IPV6_CREATE</p> <p>IPV6_DELETE</p>	<p><i>no extra</i></p>	<p><i>USE +</i></p> <p>UpdateIpv6 (also need use vnics)</p> <p>CreateIpv6, DeleteIpv6 (both also need use vnics and use subnets)</p> <p><b>Note:</b> The above operations in this cell are totally covered with just <code>manage virtual-network-family</code>.</p>

internet-gateways

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
<p>INTERNET_GATEWAY_READ</p>	<p>ListInternetGateways</p> <p>GetInternetGateway</p>	<p><i>none</i></p>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>no extra</i></p>	<p><i>none</i></p>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>no extra</i></p>	<p><i>none</i></p>

## CHAPTER 18 IAM

### MANAGE

#### Permissions

USE +

INTERNET\_GATEWAY\_ATTACH

INTERNET\_GATEWAY\_DETACH

INTERNET\_GATEWAY\_UPDATE

INTERNET\_GATEWAY\_CREATE

INTERNET\_GATEWAY\_DELETE

INTERNET\_GATEWAY\_MOVE

#### APIs Fully Covered

USE +

UpdateInternetGateway

**Note:** Ability to update an internet gateway is available only with the `manage` verb, not the `use` verb.

ChangeInternetGatewayCompartment

#### APIs Partially Covered

CreateInternetGateway,

DeleteInternetGateway (both also need `manage` `vcns`)

CreateRouteTable, DeleteRouteTable

(both also need `manage` `route-tables`, `manage` `vcns`, `manage` `drgs`, `manage` `private-ips`, `manage` `local-peering-gateways`, `use` `nat-gateways`, `use` `nat-gateways`, `use` `service-gateways`)

UpdateRouteTable (also need `manage` `route-tables`, `manage` `drgs`, `manage` `private-ips`, `manage` `local-peering-gateways`, `use` `nat-gateways`, `use` `service-gateways`)

**Note:** All of the above operations in this cell are totally covered with just `manage` `virtual-network-family`.

## nat-gateways

### INSPECT

#### Permissions

*none*

#### APIs Fully Covered

*none*

#### APIs Partially Covered

*none*

### READ

#### Permissions

NAT\_GATEWAY\_READ

#### APIs Fully Covered

ListNatGateways

GetNatGateway

#### APIs Partially Covered

*none*

## CHAPTER 18 IAM

### USE

#### Permissions

*READ* +

NAT\_GATEWAY\_ATTACH

NAT\_GATEWAY\_DETACH

#### APIs Fully Covered

*no extra*

#### APIs Partially Covered

*READ* +

CreateRouteTable, DeleteRouteTable

(both also need `manage route-tables`,  
`manage vcns`, `manage drgs`, `manage`  
`private-ips`, `manage internet-`  
`gateways`, `manage local-peering-`  
`gateways`, `use service-gateways`)

UpdateRouteTable (also need `manage`  
`route-tables`, `manage drgs`, `manage`  
`private-ips`, `manage internet-`  
`gateways`, `manage local-peering-`  
`gateways`, `use service-gateways`)

**Note:** All of the above operations in this cell are  
totally covered with just `manage virtual-`  
`network-family`.

### MANAGE

#### Permissions

*USE* +

NAT\_GATEWAY\_UPDATE

NAT\_GATEWAY\_CREATE

NAT\_GATEWAY\_DELETE

NAT\_GATEWAY\_MOVE

#### APIs Fully Covered

*USE* +

UpdateNatGateway

ChangeNatGatewayCompartment

**Note:** Ability to update a NAT gateway is  
available only with the `manage` verb, not the  
`use` verb.

#### APIs Partially Covered

CreateNatGateway, DeleteNatGateway

(both also need `manage vcns`)

**Note:** All of the above operations in this cell are  
totally covered with just `manage virtual-`  
`network-family`.

## service-gateways

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
SERVICE_GATEWAY_READ	ListServiceGateways GetServiceGateway	<i>none</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + SERVICE_GATEWAY_ATTACH SERVICE_GATEWAY_DETACH	<i>no extra</i>	READ + CreateRouteTable, DeleteRouteTable (both also need manage route-tables, manage vcn, manage internet-gateways, manage drgs, manage private-ips, manage local-peering-gateways) UpdateRouteTable (also need manage route-tables, manage drgs, manage internet-gateways, manage private-ips, manage local-peering-gateways)

## CHAPTER 18 IAM

### MANAGE

#### Permissions

USE +

SERVICE\_GATEWAY\_UPDATE

SERVICE\_GATEWAY\_CREATE

SERVICE\_GATEWAY\_DELETE

SERVICE\_GATEWAY\_ADD\_SERVICE

SERVICE\_GATEWAY\_DELETE\_SERVICE

SERVICE\_GATEWAY\_MOVE

#### APIs Fully Covered

USE +

ChangeServiceGatewayCompartment

AttachServiceId

DetachServiceId

**Note:** Ability to update a service gateway is available only with the `manage` verb, not the `use` verb.

#### APIs Partially Covered

CreateServiceGateway (also need `manage` `vcns`, and need `manage` `route-tables` if you associate a route table during creation)

UpdateServiceGateway (also need `manage` `route-tables` if you associate a route table during the update)

DeleteServiceGateway (also need `manage` `vcns`)

**Note:** All of the above operations in this cell are totally covered with just `manage` `virtual-network-family`.

## local-peering-gateways

### INSPECT

#### Permissions

LOCAL\_PEERING\_GATEWAY\_READ

#### APIs Fully Covered

ListLocalPeeringGateways

GetLocalPeeringGateway

#### APIs Partially Covered

none

### READ

#### Permissions

no extra

#### APIs Fully Covered

no extra

#### APIs Partially Covered

none

### USE

#### Permissions

no extra

#### APIs Fully Covered

no extra

#### APIs Partially Covered

none

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>no extra</i>	
LOCAL_PEERING_GATEWAY_UPDATE		CreateLocalPeeringGateway (also need manage vcn, and need manage route-tables if you associate a route table during creation)
LOCAL_PEERING_GATEWAY_ATTACH		
LOCAL_PEERING_GATEWAY_DETACH		
LOCAL_PEERING_GATEWAY_CREATE		UpdateLocalPeeringGateway (also need manage route-tables if you associate a route table during the update)
LOCAL_PEERING_GATEWAY_DELETE		DeleteLocalPeeringGateway (also need manage vcn)
LOCAL_PEERING_GATEWAY_MOVE		CreateRouteTable, DeleteRouteTable (both also need manage route-tables, manage vcn, manage internet-gateways, manage drgs, manage private-ips, use nat-gateways, use service-gateways)
		UpdateRouteTable (also need manage route-tables, manage internet-gateways, manage drgs, manage private-ips, use nat-gateways, use service-gateways)
		ChangeLocalPeeringGatewayCompartment
		<b>Note:</b> The above operations in this cell are totally covered with just manage virtual-network-family.

### local-peering-from

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
LOCAL_PEERING_GATEWAY_READ	<i>none</i>	<i>none</i>

## CHAPTER 18 IAM

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*none*

**APIs Partially Covered**

*none*

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*none*

**APIs Partially Covered**

*none*

### MANAGE

**Permissions**

*USE +*

LOCAL\_PEERING\_GATEWAY\_CONNECT\_FROM

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

`ConnectLocalPeeringGateways` (acceptor in the peering relationship must also grant the requestor `manage local-peering-to` in the compartment where the acceptor's LPG resides. See [Local VCN Peering \(Within Region\)](#).)

**Note:** The above operation in this cell is totally covered with just `manage virtual-network-family`.

## local-peering-to

### INSPECT

**Permissions**

LOCAL\_PEERING\_GATEWAY\_READ

**APIs Fully Covered**

*none*

**APIs Partially Covered**

*none*

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*none*

**APIs Partially Covered**

*none*

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*none*

**APIs Partially Covered**

*none*

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i> LOCAL_PEERING_GATEWAY_CONNECT_TO	<i>no extra</i>	ConnectLocalPeeringGateways (requestor in the peering relationship must also have manage local-peering-from in the compartment where the requestor's LPG resides. See <a href="#">Local VCN Peering (Within Region).</a> ) <b>Note:</b> The above operation in this cell is totally covered with just manage virtual-network-family.

### remote-peering-connections

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
REMOTE_PEERING_CONNECTION_READ	ListRemotePeeringConnections GetRemotePeeringConnection	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## CHAPTER 18 IAM

### MANAGE

#### Permissions

USE +

REMOTE\_PEERING\_CONNECTION\_UPDATE

REMOTE\_PEERING\_CONNECTION\_CREATE

REMOTE\_PEERING\_CONNECTION\_DELETE

REMOTE\_PEERING\_CONNECTION\_

RESOURCE\_MOVE

#### APIs Fully Covered

UpdateRemotePeeringConnection

#### APIs Partially Covered

CreateRemotePeeringConnection,

DeleteRemotePeeringConnection (both

also need manage drgs)

ChangeRemotePeeringConnectionCompartment

ent

**Note:** The above operations in this cell are totally covered with just `manage virtual-network-family`.

### remote-peering-from

### INSPECT

#### Permissions

REMOTE\_PEERING\_CONNECTION\_READ

#### APIs Fully Covered

none

#### APIs Partially Covered

none

### READ

#### Permissions

no extra

#### APIs Fully Covered

none

#### APIs Partially Covered

none

### USE

#### Permissions

no extra

#### APIs Fully Covered

none

#### APIs Partially Covered

none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>USE +</i></p> <p>REMOTE_PEERING_CONNECTION_CONNECT_</p> <p>FROM</p>	<p><i>no extra</i></p>	<p>ConnectRemotePeeringConnections</p> <p>(acceptor in the peering relationship must also grant the requestor <code>manage remote-peering-to</code> in the compartment where the acceptor's RPC resides. See <a href="#">Remote VCN Peering (Across Regions)</a>.)</p> <p><b>Note:</b> The above operation in this cell is totally covered with just <code>manage virtual-network-family</code>.</p>

remote-peering-to

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
<p>REMOTE_PEERING_CONNECTION_READ</p>	<p><i>none</i></p>	<p><i>none</i></p>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>none</i></p>	<p><i>none</i></p>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<p><i>no extra</i></p>	<p><i>none</i></p>	<p><i>none</i></p>

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<code>USE +</code> <code>REMOTE_PEERING_CONNECTION_CONNECT_</code> <code>TO</code>	<i>no extra</i>	<code>ConnectRemotePeeringConnections</code>  (requestor in the peering relationship must also have <code>manage remote-peering-from</code> in the compartment where the requestor's RPC resides. See <a href="#">Remote VCN Peering (Across Regions)</a> .)  <b>Note:</b> The above operation in this cell is totally covered with just <code>manage virtual-network-family</code> .

### drgs

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>DRG_READ</code> <code>DRG_ATTACHMENT_READ</code>	<code>ListDrgs</code> <code>GetDrg</code> <code>ListDrgAttachments</code>	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>



## CHAPTER 18 IAM

---

### **Permissions**

*USE +*

DRG\_ATTACH

DRG\_DETACH

DRG\_UPDATE

DRG\_ATTACHMENT\_UPDATE

DRG\_CREATE

DRG\_DELETE

DRG\_MOVE

### **APIs Fully Covered**

*USE +*

CreateDrg

UpdateDrg

DeleteDrg

ChangeDrgCompartment



## cpes

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CPE_READ	ListCpes GetCpe	none

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	USE +	CreateIPSecConnection,
CPE_ATTACH	CreateCpe	DeleteIPSecConnection (both also need
CPE_DETACH	UpdateCpe	manage ipsec-connections and manage
CPE_UPDATE	DeleteCpe	drgs)
CPE_CREATE	ChangeCpeCompartment	<b>Note:</b> All of the above operations in this cell are
CPE_DELETE		totally covered with just <code>manage virtual-</code>
CPE_RESOURCE_MOVE		<code>network-family</code> .

## ipsec

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
IPSEC_CONNECTION_READ	ListIPSecConnections GetIPSecConnection GetIPSecConnectionStatus ListIPSecConnectionTunnels GetIPSecConnectionTunnel	<i>none</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i> IPSEC_CONNECTION_DEVICE_CONFIG_READ	<i>INSPECT +</i> GetIPSecConnectionDeviceConfig GetIPSecConnectionTunnelSharedSecret	<i>none</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i> IPSEC_CONNECTION_CREATE IPSEC_CONNECTION_UPDATE IPSEC_CONNECTION_DELETE IPSEC_CONNECTION_DEVICE_CONFIG_UPDATE	<i>USE +</i> UpdateIPSecConnection UpdateIPSecConnectionTunnel	CreateIPSecConnection, DeleteIPSecConnection (both also need manage cpes and manage drgs) <b>Note:</b> All of the above operations in this cell are totally covered with just <code>manage virtual-</code> <code>network-family</code> .

## CHAPTER 18 IAM

### cross-connects

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CROSS_CONNECT_READ	ListCrossConnects GetCrossConnect	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	UpdateCrossConnect	UpdateVirtualCircuit (also need use
CROSS_CONNECT_UPDATE	CreateCrossConnect	virtual-circuits)
CROSS_CONNECT_CREATE	DeleteCrossConnect	CreateVirtualCircuit,
CROSS_CONNECT_DELETE	ChangeCrossConnectCompartment	DeleteVirtualCircuit (also need manage
CROSS_CONNECT_RESOURCE_MOVE		virtual-circuits)
CROSS_CONNECT_ATTACH		
CROSS_CONNECT_DETACH		

### cross-connect-groups

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CROSS_CONNECT_GROUP_READ	ListCrossConnectGroups GetCrossConnectGroup	<i>none</i>

## CHAPTER 18 IAM

### READ

**Permissions***no extra***APIs Fully Covered***no extra***APIs Partially Covered***none*

### USE

**Permissions***no extra***APIs Fully Covered***no extra***APIs Partially Covered***no extra*

### MANAGE

**Permissions***USE +*

CROSS\_CONNECT\_GROUP\_UPDATE

CROSS\_CONNECT\_GROUP\_CREATE

CROSS\_CONNECT\_GROUP\_DELETE

CROSS\_CONNECT\_GROUP\_RESOURCE\_MOVE

**APIs Fully Covered**

UpdateCrossConnectGroup

CreateCrossConnectGroup

DeleteCrossConnectGroup

ChangeCrossConnectGroupCompartment

**APIs Partially Covered***no extra*

## virtual-circuits

### INSPECT

**Permissions**

VIRTUAL\_CIRCUIT\_READ

**APIs Fully Covered**

ListVirtualCircuits

GetVirtualCircuit

**APIs Partially Covered***none*

### READ

**Permissions***no extra***APIs Fully Covered***no extra***APIs Partially Covered***none*

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<code>READ +</code> <code>VIRTUAL_CIRCUIT_UPDATE</code>	<code>no extra</code>	<code>UpdateVirtualCircuit</code> (also need <code>manage drgs</code> , and if you're also changing which cross-connect or cross-connect group the virtual circuit uses, also need <code>manage cross-connects</code> )

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<code>USE +</code> <code>VIRTUAL_CIRCUIT_CREATE</code> <code>VIRTUAL_CIRCUIT_DELETE</code> <code>VIRTUAL_CIRCUIT_RESOURCE_MOVE</code>	<code>ChangeVirtualCircuitCompartment</code>	<code>USE +</code> <code>CreateVirtualCircuit</code> , <code>DeleteVirtualCircuit</code> (both also need <code>manage drgs</code> , and if you're also creating/deleting the virtual circuit with a mapping to a specific cross-connect or cross-connect group, also need <code>manage cross-connects</code> ) <b>Note:</b> All of the above operations in this cell are totally covered with just <code>manage virtual-network-family</code> .

## vnics

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>VNIC_READ</code>	<code>GetVnic</code>	<code>CreateInstanceConfiguration</code> (if using the <code>CreateInstanceConfigurationFromInstanceDetails</code> subtype. Also need <code>read instances</code> , <code>inspect vnic-attachments</code> , <code>inspect volumes</code> , and <code>inspect volume-attachments</code> .)

## CHAPTER 18 IAM

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*none*

### USE

**Permissions**

*READ +*

VNIC\_ATTACH

VNIC\_DETACH

VNIC\_CREATE

VNIC\_DELETE

VNIC\_UPDATE

VNIC\_ASSOCIATE\_NETWORK\_SECURITY\_

GROUP

VNIC\_DISASSOCIATE\_NETWORK\_SECURITY\_

GROUP

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*READ +*

LaunchInstance (also need use subnets, use network-security-groups, and manage instance-family)

AttachVnic (also need manage instances, use subnets, and use network-security-groups)

UpdateVnic (also need use network-security-groups)

DetachVnic (also need manage instances and use subnets)

CreatePrivateIp, DeletePrivateIp (both also need use subnets and use private-ips)

### MANAGE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

## vnic-attachments

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
VNIC_ATTACHMENT_READ	GetVnicAttachment	ListVnicAttachments (also need inspect instances) CreateInstanceConfiguration (if using the CreateInstanceConfigurationFromInstanceDetails subtype. Also need read instances, inspect vnics, inspect volumes, and inspect volume-attachments.)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>none</i>	<i>no extra</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>none</i>	<i>no extra</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>none</i>	<i>no extra</i>

**FOR INSTANCE-FAMILY RESOURCE TYPES**

The `instance-family` aggregate resource-type includes extra permissions beyond the sum of the permissions for the individual resource-types included in `instance-family`. For example: It includes a few permissions for `vnics` and `volumes`, even though those resource-types aren't generally considered part of the `instance-family`. Why are there extras included? So you can write fewer policy statements to cover general use cases, like working with an instance that has an attached block volume. You can write one statement for `instance-family` instead of multiple statements covering `instances`, `vnics`, and `volumes`.

Here's a list of the extra permissions:

For `inspect instance-family`:

- `VNIC_READ`
- `VNIC_ATTACHMENT_READ`
- `VOLUME_ATTACHMENT_INSPECT`

For `read instance-family`:

- `VOLUME_ATTACHMENT_READ`

For `use instance-family`:

- `VNIC_ATTACH`
- `VNIC_DETACH`
- `VOLUME_ATTACHMENT_UPDATE`

For `manage instance-family`:

- `VOLUME_ATTACHMENT_CREATE`
- `VOLUME_ATTACHMENT_DELETE`

The following tables list the permissions and API operations covered by each of the individual resource-types included in `instance-family`.

## instances

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
INSTANCE_INSPECT	<i>none</i>	GetConsoleHistory, ListConsoleHistories (both also need inspect console-histories) ListVnicAttachments (also need inspect vnic-attachments) ListVolumeAttachments (also need inspect volumes and inspect volume-attachments) GetVolumeAttachments (also need inspect volumes and inspect volume-attachments)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + INSTANCE_READ	ListInstances GetInstance <b>Note:</b> ListInstances and GetInstance include any user-provided metadata added to the instance	INSPECT + CaptureConsoleHistory (also need manage console-histories and read instance-images) ShowConsoleHistoryData (also need read console-histories and read instance-images) CreateInstanceConfiguration (if using the CreateInstanceConfigurationFromInstanceDetails subtype. Also need inspect vnics, inspect vnic-attachments, inspect volumes, and inspect volume-attachments.)

## CHAPTER 18 IAM

### USE

#### Permissions

*READ +*

INSTANCE\_UPDATE

INSTANCE\_CREATE\_IMAGE

INSTANCE\_POWER\_ACTIONS

INSTANCE\_ATTACH\_VOLUME

INSTANCE\_DETACH\_VOLUME

#### APIs Fully Covered

*READ +*

UpdateInstance

InstanceAction

#### APIs Partially Covered

*READ +*

CreateImage (also need `manage instance-images`)

AttachVolume (also need `manage volume-attachments` and `use volumes`)

DetachVolume (also need `manage volume-attachments` and `use volumes`)

### MANAGE

#### Permissions

*USE +*

INSTANCE\_CREATE

INSTANCE\_DELETE

INSTANCE\_ATTACH\_SECONDARY\_VNIC

INSTANCE\_DETACH\_SECONDARY\_VNIC

INSTANCE\_MOVE

#### APIs Fully Covered

ChangeInstanceCompartment

#### APIs Partially Covered

*USE +*

LaunchInstance (also need `read instance-images`, `use vnics`, `use subnets`, `use network-security-groups`, and `read app-catalog-listing`)

TerminateInstance (also need `use vnics` and `use subnets`; also need `manage volume-attachments` and `use volumes` if a volume is attached)

AttachVnic (also need `use subnets`, `use network-security-groups`, and either `use vnics` or `use instance-family`)

DetachVnic (also need `use subnets` and either `use vnics` or `use instance-family`)

GetWorkRequest, ListWorkRequestErrors, and ListWorkRequestLogs (for work requests related to instances resource types.

All also need the permissions for

LaunchInstance)

## console-histories

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
CONSOLE_HISTORY_INSPECT	<i>none</i>	ListConsoleHistories, GetConsoleHistory (both also need inspect instances)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + CONSOLE_HISTORY_READ	<i>none</i>	INSPECT + ShowConsoleHistoryData (also need read instances and read instance-images)

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>none</i>	<i>no extra</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + CONSOLE_HISTORY_CREATE CONSOLE_HISTORY_DELETE	DeleteConsoleHistory	USE + CaptureConsoleHistory (also need read instances and read instance-images)

## instance-console-connection

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
INSTANCE_CONSOLE_CONNECTION_INSPECT	<i>none</i>	ListInstanceConsoleConnections (also need inspect instances and read instances)

## CHAPTER 18 IAM

---

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + INSTANCE_CONSOLE_CONNECTION_READ	<i>none</i>	<i>INSPECT</i> + GetInstanceConsoleConnection (also need read instances)

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	<i>none</i>	<i>no extra</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + INSTANCE_CONSOLE_CONNECTION_CREATE INSTANCE_CONSOLE_CONNECTION_DELETE	DeleteInstanceConsoleConnection	CreateInstanceConsoleConnection (also need read instances)

## instance-images

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
INSTANCE_IMAGE_INSPECT	ListImages GetImage	<i>none</i>

## CHAPTER 18 IAM

### READ

#### Permissions

*INSPECT* +  
INSTANCE\_IMAGE\_READ

#### APIs Fully Covered

*no extra*

#### APIs Partially Covered

*INSPECT* +  
LaunchInstance (also need *manage instances, use vnics, use subnets, and use network-security-groups*)  
CaptureConsoleHistory (also need *read instances and manage console-histories*)  
ShowConsoleHistoryData (also need *read instances and read console-histories*)

### USE

#### Permissions

*READ* +  
INSTANCE\_IMAGE\_UPDATE

#### APIs Fully Covered

UpdateImage

#### APIs Partially Covered

*no extra*

### MANAGE

#### Permissions

*USE* +  
INSTANCE\_IMAGE\_CREATE  
INSTANCE\_IMAGE\_DELETE  
INSTANCE\_IMAGE\_MOVE

#### APIs Fully Covered

DeleteImage  
ChangeImageCompartment

#### APIs Partially Covered

*USE* +  
CreateImage (also need *use instances*)  
GetWorkRequest, ListWorkRequestErrors, and ListWorkRequestLogs (for work requests related to *instance-images* resource types. All also need the permissions for *CreateImage*)

## app-catalog-listing

### INSPECT

#### Permissions

APP\_CATALOG\_LISTING\_INSPECT

#### APIs Fully Covered

ListAppCatalogSubscriptions

#### APIs Partially Covered

*none*

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + APP_CATALOG_LISTING_READ	<i>no extra</i>	<i>INSPECT</i> + LaunchInstance (Also need use instances, read instance-images, use vnics, use subnets, and use network-security- groups)

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + APP_CATALOG_LISTING_SUBSCRIBE	<i>READ</i> + CreateAppCatalogSubscription DeleteAppCatalogSubscription	<i>none</i>

### FOR COMPUTE-MANAGEMENT-FAMILY RESOURCE TYPES

The following tables list the permissions and API operations covered by each of the individual resource-types included in `compute-management-family`.

### instance-configurations

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
INSTANCE_CONFIGURATION_INSPECT	ListInstanceConfigurations	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + INSTANCE_CONFIGURATION_READ	<i>INSPECT</i> + GetInstanceConfiguration	<i>none</i>

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**MANAGE**

<b>Permissions</b>	<b>APIs Fully Covered</b>	<b>APIs Partially Covered</b>
<i>USE</i> +	<i>USE</i> +	<i>none</i>
INSTANCE_CONFIGURATION_CREATE	CreateInstanceConfiguration (if using the	
INSTANCE_CONFIGURATION_UPDATE	CreateInstanceConfigurationDetails	
INSTANCE_CONFIGURATION_LAUNCH	subtype)	
INSTANCE_CONFIGURATION_DELETE	UpdateInstanceConfiguration	
INSTANCE_CONFIGURATION_MOVE	LaunchInstanceConfiguration	
	DeleteInstanceConfiguration	
	ChangeInstanceConfigurationCompartment	
	t	

## instance-pools

**INSPECT**

<b>Permissions</b>	<b>APIs Fully Covered</b>	<b>APIs Partially Covered</b>
INSTANCE_POOL_INSPECT	ListInstancePools	<i>none</i>

**READ**

<b>Permissions</b>	<b>APIs Fully Covered</b>	<b>APIs Partially Covered</b>
<i>INSPECT</i> +	<i>INSPECT</i> +	<i>none</i>
INSTANCE_POOL_READ	GetInstancePool	
	ListInstancePoolInstances	

**USE**

<b>Permissions</b>	<b>APIs Fully Covered</b>	<b>APIs Partially Covered</b>
<i>READ</i> +	<i>no extra</i>	ResetInstancePool
INSTANCE_POOL_POWER_ACTIONS		SoftresetInstancePool
		StartInstancePool
		StopInstancePool
		All also need use instances.

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>USE +</i>
INSTANCE_POOL_CREATE	UpdateInstancePool	CreateInstancePool (also need manage instances, read instance-images, use vnics, and use subnets)
INSTANCE_POOL_UPDATE	ChangeInstancePoolCompartment	TerminateInstancePool (also need manage instances, use vnics, use subnets, manage volume-attachments, and use volumes)
INSTANCE_POOL_DELETE		GetWorkRequest, ListWorkRequestErrors, and ListWorkRequestLogs (for work requests related to instance-pools resource types. All also need the permissions for CreateInstancePool or TerminateInstancePool, depending on the operation that spawns the work request)
INSTANCE_POOL_MOVE		

## cluster-networks

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CLUSTER_NETWORK_INSPECT	ListClusterNetworks	none

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i>	<i>INSPECT +</i>	ListClusterNetworkInstances (also need read instance-pools)
CLUSTER_NETWORK_READ	GetClusterNetwork	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>USE +</i>
CLUSTER_NETWORK_CREATE	UpdateClusterNetwork	CreateClusterNetwork (also need manage instances, manage instance-pools, read instance-images, use vnics, and use subnets)
CLUSTER_NETWORK_UPDATE	ChangeClusterNetworkCompartment	TerminateClusterNetwork (also need manage instances, manage instance-pools, use vnics, use subnets, manage volume-attachments, and use volumes)
CLUSTER_NETWORK_DELETE		GetWorkRequest, ListWorkRequestErrors, and ListWorkRequestLogs (for work requests related to cluster-networks resource types. All also need the permissions for CreateClusterNetwork or TerminateClusterNetwork, depending on the operation that spawns the work request)
CLUSTER_NETWORK_MOVE		

### FOR ADDITIONAL COMPUTE INDIVIDUAL RESOURCE TYPES

The following tables list the permissions and API operations covered by other Compute resource-types that aren't included in any aggregate resource-types.

### auto-scaling-configurations

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
AUTO_SCALING_CONFIGURATION_INSPECT	ListAutoScalingConfigurations ListAutoScalingPolicies	none

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + AUTO_SCALING_CONFIGURATION_READ	<i>INSPECT</i> + GetAutoScalingConfiguration GetAutoScalingPolicy	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + AUTO_SCALING_CONFIGURATION_CREATE AUTO_SCALING_CONFIGURATION_UPDATE AUTO_SCALING_CONFIGURATION_DELETE AUTO_SCALING_CONFIGURATION_MOVE	<i>USE</i> + ChangeAutoScalingConfigurationCompart ment	<i>USE</i> + CreateAutoScalingConfiguration UpdateAutoScalingConfiguration DeleteAutoScalingConfiguration CreateAutoScalingPolicy UpdateAutoScalingPolicy DeleteAutoScalingPolicy  All also need <code>manage instance-pools</code> .

## dedicated-vm-hosts

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DEDICATED_VM_HOST_INSPECT	ListDedicatedVmHosts	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + DEDICATED_VM_HOST_READ	<i>INSPECT</i> + GetDedicatedVmHost ListDedicatedVmHostInstances	<i>none</i>

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + DEDICATED_VM_HOST_LAUNCH_INSTANCE DEDICATED_VM_HOST_UPDATE	<i>INSPECT</i> + UpdateDedicatedVmHost	<i>INSPECT</i> + LaunchInstance  All also need <code>create_instance</code> in the compartment to launch the instance in and dedicated vm host launch instance in the compartment for the dedicated virtual machine host.

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + DEDICATED_VM_HOST_CREATE DEDICATED_VM_HOST_MOVE DEDICATED_VM_HOST_DELETE	<i>USE</i> + CreateDedicatedVmHost DeleteDedicatedVmHost ChangeDedicatedVmHostCompartment	<i>USE</i> + <i>none</i>

## work-requests

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
WORKREQUEST_INSPECT	ListWorkRequests	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## CHAPTER 18 IAM

### MANAGE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*none*

### FOR VOLUME-FAMILY RESOURCE TYPES

The following tables list the permissions and API operations covered by each of the individual resource-types included in `volume-family`.

## volumes

### INSPECT

**Permissions**

VOLUME\_INSPECT

**APIs Fully Covered**

ListVolumes

GetVolume

**APIs Partially Covered**

ListVolumeBackups, GetVolumeBackup

(these also need `inspect volume-backups`)

UpdateVolumeBackup (also need `read volume-backups`)

DeleteVolumeBackup (also need `manage volume-backups`)

GetVolumeAttachment (also need `inspect instances` and `inspect volume-attachments`). If you need to get the CHAP secret if it exists, `read volume-attachments` is required.

CreateInstanceConfiguration (if using the `CreateInstanceConfigurationFromInstanceDetails` subtype. Also need `read instances`, `inspect vnics`, `inspect vnic-attachments`, and `inspect volume-attachments`.)

## CHAPTER 18 IAM

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

### USE

**Permissions**

*READ +*

VOLUME\_UPDATE

VOLUME\_WRITE

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*READ +*

AttachVolume and DetachVolume (both also need manage volume-attachments, use instances)

CreateVolumeBackup (also need manage volume-backups)

### MANAGE

**Permissions**

*USE +*

VOLUME\_CREATE

VOLUME\_DELETE

VOLUME\_MOVE

**APIs Fully Covered**

*USE +*

CreateVolume

DeleteVolume

ChangeVolumeCompartment

When moving volumes between compartments, the `move volume` permission is needed for both source and destination compartments.

**APIs Partially Covered**

*USE +*

If creating a volume *from a backup*, also need `read volume-backups`.

If creating a volume *encrypted with a Key Management master encryption key*, also need `use key-delegate` (for the caller) and `read keys` (for the service principal). For more information, see [Details for the Key Management Service](#).

## volume-attachments

## INSPECT

**Permissions**

VOLUME\_ATTACHMENT\_INSPECT

**APIs Fully Covered**

ListVolumeAttachments

**APIs Partially Covered**

GetVolumeAttachment (also need inspect volumes and inspect instances)

**Note:** The CHAP secret (if it exists) is NOT included with inspect volume-attachments.

CreateInstanceConfiguration (if using the CreateInstanceConfigurationFromInstanceDetails subtype. Also need read instances, inspect vnics, inspect vnic-attachments, and inspect volumes.)

## READ

**Permissions**

INSPECT +

VOLUME\_ATTACHMENT\_READ

**APIs Fully Covered**

no extra

**APIs Partially Covered**

Same as for inspect volume-attachments, except that GetVolumeAttachment also includes the CHAP secret, if it exists.

## USE

**Permissions**

READ +

VOLUME\_ATTACHMENT\_UPDATE

**APIs Fully Covered**

no extra

**APIs Partially Covered**

no extra

## MANAGE

**Permissions**

USE +

VOLUME\_ATTACHMENT\_CREATE

VOLUME\_ATTACHMENT\_DELETE

**APIs Fully Covered**

no extra

**APIs Partially Covered**

USE +

AttachVolume, DetachVolume (both also need use volumes and use instances)

## volume-backups

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
VOLUME_BACKUP_INSPECT	<i>none</i>	ListVolumeBackups, GetVolumeBackup (both also need inspect volumes)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + VOLUME_BACKUP_READ	<i>none</i>	INSPECT + CreateVolume when creating volume from an backup (also need manage volumes)

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + VOLUME_BACKUP_COPY VOLUME_BACKUP_UPDATE	<i>none</i>	READ + UpdateVolumeBackup (also need inspect volumes) CopyVolumeBackup (also need create volume backups in destination region)

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + VOLUME_BACKUP_CREATE VOLUME_BACKUP_DELETE VOLUME_BACKUP_MOVE	ChangeVolumeBackupCompartment  When moving volume backups between compartments, the move volume backup permission is needed for both source and destination compartments.	USE + CreateVolumeBackup (also need use volumes) DeleteVolumeBackup (also need inspect volumes)

## boot-volume-backups

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
BOOT_VOLUME_BACKUP_INSPECT	<i>none</i>	ListBootVolumeBackups, GetBootVolumeBackup (both also need inspect volumes)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + BOOT_VOLUME_BACKUP_READ	<i>none</i>	INSPECT + CreateBootVolume when creating volume from a backup (also need manage volumes)

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + BOOT_VOLUME_BACKUP_UPDATE	<i>none</i>	READ + UpdateBootVolumeBackup (also need inspect volumes)

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + BOOT_VOLUME_BACKUP_CREATE BOOT_VOLUME_BACKUP_DELETE BOOT_VOLUME_BACKUP_MOVE	ChangeVolumeBackupCompartment When moving boot volume backups between compartments, the <code>move boot volume backup</code> permission is needed for both source and destination compartments.	USE + CreateBootVolumeBackup (also need use volumes) DeleteBootVolumeBackup (also need inspect volumes)

## backup-policies

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
BACKUP_POLICY_INSPECT	ListVolumeBackupPolicies GetVolumeBackupPolicy	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + BACKUP_POLICIES_UPDATE	READ + UpdateVolumeBackupPolicy	none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + BACKUP_POLICIES_CREATE BACKUP_POLICIES_DELETE	USE + CreateVolumeBackupPolicy DeleteVolumeBackupPolicy	none

## backup-policy-assignments

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
BACKUP_POLICY_ASSIGNMENT_INSPECT	GetVolumeBackupPolicyAssignment	GetVolumeBackupPolicyAssetAssignment (also need inspect volumes)

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

## CHAPTER 18 IAM

---

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

### MANAGE

**Permissions**

*USE +*

BACKUP\_POLICY\_ASSIGNMENT\_CREATE

BACKUP\_POLICY\_ASSIGNMENT\_DELETE

**APIs Fully Covered**

*USE +*

CreateVolumeBackupPolicyAssignment

DeleteVolumeBackupPolicyAssignment

**APIs Partially Covered**

*none*

## volume-groups

### INSPECT

**Permissions**

VOLUME\_GROUP\_INSPECT

**APIs Fully Covered**

ListVolumeGroups

GetVolumeGroup

**APIs Partially Covered**

*no extra*

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*



## CHAPTER 18 IAM

---

### Permissions

*USE +*

VOLUME\_GROUP\_UPDATE

VOLUME\_GROUP\_CREATE

VOLUME\_GROUP\_DELETE

VOLUME\_GROUP\_MOVE

### APIs Fully Covered

*USE +*

DeleteVolumeGroup



## volume-group-backups

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
VOLUME_GROUP_BACKUP_INSPECT	ListVolumeGroupBackups GetVolumeGroupBackup	<i>no extra</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>no extra</i>

**MANAGE**

**Permissions**

USE +

VOLUME\_GROUP\_BACKUP\_UPDATE

VOLUME\_GROUP\_BACKUP\_CREATE

VOLUME\_GROUP\_BACKUP\_DELETE

VOLUME\_GROUP\_BACKUP\_MOVE

**APIs Fully Covered**

USE +

UpdateVolumeGroupBackup

**APIs Partially Covered**

USE +

CreateVolumeGroupBackup also need the following:

- inspect volume group for the source volume group
- create volume group backup
- write volume for the source volumes
- create volume backup **OR** create boot volume backup for the destination volumes

DeleteVolumeGroupBackup also need delete volume backup **OR** delete boot volume backup

ChangeVolumeGroupBackupCompartment

(also need move volume backup **OR** move boot volume backup for the volumes in the request)

When moving volume group backups between compartments, the move volume group backup **and** move volume backup permissions are needed for both source and destination compartments.

### Permissions Required for Each API Operation

The following tables list the API operations grouped by resource type. The resource types are listed in alphabetical order.

For information about permissions, see [Permissions](#).

#### CORE SERVICES API OPERATIONS

API Operation	Permissions Required to Use the Operation
CreateVolumeBackupPolicy	BACKUP_POLICIES_CREATE
DeleteVolumeBackupPolicy	BACKUP_POLICIES_DELETE
GetVolumeBackupPolicy	BACKUP_POLICIES_INSPECT
ListVolumeBackupPolicies	BACKUP_POLICIES_INSPECT
CreateVolumeBackupPolicyAssignment	BACKUP_POLICY_ASSIGNMENT_CREATE
DeleteVolumeBackupPolicyAssignment	BACKUP_POLICY_ASSIGNMENT_DELETE
GetVolumeBackupPolicyAssetAssignment	BACKUP_POLICY_ASSIGNMENT_INSPECT and VOLUME_INSPECT
GetVolumeBackupPolicyAssignment	BACKUP_POLICY_ASSIGNMENT_INSPECT
ListClusterNetworks	CLUSTER_NETWORK_INSPECT and INSTANCE_POOL_INSPECT
ListClusterNetworkInstances	CLUSTER_NETWORK_READ and INSTANCE_POOL_READ
GetClusterNetwork	CLUSTER_NETWORK_READ and INSTANCE_POOL_READ
UpdateClusterNetwork	CLUSTER_NETWORK_UPDATE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
CreateClusterNetwork	CLUSTER_NETWORK_CREATE and INSTANCE_POOL_CREATE
ChangeClusterNetworkCompartment	CLUSTER_NETWORK_MOVE
TerminateClusterNetwork	CLUSTER_NETWORK_DELETE and INSTANCE_POOL_DELETE
ListConsoleHistories	CONSOLE_HISTORY_READ and INSTANCE_INSPECT
GetConsoleHistory	CONSOLE_HISTORY_READ and INSTANCE_INSPECT
ShowConsoleHistoryData	CONSOLE_HISTORY_READ and INSTANCE_READ and INSTANCE_IMAGE_READ
CaptureConsoleHistory	CONSOLE_HISTORY_CREATE and INSTANCE_READ and INSTANCE_IMAGE_READ
DeleteConsoleHistory	CONSOLE_HISTORY_DELETE
ListCpes	CPE_READ
GetCpe	CPE_READ
UpdateCpe	CPE_UPDATE
CreateCpe	CPE_CREATE
DeleteCpe	CPE_DELETE
ChangeCpeCompartment	CPE_RESOURCE_MOVE

API Operation	Permissions Required to Use the Operation
ListCrossConnects	CROSS_CONNECT_READ
GetCrossConnect	CROSS_CONNECT_READ
UpdateCrossConnect	CROSS_CONNECT_UPDATE
CreateCrossConnect	<p>CROSS_CONNECT_CREATE if not creating cross-connect in a cross-connect group.</p> <p>If creating the cross-connect in a cross-connect group, also need CROSS_CONNECT_CREATE and CROSS_CONNECT_ATTACH</p>
DeleteCrossConnect	<p>CROSS_CONNECT_DELETE if cross-connect is not in a cross-connect group.</p> <p>If the cross-connect is in a cross-connect group, also need CROSS_CONNECT_DELETE and CROSS_CONNECT_DETACH</p>
ChangeCrossConnectCompartment	CROSS_CONNECT_RESOURCE_MOVE
ListCrossConnectGroups	CROSS_CONNECT_GROUP_READ
GetCrossConnectGroup	CROSS_CONNECT_GROUP_READ
UpdateCrossConnectGroup	CROSS_CONNECT_GROUP_UPDATE
CreateCrossConnectGroup	CROSS_CONNECT_GROUP_CREATE
DeleteCrossConnectGroup	CROSS_CONNECT_GROUP_DELETE
ChangeCrossConnectGroupCompartment	CROSS_CONNECT_GROUP_RESOURCE_MOVE
ListDhcpOptions	DHCP_READ

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
GetDhcpOptions	DHCP_READ
UpdateDhcpOptions	DHCP_UPDATE
CreateDhcpOptions	DHCP_CREATE and VCN_ATTACH
DeleteDhcpOptions	DHCP_DELETE and VCN_DETACH
ChangeDhcpOptionsCompartment	DHCP_MOVE
ListDrgs	DRG_READ
GetDrg	DRG_READ
UpdateDrg	DRG_UPDATE
CreateDrg	DRG_CREATE
DeleteDrg	DRG_DELETE
ChangeDrgCompartment	DRG_MOVE
ListDrgAttachments	DRG_ATTACHMENT_READ
GetDrgAttachment	DRG_ATTACHMENT_READ
UpdateDrgAttachment	DRG_ATTACHMENT_UPDATE  ROUTE_TABLE_ATTACH is necessary to associate a route table with the DRG attachment during the update.

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
CreateDrgAttachment	DRG_ATTACH and VCN_ATTACH  ROUTE_TABLE_ATTACH is necessary to associate a route table with the DRG attachment during creation.
DeleteDrgAttachment	DRG_DETACH and VCN_DETACH
CreateInstanceConsoleConnection	INSTANCE_CONSOLE_CONNECTION_CREATE and INSTANCE_READ
DeleteInstanceConsoleConnection	INSTANCE_CONSOLE_CONNECTION_DELETE
GetInstanceConsoleConnection	INSTANCE_CONSOLE_CONNECTION_READ and INSTANCE_READ
ListInstanceConsoleConnections	INSTANCE_CONSOLE_CONNECTION_INSPECT and INSTANCE_INSPECT and INSTANCE_READ
ListImages	INSTANCE_IMAGE_READ
GetImage	INSTANCE_IMAGE_READ
UpdateImage	INSTANCE_IMAGE_UPDATE
CreateImage	INSTANCE_IMAGE_CREATE and INSTANCE_CREATE_IMAGE  The first permission is related to the <code>instance-image</code> ; the second is related to the <code>instance</code> .
ChangeImageCompartment	INSTANCE_IMAGE_MOVE
DeleteImage	INSTANCE_IMAGE_DELETE

API Operation	Permissions Required to Use the Operation
ListInstances	INSTANCE_READ
GetInstance	INSTANCE_READ
LaunchInstance	<p>INSTANCE_CREATE and INSTANCE_IMAGE_READ and VNIC_CREATE and VNIC_ATTACH and SUBNET_ATTACH</p> <p>If putting the instance in a network security group during instance creation, also need NETWORK_SECURITY_GROUP_UPDATE_MEMBERS and VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP</p>
UpdateInstance	INSTANCE_UPDATE
InstanceAction	INSTANCE_POWER_ACTIONS
ChangeInstanceCompartment	INSTANCE_MOVE
TerminateInstance	INSTANCE_DELETE and VNIC_DELETE and SUBNET_DETACH
ListInstanceConfigurations	INSTANCE_CONFIGURATION_INSPECT
GetInstanceConfiguration	INSTANCE_CONFIGURATION_READ
LaunchInstanceConfiguration	INSTANCE_CONFIGURATION_LAUNCH
UpdateInstanceConfiguration	INSTANCE_CONFIGURATION_UPDATE

API Operation	Permissions Required to Use the Operation
CreateInstanceConfiguration	<p>INSTANCE_CONFIGURATION_CREATE (if using the CreateInstanceConfigurationDetails subtype)</p> <p>INSTANCE_READ and VNIC_READ and VNIC_ATTACHMENT_READ and VOLUME_INSPECT and VOLUME_ATTACHMENT_INSPECT (if using the CreateInstanceConfigurationFromInstanceDetails subtype)</p>
ChangeInstanceConfigurationCompartment	INSTANCE_CONFIGURATION_MOVE
DeleteInstanceConfiguration	INSTANCE_CONFIGURATION_DELETE
ListInstancePools	INSTANCE_POOL_INSPECT
ListInstancePoolInstances	INSTANCE_POOL_READ
GetInstancePool	INSTANCE_POOL_READ
UpdateInstancePool	INSTANCE_POOL_UPDATE
ResetInstancePool	INSTANCE_POOL_POWER_ACTIONS
SoftresetInstancePool	INSTANCE_POOL_POWER_ACTIONS
StartInstancePool	INSTANCE_POOL_POWER_ACTIONS
StopInstancePool	INSTANCE_POOL_POWER_ACTIONS
CreateInstancePool	INSTANCE_POOL_CREATE and INSTANCE_CREATE and IMAGE_READ and VNIC_CREATE and SUBNET_ATTACH

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ChangeInstancePoolCompartment	INSTANCE_POOL_MOVE
TerminateInstancePool	INSTANCE_POOL_DELETE and INSTANCE_DELETE and VNIC_DELETE and SUBNET_DETACH and VOLUME_ATTACHMENT_DELETE and VOLUME_WRITE
ListInternetGateways	INTERNET_GATEWAY_READ
GetInternetGateway	INTERNET_GATEWAY_READ
UpdateInternetGateway	INTERNET_GATEWAY_UPDATE
CreateInternetGateway	INTERNET_GATEWAY_CREATE and VCN_ATTACH
DeleteInternetGateway	INTERNET_GATEWAY_DELETE and VCN_DETACH
ChangeInternetGatewayCompartment	INTERNET_GATEWAY_MOVE
ListIPSecConnections	IPSEC_CONNECTION_READ
GetIPSecConnection	IPSEC_CONNECTION_READ
UpdateIpSecConnection	IPSEC_CONNECTION_UPDATE
CreateIPSecConnection	DRG_ATTACH and CPE_ATTACH and IPSEC_CONNECTION_CREATE
DeleteIPSecConnection	DRG_DETACH and CPE_DETACH and IPSEC_CONNECTION_DELETE
GetIPSecConnectionDeviceConfig	IPSEC_CONNECTION_DEVICE_CONFIG_READ

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
GetIPSecConnectionDeviceStatus	IPSEC_CONNECTION_READ
ListIPSecConnectionTunnels	IPSEC_CONNECTION_READ
GetIPSecConnectionTunnel	IPSEC_CONNECTION_READ
UpdateIPSecConnectionTunnel	IPSEC_CONNECTION_UPDATE
GetIPSecConnectionTunnelSharedSecret	IPSEC_CONNECTION_DEVICE_CONFIG_READ
UpdateIPSecConnectionTunnelSharedSecret	IPSEC_CONNECTION_DEVICE_CONFIG_UPDATE
ListIpv6s	IPV6_READ and SUBNET_READ (if listing by subnet) and VNIC_READ (if listing by VNIC)
GetIpv6	IPV6_READ
UpdateIpv6	IPV6_UPDATE and VNIC_UNASSIGN and VNIC_ASSIGN (if moving IPv6 to a different VNIC)
CreateIpv6	IPV6_CREATE and SUBNET_ATTACH and VNIC_ASSIGN
DeleteIpv6	IPV6_DELETE and SUBNET_DETACH and VNIC_UNASSIGN
ListLocalPeeringGateways	LOCAL_PEERING_GATEWAY_READ
GetLocalPeeringGateway	LOCAL_PEERING_GATEWAY_READ

API Operation	Permissions Required to Use the Operation
UpdateLocalPeeringGateway	LOCAL_PEERING_GATEWAY_UPDATE  ROUTE_TABLE_ATTACH is necessary to associate a route table with the LPG during the update.
CreateLocalPeeringGateway	LOCAL_PEERING_GATEWAY_CREATE and VCN_ATTACH  ROUTE_TABLE_ATTACH is necessary to associate a route table with the LPG during creation.
DeleteLocalPeeringGateway	LOCAL_PEERING_GATEWAY_DELETE and VCN_DETACH
ConnectLocalPeeringGateway	LOCAL_PEERING_GATEWAY_CONNECT_FROM and LOCAL_PEERING_GATEWAY_CONNECT_TO
ChangeLocalPeeringGatewayCompartment	LOCAL_PEERING_GATEWAY_MOVE
ListNatGateways	NAT_GATEWAY_READ
GetNatGateway	NAT_GATEWAY_READ
UpdateNatGateway	NAT_GATEWAY_UPDATE
CreateNatGateway	NAT_GATEWAY_CREATE and VCN_READ and VCN_ATTACH

API Operation	Permissions Required to Use the Operation
DeleteNatGateway	NAT_GATEWAY_DELETE and VCN_READ and VCN_DETACH
ChangeNatGatewayCompartment	NAT_GATEWAY_MOVE
ListNetworkSecurityGroups	NETWORK_SECURITY_GROUP_READ
GetNetworkSecurityGroup	NETWORK_SECURITY_GROUP_READ
UpdateNetworkSecurityGroup	NETWORK_SECURITY_GROUP_UPDATE
CreateNetworkSecurityGroup	NETWORK_SECURITY_GROUP_CREATE and VCN_ATTACH
DeleteNetworkSecurityGroup	NETWORK_SECURITY_GROUP_DELETE and VCN_DETACH
ChangeNetworkSecurityGroupCompartment	NETWORK_SECURITY_GROUP_MOVE
ListNetworkSecurityGroupSecurityRules	NETWORK_SECURITY_GROUP_LIST_SECURITY_RULES
UpdateNetworkSecurityGroupSecurityRules	NETWORK_SECURITY_GROUP_UPDATE_SECURITY_RULES and NETWORK_SECURITY_GROUP_INSPECT if writing a rule that specifies a network security group as the source (for ingress rules) or destination (for egress rules)

API Operation	Permissions Required to Use the Operation
AddNetworkSecurityGroupSecurityRules	NETWORK_SECURITY_GROUP_UPDATE_SECURITY_RULES and  NETWORK_SECURITY_GROUP_INSPECT if writing a rule that specifies a network security group as the source (for ingress rules) or destination (for egress rules)
RemoveNetworkSecurityGroupSecurityRules	NETWORK_SECURITY_GROUP_UPDATE_SECURITY_RULES
ListPrivateIps	PRIVATE_IP_READ
GetPrivateIp	PRIVATE_IP_READ
UpdatePrivateIp	PRIVATE_IP_UPDATE
CreatePrivateIp	PRIVATE_IP_CREATE and PRIVATE_IP_ASSIGN and VNIC_ASSIGN and SUBNET_ATTACH
DeletePrivateIp	PRIVATE_IP_DELETE and PRIVATE_IP_UNASSIGN and VNIC_UNASSIGN and SUBNET_DETACH
ListRemotePeeringConnections	REMOTE_PEERING_CONNECTION_READ
GetRemotePeeringConnection	REMOTE_PEERING_CONNECTION_READ
UpdateRemotePeeringConnection	REMOTE_PEERING_CONNECTION_UPDATE
CreateRemotePeeringConnection	REMOTE_PEERING_CONNECTION_CREATE and DRG_ATTACH

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
DeleteRemotePeeringConnection	REMOTE_PEERING_CONNECTION_DELETE and DRG_DETACH
ChangeRemotePeeringConnectionCompartment	REMOTE_PEERING_CONNECTION_RESOURCE_MOVE
ConnectRemotePeeringConnections	REMOTE_PEERING_CONNECTION_CONNECT_FROM and REMOTE_PEERING_CONNECTION_CONNECT_TO
ListPublicIps	For ephemeral public IPs: PRIVATE_IP_READ For reserved public IPs: PUBLIC_IP_READ
GetPublicIp	For ephemeral public IPs: PRIVATE_IP_READ For reserved public IPs: PUBLIC_IP_READ
GetPublicIpByPrivateIpId	For ephemeral public IPs: PRIVATE_IP_READ For reserved public IPs: PUBLIC_IP_READ
GetPublicIpByIpAddress	For ephemeral public IPs: PRIVATE_IP_READ For reserved public IPs: PUBLIC_IP_READ
UpdatePublicIP	For ephemeral public IPs: PRIVATE_IP_UPDATE For reserved public IPs: PUBLIC_IP_UPDATE and PRIVATE_IP_ASSIGN_PUBLIC_IP and PUBLIC_IP_ASSIGN_PRIVATE_IP and PRIVATE_IP_UNASSIGN_PUBLIC_IP and PUBLIC_IP_UNASSIGN_PRIVATE_IP

API Operation	Permissions Required to Use the Operation
CreatePublicIp	<p>For ephemeral public IPs: PRIVATE_IP_ASSIGN_PUBLIC_IP</p> <p>For reserved public IPs: PUBLIC_IP_CREATE and PUBLIC_IP_ASSIGN_PRIVATE_IP and PRIVATE_IP_ASSIGN_PUBLIC_IP</p>
DeletePublicIp	<p>For ephemeral public IPs: PRIVATE_IP_UNASSIGN_PUBLIC_IP</p> <p>For reserved public IPs: PUBLIC_IP_DELETE and PUBLIC_IP_UNASSIGN_PRIVATE_IP and PRIVATE_IP_UNASSIGN_PUBLIC_IP</p>
ChangePublicIpCompartment	<p>PUBLIC_IP_MOVE</p> <p>Note: This operation applies only to reserved public IPs.</p>
ListRouteTables	ROUTE_TABLE_READ
GetRouteTable	ROUTE_TABLE_READ

API Operation	Permissions Required to Use the Operation
UpdateRouteTable	<p>ROUTE_TABLE_UPDATE and</p> <p>INTERNET_GATEWAY_ATTACH (if creating a route rule that uses an internet gateway as a target) and</p> <p>INTERNET_GATEWAY_DETACH (if deleting a route rule that uses an internet gateway as a target) and</p> <p>DRG_ATTACH (if creating a route rule that uses a DRG as a target) and</p> <p>DRG_DETACH (if deleting a route rule that uses a DRG as a target) and</p> <p>PRIVATE_IP_ROUTE_TABLE_ATTACH (if creating a route rule that uses a private IP as a target) and</p> <p>PRIVATE_IP_ROUTE_TABLE_DETACH (if deleting a route rule that uses a private IP as a target) and</p> <p>LOCAL_PEERING_GATEWAY_ATTACH (if creating a route rule that uses an LPG as a target) and</p> <p>LOCAL_PEERING_GATEWAY_DETACH (if deleting a route rule that uses an LPG as a target) and</p> <p>NAT_GATEWAY_ATTACH (if creating a route rule that uses a NAT gateway as a target) and</p> <p>NAT_GATEWAY_DETACH (if deleting a route</p>

API Operation	Permissions Required to Use the Operation
	<p>rule that uses a NAT gateway as a target) and</p> <p>SERVICE_GATEWAY_ATTACH (if creating a route rule that uses a service gateway as a target) and</p> <p>SERVICE_GATEWAY_DETACH (if deleting a route rule that uses a service gateway as a target)</p>
CreateRouteTable	<p>ROUTE_TABLE_CREATE and VCN_ATTACH and</p> <p>INTERNET_GATEWAY_ATTACH (if creating a route rule that uses an internet gateway as a target) and</p> <p>DRG_ATTACH (if creating a route rule that uses a DRG as a target) and</p> <p>PRIVATE_IP_ROUTE_TABLE_ATTACH (if creating a route rule that uses a private IP as a target) and</p> <p>LOCAL_PEERING_GATEWAY_ATTACH (if creating a route rule that uses an LPG as a target) and</p> <p>NAT_GATEWAY_ATTACH (if creating a route rule that uses a NAT gateway as a target) and</p> <p>SERVICE_GATEWAY_ATTACH (if creating a route rule that uses a service gateway as a target)</p>

API Operation	Permissions Required to Use the Operation
DeleteRouteTable	ROUTE_TABLE_DELETE and VCN_DETACH and INTERNET_GATEWAY_DETACH (if deleting a route rule that uses an internet gateway as a target) and DRG_DETACH (if deleting a route rule that uses a DRG as a target) and PRIVATE_IP_ROUTE_TABLE_DETACH (if deleting a route rule that uses a private IP as a target) and LOCAL_PEERING_GATEWAY_DETACH (if deleting a route rule that uses an LPG as a target) and NAT_GATEWAY_DETACH (if deleting a route rule that uses a NAT gateway as a target) and SERVICE_GATEWAY_DETACH (if deleting a route rule that uses a service gateway as a target)
ChangeRouteTableCompartment	ROUTE_TABLE_MOVE
ListSecurityLists	SECURITY_LIST_READ
GetSecurityList	SECURITY_LIST_READ
UpdateSecurityList	SECURITY_LIST_UPDATE
ChangeSecurityListCompartment	SECURITY_LIST_MOVE
CreateSecurityList	SECURITY_LIST_CREATE and VCN_ATTACH

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
DeleteSecurityList	SECURITY_LIST_DELETE and VCN_DETACH
ListServiceGateways	SERVICE_GATEWAY_READ
GetServiceGateway	SERVICE_GATEWAY_READ
UpdateServiceGateway	SERVICE_GATEWAY_UPDATE  ROUTE_TABLE_ATTACH is necessary to associate a route table with the service gateway during the update.
ChangeServiceGatewayCompartment	SERVICE_GATEWAY_MOVE
CreateServiceGateway	SERVICE_GATEWAY_CREATE and VCN_READ and VCN_ATTACH  ROUTE_TABLE_ATTACH is necessary to associate a route table with the service gateway during creation.
DeleteServiceGateway	SERVICE_GATEWAY_DELETE and VCN_READ and VCN_DETACH
AttachServiceId	SERVICE_GATEWAY_ADD_SERVICE
DetachServiceId	SERVICE_GATEWAY_DELETE_SERVICE
ListShapes	MACHINE_SHAPE_READ
ListSubnets	SUBNET_READ
GetSubnet	SUBNET_READ

API Operation	Permissions Required to Use the Operation
UpdateSubnet	<p>SUBNET_UPDATE</p> <p>If changing which route table is associated with the subnet, also need ROUTE_TABLE_ATTACH and ROUTE_TABLE_DETACH</p> <p>If changing which security lists are associated with the subnet, also need SECURITY_LIST_ATTACH and SECURITY_LIST_DETACH</p> <p>If changing which set of DHCP options are associated with the subnet, also need DHCP_ATTACH and DHCP_DETACH</p>
CreateSubnet	SUBNET_CREATE and VCN_ATTACH and ROUTE_TABLE_ATTACH and SECURITY_LIST_ATTACH and DHCP_ATTACH
DeleteSubnet	SUBNET_DELETE and VCN_DETACH and ROUTE_TABLE_DETACH and SECURITY_LIST_DETACH and DHCP_DETACH
ChangeSubnetCompartment	SUBNET_MOVE
ListVcns	VCN_READ
GetVcn	VCN_READ
UpdateVcn	VCN_UPDATE
CreateVcn	VCN_CREATE
DeleteVcn	VCN_DELETE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ChangeVcnCompartment	VCN_MOVE
ListVirtualCircuits	VIRTUAL_CIRCUIT_READ
GetVirtualCircuit	VIRTUAL_CIRCUIT_READ
UpdateVirtualCircuit	VIRTUAL_CIRCUIT_UPDATE and DRG_ATTACH and DRG_DETACH  If updating which cross-connect or cross-connect group the virtual circuit is using, also need CROSS_CONNECT_DETACH and CROSS_CONNECT_ATTACH
CreateVirtualCircuit	VIRTUAL_CIRCUIT_CREATE and DRG_ATTACH  If creating the virtual circuit with a mapping to a specific cross-connect or cross-connect group, also need CROSS_CONNECT_ATTACH
DeleteVirtualCircuit	VIRTUAL_CIRCUIT_DELETE and DRG_DETACH  If deleting a virtual circuit that's currently using a cross-connect or cross-connect group, also need CROSS_CONNECT_DETACH
changeVirtualCircuitCompartment	VIRTUAL_CIRCUIT_RESOURCE_MOVE
GetVnic	VNIC_READ

API Operation	Permissions Required to Use the Operation
AttachVnic	<p>INSTANCE_ATTACH_SECONDARY_VNIC and VNIC_ATTACH and VNIC_CREATE and SUBNET_ATTACH</p> <p>If putting the secondary VNIC in a network security group during VNIC creation, also need NETWORK_SECURITY_GROUP_UPDATE_MEMBERS and VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP</p>
DetachVnic	<p>INSTANCE_DETACH_SECONDARY_VNIC and VNIC_DETACH and VNIC_DELETE and SUBNET_DETACH</p>
UpdateVnic	<p>VNIC_UPDATE</p> <p>If adding or removing the VNIC from a network security group, also need NETWORK_SECURITY_GROUP_UPDATE_MEMBERS and VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP</p>
ListVnicAttachments	<p>VNIC_ATTACHMENT_READ and INSTANCE_INSPECT</p>
GetVnicAttachment	<p>VNIC_ATTACHMENT_READ</p>
ListVolumes	<p>VOLUME_INSPECT</p>
GetVolume	<p>VOLUME_INSPECT</p>
UpdateVolume	<p>VOLUME_UPDATE</p>

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
CreateVolume	VOLUME_CREATE (and VOLUME_BACKUP_READ if creating volume from a backup)
DeleteVolume	VOLUME_DELETE
ChangeVolumeCompartment	VOLUME_MOVE
ListVolumeAttachments	VOLUME_ATTACHMENT_INSPECT and VOLUME_INSPECT and INSTANCE_INSPECT
GetVolumeAttachment	VOLUME_ATTACHMENT_INSPECT and VOLUME_INSPECT and INSTANCE_INSPECT  <b>Note:</b> To also get the CHAP secret for the volume, then VOLUME_ATTACHMENT_READ is required instead of VOLUME_ATTACHMENT_INSPECT
AttachVolume	VOLUME_ATTACHMENT_CREATE and VOLUME_WRITE and INSTANCE_ATTACH_VOLUME
DetachVolume	VOLUME_ATTACHMENT_DELETE and VOLUME_WRITE and INSTANCE_DETACH_VOLUME
ListVolumeBackups	VOLUME_BACKUP_INSPECT and VOLUME_INSPECT
GetVolumeBackup	VOLUME_BACKUP_INSPECT and VOLUME_INSPECT
UpdateVolumeBackup	VOLUME_BACKUP_UPDATE and VOLUME_INSPECT
CreateVolumeBackup	VOLUME_BACKUP_CREATE and VOLUME_WRITE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
DeleteVolumeBackup	VOLUME_BACKUP_DELETE and VOLUME_INSPECT
ChangeVolumeBackupCompartment	VOLUME_BACKUP_MOVE
GetBootVolume	VOLUME_INSPECT
ListBootVolumes	VOLUME_INSPECT
UpdateBootVolume	VOLUME_UPDATE
DeleteBootVolume	VOLUME_DELETE
ChangeBootVolumeCompartment	BOOT_VOLUME_MOVE
CreateBootVolumeBackup	BOOT_VOLUME_BACKUP_CREATE, VOLUME_WRITE
ListBootVolumeBackups	BOOT_VOLUME_BACKUP_INSPECT, VOLUME_INSPECT
GetBootVolumeBackup	BOOT_VOLUME_BACKUP_INSPECT, VOLUME_INSPECT
UpdateBootVolumeBackup	BOOT_VOLUME_BACKUP_UPDATE, VOLUME_INSPECT
DeleteBootVolumeBackup	BOOT_VOLUME_BACKUP_DELETE, VOLUME_INSPECT
ChangeBootVolumeBackupCompartment	BOOT_VOLUME_BACKUP_MOVE

API Operation	Permissions Required to Use the Operation
CreateVolumeGroup	<p>VOLUME_GROUP_CREATE, VOLUME_INSPECT if creating the volume group from a list of volumes.</p> <p>VOLUME_GROUP_CREATE, VOLUME_GROUP_INSPECT, VOLUME_CREATE, VOLUME_WRITE if cloning a volume group.</p> <p>VOLUME_GROUP_CREATE, VOLUME_GROUP_BACKUP_INSPECT, VOLUME_BACKUP_READ/BOOT_VOLUME_BACKUP_READ, VOLUME_CREATE, VOLUME_WRITE if restoring from a volume group backup.</p>
DeleteVolumeGroup	VOLUME_GROUP_DELETE
GetVolumeGroup	VOLUME_GROUP_INSPECT
ListVolumeGroups	VOLUME_GROUP_INSPECT
UpdateVolumeGroup	VOLUME_GROUP_UPDATE, VOLUME_INSPECT
ChangeVolumeGroupCompartment	VOLUME_GROUP_MOVE, VOLUME_MOVE/BOOT_VOLUME_MOVE
CreateVolumeGroupBackup	VOLUME_GROUP_BACKUP_CREATE, VOLUME_GROUP_INSPECT, VOLUME_WRITE, VOLUME_BACKUP_CREATE/BOOT_VOLUME_BACKUP_CREATE
DeleteVolumeGroupBackup	VOLUME_GROUP_BACKUP_DELETE, VOLUME_BACKUP_DELETE/BOOT_VOLUME_BACKUP_DELETE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
GetVolumeGroupBackup	VOLUME_GROUP_BACKUP_INSPECT
ListVolumeGroupBackups	VOLUME_GROUP_BACKUP_INSPECT
UpdateVolumeGroupBackup	VOLUME_GROUP_BACKUP_UPDATE
ChangeVolumeGroupBackupCompartment	VOLUME_GROUP_BACKUP_MOVE, VOLUME_BACKUP_MOVE/BOOT_VOLUME_BACKUP_MOVE

### DEDICATED VIRTUAL MACHINE HOST API OPERATIONS

API Operation	Permissions Required to Use the Operation
CreateDedicatedVmHost	DEDICATED_VM_HOST_CREATE
ChangeDedicatedVmHostCompartment	DEDICATED_VM_HOST_MOVE
DeleteDedicatedVmHost	DEDICATED_VM_HOST_DELETE
GetDedicatedVmHost	DEDICATED_VM_HOST_READ
ListDedicatedVmHosts	DEDICATED_VM_HOST_INSPECT
ListDedicatedVmHostInstances	DEDICATED_VM_HOST_READ
ListDedicatedVmHostInstanceShapes	None
ListDedicatedVmHostShapes	None

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
LaunchInstance	DEDICATED_VM_HOST_LAUNCH_INSTANCE in dedicated virtual machine host compartment INSTANCE_CREATE in compartment for the instance launched on the dedicated virtual machine host
UpdateDedicatedVmHost	AUTO_SCALING_CONFIGURATION_CREATE and INSTANCE_POOL_UPDATE

### AUTOSCALING API OPERATIONS

API Operation	Permissions Required to Use the Operation
ListAutoScalingConfigurations	AUTO_SCALING_CONFIGURATION_INSPECT
GetAutoScalingConfiguration	AUTO_SCALING_CONFIGURATION_READ
UpdateAutoScalingConfiguration	AUTO_SCALING_CONFIGURATION_UPDATE and INSTANCE_POOL_UPDATE
CreateAutoScalingConfiguration	AUTO_SCALING_CONFIGURATION_CREATE and INSTANCE_POOL_UPDATE
ChangeAutoScalingConfigurationCompartment	AUTO_SCALING_CONFIGURATION_MOVE
DeleteAutoScalingConfiguration	AUTO_SCALING_CONFIGURATION_DELETE and INSTANCE_POOL_UPDATE

API Operation	Permissions Required to Use the Operation
ListAutoScalingPolicies	AUTO_SCALING_CONFIGURATION_READ
GetAutoScalingPolicy	AUTO_SCALING_CONFIGURATION_READ
UpdateAutoScalingPolicy	AUTO_SCALING_CONFIGURATION_UPDATE and INSTANCE_POOL_UPDATE
CreateAutoScalingPolicy	AUTO_SCALING_CONFIGURATION_CREATE and INSTANCE_POOL_UPDATE
DeleteAutoScalingPolicy	AUTO_SCALING_CONFIGURATION_DELETE and INSTANCE_POOL_UPDATE

**WORK REQUESTS API OPERATIONS**

API Operation	Permissions Required to Use the Operation
ListWorkRequests	WORKREQUEST_INSPECT
GetWorkRequests	Work requests inherit the permissions of the operation that spawns the work request. Generally, <i>&lt;RESOURCE&gt;</i> _CREATE permissions for the associated resource are required.
ListWorkRequestLogs	Work requests inherit the permissions of the operation that spawns the work request. Generally, <i>&lt;RESOURCE&gt;</i> _CREATE permissions for the associated resource are required.
ListWorkRequestErrors	Work requests inherit the permissions of the operation that spawns the work request. Generally, <i>&lt;RESOURCE&gt;</i> _CREATE permissions for the associated resource are required.

## Details for Container Engine for Kubernetes

This topic covers details for writing policies to control access to Container Engine for Kubernetes.

### Resource-Types

#### AGGREGATE RESOURCE-TYPE

- `cluster-family`

#### INDIVIDUAL RESOURCE-TYPES

- `clusters`
- `cluster-node-pools`
- `cluster-work-requests`

#### COMMENTS

A policy that uses `<verb> cluster-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `cluster-family`.

### Supported Variables

Container Engine for Kubernetes supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

The `clusters` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.cluster.id</code>	Entity (OCID)	

The `cluster-node-pools` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.nodepool.id</code>	Entity (OCID)	

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `clusters` resource-type includes the same permissions and API operations as the `inspect` verb, plus the `CLUSTER_READ` permission and a number of API operations (e.g., `GetCluster`, etc.). The `use` verb covers still another permission and API operation compared to `read`. Lastly, `manage` covers more permissions and operations compared to `use`.

## clusters

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>CLUSTER_INSPECT</code>	<code>ListClusters</code> <code>ListWorkRequests</code>	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + <code>CLUSTER_READ</code>	<i>INSPECT</i> + <code>GetCluster</code> <code>GetWorkRequest</code> <code>ListWorkRequestErrors</code> <code>ListWorkRequestLogs</code>	<i>none</i>

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + CLUSTER_USE	<i>READ</i> + GetClusterKubeconfig	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + CLUSTER_CREATE CLUSTER_DELETE CLUSTER_UPDATE CLUSTER_MANAGE	<i>USE</i> + UpdateCluster AdministerK8s	CreateCluster (also need use subnets, read virtual-network-family, and inspect compartments) DeleteCluster (also need manage instance-family, use subnets, and use vnics)

## cluster-node-pools

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CLUSTER_NODE_POOL_INSPECT	ListNodePools ListWorkRequests	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + CLUSTER_NODE_POOL_READ	<i>INSPECT</i> + GetNodePool GetWorkRequest ListWorkRequestErrors ListWorkRequestLogs	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + CLUSTER_NODE_POOL_CREATE CLUSTER_NODE_POOL_DELETE CLUSTER_NODE_POOL_UPDATE	<i>no extra</i>	CreateNodePool, DeleteNodePool, and UpdateNodePool (also need manage instance-family, use subnets, use vnics, and inspect compartments)

## cluster-work-requests

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
CLUSTER_WORK_REQUEST_INSPECT	ListWorkRequests	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + CLUSTER_WORK_REQUEST_READ	<i>INSPECT</i> + GetWorkRequest ListWorkRequestErrors ListWorkRequestLogs	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + CLUSTER_WORK_REQUEST_DELETE	<i>USE</i> + DeleteWorkRequest	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type. For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListClusters	CLUSTER_INSPECT
CreateCluster	CLUSTER_CREATE
GetClusterKubeconfig	CLUSTER_USE
GetCluster	CLUSTER_READ
UpdateCluster	CLUSTER_UPDATE
DeleteCluster	CLUSTER_DELETE, CLUSTER_NODE_POOL_DELETE
AdministerK8s	CLUSTER_MANAGE
ListNodePools	CLUSTER_NODE_POOL_INSPECT
CreateNodePool	CLUSTER_NODE_POOL_CREATE
GetNodePool	CLUSTER_NODE_POOL_READ
GetNodePoolOptions	CLUSTER_READ
UpdateNodePool	CLUSTER_NODE_POOL_UPDATE
DeleteNodePool	CLUSTER_NODE_POOL_DELETE
ListWorkRequests	CLUSTER_WORK_REQUEST_INSPECT, CLUSTER_NODE_POOL_INSPECT, CLUSTER_INSPECT
GetWorkRequest	CLUSTER_WORK_REQUEST_READ, CLUSTER_NODE_POOL_READ, CLUSTER_READ

API Operation	Permissions Required to Use the Operation
ListWorkRequestErrors	CLUSTER_WORK_REQUEST_READ, CLUSTER_NODE_POOL_READ, CLUSTER_READ
ListWorkRequestLogs	CLUSTER_WORK_REQUEST_READ, CLUSTER_NODE_POOL_READ, CLUSTER_READ
DeleteWorkRequest	CLUSTER_WORK_REQUEST_DELETE

## Details for the Database Service

See the following topics for details for writing [policies](#) to control access to Oracle Cloud Infrastructure Database service resources:

- [Policy Details for Autonomous Database](#)
- [Policy Details for DB Systems](#)
- [Policy Details for Exadata Cloud at Customer](#)

## Details for the DNS Service

This topic covers details for writing policies to control access to the DNS service.

### Aggregate Resource-Type

dns

### Individual Resource-Types

dns-zones

dns-records

dns-steering-policies

dns-steering-policy-attachments

## CHAPTER 18 IAM

---

### COMMENTS

A policy that uses `<verb> dns` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `dns`.

### Supported Variables

The DNS Service supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

The `dns-zones` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.dns-zone.id</code>	Entity (OCID)	Use this variable to control access to specific DNS zones by OCID.
<code>target.dns-zone.name</code>	String	Use this variable to control access to specific DNS zones by name.

The `dns-records` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.dns-zone.id</code>	Entity (OCID)	Use this variable to control access to specific DNS zones by OCID.
<code>target.dns-zone.name</code>	String	Use this variable to control access to specific DNS zones by name.
<code>target.dns-zone.scope</code>	String	Valid values are "public" and "private".

Variable	Variable Type	Comments
<code>target.dns-record.type</code>	List (String)	Use this variable to control access to specific DNS records by type. Valid values in the last can be any supported DNS resource type. For example, "A", "AAAA", "TXT", and so on. See .
<code>target.dns-domain.name</code>	List (String)	Use this variable to control access to specific domain names. Applicable to the following API operations: <ul style="list-style-type: none"> <li>• <code>GetDomainRecords</code></li> <li>• <code>PatchDomainRecords</code></li> <li>• <code>UpdateDomainRecords</code></li> <li>• <code>DeleteRRSet</code></li> <li>• <code>GetRRSet</code></li> <li>• <code>PatchRRSet</code></li> <li>• <code>UpdateRRSet</code></li> </ul>

The `dns-steering-policies` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.dns-steering-policy.id</code>	Entity (OCID)	Use this variable to control access to specific steering policies by OCID.
<code>target.dns-steering-policy.display-name</code>	String	Use this variable to control access to specific steering policies by name.

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `manage` verb for the `dns-records` resource-type covers no extra permissions or API operations compared to the `use` verb.

### dns-zones

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DNS_ZONE_INSPECT	ListZones	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	GetZone	GetZoneRecords
DNS_ZONE_READ		

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ +	UpdateZone	UpdateZoneRecords
DNS_ZONE_UPDATE		PatchZoneRecords
		CreateSteeringPolicyAttachment
		DeleteSteeringPolicyAttachment

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
UPDATE +	CreateZone	none
DNS_ZONE_CREATE	DeleteZone	
DNS_ZONE_DELETE		

## dns-records

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
DNS_RECORD_INSPECT	<i>none</i>	<i>none</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> +	GetDomainRecords	GetZoneRecords
DNS_RECORD_READ	GetRRSet	

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	PatchDomainRecords	UpdateZoneRecords
DNS_RECORD_UPDATE	UpdateDomainRecords	PatchZoneRecords
	DeleteRRSet	UpdateSteeringPolicyAttachment
	PatchRRSet	
	UpdateRRSet	

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>UPDATE</i> +	<i>no extra</i>	<i>none</i>
DNS_RECORD_CREATE		
DNS_RECORD_DELETE		

## dns-steering-policies

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
DNS_STEERING_POLICY_INSPECT	ListSteeringPolicies	<i>none</i>

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + DNS_STEERING_POLICY_READ	GetSteeringPolicy	CreateSteeringPolicyAttachment UpdateSteeringPolicyAttachment DeleteSteeringPolicyAttachment

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + DNS_POLICY_STEERING_UPDATE	UpdateSteeringPolicy	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>UPDATE</i> + DNS_STEERING_POLICY_CREATE DNS_STEERING_POLICY_DELETE	CreateSteeringPolicy DeleteSteeringPolicy	<i>none</i>

## dns-steering-policy-attachments

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
DNS_STEERING_ATTACHMENT_INSPECT	ListSteeringPolicyAttachments	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + DNS_STEERING_ATTACHMENT_READ	GetSteeringPolicyAttachment	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type. For information about permissions, see [Permissions](#).

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ListZones	DNS_ZONE_INSPECT
CreateZone	DNS_ZONE_CREATE
DeleteZone	DNS_ZONE_DELETE
GetZone	DNS_ZONE_READ
UpdateZone	DNS_ZONE_UPDATE
GetZoneRecords	DNS_ZONE_READ and DNS_RECORD_READ
PatchZoneRecords	DNS_ZONE_UPDATE and DNS_RECORD_UPDATE
UpdateZoneRecords	DNS_ZONE_UPDATE and DNS_RECORD_UPDATE
GetDomainRecords	DNS_RECORD_READ
PatchDomainRecords	DNS_RECORD_UPDATE
UpdateDomainRecords	DNS_RECORD_UPDATE
DeleteRRSet	DNS_RECORD_UPDATE
GetRRSet	DNS_RECORD_READ
PatchRRSet	DNS_RECORD_UPDATE
UpdateRRSet	DNS_RECORD_UPDATE
ListSteeringPolicies	DNS_STEERING_ POLICY_INSPECT

API Operation	Permissions Required to Use the Operation
CreateSteeringPolicy	DNS_STEERING_POLICY_CREATE
GetSteeringPolicy	DNS_STEERING_POLICY_READ
UpdateSteeringPolicy	DNS_STEERING_POLICY_UPDATE
DeleteSteeringPolicy	DNS_STEERING_POLICY_DELETE
ListSteeringPolicyAttachments	DNS_STEERING_ATTACHMENT_INSPECT
CreateSteeringPolicyAttachment	DNS_ZONE_UPDATE and DNS_STEERING_POLICY_READ
GetSteeringPolicyAttachment	DNS_STEERING_ATTACHMENT_READ
UpdateSteeringPolicyAttachment	DNS_ZONE_UPDATE and DNS_STEERING_POLICY_READ
DeleteSteeringPolicyAttachment	DNS_ZONE_UPDATE and DNS_STEERING_POLICY_READ

## Details for the Email Service

This topic covers details for writing policies to control access to the Email service.

**Resource-Types**

approved-senders

suppressions

**Supported Variables**

The Email Service supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

The `approved-senders` resource type can use the following variables:

Variable	Variable Type	Comments
<code>target.approved-sender.id</code>	Entity (OCID)	
<code>target.approved-sender.emailaddress</code>	String	Use this variable with the APPROVED_SENDER_USE permissions only.

**Details for Verb + Resource-Type Combinations**

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `use` verb for the `suppressions` resource-type covers no extra permissions or API operations compared to the `read` verb.

## approved-senders

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
APPROVED_SENDER_INSPECT	ListSenders	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + APPROVED_SENDER_READ	GetSender	None

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + APPROVED_SENDER_USE	SmtSend	None

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + APPROVED_SENDER_CREATE	CreateSender	none
APPROVED_SENDER_DELETE	DeleteSender	
APPROVED_SENDER_UPDATE	UpdateSender	
APPROVED_SENDER_MOVE	MoveSender	

## suppressions

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
SUPPRESSION_INSPECT	ListSuppression	none

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + SUPPRESSION_READ	GetSuppression	<i>None</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + SUPPRESSION_CREATE SUPPRESSION_DELETE	CreateSuppression DeleteSuppression	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListSenders	APPROVED_SENDER_INSPECT
GetSender	APPROVED_SENDER_READ
CreateSender	APPROVED_SENDER_CREATE
DeleteSender	APPROVED_SENDER_DELETE
MoveSender	APPROVED_SENDER_MOVE
SmtplibSend	APPROVED_SENDER_USE

API Operation	Permissions Required to Use the Operation
ListSuppression	SUPPRESSION_INSPECT
GetSuppression	SUPPRESSION_READ
CreateSuppression	SUPPRESSION_CREATE
DeleteSuppression	SUPPRESSION_DELETE

## Details for the Events Service

This topic covers details for writing user IAM policies that control access to rules for the Events service.

### Resource-Types

`cloudevents-rules`

### Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for `cloudevents-rules` includes the same permissions and API operations as the `inspect` verb, plus the `EVENTRULE_READ` permissions and the corresponding API operation `GetEventRule`. The `use` verb adds no extra permissions or API operations compared to `read`. However, `manage` adds more permissions and operations compared to `use`.

## CHAPTER 18 IAM

### CLOUDEVENTS-RULES

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
EVENTRULE_LIST	ListRules	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<b>INSPECT +</b> EVENTRULE_READ	<b>INSPECT +</b> GetRule	none

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<b>USE +</b> EVENTRULE_CREATE EVENTRULE_DELETE EVENTRULE_MODIFY	<b>USE +</b> CreateRule DeleteRule UpdateRule ChangeRuleCompartment	none

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListRules	EVENTRULE_LIST
CreateRule	EVENTRULE_CREATE
GetRule	EVENTRULE_READ

API Operation	Permissions Required to Use the Operation
DeleteRule	EVENTRULE_DELETE
UpdateRule	EVENTRULE_MODIFY
ChangeRuleCompartment	EVENTRULE_MODIFY

## Details for the File Storage Service

This topic covers details for writing policies to control access to the File Storage Service.

### Aggregate Resource-Type

- `file-family`

### Individual Resource-Types

- `file-systems`
- `mount-targets`
- `export-sets`

#### COMMENTS

A policy that uses `<verb> file-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `file-family`.

### Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `file-systems` resource-type includes the same permissions and API operations as the `inspect` verb, plus the `FILE_SYSTEM_READ` permission and a number of API operations (e.g., `GetFileSystem`, `ListMountTargets`, etc.). The `use` verb covers still another permission and set of API operations compared to `read`. Lastly, `manage` covers two more permissions and operations compared to `use`.

### export-sets

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
EXPORT_SET_INSPECT	ListExportSets	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	INSPECT +	none
EXPORT_SET_READ	GetExport GetExportSet ListExports	

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> +	<i>USE</i> +	CreateExport
EXPORT_SET_CREATE	CreateExportSet	DeleteExport
EXPORT_SET_UPDATE	UpdateExportSet	(both also need use file-systems.)
EXPORT_SET_DELETE	DeleteExportSet	

## file-systems

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
FILE_SYSTEM_INSPECT	ListFileSystems	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> +	<i>INSPECT</i> +	<i>none</i>
FILE_SYSTEM_READ	GetFileSystem	
	GetSnapshot	
	ListSnapshots	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	<i>no extra</i>	CreateExport
FILE_SYSTEM_NFSv3_EXPORT		DeleteExport
FILE_SYSTEM_NFSv3_UNEXPORT		(both also need manage export-sets.)

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	If creating a file system <i>encrypted with a Key Management master encryption key</i> , also
FILE_SYSTEM_CREATE	CreateFileSystem	need use <code>key-delegate</code> (for the caller) and
FILE_SYSTEM_UPDATE	UpdateFileSystem	<code>read keys</code> (for the service principal). For
FILE_SYSTEM_DELETE	DeleteFileSystem	more information, see <a href="#">Details for the Key Management Service</a> .
FILE_SYSTEM_CREATE_SNAPSHOT	CreateSnapshot	
FILE_SYSTEM_DELETE_SNAPSHOT	DeleteSnapshot	

## mount-targets

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
MOUNT_TARGET_INSPECT	ListMountTargets	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i>	<i>INSPECT +</i>	<i>none</i>
MOUNT_TARGET_READ	GetMountTarget	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	UpdateMountTarget	<i>USE +</i>
MOUNT_TARGET_CREATE		CreateMountTarget,
MOUNT_TARGET_UPDATE		DeleteMountTarget
MOUNT_TARGET_DELETE		(both also need use <code>vnics</code> , use <code>private-ips</code> , and use <code>subnets</code> .)

## Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.



### Tip

If a group uses the Console to create file systems, permissions to read mount targets is required. [See the file storage policy examples](#) for further guidance.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListExports	EXPORT_SET_READ
CreateExport	EXPORT_SET_UPDATE + FILE_SYSTEM_NFSv3_EXPORT
GetExport	EXPORT_SET_READ
DeleteExport	EXPORT_SET_UPDATE + FILE_SYSTEM_NFSv3_UNEXPORT
ListExportSets	EXPORT_SET_INSPECT
CreateExportSet	EXPORT_SET_CREATE
GetExportSet	EXPORT_SET_READ
UpdateExportSet	EXPORT_SET_UPDATE
DeleteExportSet	EXPORT_SET_DELETE
ListFileSystems	FILE_SYSTEM_INSPECT
CreateFileSystem	FILE_SYSTEM_CREATE
GetFileSystem	FILE_SYSTEM_READ

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
UpdateFileSystem	FILE_SYSTEM_UPDATE
DeleteFileSystem	FILE_SYSTEM_DELETE
ListMountTargets	MOUNT_TARGET_INSPECT
CreateMountTarget	MOUNT_TARGET_CREATE + VNIC_CREATE(vnicCompartment) + SUBNET_ATTACH(subnetCompartment) + VNIC_ATTACH(vnicCompartment) + PRIVATE_IP_CREATE(subnetCompartment) + PRIVATE_IP_ASSIGN(subnetCompartment) + VNIC_ASSIGN(subnetCompartment)
GetMountTarget	MOUNT_TARGET_READ
UpdateMountTarget	MOUNT_TARGET_UPDATE
DeleteMountTarget	MOUNT_TARGET_DELETE + VNIC_DELETE(vnicCompartment) + SUBNET_DETACH(subnetCompartment) + VNIC_DETACH(vnicCompartment) + PRIVATE_IP_DELETE(subnetCompartment) + PRIVATE_IP_UNASSIGN(subnetCompartment) + VNIC_UNASSIGN(vnicCompartment)
ListSnapshots	FILE_SYSTEM_READ

API Operation	Permissions Required to Use the Operation
CreateSnapshot	FILE_SYSTEM_CREATE_SNAPSHOT
GetSnapshot	FILE_SYSTEM_READ
DeleteSnapshot	FILE_SYSTEM_DELETE_SNAPSHOT

## Details for Functions

This topic covers details for writing policies to control access to Oracle Functions.

### Resource-Types

#### AGGREGATE RESOURCE-TYPE

- `functions-family`

#### INDIVIDUAL RESOURCE-TYPES

- `fn-app`
- `fn-function`
- `fn-invocation`

#### COMMENTS

A policy that uses `<verb> functions-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `functions-family`.

### Supported Variables

Oracle Functions supports all the general variables (see [General Variables for All Requests](#)).

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `fn-app` resource-type includes the same permissions and API operations as the `inspect` verb, plus the `FN_APP_READ` permission and the `GetApp` API operation. In the case of the `fn-app` resource-type, the `use` verb covers no additional permissions or API operations compared to `read`. Lastly, `manage` covers more permissions and operations compared to `use`.

### fn-app

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
FN_APP_LIST	ListApp	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	INSPECT +	none
FN_APP_READ	GetApp	

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	USE +	none
FN_APP_CREATE	CreateApp	
FN_APP_DELETE	DeleteApp	
FN_APP_UPDATE	UpdateApp	

## fn-function

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
FN_FUNCTION_LIST	ListFunctions	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	INSPECT +	none
FN_FUNCTION_READ	GetFunction	

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	USE +	none
FN_FUNCTION_CREATE	CreateFunction	
FN_FUNCTION_DELETE	DeleteFunction	
FN_FUNCTION_UPDATE	UpdateFunction	

## fn-invocation

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
none	none	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
none	none	none

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
FN_INVOCATION	InvokeFunction	none

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type. For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
CreateApp	FN_APP_CREATE
DeleteApp	FN_APP_DELETE
ListApp	FN_APP_LIST
GetApp	FN_APP_READ
UpdateApp	FN_APP_UPDATE
CreateFunction	FN_FUNCTION_CREATE
DeleteFunction	FN_FUNCTION_DELETE
ListFunctions	FN_FUNCTION_LIST
GetFunction	FN_FUNCTION_READ
UpdateFunction	FN_FUNCTION_UPDATE
InvokeFunction	FN_INVOCATION

## Details for the Health Checks Service

This topic covers details for writing policies to control access to the Health Checks service.

### Resource-Types

health-check-monitor

health-check-results

on-demand-probe

vantage-points

health-check-family

### Supported Variables

The Health Checks Service supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here. Values in the list can be any valid test type. For example, HTTP, HTTPS, ICMP, etc.

Variable	Variable Type	Comments
target.health-check-monitor.test-type	String	
target.on-demand-probe.test-type	String	
target.health-check-results.test-type	String	

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `use` verb for the `health-check-monitor` resource-type covers no extra permissions or API operations compared to the `read` verb.

### health-check-monitor

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
HEALTH_CHECK_MONITOR_INSPECT	ListOHCMonitors	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	GetOHCMonitor	None
HEALTH_CHECK_MONITOR_READ		

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
No extra	None	None

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	CreateOHCMonitor	None
HEALTH_CHECK_MONITOR_MANAGE	UpdateOHCMonitor	
	DeleteOHCMonitor	
	MoveOHCMonitor	

## health-check-results

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
HEALTH_CHECK_RESULTS_READ	ListOHCProbeResults ListOHCProbeResultsForTarget	<i>None</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

## vantage-points

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
VANTAGE_POINTS_INSPECT	ListVantagePoints	<i>none</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>No extra</i>	<i>None</i>	<i>None</i>

**MANAGE****Permissions***No extra***APIs Fully Covered***None***APIs Partially Covered***None*

## on-demand-probe

**INSPECT****Permissions***No extra***APIs Fully Covered***None***APIs Partially Covered***None***READ****Permissions***No extra***APIs Fully Covered***None***APIs Partially Covered***None***USE****Permissions***No extra***APIs Fully Covered***None***APIs Partially Covered***None***MANAGE****Permissions***USE +**ON\_DEMAND\_PROBE\_MANAGE***APIs Fully Covered***CreateOnDemandOHCPProbe***APIs Partially Covered***None***Permissions Required for Each API Operation**

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ListOHCMonitors	HEALTH_CHECK_MONITOR_INSPECT
CreateOHCMonitor	HEALTH_CHECK_MONITOR_MANAGE
GetOHCMonitor	HEALTH_CHECK_MONITOR_READ
UpdateOHCMonitor	HEALTH_CHECK_MONITOR_MANAGE
DeleteOHCMonitor	HEALTH_CHECK_MONITOR_MANAGE
ListOHCProbeResults	HEALTH_CHECK_RESULTS_READ
ListOHCProbeResultsForTarget	HEALTH_CHECK_RESULTS_READ
ListVantagePoints	VANTAGE_POINTS_INSPECT
CreateOnDemandOHCProbe	ON_DEMAND_PROBE_MANAGE
MoveOHCMonitor	HEALTH_CHECK_MONITOR_MOVE

### Details for IAM

This topic covers details for writing policies to control access to IAM.

#### Resource-Types

authentication-policies

compartments

users

groups

dynamic-groups

policies

## CHAPTER 18 IAM

---

identity-providers

tenancies

tag-namespaces

tagdefinitions

tag-defaults

workrequest

### Supported Variables

IAM supports all the general variables (see [General Variables for All Requests](#)), plus additional ones listed here:

Operations for This Resource-Type...	Can Use These Variables...	Variable Type	Comments
users	target .user.id	Entity (OCID)	Not available to use with <a href="#">CreateUser</a> .
	target .user.name	String	
groups	target .group.id	Entity (OCID)	Not available to use with <a href="#">CreateGroup</a> .
	target .group.name	String	
	target .group .member	Boolean	True if request.user is a member of target.group.

Operations for This Resource-Type...	Can Use These Variables...	Variable Type	Comments
policies	target.policy.id	Entity (OCID)	Not available to use with <a href="#">CreatePolicy</a> .
	target.policy.name	String	
	target.policy.autoupdate	Boolean	Whether the policy being acted upon uses "Keep policy current" as its version date (i.e., either null or an empty string for the <code>versionDate</code> parameter in <a href="#">CreatePolicy</a> and <a href="#">UpdatePolicy</a> ).
compartments	target.compartment.id	Entity (OCID)	For <a href="#">CreateCompartment</a> , this will be the value of the parent compartment (e.g., the root compartment).  This is a universal variable available to use with any request across all services (see <a href="#">General Variables for All Requests</a> ).
	target.compartment.name	String	

Operations for This Resource-Type...	Can Use These Variables...	Variable Type	Comments
tag-namespace	target.tag-namespace.id	Entity (OCID)	This variable is supported only in statements granting permissions for the tag-namespaces resource-type. For an example, see <a href="#">Required Permissions for Working with Defined Tags</a> . Not available to use with <a href="#">CreateTagNamespace</a> .
	target.tag-namespace.name	String	

### Details for Verbs + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for compartments covers no extra permissions or API operations compared to the `inspect` verb. The `use` verb includes the same ones as the `read` verb, plus the `COMPARTMENT_UPDATE` permission and `UpdateCompartment` API operation. The `manage` verb includes the same permissions and API operations as the `use` verb, plus the `COMPARTMENT_CREATE` permission and two API operations: `CreateCompartment` and `DeleteCompartment`.

## authentication-policies

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
AUTHENTICATION_POLICY_INSPECT	GetAuthenticationPolicy	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + AUTHENTICATION_POLICY_UPDATE	USE + UpdateAuthenticationPolicy	none

## compartments

To move a compartment (that is, use the `MoveCompartment` operation) you must belong to a group that has `manage all-resources` permissions on the lowest shared parent compartment of the current compartment and the destination compartment.

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
COMPARTMENT_INSPECT	ListCompartments GetCompartment ListAvailabilityDomains ListFaultDomains	none

## CHAPTER 18 IAM

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*none*

### USE

**Permissions**

*READ +*

COMPARTMENT\_UPDATE

**APIs Fully Covered**

*READ +*

UpdateCompartment

GetWorkRequest

**APIs Partially Covered**

*none*

### MANAGE

**Permissions**

*USE +*

COMPARTMENT\_CREATE

COMPARTMENT\_DELETE

**APIs Fully Covered**

*USE +*

CreateCompartment

DeleteCompartment

**APIs Partially Covered**

*none*

## USERS

### INSPECT

**Permissions**

USER\_INSPECT

**APIs Fully Covered**

ListUsers

GetUser

**APIs Partially Covered**

GetUserGroupMembership (also need

inspect groups)

### READ

**Permissions**

*INSPECT +*

USER\_READ

**APIs Fully Covered**

*INSPECT +*

ListApiKeys

ListSwiftPasswords

ListAuthTokens

ListCustomerSecretKeys

ListMfaTotpDevices

**APIs Partially Covered**

*no extra*

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>READ +</i>	<i>READ +</i>
USER_UPDATE	UpdateUser	AddUserToGroup (also need use groups) RemoveUserFromGroup (also need use groups)

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>no extra</i>
USER_CREATE	CreateUser	
USER_DELETE	DeleteUser	
USER_UNBLOCK	UpdateUserState	
USER_APIKEY_ADD	UploadApiKey	
USER_APIKEY_REMOVE	DeleteApiKey	
USER_UIPASS_SET	CreateOrResetUIPassword	
USER_UIPASS_RESET	UpdateSwiftPassword	
USER_SWIFTPASS_SET	CreateSwiftPassword	
USER_SWIFTPASS_RESET	DeleteSwiftPassword	
USER_SWIFTPASS_REMOVE	UpdateAuthToken	
USER_AUTHTOKEN_SET	CreateAuthToken	
USER_AUTHTOKEN_RESET	DeleteAuthToken	
USER_AUTHTOKEN_REMOVE	CreateSecretKey	
USER_SECRETKEY_ADD	UpdateCustomerSecretKey	
USER_SECRETKEY_UPDATE	DeleteCustomerSecretKey	
USER_SECRETKEY_REMOVE	CreateMfaTotpDevice	
USER_TOTPDEVICE_ADD	ActivateMfaTotpDevice	
USER_TOTPDEVICE_REMOVE	DeleteMfaTotpDevice	
USER_TOTPDEVICE_UPDATE		

## groups

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
GROUP_INSPECT	ListGroups GetGroup	GetUserGroupMembership (also need inspect users) ListIdpGroupMappings, GetIdpGroupMapping (both also need inspect identity-providers)

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + GROUP_UPDATE	READ + UpdateGroup	READ + AddUserToGroup (also need use users) RemoveUserFromGroup (also need use users) AddIdpGroupMapping, DeleteIdpGroupMapping (both also need manage identity-providers)

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + GROUP_CREATE GROUP_DELETE	USE + CreateGroup DeleteGroup	no extra

## dynamic-groups

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
DYNAMIC_GROUP_INSPECT	ListDynamicGroups GetDynamicGroup	No extra

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	no extra

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + DYNAMIC_GROUP_UPDATE	READ + UpdateDynamicGroup	No extra

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + DYNAMIC_GROUP_CREATE DYNAMIC_GROUP_DELETE	USE + CreateDynamicGroup DeleteDynamicGroup	no extra

## policies

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
POLICY_READ	ListPolicies GetPolicy	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

## CHAPTER 18 IAM

---

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**Note:** The ability to update policies is available only with `manage_policies`.

**APIs Partially Covered**

*none*

### MANAGE

**Permissions**

*USE +*

POLICY\_UPDATE

POLICY\_CREATE

POLICY\_DELETE

**APIs Fully Covered**

*USE +*

UpdatePolicy

CreatePolicy

DeletePolicy

**APIs Partially Covered**

*none*

## identity-providers

### INSPECT

**Permissions**

IDENTITY\_PROVIDER\_INSPECT

**APIs Fully Covered**

ListIdentityProviders

GetIdentityProvider

**APIs Partially Covered**

ListIdpGroupMappings,

GetIdpGroupMapping (both also need

`inspect_groups`)

### READ

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

### USE

**Permissions**

*no extra*

**APIs Fully Covered**

*no extra*

**APIs Partially Covered**

*no extra*

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>USE +</i>
IDENTITY_PROVIDER_UPDATE	UpdateIdentityProvider	AddIdpGroupMapping,
IDENTITY_PROVIDER_CREATE	CreateIdentityProvider	DeleteIdpGroupMapping (both also need use
IDENTITY_PROVIDER_DELETE	DeleteIdentityProvider	groups)

### tenancies

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
TENANCY_INSPECT	ListRegionSubscriptions GetTenancy ListRegions	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>no extra</i>	<i>none</i>
TENANCY_UPDATE		

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>none</i>
TENANCY_UPDATE	CreateRegionSubscription	

## tag-namespaces

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
TAG_NAMESPACE_INSPECT	ListTagNamespaces GetTagNamespace ListTags ListCostTrackingTags GetTag GetTaggingWorkRequest ListTaggingWorkRequest ListTaggingWorkRequestErrors ListTaggingWorkRequestLogs	<i>none</i>

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + TAG_NAMESPACE_USE	<i>no extra</i>	<i>none</i>
<p><b>Note:</b> To apply, update, or remove defined tags for a resource, a user must be granted permissions on the resource and permissions to use the tag namespace.</p>		

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>none</i>
TAG_NAMESPACE_UPDATE	UpdateTagNamespace	
TAG_NAMESPACE_CREATE	CreateTag	
TAG_NAMESPACE_MOVE	UpdateTag	
TAG_NAMESPACE_DELETE	CreateTagNamespace	
	ChangeTagNamespaceCompartment	
	DeleteTagNamespace	
	DeleteTag	

### tagdefinitions

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
TAG_NAMESPACE_INSPECT	<i>no extra</i>	<i>none</i>

**Note:** See tag-namespaces.

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>no extra</i>	<i>none</i>
TAG_NAMESPACE_UPDATE	<b>Note:</b> See tag-namespaces.	

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>no extra</i>	<i>none</i>
TAG_NAMESPACE_CREATE	<b>Note:</b> See tag-namespaces.	
TAG_NAMESPACE_MOVE		
TAG_NAMESPACE_DELETE		

## tag-defaults

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
TAG_DEFAULT_INSPECT	ListTagDefaults GetTagDefault	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT +	USE +	none
TAG_DEFAULT_CREATE	CreateTagDefault	
TAG_DEFAULT_UPDATE	UpdateTagDefault	
TAG_DEFAULT_DELETE	DeleteTagDefault	

**Permissions Required for Each API Operation**

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListRegions	TENANCY_INSPECT
ListRegionSubscriptions	TENANCY_INSPECT

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
CreateRegionSubscription	TENANCY_UPDATE
GetTenancy	TENANCY_INSPECT
GetAuthenticationPolicy	AUTHENTICATION_POLICY_INSPECT
UpdateAuthenticationPolicy	AUTHENTICATION_POLICY_UPDATE
ListAvailabilityDomains	COMPARTMENT_INSPECT
ListFaultDomains	COMPARTMENT_INSPECT
ListCompartments	COMPARTMENT_INSPECT
GetCompartment	COMPARTMENT_INSPECT
UpdateCompartment	COMPARTMENT_UPDATE
CreateCompartment	COMPARTMENT_CREATE
DeleteCompartment	COMPARTMENT_DELETE
MoveCompartment	There is not a single permission associated with the <code>MoveCompartment</code> operation. This operation requires <code>manage all-resources</code> permissions on the lowest shared parent compartment of the current compartment and the destination compartment.
GetWorkRequest	COMPARTMENT_READ
ListUsers	USER_INSPECT
GetUser	USER_INSPECT
UpdateUser	USER_UPDATE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
UpdateUserState	USER_UPDATE and USER_UNBLOCK
CreateUser	USER_CREATE
DeleteUser	USER_DELETE
CreateOrResetUIPassword	USER_UPDATE and USER_UIPASS_RESET
ListApiKeys	USER_READ
UploadApiKey	USER_UPDATE and USER_APIKEY_ADD
DeleteApiKey	USER_UPDATE and USER_APIKEY_REMOVE
ListAuthTokens	USER_READ
UpdateAuthToken	USER_UPDATE and USER_AUTHTOKEN_RESET
CreateAuthToken	USER_UPDATE and USER_AUTHTOKEN_SET
DeleteAuthToken	USER_UPDATE and USER_AUTHTOKEN_REMOVE
ListSwiftPasswords	USER_READ
UpdateSwiftPassword	USER_UPDATE and USER_SWIFTPASS_RESET
CreateSwiftPassword	USER_UPDATE and USER_SWIFTPASS_SET
DeleteSwiftPassword	USER_UPDATE and USER_SWIFTPASS_REMOVE
ListCustomerSecretKeys	USER_READ
CreateSecretKey	USER_UPDATE and USER_SECRETKEY_ADD
UpdateCustomerSecretKey	USER_UPDATE and USER_SECRETKEY_UPDATE
DeleteCustomerSecretKey	USER_UPDATE and USER_SECRETKEY_REMOVE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ListUserGroupMemberships	GROUP_INSPECT and USER_INSPECT
GetUserGroupMembership	USER_INSPECT and GROUP_INSPECT
AddUserToGroup	GROUP_UPDATE and USER_UPDATE
RemoveUserFromGroup	GROUP_UPDATE and USER_UPDATE
ListGroups	GROUP_INSPECT
GetGroup	GROUP_INSPECT
UpdateGroup	GROUP_UPDATE
CreateGroup	GROUP_CREATE
DeleteGroup	GROUP_DELETE
ListDynamicGroups	DYNAMIC_GROUP_INSPECT
GetDynamicGroup	DYNAMIC_GROUP_INSPECT
UpdateDynamicGroup	DYNAMIC_GROUP_UPDATE
CreateDynamicGroup	DYNAMIC_GROUP_CREATE
DeleteDynamicGroup	DYNAMIC_GROUP_DELETE
ListPolicies	POLICY_READ
GetPolicy	POLICY_READ
UpdatePolicy	POLICY_UPDATE
CreatePolicy	POLICY_CREATE
DeletePolicy	POLICY_DELETE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ListIdentityProviders	IDENTITY_PROVIDER_INSPECT
GetIdentityProvider	IDENTITY_PROVIDER_INSPECT
UpdateIdentityProvider	IDENTITY_PROVIDER_UPDATE
CreateIdentityProvider	IDENTITY_PROVIDER_CREATE
DeleteIdentityProvider	IDENTITY_PROVIDER_DELETE
ListIdpGroupMappings	IDENTITY_PROVIDER_INSPECT and GROUP_INSPECT
GetIdpGroupMapping	IDENTITY_PROVIDER_INSPECT and GROUP_INSPECT
AddIdpGroupMapping	IDENTITY_PROVIDER_UPDATE and GROUP_UPDATE
DeleteIdpGroupMapping	IDENTITY_PROVIDER_UPDATE and GROUP_UPDATE
ListTagNamespaces	TAG_NAMESPACE_INSPECT
ListTaggingWorkRequest	TAG_NAMESPACE_INSPECT
ListTaggingWorkRequestErrors	TAG_NAMESPACE_INSPECT
ListTaggingWorkRequestLogs	TAG_NAMESPACE_INSPECT
GetTaggingWorkRequest	TAG_NAMESPACE_INSPECT
GetTagNamespace	TAG_NAMESPACE_INSPECT
CreateTagNamespace	TAG_NAMESPACE_CREATE
UpdateTagNamespace	TAG_NAMESPACE_UPDATE
ChangeTagNamespaceCompartment	TAG_NAMESPACE_MOVE
DeleteTagNamespace	TAG_NAMESPACE_DELETE

API Operation	Permissions Required to Use the Operation
ListTags	TAG_NAMESPACE_INSPECT
ListCostTrackingTags	TAG_NAMESPACE_INSPECT
GetTag	TAG_NAMESPACE_INSPECT
CreateTag	TAG_NAMESPACE_UPDATE
UpdateTag	TAG_NAMESPACE_UPDATE
DeleteTag	TAG_NAMESPACE_DELETE
ListTagDefaults	TAG_DEFAULT_INSPECT
GetTagDefault	TAG_DEFAULT_INSPECT
CreateTagDefault	TAG_DEFAULT_MANAGE
UpdateTagDefault	TAG_DEFAULT_MANAGE
DeleteTagDefault	TAG_DEFAULT_MANAGE

## Details for the Key Management Service

This topic covers details for writing policies to control access to the Key Management service.

### Resource-Types

vaults

keys

key-delegate

## Supported Variables

Key Management supports all the general variables, plus the ones listed here. For more information about general variables supported by Oracle Cloud Infrastructure services, see [General Variables for All Requests](#).

Variable	Variable Type	Comments
<code>target.key.id</code>	Entity (OCID)	Use this variable to control access to specific keys by OCID.
<code>target.vault.id</code>	Entity (OCID)	Use this variable to control access to specific vaults by OCID.
<code>request.includePlainTextKey</code>	String	Use this variable to control whether to return the plaintext key in addition to the encrypted key in response to a request to generate a data encryption key.
<code>request.kms-key.id</code>	String	Use this variable to control whether block volumes or buckets can be created without a Key Management master encryption key.
<code>target.boot-volume.kms-key.id</code>	String	Use this variable to control whether Compute instances can be launched with boot volumes that were created without a Key Management master encryption key.

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

## CHAPTER 18 IAM

For example, the `use` verb for the `keys` resource-type includes the same permissions and API operations as the `read` verb, plus the `KEY_ENCRYPT` and `KEY_DECRYPT` permissions and a number of API operations (`Encrypt`, `Decrypt`, and `GenerateDataEncryptionKey`). The `manage` verb allows even more permissions and API operations when compared to the `use` verb.

### vaults

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
VAULT_INSPECT	ListVaults	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + VAULT_READ	INSPECT +  GetVault	none

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
READ + VAULT_CREATE_KEY	READ +  CreateKey	none

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + VAULT_CREATE VAULT_UPDATE VAULT_DELETE VAULT_MOVE	USE +  CreateVault UpdateVault ScheduleVaultDeletion CancelVaultDeletion ChangeVaultCompartment	none

## keys

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
KEY_INSPECT	ListKeys ListKeyVersions	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT + KEY_READ	INSPECT + GetKey GetKeyVersion	none

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
READ + KEY_ENCRYPT KEY_DECRYPT	READ + Encrypt GenerateDataEncryptionKey Decrypt	none

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE + KEY_CREATE KEY_UPDATE KEY_ROTATE KEY_DELETE KEY_MOVE	USE + CreateKey UpdateKey DisableKey EnableKey CreateKeyVersion ScheduleKeyDeletion CancelKeyDeletion ChangeKeyCompartment	none

## key-delegate

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
KEY_ASSOCIATE	Encrypt	<i>none</i>
KEY_DISASSOCIATE	GenerateDataEncryptionKey Decrypt	

**Permissions Required for Each API Operation**

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListVaults	VAULT_INSPECT
GetVault	VAULT_READ
CreateVault	VAULT_CREATE
UpdateVault	VAULT_UPDATE
ScheduleVaultDeletion	VAULT_DELETE
CancelVaultDeletion	VAULT_DELETE
ChangeVaultCompartment	VAULT_MOVE
ListKeys	KEY_INSPECT
ListKeyVersions	KEY_INSPECT
GetKey	KEY_READ
CreateKey	KEY_CREATE and VAULT_CREATE_KEY

API Operation	Permissions Required to Use the Operation
EnableKey	KEY_UPDATE
DisableKey	KEY_UPDATE
UpdateKey	KEY_UPDATE
ScheduleKeyDeletion	KEY_DELETE
CancelKeyDeletion	KEY_DELETE
ChangeKeyCompartment	KEY_MOVE
GetKeyVersion	KEY_READ
CreateKeyVersion	KEY_ROTATE
GenerateDataEncryptionKey	KEY_ENCRYPT
Encrypt	KEY_ENCRYPT
Decrypt	KEY_DECRYPT

## Details for Load Balancing

This topic covers details for writing policies to control access to the Load Balancing service.

### Resource-Types

`load-balancers`

### Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for load-balancers includes the same permissions and API operations as the `inspect` verb, plus the `LOAD_BALANCER_READ` permission and a number of API operations (e.g., `GetLoadBalancer`, `ListWorkRequests`, etc.). The `use` verb covers still another permission and set of API operations compared to `read`. And `manage` covers two more permissions and operations compared to `use`.

### LOAD-BALANCERS

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>LOAD_BALANCER_INSPECT</code>	<code>ListLoadBalancers</code> <code>ListShapes</code> <code>ListPolicies</code> <code>ListProtocols</code>	<i>none</i>

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<code>INSPECT</code> +	<code>INSPECT</code> +	<i>none</i>
<code>LOAD_BALANCER_READ</code>	<code>GetLoadBalancer</code> <code>ListWorkRequests</code> <code>GetWorkRequest</code> <code>ListBackendSets</code> <code>GetBackendSet</code> <code>ListBackends</code> <code>GetBackend</code> <code>GetHealthChecker</code> <code>ListCertificates</code>	

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>READ +</i>	<i>none</i>
LOAD_BALANCER_UPDATE	UpdateLoadBalancer	
LOAD_BALANCER_MOVE	ChangeLoadBalancerCompartment	
	UpdateBackendSet	
	CreateBackendSet	
	DeleteBackendSet	
	UpdateBackend	
	CreateBackend	
	DeleteBackend	
	UpdateHealthChecker	
	CreateCertificate	
	DeleteCertificate	
	UpdateListener	
	CreateListener	
	DeleteListener	

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>none</i>
LOAD_BALANCER_CREATE	CreateLoadBalancer	
LOAD_BALANCER_DELETE	DeleteLoadBalancer	

**Permissions Required for Each API Operation**

The following table lists the API operations in a logical order, grouped by resource type.



**Tip**

If a group uses the Console to manage load balancers, permissions to use the associated networking resources are required. See the [load balancing policy examples](#) for further guidance.

## CHAPTER 18 IAM

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListLoadBalancers	LOAD_BALANCER_INSPECT and
GetLoadBalancer	LOAD_BALANCER_READ
ChangeLoadBalancerCompartment	LOAD_BALANCER_MOVE
UpdateLoadBalancer	LOAD_BALANCER_UPDATE
CreateLoadBalancer	LOAD_BALANCER_CREATE
DeleteLoadBalancer	LOAD_BALANCER_DELETE
ListShapes	LOAD_BALANCER_INSPECT
ListWorkRequests	LOAD_BALANCER_READ
GetWorkRequest	LOAD_BALANCER_READ
ListBackendSets	LOAD_BALANCER_READ
GetBackendSet	LOAD_BALANCER_READ
UpdateBackendSet	LOAD_BALANCER_UPDATE
CreateBackendSet	LOAD_BALANCER_UPDATE
DeleteBackendSet	LOAD_BALANCER_UPDATE
ListBackends	LOAD_BALANCER_READ
GetBackend	LOAD_BALANCER_READ
UpdateBackend	LOAD_BALANCER_UPDATE
CreateBackend	LOAD_BALANCER_UPDATE

API Operation	Permissions Required to Use the Operation
DeleteBackend	LOAD_BALANCER_UPDATE
GetHealthChecker	LOAD_BALANCER_READ
UpdateHealthChecker	LOAD_BALANCER_UPDATE
ListCertificates	LOAD_BALANCER_READ
CreateCertificate	LOAD_BALANCER_UPDATE
DeleteCertificate	LOAD_BALANCER_UPDATE
UpdateListener	LOAD_BALANCER_UPDATE
CreateListener	LOAD_BALANCER_UPDATE
DeleteListener	LOAD_BALANCER_UPDATE
ListPolicies	LOAD_BALANCER_INSPECT
ListProtocols	LOAD_BALANCER_INSPECT

## Details for Monitoring

This topic covers details for writing policies to control access to the Monitoring service.

### Resource-Types

alarms

metrics

### Supported Variables

Monitoring supports all the general variables (see [General Variables for All Requests](#)), plus the one listed here:

Operations for This Resource-Type...	Can Use This Variable	Variable Type	Comments
metrics	target.metrics.namespace	String	<p>Use this variable to control access to specific resource types. Surround the namespace value with single quotes. For example, to control access to metrics for Compute instances, use the following phrase: <code>where target.metrics.namespace='oci_computeagent'</code></p> <p>For an example policy, see <a href="#">Restrict user access to a specific metric namespace</a>. For valid namespace values, see <a href="#">Supported Services</a>.</p>

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

## alarms

### INSPECT

#### Permissions

ALARM\_INSPECT

#### APIs Fully Covered

ListAlarms

ListAlarmsStatus

#### APIs Partially Covered

none

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + ALARM_READ	GetAlarmHistory	GetAlarm (also need METRIC_READ for the metric compartment and metric namespace)

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + <i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + ALARM_CREATE ALARM_UPDATE ALARM_DELETE ALARM_MOVE	ChangeAlarmCompartment DeleteAlarm RemoveAlarmSuppression	CreateAlarm (also need METRIC_READ for the metric compartment and metric namespace) UpdateAlarm (also need METRIC_READ for the metric compartment and metric namespace)

## metrics

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
METRIC_INSPECT	ListMetrics	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + METRIC_READ	SummarizeMetricsData	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + METRIC_WRITE	PostMetricData	<i>none</i>

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE + <i>no extra</i>	<i>no extra</i>	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListMetrics	METRIC_INSPECT or METRIC_READ
SummarizeMetricsData	METRIC_READ
PostMetricData	METRIC_WRITE
ListAlarms	ALARM_INSPECT
ListAlarmsStatus	ALARM_INSPECT
GetAlarm	ALARM_READ and METRIC_READ
GetAlarmHistory	ALARM_READ
CreateAlarm	ALARM_CREATE and METRIC_READ
ChangeAlarmCompartment	ALARM_MOVE
UpdateAlarm	ALARM_UPDATE and METRIC_READ
RemoveAlarmSuppression	ALARM_UPDATE
DeleteAlarm	ALARM_DELETE

## Details for the Notifications Service

This topic covers details for writing policies to control access to the Notifications service.

### Aggregate Resource-Type

`ons-family`

### Individual Resource-Types

`ons-topics`

`ons-subscriptions`

### Supported Variables

Only the general variables are supported (see [General Variables for All Requests](#)).

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

#### ons-topics

##### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>ONS_TOPIC_INSPECT</code>	<code>ListTopics</code>	<i>none</i>

##### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<code>INSPECT +</code>	<code>GetTopic</code>	<i>none</i>
<code>ONS_TOPIC_READ</code>		

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	CreateSubscription	<i>none</i>
ONS_TOPIC_PUBLISH	UpdateSubscription	
ONS_TOPIC_SUBSCRIBE	DeleteSubscription	
	GetSubscription	
	ResendSubscriptionConfirmation	
	PublishMessage	

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> +	CreateTopic	<i>none</i>
ONS_TOPIC_CREATE	ChangeTopicCompartment	
ONS_TOPIC_MOVE	UpdateTopic	
ONS_TOPIC_UPDATE	DeleteTopic	
ONS_TOPIC_DELETE		

## ons-subscriptions

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
ONS_SUBSCRIPTION_INSPECT	ListSubscriptions	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> +	<i>no extra</i>	<i>none</i>
<i>no extra</i>		

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	<i>no extra</i>	<i>none</i>
<i>no extra</i>		

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	ChangeSubscriptionCompartment	<i>none</i>
ONS_SUBSCRIPTION_MOVE	CreateSubscription	
ONS_TOPIC_SUBSCRIBE	UpdateSubscription	
	DeleteSubscription	
	GetSubscription	
	ResendSubscriptionConfirmation	

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListTopics	ONS_TOPIC_INSPECT
GetTopic	ONS_TOPIC_READ
CreateTopic	ONS_TOPIC_CREATE
ChangeTopicCompartment	ONS_TOPIC_MOVE
UpdateTopic	ONS_TOPIC_UPDATE
DeleteTopic	ONS_TOPIC_DELETE
ListSubscriptions	ONS_SUBSCRIPTION_INSPECT
CreateSubscription	ONS_TOPIC_SUBSCRIBE
ChangeSubscriptionCompartment	ONS_SUBSCRIPTION_MOVE
UpdateSubscription	ONS_TOPIC_SUBSCRIBE

API Operation	Permissions Required to Use the Operation
DeleteSubscription	ONS_TOPIC_SUBSCRIBE
GetSubscription	ONS_TOPIC_SUBSCRIBE
GetConfirmSubscription	(no permissions required; available to anyone)
ResendSubscriptionConfirmation	ONS_TOPIC_SUBSCRIBE
GetUnsubscription	(no permissions required; available to anyone)
PublishMessage	ONS_TOPIC_PUBLISH

## Details for Object Storage, Archive Storage, and Data Transfer

This topic covers details for writing policies to control access to Archive Storage, Object Storage, and Data Transfer.



### Tip

The object lifecycle policies feature requires that you grant permissions to the Object Storage service to archive and delete objects on your behalf. See [Using Object Lifecycle Policies](#) for more information.

## Resource-Types

### INDIVIDUAL RESOURCE-TYPES

objectstorage-namespaces

buckets

objects

### **AGGREGATE RESOURCE-TYPE**

`object-family`

A policy that uses `<verb> object-family` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `object-family`.

### **ADDITIONAL INDIVIDUAL RESOURCE-TYPE FOR DATA TRANSFER**

`data-transfer-jobs`

### **Supported Variables**

Object Storage supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here:

Operations for This Resource-Type...	Can Use This Variable	Variable Type	Comments
buckets and objects	<code>target.bucket.name</code>	String	Use this variable to control access to a specific bucket. For an example policy, see <a href="#">Let users write objects to Object Storage buckets</a> . <b>Important:</b> Condition matching is case insensitive. If you have a bucket named "BucketA" and a bucket named "bucketA", the condition where <code>target.bucket.name="BucketA"</code> applies to both. To avoid potential issues with resource names in policy, give your resources distinct names.
buckets and objects	<code>request.vcn.id</code>	String	If you're using a <a href="#">service gateway</a> with your VCN, you can use this variable to allow access to a bucket only from a specific virtual cloud network (VCN). For an example policy, see <a href="#">Task 4: (Optional) Update IAM Policies to Restrict Object Storage Bucket Access</a> .

Operations for This Resource-Type...	Can Use This Variable	Variable Type	Comments
buckets and objects	request.ipv4.ipaddress	String	If you're using a <a href="#">service gateway</a> with your VCN, you can use this variable to allow access to a bucket only from a specific CIDR range. For an example policy, see <a href="#">Task 4: (Optional) Update IAM Policies to Restrict Object Storage Bucket Access</a> .

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

#### FOR OBJECT-FAMILY RESOURCE TYPES

#### objectstorage-namespaces

##### READ

###### Permissions

None

###### Permissions

OBJECTSTORAGE\_NAMESPACE\_READ

###### APIs Fully Covered

GetNamespace

###### APIs Fully Covered

GetNamespace with optional `compartmentId` parameter  
GetNamespaceMetadata

###### APIs Partially Covered

*none*

## CHAPTER 18 IAM

---

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + OBJECTSTORAGE_NAMESPACE_UPDATE	<i>READ</i> + UpdateNamespaceMetadata	<i>none</i>

### buckets

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
BUCKET_INSPECT	HeadBucket ListBuckets	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + BUCKET_READ	<i>INSPECT</i> + GetBucket ListMultipartUploads GetObjectLifecyclePolicy	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + BUCKET_UPDATE	<i>READ</i> + UpdateBucket DeleteObjectLifecyclePolicy ReencryptBucket	PutObjectLifecyclePolicy

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	<i>none</i>
BUCKET_CREATE	CreateBucket	
BUCKET_DELETE	DeleteBucket	
PAR_MANAGE	CreatePar	
	GetPar	
	ListPar	
	DeletePar	

objects

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
OBJECT_INSPECT	HeadObject	<i>none</i>
	ListObjects	
	ListMultipartUploadParts	

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i>	<i>INSPECT +</i>	<i>none</i>
OBJECT_READ	GetObject	

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>READ +</i>	<i>READ +</i>
OBJECT_OVERWRITE	PutObject	CreateMultipartUpload, UploadPart, CommitMultipartUpload (these operations also need manage objects)

## CHAPTER 18 IAM

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i>	<i>USE +</i>	PutObjectLifecyclePolicy (also needs
OBJECT_CREATE	CreateObject	manage objects)
OBJECT_DELETE	RenameObject	
OBJECT_RESTORE	RestoreObject	
	DeleteObject	
	CreateMultipartUpload	
	UploadPart	
	CommitMultipartUpload	
	AbortMultipartUpload	

### data-transfer-jobs

Policies for data transfer jobs also require either manage objects or manage objects and manage buckets. See [Creating the Required IAM Users, Groups, and Policies](#) for details.

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
TRANSFER_JOB_INSPECT	<i>no customer-facing API</i>	<i>no customer-facing API</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i>	<i>INSPECT +</i>	<i>no customer-facing API</i>
TRANSFER_JOB_READ	<i>no customer-facing API</i>	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ +</i>	<i>READ +</i>	<i>READ +</i>
TRANSFER_JOB_UPDATE	<i>no customer-facing API</i>	<i>no customer-facing API</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	USE +	USE +
TRANSFER_JOB_CREATE	<i>no customer-facing API</i>	<i>no customer-facing API</i>
TRANSFER_JOB_DELETE		

**Permissions Required for Each API Operation**

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
GetNamespace	API requires no permissions and returns the caller's namespace. Use the API to validate your credentials.  OBJECTSTORAGE_NAMESPACE_READ permission is required if you include the optional <code>compartmentId</code> parameter. Use the <code>compartmentId</code> parameter to determine the namespace for a third-party tenancy.
GetNamespaceMetadata	OBJECTSTORAGE_NAMESPACE_READ
UpdateNamespaceMetadata	OBJECTSTORAGE_NAMESPACE_UPDATE
CreateBucket	BUCKET_CREATE
UpdateBucket	BUCKET_UPDATE
GetBucket	BUCKET_READ
HeadBucket	BUCKET_INSPECT
ListBuckets	BUCKET_INSPECT
DeleteBucket	BUCKET_DELETE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ReencryptBucket	BUCKET_UPDATE
PutObject	The permission required depends on whether or not the object already exists in the bucket: <ul style="list-style-type: none"> <li>• OBJECT_CREATE is required when an object with that name does not already exist in the bucket.</li> <li>• OBJECT_OVERWRITE is required when an object with that name already exists in the bucket.</li> </ul>
RenameObject	OBJECT_CREATE and OBJECT_OVERWRITE
GetObject	OBJECT_READ
HeadObject	OBJECT_READ or OBJECT_INSPECT
DeleteObject	OBJECT_DELETE
ListObjects	OBJECT_INSPECT
RestoreObjects	OBJECT_RESTORE
CreateMultipartUpload	OBJECT_CREATE and OBJECT_OVERWRITE
UploadPart	OBJECT_CREATE and OBJECT_OVERWRITE
CommitMultipartUpload	OBJECT_CREATE and OBJECT_OVERWRITE
ListMultipartUploadParts	OBJECT_INSPECT
ListMultipartUploads	BUCKET_READ
AbortMultipartUpload	OBJECT_DELETE
CreatePar	PAR_MANAGE

## CHAPTER 18 IAM

---

API Operation	Permissions Required to Use the Operation
GetPar	PAR_MANAGE
ListPars	PAR_MANAGE
DeletePar	PAR_MANAGE
PutObjectLifecyclePolicy	BUCKET_UPDATE, OBJECT_CREATE, and OBJECT_DELETE
GetObjectLifecyclePolicy	BUCKET_READ
DeleteObjectLifecyclePolicy	BUCKET_UPDATE
CreateCopyRequest	OBJECT_READ, OBJECT_CREATE, OBJECT_OVERWRITE, and OBJECT_INSPECT
GetWorkRequest	OBJECT_READ
ListWorkRequests	OBJECT_INSPECT
CancelWorkRequest	OBJECT_DELETE

### Details for the Quotas Service

This topic covers details for writing policies to control access to the Quotas service.

#### Resource-Types

quota

#### Supported Variables

The Quotas service supports all the general variables (see [General Variables for All Requests](#)) plus the following:

Variable	Variable Type	Source
target.quota.id	Entity (OCID)	Request
target.quota.name	String	Request/Stored

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

## quotas

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
QUOTA_INSPECT	<code>listQuotas</code>	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
QUOTA_READ	<code>getQuota</code>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	createQuota	none
QUOTA_CREATE	deleteQuota	
QUOTA_DELETE	updateQuota	
QUOTA_UPDATE		

**QUOTAS API OPERATIONS**

API Operation	Permissions Required to Use the Operation
listQuotas	QUOTA_INSPECT
createQuota	QUOTA_CREATE
getQuota	QUOTA_READ
deleteQuota	QUOTA_DELETE
updateQuota	QUOTA_UPDATE

Details for Registry

This topic covers details for writing policies to control access to the Registry.

**Resource-Types**

- repos

**Supported Variables**

Oracle Cloud Infrastructure Registry supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

The `repos` resource-type can use the following variables:

Variable	Variable Type	Comments
<code>target.repo.name</code>	String	Use this variable to control access to specific repositories. For an example policy, see <a href="#">Policies to Control Repository Access</a> .

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `repos` resource-type includes the same permissions and API operations as the `inspect` verb, plus the `REPOSITORY_READ` permission and a number of API operations (e.g., `ReadDockerRepositoryMetadata`, etc.). The `use` verb covers still another permission and API operation compared to `read`. Lastly, `manage` covers more permissions and operations compared to `use`.

Note the Registry API is not currently available.

## repos

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<code>REPOSITORY_INSPECT</code>	<code>ListDockerRepositories</code> <code>ListDockerRepositoryManifests</code>	<i>none</i>

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT +</i> REPOSITORY_READ	<i>INSPECT +</i> ReadDockerRepositoryMetadata ReadDockerRepositoryManifest PullDockerLayer	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE +</i> REPOSITORY_CREATE REPOSITORY_DELETE REPOSITORY_UPDATE REPOSITORY_MANAGE	<i>USE +</i> CreateDockerRepository DeleteDockerRepository UploadDockerImage DeleteDockerImage DeleteDockerLayer UpdateDockerRepositoryMetadata	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

Note the Registry API is not currently available.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListDockerRepositories	REPOSITORY_INSPECT
ListDockerRepositoryManifests	REPOSITORY_INSPECT

API Operation	Permissions Required to Use the Operation
ReadDockerRepositoryMetadata	REPOSITORY_READ
ReadDockerRepositoryManifest	REPOSITORY_READ
CreateDockerRepository	REPOSITORY_CREATE
DeleteDockerRepository	REPOSITORY_DELETE
DeleteDockerRepositoryContents	REPOSITORY_UPDATE
UpdateDockerRepositoryMetadata	REPOSITORY_MANAGE
UploadDockerImage	REPOSITORY_UPDATE + REPOSITORY_CREATE
DeleteDockerImage	REPOSITORY_UPDATE
DeleteDockerLayer	REPOSITORY_UPDATE
PullDockerLayer	REPOSITORY_READ
UploadDockerLayer	REPOSITORY_UPDATE + REPOSITORY_CREATE

## Details for Resource Manager

This topic covers details for writing policies to control access to the Resource Manager service.

### Aggregate Resource-Type

`orm-family`

### Individual Resource-Types

`orm-stacks`

`orm-jobs`

`orm-work-requests`

### Supported Variables

Resource Manager supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

The `orm-jobs` resource type can use the following variables.

Variable	Variable Type	Comments
<code>target.job.operation</code>	String	Use this variable to control access for running specified job types. For example, to limit access to PLAN and APPLY jobs, use the following phrase: <code>where any {target.job.operation = 'PLAN', target.job.operation = 'APPLY'}</code>

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access..

### orm-stacks

#### INSPECT

##### Permissions

`ORM_STACK_INSPECT`

##### APIs Fully Covered

`ListStacks`

`ListTerraformVersions`

##### APIs Partially Covered

*none*

## CHAPTER 18 IAM

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> +	GetStack	<i>none</i>
ORM_STACK_READ	GetStackTfConfig	

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> +	<i>no extra</i>	CreateJob (also need <code>manage_orm-jobs</code> )
ORM_STACK_USE		

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> +	CreateStack	<i>none</i>
ORM_STACK_CREATE	UpdateStack	
ORM_STACK_UPDATE	ChangeStackCompartment	
ORM_STACK_MOVE	DeleteStack	
ORM_STACK_DELETE	ListTerraformVersions	

## orm-jobs

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
ORM_JOB_INSPECT	ListJobs	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> +	GetJob	<i>none</i>
ORM_JOB_READ	GetJobTfState	
	GetJobTfConfig	
	GetJobTfExecutionPlan	
	GetJobLogs	
	GetJobLogsContent	

## CHAPTER 18 IAM

---

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + <i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + ORM_JOB_MANAGE	UpdateJob CancelJob	CreateJob (also need <i>use_orm-stacks</i> )

## orm-work-requests

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
ORM_WORK_REQUEST_INSPECT	ListWorkRequests	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>INSPECT</i> + ORM_WORK_REQUEST_READ	ListWorkRequestErrors ListWorkRequestLogs GetWorkRequest	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>READ</i> + <i>no extra</i>	<i>no extra</i>	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>USE</i> + <i>no extra</i>	<i>no extra</i>	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
ListStacks	ORM_STACK_INSPECT
CreateStack	ORM_STACK_CREATE
GetStack	ORM_STACK_READ
UpdateStack	ORM_STACK_UPDATE
ChangeStackCompartment	ORM_STACK_MOVE
DeleteStack	ORM_STACK_DELETE
GetStackTfConfig	ORM_STACK_READ
ListTerraformVersions	ORM_STACK_INSPECT
ListJobs	ORM_JOB_INSPECT
CreateJob	ORM_JOB_MANAGE and ORM_STACK_USE
GetJob	ORM_JOB_READ
UpdateJob	ORM_JOB_MANAGE
CancelJob	ORM_JOB_MANAGE
GetJobTfState	ORM_JOB_READ
GetJobTfConfig	ORM_JOB_READ
GetJobTfExecutionPlan	ORM_JOB_READ

API Operation	Permissions Required to Use the Operation
GetJobLogs	ORM_JOB_READ
GetJobLogsContent	ORM_JOB_READ
ListWorkRequestErrors	ORM_WORK_REQUEST_READ
ListWorkRequestLogs	ORM_WORK_REQUEST_READ
ListWorkRequests	ORM_WORK_REQUEST_INSPECT
GetWorkRequest	ORM_WORK_REQUEST_READ

## Details for Search

The Search service does not require permissions for its API operations. You do not need to write policies specifically to control access to Search. However, what you can see in search or query results depends on the permissions you have. If a policy exists to give you access to the `inspect` verb for a particular resource type, you have access to the permissions needed to view that resource type and its associated metadata in search results. If a service does not recognize the `inspect` verb or if the resource type's `inspect` verb does not fully cover list operations, permissions to view the service's supported resource types are granted by the `read` verb instead.

For more information about permissions, see the Permissions section of [Advanced Policy Features](#).

## Permissions Required to View Each Resource Type

The following table lists the resource types grouped by service, which are listed in alphabetical order. The Search API operations that can access the metadata for these resource types with these permissions are `GetResourceType`, `ListResourceTypes`, and `SearchResources`.

## CHAPTER 18 IAM

Service	Resource Type	Permissions Required to View in Search Results
Block Volume	volumes	VOLUME_INSPECT
Block Volume	volume-backups	VOLUME_BACKUP_INSPECT
Compute	console-histories	CONSOLE_HISTORY_INSPECT
Compute	instance-images	INSTANCE_IMAGE_READ
Compute	instances	INSTANCE_READ
Database	databases	DATABASE_INSPECT
Database	db-homes	DB_HOME_INSPECT (if you want to filter results using db-homes attributes)
Database	db-systems	DB_SYSTEM_INSPECT
IAM	compartments	COMPARTMENT_INSPECT
IAM	groups	GROUP_INSPECT
IAM	identity-providers	IDENTITY_PROVIDER_INSPECT
IAM	users	USER_INSPECT
Networking	route-tables	ROUTE_TABLE_READ
Networking	security-lists	SECURITY_LIST_READ

Service	Resource Type	Permissions Required to View in Search Results
Networking	subnets	SUBNET_READ
Networking	vcns	VCN_READ
Object Storage	buckets	BUCKET_INSPECT

## Details for the Streaming Service

This topic covers details for writing policies to control access to the Streaming service.

### Resource-Types

`streams`

`stream-pull`

`stream-push`

### Supported Variables

The Streaming service supports all the general variables (see [General Variables for All Requests](#)) plus the following:

The `streams` resource type can use the following variables:

Variable	Variable Type	Source
<code>target.stream.id</code>	Entity (OCID)	Request

### Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a

table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

## streams

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
STREAM_INSPECT	ListStreams	none

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
STREAM_READ	GetStream	none

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
no extra	no extra	none

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
USE +	CreateStream	none
STREAM_CREATE	DeleteStream	
STREAM_DELETE	UpdateStream	
STREAM_UPDATE		

## stream-pull

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
none	none	none

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
none	none	none

## CHAPTER 18 IAM

---

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
STREAM_CONSUME	GetMessages CreateCursor CreateGroupCursor GetGroup UpdateGroup ConsumerHeartbeat ConsumerCommit	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

## stream-push

### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
<i>none</i>	<i>none</i>	<i>none</i>

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
<i>none</i>	<i>none</i>	<i>none</i>

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
STREAM_PRODUCE	PutMessages	<i>none</i>

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
<i>no extra</i>	<i>no extra</i>	<i>none</i>

**STREAMING API OPERATIONS**

<b>API Operation</b>	<b>Permissions Required to Use the Operation</b>
ListStreams	STREAM_INSPECT
CreateStream	STREAM_CREATE
GetStream	STREAM_READ
DeleteStream	STREAM_DELETE
GetMessages	STREAM_CONSUME
PutMessages	STREAM_PRODUCE
UpdateStream	STREAM_UPDATE
CreateCursor	STREAM_CONSUME
CreateGroupCursor	STREAM_CONSUME
GetGroup	STREAM_CONSUME
UpdateGroup	STREAM_CONSUME
ConsumerHeartbeat	STREAM_CONSUME
ConsumerCommit	STREAM_CONSUME

**Details for the WAF Service**

This topic covers details for writing policies to control access to the WAAS service.

**Aggregate Resource-Type**

waas-family

## Individual Resource-Types

```

waas-policy
waas-certificate
waas-work-request
waas-metering

```

### COMMENTS

A policy that uses `<verb> waas` is equivalent to writing one with a separate `<verb> <individual resource-type>` statement for each of the individual resource-types.

See the table in [Details for Verb + Resource-Type Combinations](#) for a detailed breakout of the API operations covered by each verb, for each individual resource-type included in `waas`.

## Supported Variables

The WAF Service supports all the general variables (see [General Variables for All Requests](#)), plus the ones listed here.

Variable	Variable Type	Comments
<code>target.waas-policy.id</code>	Entity (OCID)	Use this variable to control access to specific WAAS policies by OCID.
<code>target.waf-rule-key</code>	String	Use this variable to control access to specific WAF rules by name.

## Details for Verb + Resource-Type Combinations

The following tables show the [permissions](#) and API operations covered by each verb. The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

## CHAPTER 18 IAM

---

For example, the `use` and `manage` verbs for the `waas-policy` resource-type cover no extra permissions or API operations compared to the `read` verb.

### waas-policy

#### INSPECT

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_POLICY_INSPECT	ListWaasPolicies ListWaasOriginRequestCidrs	none

#### READ

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_POLICY_INSPECT	GetWaasPolicy	none
WAAS_POLICY_READ		

#### USE

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_POLICY_INSPECT	UpdateWaasPolicy	none
WAAS_POLICY_READ		
WAAS_POLICY_UPDATE		

#### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_POLICY_INSPECT	CreateWaasPolicy	none
WAAS_POLICY_READ	DeleteWaasPolicy	
WAAS_POLICY_UPDATE	ChangeWaasPolicyCompartment	
WAAS_POLICY_CREATE		
WAAS_POLICY_DELETE		
WAAS_POLICY_MOVE		

## waas-certificate

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_CERTIFICATE_INSPECT	ListCertificates	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_CERTIFICATE_INSPECT	GetCertificate	none
WAAS_CERTIFICATE_READ		

**USE**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_CERTIFICATE_INSPECT	CreateCertificate	none
WAAS_CERTIFICATE_READ		

**MANAGE**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_CERTIFICATE_INSPECT	DeleteCertificate	none
WAAS_CERTIFICATE_READ		
WAAS_CERTIFICATE_CREATE		
WAAS_CERTIFICATE_DELETE		

## waas-work-request

**INSPECT**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_WORK_REQUEST_INSPECT	ListWorkRequests	none

**READ**

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_WORK_REQUEST_INSPECT	GetWorkRequestDetails	none
WAAS_WORK_REQUEST_READ		

## CHAPTER 18 IAM

### USE

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_WORK_REQUEST_INSPECT	<i>no extra</i>	<i>none</i>
WAAS_WORK_REQUEST_READ		

### MANAGE

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_WORK_REQUEST_INSPECT	DeleteWorkRequest	<i>none</i>
WAAS_WORK_REQUEST_READ		
WAAS_WORK_REQUEST_DELETE		

## waas-metering

### READ

Permissions	APIs Fully Covered	APIs Partially Covered
WAAS_METERING_READ	GetWafReport	<i>none</i>

### Permissions Required for Each API Operation

The following table lists the API operations in a logical order, grouped by resource type.

For information about permissions, see [Permissions](#).

API Operation	Permissions Required to Use the Operation
CreateWaasPolicy	WAAS_POLICY_CREATE
ListWaasPolcies	WAAS_POLICY_INSPECT
GetWaasPolicy	WAAS_POLICY_READ
UpdateWaasPolicy	WAAS_POLICY_UPDATE
DeleteWaasPolicy	WAAS_POLICY_DELETE

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
ListReports	WAAS_POLICY_INSPECT
ListWafReports	WAAS_POLICY_INSPECT
GetWafTraffic	WAAS_POLICY_READ
GetWafBlocked	WAAS_POLICY_READ
GetWafRequests	WAAS_POLICY_READ
GetWafSettings	WAAS_POLICY_READ
UpdateWafSettings	WAAS_POLICY_UPDATE
GetAccessRules	WAAS_POLICY_READ
UpdateAccessRules	WAAS_POLICY_UPDATE
GetCaptchas	WAAS_POLICY_READ
GetHumanInteractionChallenge	WAAS_POLICY_READ
UpdateHumanInteractionChallenge	WAAS_POLICY_UPDATE
GetJsChallenge	WAAS_POLICY_READ
UpdateJsChallenge	WAAS_POLICY_UPDATE
UpdateIpRateLimiting	WAAS_POLICY_UPDATE
GetGoodBots	WAAS_POLICY_READ
UpdateGoodBots	WAAS_POLICY_UPDATE
GetWafWhitelists	WAAS_POLICY_READ
GetWafRecommendations	WAAS_POLICY_READ

## CHAPTER 18 IAM

API Operation	Permissions Required to Use the Operation
AcceptWafRecommendations	WAAS_POLICY_UPDATE
ListWafRules	WAAS_POLICY_INSPECT
UpdateWafRuleActions	WAAS_POLICY_UPDATE
GetWafRule	WAAS_POLICY_READ
GetThreatFeeds	WAAS_POLICY_READ
UpdateThreatFeedAction	WAAS_POLICY_UPDATE
ChangeWaasPolicyCompartment	WAAS_POLICY_MOVE
GetAlerts	WAAS_POLICY_READ
ListWorkRequests	WAAS_WORK_REQUEST_READ
ListWaasOriginRequestCidrs	WAAS_POLICY_INSPECT
GetWorkRequestDetails	WAAS_WORK_REQUEST_READ
DeleteWorkRequest	WAAS_WORK_REQUEST_DELETE
CreateCertificate	WAAS_CERTIFICATE_CREATE
ListCertificates	WAAS_CERTIFICATE_INSPECT
GetCertificate	WAAS_CERTIFICATE_READ
DeleteCertificate	WAAS_CERTIFICATE_DELETE
GetWafReport	WAAS_METERING_READ

## User Credentials

There are several types of credentials that you manage with Oracle Cloud Infrastructure Identity and Access Management (IAM):

- **Console password:** For signing in to the Console, the user interface for interacting with Oracle Cloud Infrastructure. Note that federated users can't have Console passwords because they sign in through their identity provider. See [Federating with Identity Providers](#).
- **API signing key (in PEM format):** For sending API requests, which require authentication.
- **Auth token:** An Oracle-generated token that you can use to authenticate with third-party APIs. For example, use an auth token to authenticate with a Swift client when using Recovery Manager (RMAN) to back up an Oracle Database System (DB System) database to Object Storage.
- **Customer Secret Keys:** For using the Amazon S3 Compatibility API with Object Storage. See [Amazon S3 Compatibility API](#).
- **SMTP Credentials:** For using the [Email Delivery service](#).



### Important

API signing keys are different from the SSH keys you use to access a compute instance (see [Security Credentials](#)). For more information about API signing keys, see [Required Keys and OCIDs](#). For more information about instance SSH keys, see [Managing Key Pairs](#).

## User Password

The administrator who creates a new user in IAM also needs to generate a one-time Console password for the user (see [To create or reset another user's Console password](#)). The

administrator needs to securely deliver the password to the user by providing it verbally, printing it out, or sending it through a secure email service.

When the user signs in to the Console the first time, they'll be immediately prompted to change the password. If the user waits more than 7 days to initially sign in and change the password, it will expire and an administrator will need to create a new one-time password for the user.

Once the user successfully signs in to the Console, they can use Oracle Cloud Infrastructure resources according to permissions they've been granted through policies.



### Note

A user automatically has the ability to change their password in the Console. An administrator does not need to create a policy to give a user that ability.

### Changing a Password

If a user wants to change their own password *sometime after* they change their initial one-time password, they can do it in the Console. Remember that a user can automatically change *their own* password; an administrator does not need to create a policy to give the user that ability.

For more information, see [To change your Console password](#).

### If a User Needs Their Console Password Reset

If a user forgets their Console password and also has no access to the API, they can use the Console's **Forgot Password** link to have a temporary password sent to them. This option is available if the user has an email address in their user profile.

If the user does not have an email address in their user profile, then they need to ask an administrator to reset their password for them. All administrators (and anyone else who has permission to the tenancy) can reset Console passwords. The process of resetting the

password generates a new one-time password that the administrator needs to deliver to the user. The user will need to change their password the next time they sign in to the Console.

If you're an administrator who needs to reset a user's Console password, see [To create or reset another user's Console password](#).

### **If a User Is Blocked from Signing In to the Console**

If a user tries 10 times in a row to sign in to the Console unsuccessfully, they will be automatically blocked from further attempts. They'll need to contact an administrator to get unblocked (see [To unblock a user](#)).

## API Signing Keys

A user who needs to make API requests must upload an **RSA public key in PEM format (minimum 2048 bits)** to IAM and sign the API requests with the corresponding private key (see [Required Keys and OCIDs](#)).



### **Important**

A user automatically has the ability to upload and manage *their own* API keys in the Console or API. An administrator does not need to write a policy to give the user that ability. Remember that a user can't use the API to change or delete their own credentials until they themselves upload a key in the Console, or an administrator uploads a key for that user in the Console or the API.

If you have a non-human system that needs to make API requests, an administrator needs to create a user for that system and then upload a public key to the IAM service for the system. There's no need to generate a Console password for the user.

For instructions on uploading an API key, see [To upload an API signing key](#).

### Auth Tokens

Auth tokens are authentication tokens generated by Oracle. You use auth tokens to authenticate with third-party APIs that do not support the Oracle Cloud Infrastructure signature-based authentication, for example, the Swift API. If your service requires an auth token, the service-specific documentation instructs you to generate one and how to use it.

## Federating with Identity Providers

This topic describes identity federation concepts. Oracle Cloud Infrastructure supports federation with [Oracle Identity Cloud Service](#), and Microsoft Active Directory (via Active Directory Federation Services (AD FS)), Microsoft Azure Active Directory, Okta, and other identity providers that supports the Security Assertion Markup Language (SAML) 2.0 protocol.

### Overview

Enterprise companies commonly use an *identity provider (IdP)* to manage user login/passwords and to authenticate users for access to secure websites, services, and resources.

When someone in your company wants to use Oracle Cloud Infrastructure resources in the Console, they must sign in with a user login and password. Your administrators can federate with a supported IdP so that each employee can use an existing login and password and not have to create a new set to use Oracle Cloud Infrastructure resources.

To federate, an administrator goes through a short process to set up a relationship between the IdP and Oracle Cloud Infrastructure (commonly referred to as a *federation trust*). After an administrator sets up that relationship, any person in your company who goes to the Oracle Cloud Infrastructure Console is prompted with a "single sign-on" experience provided by the IdP. The user signs in with the login/password that they've already set up with the IdP. The IdP authenticates the user, and then that user can access Oracle Cloud Infrastructure.

When working with your IdP, your administrator defines groups and assigns each user to one or more groups according to the type of access the user needs. Oracle Cloud Infrastructure also uses the concept of groups (in conjunction with IAM policies) to define the type of access

a user has. As part of setting up the relationship with the IdP, your administrator can map each IdP group to a similarly defined IAM group, so that your company can re-use the IdP group definitions when authorizing user access to Oracle Cloud Infrastructure resources. Here's a screenshot from the mapping process:

The screenshot shows the 'Edit Identity Provider' interface. At the top right is a 'cancel' link. Below the title is an explanatory text: 'Here you'll map groups defined in your Identity Provider to groups defined in Oracle Cloud Infrastructure. Each group can be mapped to one or more groups of the other kind.' The main area is a 'Mapping' dialog box with a red close button. It contains three columns: 'IDENTITY PROVIDER GROUP' with a text input 'IdP\_Group\_Name', 'ORACLE CLOUD INFRASTRUCTURE GROUP' with a dropdown menu 'New Oracle Cloud Infrastructure Group', and 'NEW OCI GROUP' with an empty text input. An arrow points from the first input to the second. Below the dialog is a '+ Add Mapping' button and a 'Submit' button at the bottom left.

For information about the number of federations and group mappings you can have, see [Service Limits](#). There's no limit on the number of federated users.



### Note

Any users who are in more than 50 IdP groups cannot be authenticated to use the Oracle Cloud Infrastructure Console.

### Automated User Provisioning and Synchronization with SCIM

Tenancies federated with Oracle Identity Cloud Service or the [third-party provider Okta](#), can also leverage [SCIM \(System for Cross-domain Identity Management\)](#) to enable provisioning of federated users in Oracle Cloud Infrastructure. Federated users that have been provisioned in Oracle Cloud Infrastructure through this process can have the additional user credentials such as API keys and auth tokens that are managed in the **User Settings** page. This enables federated users to use the SDK and CLI, and other features that require the additional user credentials. For more information, see [User Provisioning for Federated Users](#).

### General Concepts

Here's a list of the basic concepts you need to be familiar with.

#### **IdP**

IdP is short for *identity provider*, which is a service that provides identifying credentials and authentication for users.

Tenancies created after December 18, 2017 are automatically federated with [Oracle Identity Cloud Service](#) as the IdP. Oracle Cloud Infrastructure can be federated with any IdP that supports the Security Assertion Markup Language (SAML) 2.0 protocol.

#### **SERVICE PROVIDER (SP)**

A service (such as an application, website, and so on) that calls upon an IdP to authenticate users. In this case, Oracle Cloud Infrastructure is the SP.

#### **FEDERATION TRUST**

A relationship that an administrator configures between an IdP and SP. You can use the Oracle Cloud Infrastructure Console or API to set up that relationship. Then, the specific IdP is "federated" to that SP. In the Console and API, the process of federating is thought of as *adding an identity provider to the tenancy*.

### **SAML METADATA DOCUMENT**

An IdP-provided XML-based document that provides the required information to an SP to federate with that IdP. Oracle Cloud Infrastructure supports the SAML 2.0 protocol, which is an XML-based standard for sharing required information between the IdP and SP. Depending on which IdP you are federating with, you must either provide the metadata URL (see below) to this document or upload the document to Oracle Cloud Infrastructure.

### **METADATA URL**

An IdP-provided URL that enables an SP to get required information to federate with that IdP. Oracle Cloud Infrastructure supports the SAML 2.0 protocol, which is an XML-based standard for sharing required information between the IdP and SP. The metadata URL points to the SAML metadata document the SP needs.

### **FEDERATED USER**

Someone who signs in to use the Oracle Cloud Infrastructure Console by way of a federated IdP.

### **LOCAL USER**

A non-federated user. In other words, someone who signs in to use the Oracle Cloud Infrastructure Console with a login and password created in Oracle Cloud Infrastructure.

### **GROUP MAPPING**

A mapping between an IdP group and an Oracle Cloud Infrastructure group, used for the purposes of user authorization.

### **SCIM**

[SCIM \(System for Cross-domain Identity Management\)](#) is an IETF standard protocol that enables user provisioning across identity systems. Oracle Cloud Infrastructure hosts a SCIM endpoint for provisioning federated users into Oracle Cloud Infrastructure. Using a SCIM client to provision users in Oracle Cloud Infrastructure enables you to assign credentials to the users in Oracle Cloud Infrastructure.

### **PROVISIONED (OR SYNCHRONIZED) USER**

A user provisioned by the identity provider's SCIM client in Oracle Cloud Infrastructure. These users can be listed in the Oracle Cloud Infrastructure Console and can have all the Oracle Cloud Infrastructure user credentials except for a Console password.

### **ENCRYPT ASSERTION**

Some IdPs support the encryption of the SAML assertion. When enabled, the service provider expects the SAML assertion to be encrypted by the identity provider, using the service provider's encryption key. In this case, the service provider is Oracle Cloud Infrastructure authentication service. If you choose to enable this feature of your IdP, you must also enable the feature when you set up your Federation provider in the IAM service. Note that Microsoft AD FS enables the encryption of the SAML assertion by default. If your IdP is Microsoft AD FS, you must either enable this feature in IAM or disable it for Microsoft AD FS.

## Experience for Federated Users

Federated users can use the Console to access Oracle Cloud Infrastructure (according to IAM policies for the groups the users are in).

They'll be prompted to enter their Oracle Cloud Infrastructure tenant (for example, ABCCorp).

They then see a page with two sets of sign-in instructions: one for federated users and one for non-federated (Oracle Cloud Infrastructure) users. See the following screenshot.

Signing in to cloud tenant:  
**ABCCorp**  
[Change tenant](#)

**Single Sign-On (SSO)**

We have detected that your tenancy has been federated to another Identity Provider.

Select your Identity Provider below and log in.

IDENTITY PROVIDER  
OracleIdentityCloudService ▼

[Continue](#)

**Oracle Cloud Infrastructure**

The login is uncommon for federated accounts. If you have questions, please contact your tenancy administrator.

or

USER NAME  
[Redacted]

PASSWORD  
[Redacted]

[Sign In](#) [Forgot password?](#)

The tenant name is shown on the left. Directly below is the sign-in area for federated users. On the right is the sign-in area for non-federated users.

Federated users choose which identity provider to use for sign-in, and then they're redirected to that identity provider's sign-in experience for authentication. After entering their login and password, they are authenticated by the IdP and redirected back to the Oracle Cloud Infrastructure Console.

The federated users (without SCIM configuration) cannot access the "User Settings" page in the Console. This page is where a user can change or reset their Console password and manage other Oracle Cloud Infrastructure credentials such as [API signing keys](#) and [auth tokens](#).

### Experience for Federated Users with SCIM Configuration

If your IdP has also been configured with a SCIM client, a user signed in through their identity provider can access the User Settings page and have user capabilities such as API keys, auth tokens, and other user credentials. (**Note:** This is currently available for Oracle Identity Cloud Service and Okta federations only.)

### Required IAM Policy

To add and manage identity providers in your tenancy, you must be authorized by an IAM policy. If you're in the [Administrators group](#), then you have the required access.

Here's a more limited policy that restricts access to only the resources related to identity providers and group mappings:

```
Allow group IdPAdmins to manage identity-providers in tenancy
```

```
Allow group IdPAdmins to manage groups in tenancy
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

### Supported Identity Providers



#### Important

Oracle Cloud Infrastructure tenancies created December 18, 2017 or later are automatically federated with Oracle Identity Cloud Service.

If your tenancy was created before December 18, 2017, and you want to set up a federation with Oracle Identity Cloud Service, see [Federating with Oracle Identity Cloud Service](#).

For instructions for federating with other identity providers, see the following:

[Federating with Microsoft Active Directory](#)

[Federating with Microsoft Azure Active Directory](#)

[Cloud Infrastructure Okta Configuration for Federation and Provisioning](#) (white paper)

[Federating with SAML 2.0 Identity Providers](#)

### Federating with Oracle Identity Cloud Service

This topic points to the appropriate topics for federating Oracle Cloud Infrastructure with [Oracle Identity Cloud Service](#) depending on when you activated your tenancy.

#### **Tenancies created December 21, 2018 and after**

These tenancies are automatically federated with Oracle Identity Cloud Service and configured to provision federated users in Oracle Cloud Infrastructure.

To manage your federated users and groups, see [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#).

For information about the federation, see [Frequently Asked Questions for Oracle Identity Cloud Service Federated Users](#).

#### **Tenancies created between December 18, 2017 and December 20, 2018**

These tenancies are automatically federated with Oracle Identity Cloud Service but are not configured to provision federated users in Oracle Cloud Infrastructure to allow these users to have additional credentials (API keys, auth tokens, etc.).

To enable this feature for users, you need to perform a one-time upgrade, see: [User Provisioning for Federated Users](#).

After you have performed this upgrade, see [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#) to manage your federated users and groups.

### Tenancies created before December 18, 2017

These tenancies must be manually federated with Oracle Identity Cloud Service. See [Federating with Oracle Identity Cloud Service](#) described below.

### Manually Federating with Oracle Identity Cloud Service

Your organization can have multiple Oracle Identity Cloud Service accounts (e.g., one for each division of the organization). You can federate multiple Identity Cloud Service accounts with Oracle Cloud Infrastructure, but each federation trust that you set up must be for a single Identity Cloud Service account.



#### Note

Before following the steps in this topic, see [Federating with Identity Providers](#) to ensure that you understand general federation concepts.

### WEB APPLICATION AND CLIENT CREDENTIALS

For each trust, you must set up a *web application* in Oracle Identity Cloud Service (also called a *trusted application*); instructions are in [Instructions for Federating with Oracle Identity Cloud Service](#). The resulting application has a set of client credentials (a client ID and client secret). When you federate your Identity Cloud Service account with Oracle Cloud Infrastructure, you must provide these credentials.

### COMPUTE BAREMETAL APPLICATION

A *trusted application* in Oracle Identity Cloud Service that contains the set of client credentials (a client ID and client secret) you'll need to provide when you federate your Identity Cloud Service account with Oracle Cloud Infrastructure.

### REQUIRED URLS

The easiest way to federate with Oracle Identity Cloud Service is through the Oracle Cloud Infrastructure Console, although you could do it programmatically with the API. If you're

using the Console, you're asked to provide a *base URL* instead of the metadata URL. The base URL is the left-most part of the URL in the browser window when you're signed in to the Identity Cloud Service console:

- **Base URL:** *<Identity Cloud Service account name>*.identity.oraclecloud.com

If you're using the API to federate, you need to provide the metadata URL, which is the base URL with /fed/v1/metadata appended, like so:

- **Metadata URL:** *<Identity Cloud Service account name>*.identity.oraclecloud.com/fed/v1/metadata

The metadata URL links directly to the IdP-provided XML required to federate. If you're using the API, you need to provide both the metadata URL and the metadata itself when federating. For more information, see [Managing Identity Providers in the API](#).

### **OCI-V2-<TENANCY\_NAME> APP**

When you manually federate an Oracle Identity Cloud Service account with Oracle Cloud Infrastructure, a new SAML application called *OCI-V2- <tenancy\_name>* is automatically created in that Oracle Identity Cloud Service account. If you later need to delete the Oracle Identity Cloud Service identity provider from your Oracle Cloud Infrastructure tenancy, make sure to also delete the *OCI-V2- <tenancy\_name>* from Oracle Identity Cloud Service. If you don't, and you later try to federate the same Oracle Identity Cloud Service account again, you'll get a 409 error saying that an application with the same name already exists (that is, *OCI-V2- <tenancy\_name>*).

### **PROVISIONED USER**

A provisioned user is provisioned by Oracle Identity Cloud Service in Oracle Cloud Infrastructure and is synched to a federated user that is managed in Oracle Identity Cloud Service. The provisioned user can have the special Oracle Cloud Infrastructure credentials like API keys and auth tokens to enable programmatic access. Provisioned users cannot have Console passwords.

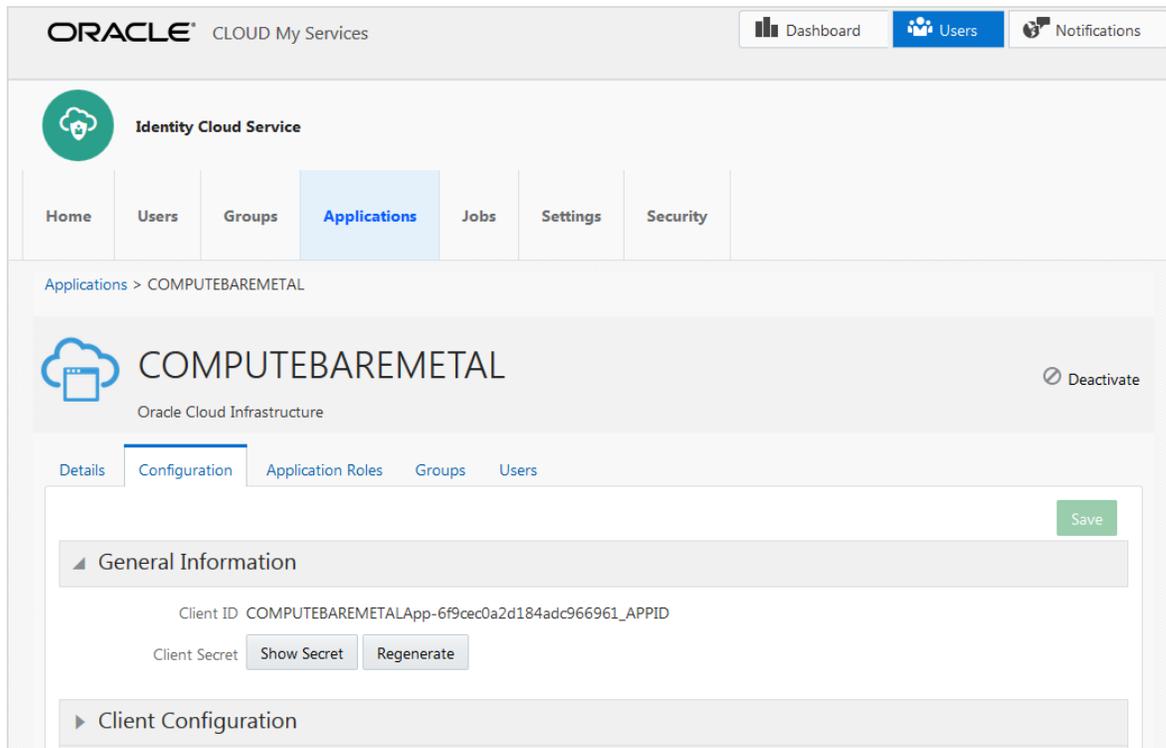
### Instructions for Federating with Oracle Identity Cloud Service

Following is the general process an administrator goes through to set up the identity provider, and below are instructions for each step. It's assumed that the administrator is an Oracle Cloud Infrastructure user with the required credentials and access.

1. In Oracle Identity Cloud Service, get the required information you'll need to perform the set up steps in Oracle Cloud Infrastructure.
2. In Oracle Cloud Infrastructure, set up the federation:
  - a. Set up Oracle Identity Cloud Service as an identity provider.
  - b. Map Oracle Identity Cloud Service groups to IAM groups.
3. In Oracle Cloud Infrastructure, set up the IAM policies for the IAM groups to define the access you want the members of the mapped groups to have.
4. Inform your users of the name of your Oracle Cloud Infrastructure tenant and the URL for the Console (for example, <https://console.us-ashburn-1.oraclecloud.com>).

### Step 1: Get required information from Oracle Identity Cloud Service

1. Go to the Oracle Identity Cloud Service console and sign in with admin privileges. Make sure you're viewing the Admin Console.
2. In the Identity Cloud Service console, click **Applications**. The list of trusted applications is displayed.
3. Click COMPUTEBAREMETAL.
4. Click **Configuration**.
5. Expand **General Information**. The client ID is displayed. Click **Show Secret** to display the client secret.



6. Record the Client ID and Client Secret. They look similar to this:

- Client ID: de06b81cb45a45a8acdcde923402a9389d8
- Client Secret: 8a297afd-66df-49ee-c67d-39cdf3d1c31

## Step 2: Add Oracle Identity Cloud Service as an identity provider in Oracle Cloud Infrastructure

1. Go to the [Console](#) and sign in with your Oracle Cloud Infrastructure login and password.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
3. Click **Add identity provider**.

4. Enter the following:
  - a. **Name:** A unique name for this federation trust. This is the name federated users see when choosing which identity provider to use when signing in to the Console (for example., *ABCCorp\_IDCS* as shown in the screenshot in [Experience for Federated Users](#)). The name must be unique across all identity providers you add to the tenancy. You cannot change this later.
  - b. **Description:** A friendly description.
  - c. **IDCS Base URL:** See [Required URLs](#).
  - d. **Client ID:** From [Step 1: Get required information from Oracle Identity Cloud Service](#).
  - e. **Client secret:** From [Step 1: Get required information from Oracle Identity Cloud Service](#).
  - f. **Encrypt Assertion:** Selecting the check box lets the IAM service know to expect the encryption from the IdP. If you select this check box, you must also set up encryption of the assertion in IDCS. For more information, see [Encrypt Assertion](#). For information about setting this feature up in the IDCS, see [Managing Oracle Identity Cloud Service Applications](#).
  - g. **Force Authentication:** Selected by default. When selected, users are required to provide their credentials to the IdP (re-authenticate) even when they are already signed in to another session.
  - h. **Authentication Context Class References:** This field is required for Government Cloud customers. When one or more values are specified, Oracle Cloud Infrastructure (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the IdP must contain an authentication statement with that authentication context class reference. If the SAML response authentication context does not match what is specified here, the Oracle Cloud Infrastructure auth service rejects the SAML response with a 400.

Several common authentication context class references are listed in the menu. To use a different context class, select **Custom**, then manually enter the class reference.

- i. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Continue**.
  6. Set up the mappings between Oracle Identity Cloud Service groups and IAM groups in Oracle Cloud Infrastructure. A given Oracle Identity Cloud Service group can be mapped to zero, one, or multiple IAM groups, and vice versa. However, each individual mapping is between only a single Oracle Identity Cloud Service group and a single IAM group. Changes to group mappings take effect typically within seconds.



### Note

If you don't want to set up the group mappings now, you can simply click **Create** and come back to add the mappings later.

To create a group mapping:

- a. Select the Oracle Identity Cloud Service group from the list under **Identity Provider Group**.
- b. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".



### Tip

Requirements for IAM group name: No spaces. Allowed characters: letters, numerals, hyphens, periods, underscores, and plus signs (+). The name cannot be changed later.

- c. Repeat the above sub-steps for each mapping you want to create, and then click **Create**.

### AFTER THE FEDERATION SET UP

The identity provider is now added to your tenancy and appears in the list on the **Federation** page. Click the identity provider to view its details and the group mappings you just set up.

Oracle assigns the identity provider and each group mapping a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

In the future, come to the **Federation** page if you want to edit the group mappings or delete the identity provider from your tenancy.

Users that are members of the Oracle Identity Cloud Service groups mapped to the Oracle Cloud Infrastructure groups are now listed in the Console on the Users page. See [Managing User Capabilities for Federated Users](#) for more information on assigning these users additional credentials.

### Step 3: Set up IAM policies for the groups

If you haven't already, set up IAM policies to control the access the federated users have to your organization's Oracle Cloud Infrastructure resources. For more information, see [Getting Started with Policies](#) and [Common Policies](#).

### Step 4: Give your federated users the name of the tenant and URL to sign in

The federated users need the URL for the Oracle Cloud Infrastructure Console (for example, <https://console.us-ashburn-1.oraclecloud.com>) and the name of your tenant. They'll be prompted to provide the tenant name when they sign in to the Console.

### Managing Identity Providers in the Console



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To add an Oracle Identity Cloud Service as an identity provider

See [Instructions for Federating with Oracle Identity Cloud Service](#).

### To delete the identity provider

All the group mappings will also be deleted.

1. Delete the identity provider from your tenancy:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
  - b. Click the identity provider to view its details.
  - c. Click **Delete**.
  - d. Confirm when prompted.

2. Delete the [OCI-V2-\*<tenancy\\_name>\*](#) from your Oracle Identity Cloud Service account:
  - a. Go to Oracle Identity Cloud Service and sign in to the federated account.
  - b. Click **Applications**. The list of applications is displayed.
  - c. Locate the [OCI-V2-\*<tenancy\\_name>\*](#) and click its name to view its details page.
  - d. In the upper right of the page, click **Deactivate**. Confirm when prompted.
  - e. Click **Remove**. Confirm when prompted.

### To add group mappings for Oracle Identity Cloud Service

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

A list of the identity providers in your tenancy is displayed.
2. Click the name you chose for your Oracle Identity Cloud Service federation to view its details.
3. Click **Edit Provider Details**.
4. Add at least one mapping:
  - a. Click **+ Add Mapping**.
  - b. Select the Oracle Identity Cloud Service group from the list under **Identity Provider Group**.
  - c. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
  - d. Repeat the above sub-steps for each mapping you want to create, and then click **Submit**.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions.

Users that are members of the Oracle Identity Cloud Service groups mapped to the Oracle Cloud Infrastructure groups are now listed in the Console on the Users page. See [Managing User Capabilities for Federated Users](#) for more information on assigning these users additional credentials.

### To update or delete a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Update the mappings (or click the X to delete a mapping), and then click **Submit**.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions.

If this action results in federated users no longer having membership in any group that is mapped to Oracle Cloud Infrastructure, the federated users' provisioned users' will also be removed from Oracle Cloud Infrastructure. Typically, this process takes several minutes.

### Managing Identity Providers in the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations:

#### Identity providers:

- [CreateIdentityProvider](#)
- [ListIdentityProviders](#)
- [GetIdentityProvider](#)
- [UpdateIdentityProvider](#)
- [DeleteIdentityProvider](#): Before you can use this operation, you must first use [DeleteIdpGroupMapping](#) to remove all the group mappings for the identity provider.

### **Group mappings:**

- [CreateIdpGroupMapping](#): Each group mapping is a separate entity with its own OCID.
- [ListIdpGroupMappings](#)
- [GetIdpGroupMapping](#)
- [UpdateIdpGroupMapping](#)
- [DeleteIdpGroupMapping](#)

### **Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console**

This topic describes how to use the Oracle Cloud Infrastructure Console to manage your Oracle Identity Cloud Service users and groups. Before you get started, understand basic federation concepts. See [Federating with Identity Providers](#).

#### **OVERVIEW OF WORKING WITH ORACLE IDENTITY CLOUD SERVICE USERS AND GROUPS IN THE CONSOLE**

The Oracle Cloud Infrastructure Console provides an integration with Oracle Identity Cloud Service (IDCS) that lets you perform many management tasks for your IDCS users and groups in the Console.

#### **User Management Tasks**

In the Console, you can do the following user management tasks:

- Add users
- Remove users

- Add users to groups
- Assign roles to users to access services and instances
- Reset user password

For information on more user management tasks, see [Managing Oracle Identity Cloud Service Users](#) in *Administering Oracle Identity Cloud Service*.

### Group Management Tasks

In the Console, you can do the following group management tasks:

- Add groups
- Remove groups
- Add users to groups
- Map IDCS groups to IAM groups

For information on more group management tasks, see [Managing Oracle Identity Cloud Service Groups](#) in *Administering Oracle Identity Cloud Service*.

### REQUIRED POLICIES AND PERMISSIONS

To manage Oracle Identity Cloud Service users and groups in the Console, you'll need to be granted permissions in both the Oracle Cloud Infrastructure IAM service and in Oracle Identity Cloud Service.

Members of the OCI\_Admistrators group have the required permissions to create groups and policies in Oracle Cloud Infrastructure.

**Important:** To create users and groups in the Oracle Identity Cloud Service federation, you'll need the Identity Domain Administrator role, or be a member of a group that has been granted that role. For information on Oracle Identity Cloud Service roles, see [Administering Oracle Identity Cloud Service](#).

To quickly create a user with the required permissions, see [Add a User with Oracle Cloud Administrator Permissions](#).

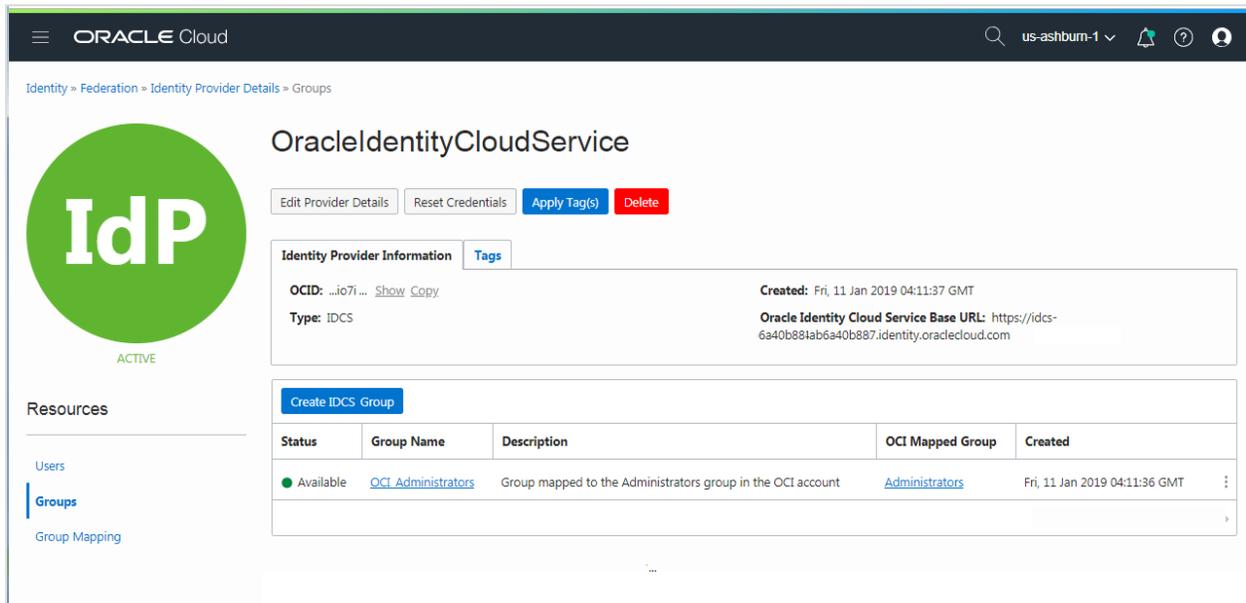
**NAVIGATING TO YOUR ORACLE IDENTITY CLOUD SERVICE USERS AND GROUPS IN THE CONSOLE**

In the Console, you can add users and groups to Oracle Identity Cloud Service from the Identity Provider Details page.

To view your identity provider details:

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.

The identity provider details page is displayed.



From the Identity Provider Details page, click **Users** to display the users created in Oracle Identity Cloud Service. Click **Groups** to display the groups created in Oracle Identity Cloud Service.

### **WORKING WITH ORACLE IDENTITY CLOUD SERVICE GROUPS**

The Console lets you perform the following tasks to manage groups in Oracle Identity Cloud Service:

- Add groups
- Delete groups
- Edit the name and description
- Add users to groups
- Remove users from groups
- Map groups to Oracle Cloud Infrastructure groups

Some tasks you can't perform in the Oracle Cloud Infrastructure Console. To add the predefined application roles for some Oracle Cloud products, you need to assign roles in the Identity Cloud Service console. For more information about using Oracle Identity Cloud Service, see [Administering Oracle Identity Cloud Service](#).

For the members of a group in Oracle Identity Cloud Service to have permissions in Oracle Cloud Infrastructure, you must map the IDCS group to a group in IAM. Before you set up any new groups in IDCS, ensure that you understand how to assign permissions to groups in Oracle Cloud Infrastructure. See [Overview of Oracle Cloud Infrastructure Identity and Access Management](#).

### **WORKING WITH ORACLE IDENTITY CLOUD SERVICE USERS**

The Console lets you perform the following tasks to manage users in Oracle Identity Cloud Service:

- Add users
- Delete users
- Edit user details
- Add users to groups
- Add roles to users

- Remove users from groups
- Reset user passwords

### *USER MANAGEMENT TASKS YOU CAN'T PERFORM IN THE CONSOLE*

The Oracle Cloud Console does not support management of the following Oracle Identity Cloud Service user features and tasks:

- Manage multi-factor authentication

For information about managing these tasks, see [Administering Oracle Identity Cloud Service](#).

### MANAGING ORACLE IDENTITY CLOUD SERVICE GROUPS IN THE CONSOLE



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## To create a group in Oracle Identity Cloud Service

This procedure creates a new group in Oracle Identity Cloud Service. Optionally, you can add users to the group at the time you create it. This group will not have any permissions in Oracle Cloud Infrastructure until you map it to an Oracle Cloud Infrastructure group.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the federations in your tenancy is displayed.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.

3. Under **Resources**, click **Groups**.  
The list of existing groups is displayed.
4. Click **Create IDCS Group**.
5. Enter the following:
  - **Name:** A unique name for the group.
  - **Description:** A friendly description. You can change this later if you want to.
  - **Users:** Add Oracle Identity Cloud Service users to this group. You can add users when you create the group, or later. Select users from the list. To find a specific user, you can start typing the user name to filter the list as you type.
6. Click **Create**.

After you create a group in Oracle Identity Cloud Service, you'll want to give the group permissions to user services:

- To grant the group access to map it to an Oracle Cloud Infrastructure group as described in the next procedure.
- To add roles to this group, see [Managing Oracle Identity Cloud Service Roles for Groups](#).

### To map an Oracle Identity Cloud Service group to an IAM group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Click **Edit Mapping**.
4. In the **Edit Identity Provider** dialog, click **+ Add Mapping**.
5. Select the **Identity Provider Group** you want to map from the list. To find a group without scrolling through the list, you can start typing the group name to filter the list as

you type.

6. Select the **OCI Group** you want to map this Identity Cloud Service group to. To find a group without scrolling through the list, you can start typing the group name to filter the list as you type.
7. To add more mappings, click **+ Add Mapping** and continue adding the mappings.
8. Select the group you want to map this group to from the list under **OCI Mapped User Group**.

Members of this group now have the permissions granted to the OCI Mapped User Group.

### To add roles to a group

Oracle Cloud Infrastructure services use policies to control access to services. However, some Oracle Cloud services use roles to manage access. This procedure describes how to add roles to an IDCS group.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the **Oracle Identity Cloud Service Console** link.  
The Identity Cloud Service console is displayed.
3. In the Identity Cloud Service console, expand the **Navigation Drawer**, and then click **Applications**.  
The list of applications is displayed. Notice that the service that the application corresponds to is displayed underneath the application name. For example, underneath the JAAS application entry, you'll see Oracle Java Cloud Service.
4. Click the name of the service that you are interested in.  
The **Details** page is displayed.
5. Click **Application Roles**.  
The roles are displayed.
6. Click the menu for the role you want to assign and select **Assign Groups**.

7. Select the group you want to assign to the role, and click **OK**.
8. Click the **Applications** breadcrumb to return to the list of applications.
9. Repeat steps 4 through 7 for each role you want to assign to this group.

### To remove roles from a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the **Oracle Identity Cloud Console** link.  
The Identity Cloud Service console is displayed.
3. In the Identity Cloud Service console, expand the **Navigation Drawer**, and then click **Applications**.  
The list of applications is displayed. Notice that the service that the application corresponds to is displayed underneath the application name. For example, underneath the JAAS application entry, you'll see Oracle Java Cloud Service.
4. Click the name of the service that you are interested in.  
The **Details** page is displayed.
5. Click **Application Roles**.  
The roles are displayed.
6. Click the menu for the role you want to remove from the group and select **Revoke Groups**.
7. Select the group you want to remove the role from, and click **OK**.
8. Click the **Applications** breadcrumb to return to the list of applications.
9. Repeat steps 4 through 7 for each role you want to remove from this group.

### To edit details for an Oracle Identity Cloud Service group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Under **Resources**, click **Groups**.  
The list of existing groups in the federation is displayed.
4. Find the group you want to edit and click its name.  
The **Group Details** page is displayed.
5. Click **Edit**.
6. You can update the **Group Name** or the **Description**.
7. Click **Update** to save your changes.



#### Warning

Changing the group name will break mappings to Oracle Cloud Infrastructure (OCI) groups. If you change the group name, ensure that you delete any existing group mappings and add new mappings with the new name. See the previous task on editing mappings.

### To add users to a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.

The identity provider details page is displayed.

3. Under **Resources**, click **Groups**.

The list of existing groups is displayed.

4. Find the group you want add a user to.

The **User Group Details** page is displayed.

5. Click **Add IDCS User**.

6. Select the user you want to add to this group from the **Users** list.

7. Click **Add**.

### To remove users from a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.

The identity provider details page is displayed.

3. Under **Resources**, click **Groups**.

The list of existing groups is displayed.

4. Find the group you want to remove the user from.

The list of users is displayed in the **Group Details** page.

5. Find the user you want to remove, and then click the the Actions icon (three dots).

6. Click **Remove User**.

7. Confirm when prompted.

### To delete a group

Note: To delete a group, you must first remove all the users.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity**

and click **Federation**.

2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.

The identity provider details page is displayed.

3. Under **Resources**, click **Groups**.

The list of existing groups is displayed.

4. Find the group you want to edit and click its name.

The **Group Details** page is displayed.

5. Click **Delete**.

6. Confirm when prompted.

### Create a policy to grant the group permissions on Oracle Cloud Infrastructure resources

The group you created in Oracle Identity Cloud Service gets permissions to access resources in Oracle Cloud Infrastructure through the policy you assign to the Oracle Cloud Infrastructure group. Before you complete this step, you need to decide what permissions you want to give your new group. For more information, see [Getting Started with Policies](#) and [Common Policies](#).

Prerequisite: The group and compartment that you're writing the policy for must already exist.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.

A list of the policies in the compartment you're viewing is displayed.

2. If you want to attach the policy to a compartment other than the one you're viewing, select the desired compartment from the list on the left. Where the policy is attached controls who can later modify or delete it (see [Policy Attachment](#)).

3. Click **Create Policy**.

4. Enter the following:

- **Name:** A unique name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later.
- **Description:** A friendly description. You can change this later if you want to.
- **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in format YYYY-MM-DD format. For more information, see [Policy Language Version](#).
- **Statement:** A policy statement. For the correct format to use, see [Policy Basics](#) and also [Policy Syntax](#). If you want to add more than one statement, click **+**.

For example:

To allow your group to manage all resources within a specified compartment enter a statement like the following:

```
Allow group <OCI_group_name> to manage all-resources in compartment <compartment_name>
```

For more policy examples, see [Common Policies](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

### MANAGING ORACLE IDENTITY CLOUD SERVICE USERS IN THE CONSOLE

After you add a user in Oracle Identity Cloud Service, a user is also automatically provisioned in Oracle Cloud Infrastructure. This provisioned user can have the Oracle Cloud Infrastructure credentials, such as API keys and auth tokens. To understand this provisioning, see [User Provisioning for Federated Users](#).

### To create a user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Click **Create IDCS User**.
4. In the **Create IDCS User** dialog enter the following:
  - **User Name**: Enter a unique name or email address for the new user. The value will be the user's login to the Console and must be unique across all other users in your tenancy.
  - **Email**: Enter an email address for this user. The initial sign-in credentials will be sent to this email address.
  - **First Name**: Enter the user's first name.
  - **Last Name**: Enter the user's last name.
  - **Phone Number**: Optionally, enter a phone number.
  - **Groups**: Optionally, select groups to add this user to.
5. Click **Create User**.



#### **Important**

For the user to have permissions in Oracle Cloud Infrastructure, you must assign the user to a group that is mapped to an Oracle Cloud Infrastructure group. Or, if you are also creating a new group, you can perform this mapping later. The user will not be able to sign in to the Console until the mapping is accomplished.

The user creation process generates an email that is sent to the address provided that you entered. The email includes the new user's username and password to use with the Oracle Cloud Infrastructure Console.

To add API keys, auth tokens, customer secret keys, or SMTP credentials for this user, see [Managing User Capabilities for Federated Users](#).

### To edit a user

You can update the following fields:

- Email address
  - First and last name
  - Phone number
1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
  2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
  3. Under **Resources**, click **Users**.  
The list of existing users is displayed.
  4. Find the user you want to edit and click its name.  
The **User Details** page is displayed.
  5. Click **Edit**.
  6. Update the fields.
  7. Click **Save** when finished.

### To reset a user's password

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Under **Resources**, click **Users**.  
The list of existing user groups in the federation is displayed.
4. Find the user you want to reset the password for and click the name.  
The **User Details** page is displayed.
5. Click **Reset Password**.  
The user's password is reset. This user can't access their account until they complete the password reset steps.
6. Click **Email Password Instructions** to send the password link and instructions to the user.  
The password link is good for 24 hours. If the user does not reset their password in time, you can generate a new password link by clicking **Reset Password** for the user again.

### To manage roles for services managed through IDCS

See see [Managing Oracle Identity Cloud Service Roles for Users](#).

### To add API keys, auth tokens, or other Oracle Cloud Infrastructure credentials

1. View the user's details:

- If you're adding credentials for *yourself*: Open the **Profile** menu () and click **User Settings**.
- If you're an administrator adding credentials for *another user*: Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.

The identity provider details page is displayed.

Find the user in the list and click the **OCI Synched User** link.

2. Add the credentials for the user.

For more details about these credentials, see [Managing User Credentials](#).

### To delete a user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Under **Resources**, click **Users**.  
The list of existing user groups in the federation is displayed.
4. Find the user you want to delete and click the name.  
The **User Details** page is displayed.
5. Click **Delete**.

### MANAGING GROUP MAPPINGS

#### To add group mappings for Oracle Identity Cloud Service

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**.  
The identity provider details page is displayed.
3. Click **Edit Provider Details**.
4. Add at least one mapping:
  - a. Click **+ Add Mapping**.
  - b. Select the Oracle Identity Cloud Service group from the list under **Identity Provider Group**.
  - c. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
  - d. Repeat the above sub-steps for each mapping you want to create, and then click **Submit**.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions.

Users that are members of the Oracle Identity Cloud Service groups mapped to the Oracle Cloud Infrastructure groups are now listed in the Console on the Users page. See [Managing User Capabilities for Federated Users](#) for more information on assigning these users additional credentials.

### To update or delete a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Update the mappings (or click the X to delete a mapping), and then click **Submit**.

If this action results in federated users no longer having membership in any group that is mapped to Oracle Cloud Infrastructure, the federated users' provisioned users' will also be removed from Oracle Cloud Infrastructure. Typically, this process takes several minutes.

### Managing Oracle Identity Cloud Service Roles for Users

This topic describes managing user roles for users created in Oracle Identity Cloud Service.

#### ABOUT USER ROLES IN ORACLE IDENTITY CLOUD SERVICE

You can assign roles to a user to allow access to those Oracle Cloud services that have predefined roles defined in Oracle Identity Cloud Service. You can also grant access just to service instances.

Services managed through Identity Cloud Service can have two types of predefined roles:

- Service access roles - grant access to use the service.
- Instance access roles - grant access to specific instances of a service. These can only be granted after the instances are created

For information about more complex role management including assigning other administrative privileges, see [Managing Oracle Identity Cloud Service Users](#).

#### AVAILABLE ROLES FOR EACH SERVICE

Service-specific roles vary from one Oracle Cloud service to another, but they typically include at least one administrator role. See [About Service Administrator Roles](#) for more

information about administrator roles. See your service-specific documentation for a description of the predefined roles for that service.

### REQUIRED PERMISSIONS TO MANAGE ROLES

Before you can manage roles using the Oracle Cloud Infrastructure Console, you must be allowed to access the Identity Provider Details page. To access this page, you must belong to a group that is allowed to inspect identity providers. If you are a Cloud Administrator or if you belong to the OCI Administrators group, this permission is included. To give this permission to non-administrators, you'll need to write a policy like the following:

```
Allow group GroupA to inspect identity-providers in tenancy
```

where you replace GroupA with the name of the group you want to grant the permission to.

To manage the service roles for another user, you must be assigned the appropriate role in Oracle Identity Cloud Service. See [Understanding Administrator Roles](#).

### MANAGING ROLES USER ROLES IN THE CONSOLE

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.  
By default, users belonging to all identity providers are displayed. To view only users that belong to your Identity Cloud Service federation, clear the check boxes for any other identity providers.
2. Click the name of the user you want to edit.
3. On the user details page, click **Manage Service Roles**. The **Manage Service Roles** page displays the list of services for which you have Administrator access. The service roles that this user has already been granted are also displayed.  
Note that you won't see services that you don't have Administrator access for.
4. Find the service you want to edit this user's access to, click the Actions icon (three dots), and then click **Manage service access**. The list of roles for the selected service is displayed.
5. Edit the user's access as follows:

- Select the check box for each role you want to give to the user.
- Clear the check box for each role you want remove from the user. Note that you can't remove a role that has been granted through a group. These roles are read only.



### Note

If a user is assigned the Cloud Account Administrator role, then you can't remove the individual entitlement roles for the user.

6. Click **Save Role Selections**.
7. Click **Apply Service Role Settings**.
8. If you are granting roles to a user, in the confirmation dialog, click **Send Email to User** to send an email to the user to notify them of this change.
9. Your email client launches with a default email message you can send to the user. You can send the email as shown, or make modifications before sending.
10. Return to the Console and click **Close**.

### MANAGING INSTANCE ROLES IN THE CONSOLE

Some services allow you to grant access to instances of the service. After you (or someone in your organization) creates an instance, use this procedure to manage individual user access to the instance.

#### Managing User Access to an Instance

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.  
By default, users belonging to all identity providers are displayed. To view only users that belong to your Identity Cloud Service federation, clear the check boxes for any other identity providers.

2. Click the name of the user you want to edit.
3. On the user details page, click **Manage Service Roles**. The **Manage Service Roles** page displays the list of services for which you have Administrator access. The service roles that this user has already been granted are also displayed.  
Note that you won't see services that you don't have Administrator access for.
4. Find the service with instances that you want to edit this user's access to, click the Actions icon (three dots), and then click **Manage instance access**. The list of instances for the selected service is displayed.
5. On the **Manage Access to Instances** page, find the name of the instance you want to edit this user's access to.  
To grant access to this instance:  
In the **Instance Role** column, select the role you want to grant to the user. You can select multiple roles from the list.  
To remove access to this instance:  
In the **Instance Role** column, click the **x** next to the role you want to remove from the user.
6. When you are finished editing roles, click **Save Instance Settings**.
7. On the **Manage Service Roles** page, click **Apply Service Role Settings**.
8. If you are granting roles to a user, in the confirmation dialog, click **Send Email to User** to send an email to the user to notify them of this change.
9. Your email client launches with a default email message you can send to the user. You can send the email as shown, or make modifications before sending.
10. Return to the Console and click **Close**.

### Managing Oracle Identity Cloud Service Roles for Groups

This topic describes managing roles for groups created in Oracle Identity Cloud Service.

#### ABOUT GROUP ROLES IN ORACLE IDENTITY CLOUD SERVICE

You can assign roles to groups to allow access to those Oracle Cloud services that have predefined roles defined in Oracle Identity Cloud Service. You can also grant access just to

service instances.

Services managed through Identity Cloud Service can have two types of predefined roles:

- Service access roles - grant access to use the service.
- Instance access roles - grant access to specific instances of a service. These can only be granted after the instances are created.

For information about more complex role management, see [Manage Oracle Identity Cloud Service Groups](#).

### AVAILABLE ROLES FOR EACH SERVICE

Service-specific roles vary from one Oracle Cloud service to another, but they typically include at least one administrator role. See [About Service Administrator Roles](#) for more information about administrator roles. See your service-specific documentation for a description of the predefined roles for that service.

### REQUIRED PERMISSIONS TO MANAGE ROLES

Before you can manage roles using the Oracle Cloud Infrastructure Console, you must be allowed to access the Identity Provider Details page. To access this page, you must belong to a group that is allowed to inspect identity providers. If you are a Cloud Administrator or if you belong to the OCI\_Admistrators group, this permission is included. To give this permission to non-administrators, you'll need to write a policy like the following:

```
Allow group GroupA to inspect identity-providers in tenancy
```

where you replace GroupA with the name of the group you want to grant the permission to.

To manage service roles, you must be assigned the Administrator role for that service.

### MANAGING GROUP ROLES IN THE CONSOLE

#### To add roles to a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**. The identity provider details page is displayed.
3. Click **Groups**.  
The list of groups is displayed.
4. Click the name of the group you want to add roles to.
5. On the group details page, click **Manage Service Roles**. The **Manage Service Roles** page displays the list of services for which you have Administrator access. The service roles that this group has already been granted are also displayed.  
Note that you won't see services that you don't have Administrator access for.
6. Find the service you want to edit this group's access to, click the Actions icon (three dots), and then click **Manage service access**. The list of roles for the selected service is displayed.
7. Select the check box for each role you want to assign to the group.
8. Click **Save Role Selections**.
9. To add more service roles to this group, repeat steps 6 - 8.
10. Click **Apply Service Role Settings**.
11. In the confirmation dialog, click **Send Email to Group** to send an email to each member of the group to notify them of this change.  
Your email client launches with a default email message to the affected users with information about the access changes. You can send the email as written, or make modifications before sending.
12. Return to the Console and click **Close**.

### To revoke roles from a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation

is named **OracleIdentityCloudService**. The identity provider details page is displayed.

3. Click **Groups**.

The list of groups is displayed.

4. Click the name of the group you want to remove roles from.

5. On the group details page, click **Manage Service Roles**. The **Manage Service Roles** page displays the list of services for which you have Administrator access. The service roles that this group has already been granted are also displayed.

Note that you won't see services that you don't have Administrator access for.

6. Find the service you want to edit this group's access to, click the Actions icon (three dots), and then click **Manage service access**. The list of roles for the selected service is displayed.

7. Clear the check box for each role you want remove from the group.

8. Click **Save Role Selections**.

9. To revoke more service roles from this group, repeat steps 6 - 8.

10. Click **Apply Service Role Settings**.

11. A confirmation dialog displays the services that you modified access to in this session. Click **Close**.

### MANAGING INSTANCE ROLES IN THE CONSOLE

Some services allow you to grant access to instances of the service. After you (or someone in your organization) creates an instance, use this procedure to manage group access to the instance.

#### Managing Group Access to an Instance

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

2. Click your Oracle Identity Cloud Service federation. For most tenancies, the federation is named **OracleIdentityCloudService**. The identity provider details page is displayed.

3. Click **Groups**.  
The list of groups is displayed.
4. On the group details page, click **Manage Service Roles**. The **Manage Service Roles** page displays the list of services for which you have Administrator access. The service roles that this group has already been granted are also displayed.  
Note that you won't see services that you don't have Administrator access for.
5. Find the service with instances that you want to edit this group's access to, click the Actions icon (three dots), and then click **Manage instance access**. The list of instances for the selected service is displayed.
6. On the **Manage Access to Instances** page, find the name of the instance you want to edit this group's access to.
  - To grant access to this instance: In the **Instance Role** column, select the role you want to grant to the group. You can select multiple roles from the list.
  - To remove access to this instance: In the **Instance Role** column, click the **x** next to the role you want to remove from the group.
7. When you are finished editing roles for this service, click **Save Instance Settings**.
8. To edit more instance roles for this group, repeat steps 6 - 7.
9. On the **Manage Service Roles** page, click **Apply Service Role Settings**.
10. If you added roles, in the confirmation dialog, click **Send Email to Group** to send an email to each member of the group to notify them of this change. Your email client launches with a default email message to the affected users with information about the access changes. You can send the email as written, or make modifications before sending. Return to the Console and click **Close**.  
If you revoked roles, a confirmation dialog displays the services that you modified access to in this session. Click **Close**.

### Frequently Asked Questions for Oracle Identity Cloud Service Federated Users

When you sign up for Oracle Cloud Infrastructure, your account is automatically federated with Oracle Identity Cloud Service as your identity provider. This topic answers some frequently asked questions about the federation.

### What resources are created in Oracle Identity Cloud Service?

*THE FOLLOWING RESOURCES ARE CREATED IN IDENTITY CLOUD SERVICE:*

- **Applications:**

- *OCI-V2-**<tenancy\_name>***

This SAML application that creates the federation with Oracle Cloud Infrastructure.

- COMPUTEBAREMETAL application

A supporting application for the federation.



**Important**

Do not delete these applications.

- **Group:**

OCI\_Administrators group

This group is mapped to the Administrators group in Oracle Cloud Infrastructure.

Members of this group have full administrator privileges in Oracle Cloud Infrastructure.

- **User:**

A default administrator user (e.g., user@example.com) who is a member of the OCI\_Administrators group.

### What resources are created in Oracle Cloud Infrastructure?

*THE FOLLOWING RESOURCES ARE CREATED IN ORACLE CLOUD INFRASTRUCTURE:*

- **Identity Provider:** OracleIdentityCloudService

- **Group Mappings:** The federation is created with one group mapping:

OCI\_Administrators group (from Oracle Identity Cloud Service) is mapped to the Administrators group (In Oracle Cloud Infrastructure).

- **Users:**

- The default administrator user created in Oracle Identity Cloud Service is provisioned in Oracle Cloud Infrastructure. This user can have the Oracle Cloud Infrastructure credentials, but not a Console password.
- A default administrator local-user with the same user name (user@example.com) is also created in Oracle Cloud Infrastructure's IAM service. Customers who choose **not** to use the Oracle Identity Cloud Service federation can use this user to administer Oracle Cloud Infrastructure.



### **Important**

The default administrator created in Oracle Identity Cloud Service and the local default administrator created in Oracle Cloud Infrastructure exist independently in their respective identity systems. Ensure that you manage passwords for them separately.

### Why is my account federated with Oracle Identity Cloud Service?

Oracle Identity Cloud Service is the identity provider for multiple Oracle services. Federating Oracle Cloud Infrastructure with Oracle Identity Cloud Service allows you to have a seamless connection between services, without having to create a separate username and password for each one.

### How do I know if I am signed in through Oracle Identity Cloud Service?

Click the **Profile** menu () to display your username. Users signed in through an identity provider will see their username prefaced with their identity provider name, for example: oracleidentitycloudservice/user@example.com

How do I add a user to Oracle Identity Cloud Service (a federated user)?

See [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#).

Can I add a user just for Oracle Cloud Infrastructure?

Yes. If you don't want to manage the user in Oracle Identity Cloud Service, you can add a user directly to the Oracle Cloud Infrastructure IAM service. See [Adding Users](#). Using this procedure, you can create users who can sign in directly to the Oracle Cloud Infrastructure Console. Users created with this procedure do not have access to any other Oracle services.

How do I manage groups?

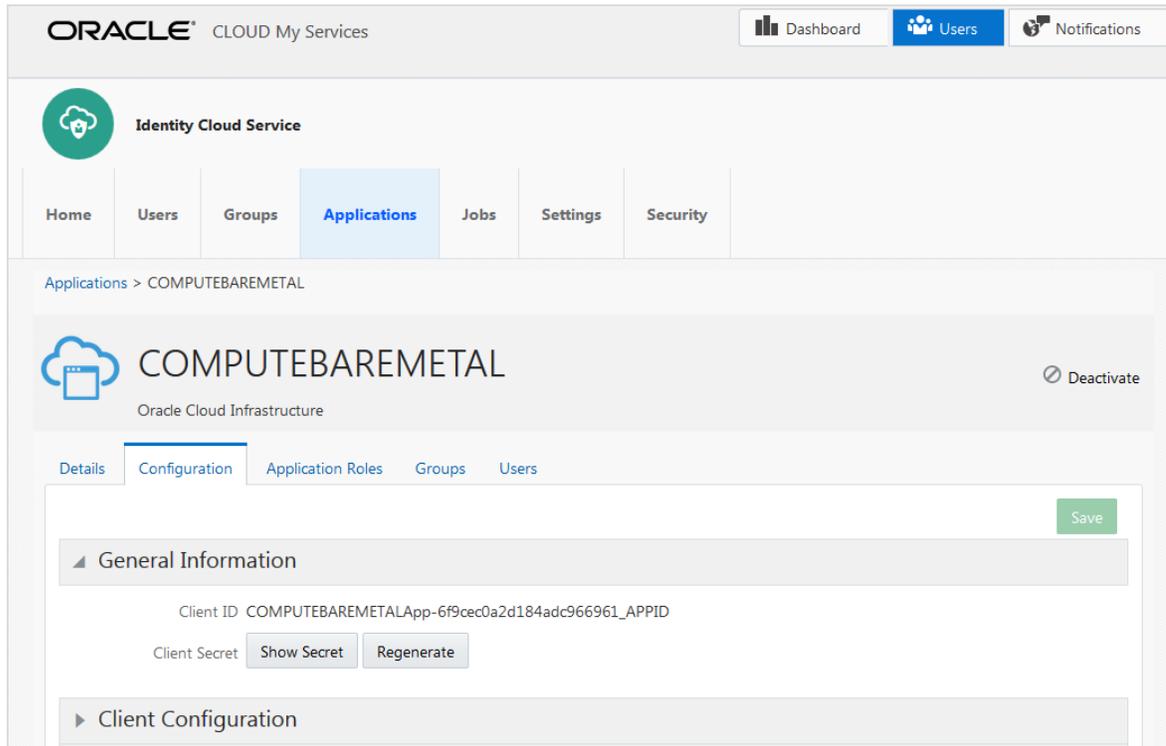
In short, managing groups requires actions in both Oracle Identity Cloud Service and Oracle Cloud Infrastructure. Groups you create in Oracle Identity Cloud Service have no privileges in Oracle Cloud Infrastructure until you map them to a group in Oracle Cloud Infrastructure. You define the policies that permit access to Oracle Cloud Infrastructure resources in the IAM service in Oracle Cloud Infrastructure. For more information, see [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#).

How do I find the client ID and client secret?

To edit mappings of your user groups in Oracle Identity Cloud Service to user groups in Oracle Cloud Infrastructure, you'll need to supply the client ID and client secret. The client ID and client secret are stored in Oracle Identity Cloud Service. To get this information:

1. Sign in to the Oracle Identity Cloud Service console.
2. In the Identity Cloud Service console, click **Applications**. The list of trusted applications is displayed.
3. Click COMPUTEBAREMETAL.
4. Click **Configuration**.

5. Expand **General Information**. The client ID is displayed. Click **Show Secret** to display the client secret.



If I delete the federation, can I later recreate it?

Yes. To recreate the federation with Oracle Identity Cloud Service, follow the instructions in the topic [Federating with Oracle Identity Cloud Service](#).

## Federating with Microsoft Active Directory

This topic describes how to federate with Microsoft Active Directory using Microsoft Active Federation Services (AD FS).



### Note

Before following the steps in this topic, see [Federating with Identity Providers](#) to ensure that you understand general federation concepts.

### About Federating with Microsoft Active Directory

Your organization can have multiple Active Directory accounts (e.g., one for each division of the organization). You can federate multiple Active Directory accounts with Oracle Cloud Infrastructure, but each federation trust that you set up must be for a single Active Directory account.

To federate with Active Directory, you set up a trust between Active Directory and Oracle Cloud Infrastructure. To set up this trust, you perform some steps in the Oracle Cloud Infrastructure Console and some steps in Active Directory Federation Services.

Following is the general process an administrator goes through to set up federation with Active Directory. Details for each step are given in the sections below.

1. Get required information from Active Directory Federation Services.
2. Federate Active Directory with Oracle Cloud Infrastructure:
  - a. Add the identity provider (AD FS) to your tenancy and provide the required information.
  - b. Map Active Directory groups to IAM groups.
3. In Active Directory Federation Services, add Oracle Cloud Infrastructure as a trusted, relying party.
4. In Active Directory Federation Services, add the claim rules required in the authentication response by Oracle Cloud Infrastructure.
5. Test your configuration by logging in to Oracle Cloud Infrastructure with your Active Directory credentials.

### Federating with Active Directory

#### Prerequisites

You have installed and configured Microsoft Active Directory Federation Services for your organization.

You have set up groups in Active Directory to map to groups in Oracle Cloud Infrastructure.



#### Tip

Consider naming Active Directory groups that you intend to map to Oracle Cloud Infrastructure groups with a common prefix, to make it easy to apply a filter rule. For example, OCI\_Admistrators, OCI\_NetworkAdmins, OCI\_InstanceLaunchers.

#### STEP 1: GET REQUIRED INFORMATION FROM ACTIVE DIRECTORY FEDERATION SERVICES

**Summary:** Get the SAML metadata document and the names of the Active Directory groups that you want to map to Oracle Cloud Infrastructure Identity and Access Management groups.

1. Locate the SAML metadata document for your AD FS federation server. By default, it is located at this URL:

```
https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml
```

Download this document and make a note of where you save it. You will upload this document to the Console in the next step.

2. Note all the Active Directory groups that you want to map to Oracle Cloud Infrastructure IAM groups. You will need to enter these in the Console in the next step.

#### STEP 2: ADD ACTIVE DIRECTORY AS AN IDENTITY PROVIDER IN ORACLE CLOUD INFRASTRUCTURE

**Summary:** Add the identity provider to your tenancy. You can set up the group mappings at the same time, or set them up later.

1. Go to the Console and sign in with your Oracle Cloud Infrastructure login and password.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
3. Click **Add identity provider**.
4. Enter the following:
  - a. **Display Name:** A unique name for this federation trust. This is the name federated users see when choosing which identity provider to use when signing in to the Console. The name must be unique across all identity providers you add to the tenancy. You cannot change this later.
  - b. **Description:** A friendly description.
  - c. **Type:** Select **Microsoft Active Directory Federation Services (ADFS)** or **SAML 2.0 compliant identity provider**.
  - d. **XML:** Upload the FederationMetadata.xml file you downloaded from Azure AD.
  - e. Click **Show Advanced Options**.
  - f. **Encrypt Assertion:** Selecting the check box lets the IAM service know to expect the encryption from IdP. Do not select this check box unless you have enabled assertion encryption in Azure AD.

To enable assertion encryption for this single sign-on application in Azure AD, set up the SAML Signing Certificate in Azure AD to sign the SAML response and assertion. For more information, see the [Azure AD documentation](#).
  - g. **Force Authentication:** Selected by default. When selected, users are required to provide their credentials to the IdP (re-authenticate) even when they are already signed in to another session.
  - h. **Authentication Context Class References:** This field is required for Government Cloud customers. When one or more values are specified, Oracle Cloud Infrastructure (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the IdP must contain an authentication statement with that authentication context class reference. If the SAML response

authentication context does not match what is specified here, the Oracle Cloud Infrastructure auth service rejects the SAML response with a 400.

Several common authentication context class references are listed in the menu. To use a different context class, select **Custom**, then manually enter the class reference.

- i. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Continue**.
  6. Set up the mappings between Active Directory groups and IAM groups in Oracle Cloud Infrastructure. A given Active Directory group can be mapped to zero, one, or multiple IAM groups, and vice versa. However, each individual mapping is between only a single Active Directory group and a single IAM group. Changes to group mappings take effect typically within seconds in your home region, but may take several minutes to propagate to all regions.



### Note

If you don't want to set up the group mappings now, you can simply click **Create** and come back to add the mappings later.

To create a group mapping:

- a. Under **Identity Provider Group**, enter the Active Directory group name. You must enter the name exactly, including the correct case.  
Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also

automatically be given this description, which you can't change: "Group created during federation".



### Tip

Requirements for IAM group name: No spaces. Allowed characters: letters, numerals, hyphens, periods, underscores, and plus signs (+). The name cannot be changed later.

- b. Repeat the above sub-steps for each mapping you want to create, and then click **Create**.

The identity provider is now added to your tenancy and appears in the list on the **Federation** page. Click the identity provider to view its details and the group mappings you just set up.

Oracle assigns the identity provider and each group mapping a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

In the future, come to the **Federation** page if you want to edit the group mappings or delete the identity provider from your tenancy.

### STEP 3: COPY THE URL FOR THE ORACLE CLOUD INFRASTRUCTURE FEDERATION METADATA DOCUMENT

**Summary:** The Federation page displays a link to the Oracle Cloud Infrastructure Federation Metadata document. Before you move on to configuring Active Directory Federation Services, you need to copy the URL.

1. On the Federation page, click **Download this document**.
2. Copy the URL. The URL looks similar to:

```
https://auth.r2.oracleiaas.com/v1/saml/ocid1.tenancy.oc1..aaaaaaaaqdt2tvdmhsa3jmvc5dzulgs3pcv6imfwfgdya4aq/metadata.xml
```

### STEP 4: IN ACTIVE DIRECTORY FEDERATION SERVICES, ADD ORACLE CLOUD INFRASTRUCTURE AS A TRUSTED RELYING PARTY

1. Go to the AD FS Management Console and sign in to the account you want to federate.
2. Add Oracle Cloud Infrastructure as a **trusted relying party**:
  - a. From the AD FS Management Console, right-click AD FS and select **Add Relying Party Trust**.
  - b. In the **Add Relying Party Trust Wizard**, click **Start**.
  - c. Select **Import data about the relying party published online or on a local network**.

Paste the Oracle Cloud Infrastructure Federation Metadata URL that you copied in Step 3. Click **Next**.

AD FS will connect to the URL. If you get an error during the attempt to read the federation metadata, you can alternatively upload the Oracle Cloud Infrastructure Federation Metadata XML document.

#### To upload the federation metadata document

- i. In a web browser, paste the Oracle Cloud Infrastructure Federation Metadata URL in the address bar.
  - ii. Save the XML document to a location that is accessible by your AD FS Management Console.
  - iii. In the **Select Data Source** step of the **Add Relying Party Trust Wizard**, select **Import data about the relying party from a file**.
  - iv. Click **Browse** and select the metadata.xml file that you saved.
  - v. Click **Next**.
- d. Set the display name for the relying party (e.g., Oracle Cloud Infrastructure) and then click **Next**.
  - e. Select **I do not want to configure multi-factor authentication settings for this relying party trust at this time**.

- f. Choose the appropriate Issuance Authorization Rules to either permit or deny all users access to the relying party. Note that if you choose "Deny", then you must later add the authorization rules to enable access for the appropriate users. Click **Next**.
- g. Review the settings and click **Next**.
- h. Check **Open the Edit Claim Rules** dialog for this relying part trust when the wizard closes and then click **Close**.

### **STEP 5: ADD THE CLAIM RULES FOR THE ORACLE CLOUD INFRASTRUCTURE RELYING PARTY**

Summary: Add the claim rules so that the elements that Oracle Cloud Infrastructure requires (Name ID and groups) are added to the SAML authentication response.

#### **Add the Name ID rule:**

1. In the **Add Transform Claim Rule Wizard**, select **Transform an Incoming Claim**, and click **Next**.
2. Enter the following:
  - **Claim rule name:** Enter a name for this rule, e.g., nameid.
  - **Incoming claim type:** Select Windows account name.
  - **Outgoing claim type:** Select Name ID.
  - **Outgoing name ID format:** Select Persistent Identifier.
  - Select **Pass through all claim value**.
  - Click **Finish**.
3. The rule is displayed in the rules list. Click **Add Rule**.

#### **Add the groups rule:**

**Important**

Any users who are in more than 100 IdP groups cannot be authenticated to use the Oracle Cloud Infrastructure Console. To enable authentication, apply a filter to the groups rule, as described below.

If your Active Directory users are in no more than 100 groups

**Add the groups rule:**

1. Under Claim rule template, select **Send Claims Using a Custom Rule**. Click **Next**.
2. In the **Add Transform Claim Rule Wizard**, enter the following:
  - a. **Claim rule name:** Enter groups.
  - b. **Custom rule:** Enter the following custom rule:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
Issuer == "AD AUTHORITY"] => issue(store = "Active Directory", types =
("https://auth.oraclecloud.com/saml/claims/groupname"), query = ";tokenGroups;{0}", param
= c.Value);
```

- c. Click **Finish**.

If your Active Directory users are in more than 100 groups

**Add the groups rule with a filter:**

To limit the groups sent to Oracle Cloud Infrastructure, create two custom claim rules. The first one retrieves all groups the user belongs to directly and indirectly. The second rule applies a filter to limit the groups passed to the service provider to only those that match the filter criteria.

Add the first rule:

1. In the Edit Claim Rules dialog, click **Add Rule**.
2. Under Claim rule template, select **Send Claims Using a Custom Rule**. Click **Next**.
3. In the **Add Transform Claim Rule Wizard**, enter the following:

- a. **Claim rule name:** Enter a name, for example, groups.
- b. **Custom rule:** Enter the following custom rule:

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname",
Issuer == "AD AUTHORITY"] => add(store = "Active Directory", types =
("https://auth.oraclecloud.com/saml/claims/groupName"), query = ";tokenGroups;{0}", param
= c.Value);
```

Note that in this custom rule you use `add` instead of `issue`. This command passes the results of the rule to the next rule, instead of sending the results to the service provider.

- c. Click **Finish**.
4. Now add the filter rule.
    - a. In the Edit Claim Rules dialog, click **Add Rule**.
    - b. Under Claim rule template, select **Send Claims Using a Custom Rule**. Click **Next**.
    - c. In the **Add Transform Claim Rule Wizard**, enter the following:
      - i. **Claim rule name:** Enter groups.
      - ii. **Custom rule:** Enter an appropriate filter rule. For example to send only groups that begin with the string "OCI", enter the following:

```
c:[Type == "https://auth.oraclecloud.com/saml/claims/groupName", Value =~ "(?i)OCI"]
=> issue(claim = c);
```

This rule filters the list from the first rule to only those groups that begin with the string `OCI`. The `issue` command, sends the results of the rule to the service provider.

You can create filters with the appropriate criteria for your organization.

For information on AD FS syntax for custom rules, see the Microsoft document: [Understanding Claim Rule Language in AD FS 2.0 and Higher](#).

- iii. Click **Finish**.

### **STEP 6: SET UP IAM POLICIES FOR THE GROUPS**

If you haven't already, set up IAM policies to control the access the federated users have to your organization's Oracle Cloud Infrastructure resources. For more information, see [Getting Started with Policies](#) and [Common Policies](#).

### **STEP 7: GIVE YOUR FEDERATED USERS THE NAME OF THE TENANT AND URL TO SIGN IN**

The federated users need the URL for the Oracle Cloud Infrastructure Console (for example, <https://console.us-ashburn-1.oraclecloud.com>) and the name of your tenant. They'll be prompted to provide the tenant name when they sign in to the Console.

## **Managing Identity Providers in the Console**



### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

To add an identity provider

See [Federating with Microsoft Active Directory](#).

To delete an identity provider

All the group mappings for the identity provider will also be deleted.

1. Delete the identity provider from your tenancy:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
  - b. Click the identity provider to view its details.
  - c. Click **Delete**.
  - d. Confirm when prompted.

### To add group mappings for an identity provider

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Add at least one mapping:
  - a. Click **+ Add Mapping**.
  - b. Under **Identity Provider Group**, enter the Active Directory group name. The name you enter here must match exactly the name in Active Directory.
  - c. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
  - d. Repeat the above sub-steps for each mapping you want to create, and then click **Submit**.

Your changes take effect typically within seconds.

### To update a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Update the mappings (or click the X to delete a mapping), and then click **Submit**.

Your changes take effect typically within seconds.

### To delete a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. For the mapping you want to delete, click **Delete** next to it.
4. Confirm when prompted.

Your changes take effect typically within seconds.

### Managing Identity Providers in the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations:

#### Identity providers:

- [CreateIdentityProvider](#)
- [ListIdentityProviders](#)

- [GetIdentityProvider](#)
- [UpdateIdentityProvider](#)
- [DeleteIdentityProvider](#): Before you can use this operation, you must first use [DeleteIdpGroupMapping](#) to remove all the group mappings for the identity provider.

### Group mappings:

- [CreateIdpGroupMapping](#): Each group mapping is a separate entity with its own OCID.
- [ListIdpGroupMappings](#)
- [GetIdpGroupMapping](#)
- [UpdateIdpGroupMapping](#)
- [DeleteIdpGroupMapping](#)

## Federating with Microsoft Azure Active Directory

This topic describes how to federate with Microsoft Azure Active Directory (AD).



### Note

Before following the steps in this topic, see [Federating with Identity Providers](#) to ensure that you understand general federation concepts.

### About Federating with Azure AD

To federate with Azure AD, you set up Oracle Cloud Infrastructure as a basic SAML single sign-on application in Azure AD. To set up this application, you perform some steps in the Oracle Cloud Infrastructure Console and some steps in Azure AD.

Following is the general process an administrator goes through to set up the federation. Details for each step are given in the next section.

1. In Oracle Cloud Infrastructure, download the federation metadata document.
2. In Azure AD, set up Oracle Cloud Infrastructure Console as an enterprise application.
3. In Azure AD, configure the Oracle Cloud Infrastructure enterprise application for single sign-on.
4. In Azure AD, set up the user attributes and claims.
5. In Azure AD, download the Azure AD SAML metadata document.
6. In Azure AD, assign user groups to the application.
7. In Oracle Cloud Infrastructure, set up Azure AD as an identity provider.
8. In Oracle Cloud Infrastructure, map your Azure AD groups to Oracle Cloud Infrastructure groups.
9. In Oracle Cloud Infrastructure, set up the IAM policies to govern access for your Azure AD groups.
10. Share the Oracle Cloud Infrastructure sign-in URL with your users.

### Steps to Federate with Azure AD

#### Prerequisites

You have an Azure tenancy with groups and users set up in Azure AD.

#### STEP 1: IN ORACLE CLOUD INFRASTRUCTURE, DOWNLOAD THE FEDERATION METADATA DOCUMENT

**Summary:** The Oracle Cloud Infrastructure Console Federation page displays a link to the Oracle Cloud Infrastructure federation metadata document. Before you set up the application in Azure AD, you need to download the document.

1. Go to the **Federation** page: Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

On the Federation page, click **Download this document**.

 You need the Oracle Cloud Infrastructure Federation Metadata document when setting up a trust with Microsoft Active Directory Federation Services or with other [SAML 2.0-compliant identity providers](#). This is an XML document that describes Oracle Cloud Infrastructure endpoint and certificate information. [Download this document](#) or [Learn more](#).

- 2.

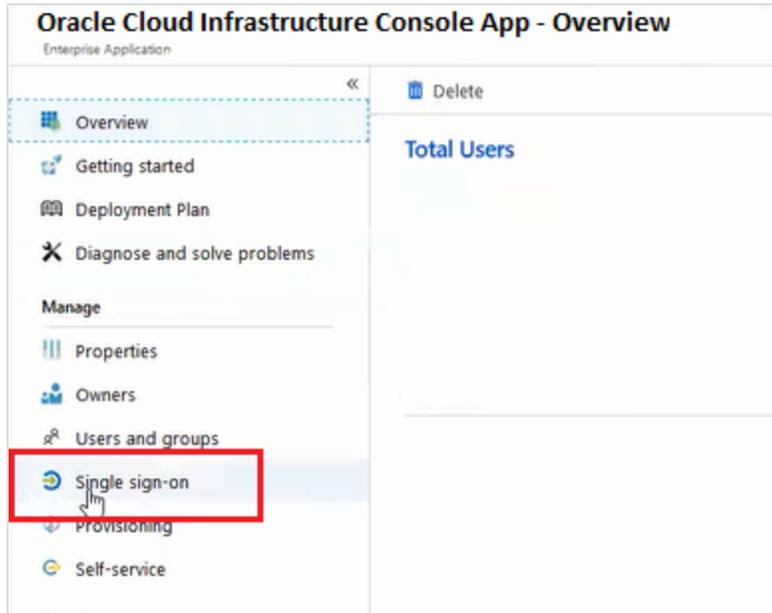
After you click the link, the metadata.xml document opens in your browser window. Use your browser's **Save page as** command to save the xml document locally where you can access it later.

### **STEP 2: IN AZURE AD, ADD ORACLE CLOUD INFRASTRUCTURE AS AN ENTERPRISE APPLICATION**

1. In the Azure portal, on the left navigation panel, select **Azure Active Directory**.
2. In the **Azure Active Directory** pane, select **Enterprise applications**. A sample of the applications in your Azure AD tenant is displayed.
3. At the top of the **All applications** pane, click **New application**.
4. In the **Add from gallery** region, enter **Oracle Cloud Infrastructure Console** in the search box.
5. Select the Oracle Cloud Infrastructure Console application from the results and select **Add**.
6. In the application-specific form, you can edit information about the application. For example, you can edit the name of the application.
7. When you are finished editing the properties, select **Add**.  
The getting started page is displayed with the options for configuring the application for your organization.

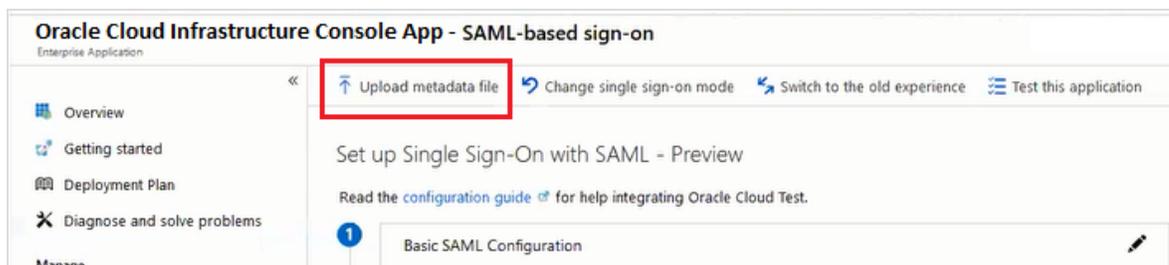
**STEP 3: IN AZURE AD, CONFIGURE ORACLE CLOUD INFRASTRUCTURE AS AN ENTERPRISE APPLICATION**

Under the **Manage** section, select **Single sign-on**.



- 1.
2. Select **SAML** to configure single sign-on. The **Set up Single Sign-On with SAML - Preview** page is displayed.

At the top of the page, click **Upload metadata file**.

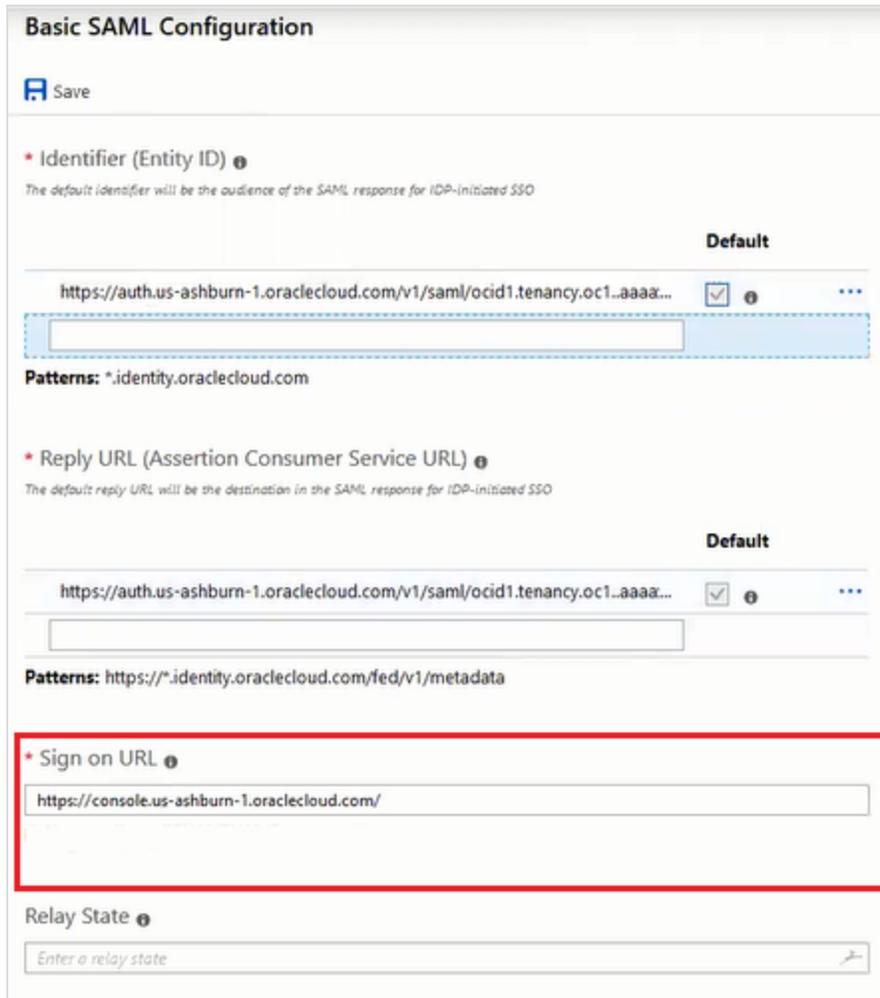


- 3.
4. Locate the federation metadata file (metadata.xml) you downloaded from Oracle Cloud

Infrastructure in Step 1, and upload it here. After you upload the file, these **Basic SAML Configuration** fields are automatically populated:

- Identifier (Entity ID)
  - Reply URL (Assertion Consumer Service URL)
5. In the **Basic SAML Configuration** section, click **Edit**. On the **Basic SAML Configuration** pane, enter the following required field:
- **Sign on URL:** Enter the URL in the following format:  
`https://console.<oci_home_region>.oraclecloud.com`  
where *oci\_home\_region* is your tenancy's home region. For example, if your home region is Ashburn, enter:  
`https://console.us-ashburn-1.oraclecloud.com`

[How do I find my home region?](#)



**Basic SAML Configuration**

[Save](#)

**\* Identifier (Entity ID)** ⓘ  
The default identifier will be the audience of the SAML response for IDP-initiated SSO

**Default**

<https://auth.us-ashburn-1.oraclecloud.com/v1/saml/ocid1.tenancy.oc1..aaaa...> ⓘ ...

**Patterns:** \*.identity.oraclecloud.com

**\* Reply URL (Assertion Consumer Service URL)** ⓘ  
The default reply URL will be the destination in the SAML response for IDP-initiated SSO

**Default**

<https://auth.us-ashburn-1.oraclecloud.com/v1/saml/ocid1.tenancy.oc1..aaaa...> ⓘ ...

**Patterns:** https://\*.identity.oraclecloud.com/fed/v1/metadata

**\* Sign on URL** ⓘ

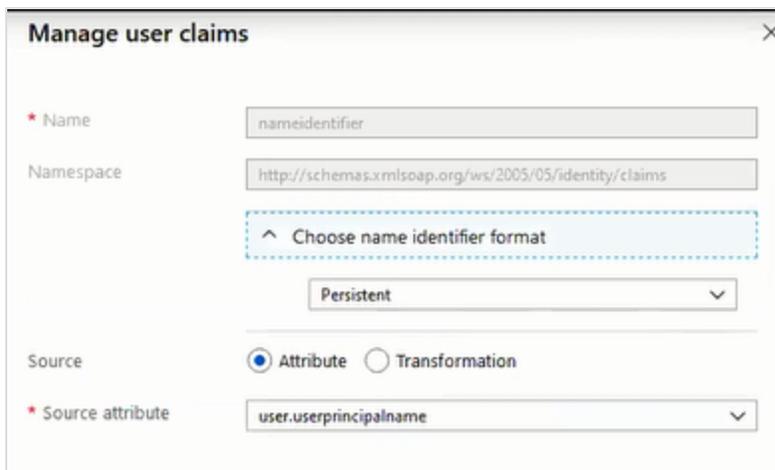
**Relay State** ⓘ

6. Click **Save**.

#### STEP 4. CONFIGURE USER ATTRIBUTES & CLAIMS

The Oracle Cloud Infrastructure Console enterprise application template is seeded with the required attributes, so you don't need to add any. However, you do need to make the following customizations:

1. In the **User Attributes & Claims** section, click **Edit** in the upper-right corner. The **Manage user claims** panel is displayed.
2. Next to the **Name identifier value** field, click **Edit**.
  - Under **Choose name identifier format**, select **Persistent**.
  - For **Source**, select **Attribute**.
  - For **Source attribute**, select **user.userprincipalname**.



The screenshot shows a 'Manage user claims' dialog box with the following fields and values:

- Name**: nameidentifier
- Namespace**: http://schemas.xmlsoap.org/ws/2005/05/identity/claims
- Choose name identifier format**: Persistent
- Source**: Attribute (selected), Transformation
- Source attribute**: user.userprincipalname

- - Click **Save**.
3. Next to the **Groups returned in claim** field, click **Edit**.
  4. In the **Group Claims (Preview)** panel, configure the following:
    - Select **Security groups**.
    - **Source attribute**: Select **Group ID**.
    - Under **Advanced Options**, select **Customize the name of the group claim**.
    - In the **Name** field, enter: **groupName**.  
Ensure that you enter **groupName** with spelling and case exactly as given.

In the **Namespace** field, enter: `https://auth.oraclecloud.com/saml/claims`

**Group Claims (Preview)**  
Manage the group claims used by Azure AD to populate SAML tokens issued to your app

Which groups associated with the user should be returned in the claim?

None  
 All groups  
 Security groups  
 Distribution lists  
 Directory roles

\* Source attribute  
Group ID

**Advanced options**

Customize the name of the group claim

Name (required)  
groupName

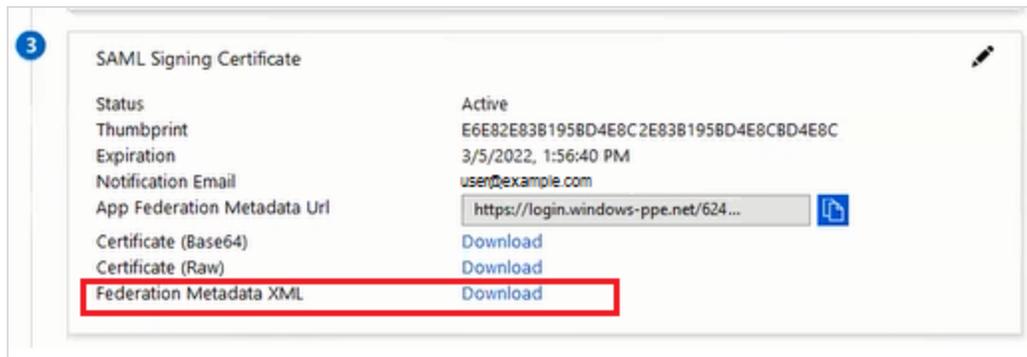
Name space (optional)  
https://auth.oraclecloud.com/saml/claims

Emit groups as role claims

- 
- Click **Save**.

**STEP 5: DOWNLOAD THE SAML METADATA DOCUMENT**

1. On the **SAML Signing Certificate** section, click the download link next to **Federation Metadata XML**.



2. Download this document and make a note of where you save it. You will upload this document to the Console in the next step.

#### STEP 6. ASSIGN USER GROUPS TO THE APPLICATION

To enable Azure AD users to sign in to Oracle Cloud Infrastructure, you need to assign the appropriate user groups to your new enterprise application.

1. On the left navigation pane, under **Manage**, select **Users and Groups**.
2. Click **Add** at the top of the **Users and Groups** list to open the **Add Assignment** pane.
3. Click the **Users and groups** selector.
4. Enter the name of the group you want to assign to the application into the **Search by name or email address** search box.
5. Hover over the group in the results list to display a check box. Select the check box to add the group to the **Selected** list.
6. When you are finished selecting groups, click **Select** to add them to the list of users and groups to be assigned to the application.
7. Click **Assign** to assign the application to the selected groups.

#### STEP 7: ADD AZURE AD AS AN IDENTITY PROVIDER IN ORACLE CLOUD INFRASTRUCTURE

**Summary:** Add the identity provider to your tenancy. You can set up the group mappings at the same time, or set them up later.

1. Go to the Console and sign in with your Oracle Cloud Infrastructure username and password.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
3. Click **Add identity provider**.
4. Enter the following:
  - a. **Display Name:** A unique name for this federation trust. This is the name federated users see when choosing which identity provider to use when signing in to the Console. The name must be unique across all identity providers you add to the tenancy. You cannot change this later.
  - b. **Description:** A friendly description.
  - c. **Type:** Select **Microsoft Active Directory Federation Services (ADFS) or SAML 2.0 compliant identity provider**.
  - d. **XML:** Upload the FederationMetadata.xml file you downloaded from Azure AD.
  - e. Click **Show Advanced Options**.
  - f. **Encrypt Assertion:** Selecting the check box lets the IAM service know to expect the encryption from IdP. Do not select this check box unless you have enabled assertion encryption in Azure AD.  
To enable assertion encryption for this single sign-on application in Azure AD, set up the SAML Signing Certificate in Azure AD to sign the SAML response and assertion. For more information, see the [Azure AD documentation](#).
  - g. **Force Authentication:** Selected by default. When selected, users are required to provide their credentials to the IdP (re-authenticate) even when they are already signed in to another session.
  - h. **Authentication Context Class References:** This field is required for Government Cloud customers. When one or more values are specified, Oracle Cloud Infrastructure (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the IdP must contain an authentication statement

with that authentication context class reference. If the SAML response authentication context does not match what is specified here, the Oracle Cloud Infrastructure auth service rejects the SAML response with a 400.

Several common authentication context class references are listed in the menu. To use a different context class, select **Custom**, then manually enter the class reference.

- i. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Continue**.



### Note

If you don't want to set up the group mappings now, you can simply click **Create** and come back to add the mappings later.

## STEP 8. ADD GROUP MAPPINGS

**Summary:** Set up the mappings between Azure AD groups and IAM groups in Oracle Cloud Infrastructure. A given Azure AD group can be mapped to zero, one, or multiple IAM groups, and vice versa. However, each individual mapping is between only a single Azure AD group and a single IAM group. Changes to group mappings take effect typically within seconds in your home region, but may take several minutes to propagate to all regions. Note that the Azure AD groups that you choose to map must also be assigned to the enterprise application in Azure AD. See [Step 6. Assign user groups to the application](#).

**Before you begin:** Have your Azure AD groups page open. From the Azure Dashboard, under **Manage**, select **Groups**. From the list of groups, select the group you want to map to an Oracle Cloud Infrastructure group. In the group's details page, click the **Copy** icon next to the Object ID for the group.

To create a group mapping:

For **Identity Provider Group**, select **Custom Group**. Enter (or paste) the Object ID of the Azure AD group. You must enter the Object ID exactly, including the correct case. An example Object ID looks like: aa0e7d64-5b2c-623g-at32-65058526179c

The screenshot shows a web interface titled "Edit Identity Provider" with a "cancel" link in the top right. Below the title is a descriptive paragraph: "Here you'll map groups defined in your Identity Provider to groups defined in Oracle Cloud Infrastructure (OCI). Each group can be mapped to one or more groups of the other kind." A section titled "Mapping 1" contains two input fields: "IDENTITY PROVIDER GROUP" with the value "Custom Group" and "OCI GROUP" with the value "Administrators". Below these is a text input field containing the Object ID "aa0e7d64-5b2c-623g-at32-65058526179c". At the bottom of the mapping section are buttons for "+ Add Mapping" and "Submit".

- 1.
2. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
3. Repeat the preceding steps for each mapping you want to create, and then click **Create**.



### Tip

Requirements for IAM group name: No spaces. Allowed characters: letters, numerals, hyphens, periods, underscores, and plus signs (+). The name cannot be changed later.

The identity provider is now added to your tenancy and appears in the list on the **Federation** page. Click the identity provider to view its details and the group mappings you just set up.

Oracle assigns the identity provider and each group mapping a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

In the future, come to the **Federation** page if you want to edit the group mappings or delete the identity provider from your tenancy.

### **STEP 9: SET UP IAM POLICIES FOR THE GROUPS**

If you haven't already, set up IAM policies to control the access the federated users have to your organization's Oracle Cloud Infrastructure resources. For more information, see [Getting Started with Policies](#) and [Common Policies](#).

### **STEP 10: GIVE YOUR FEDERATED USERS THE NAME OF THE TENANT AND URL TO SIGN IN**

The federated users need the URL for the Oracle Cloud Infrastructure Console (for example, <https://console.us-ashburn-1.oraclecloud.com>) and the name of your tenant. They'll be prompted to provide the tenant name when they sign in to the Console.

## **Managing Identity Providers in the Console**

### To add an identity provider

See [Federating with Microsoft Azure Active Directory](#).

### To delete an identity provider

All the group mappings for the identity provider will also be deleted.

1. Delete the identity provider from your tenancy:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.

- b. Click the identity provider to view its details.
- c. Click **Delete**.
- d. Confirm when prompted.

### To add group mappings for an identity provider

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Add at least one mapping:
  - a. Click **+ Add Mapping**.
  - b. Under **Identity Provider Group**, select **Custom Group**. Enter (or paste) the Object ID of the Azure AD group. You must enter the Object ID exactly, including the correct case. An example Object ID looks like: aa0e7d64-5b2c-623g-at32-65058526179c. Note that for groups to be able to sign in to Oracle Cloud Infrastructure, they must also be assigned to the enterprise application in Azure AD. See [Step 6. Assign user groups to the application](#).
  - c. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
  - d. Repeat the preceding steps for each mapping you want to create, and then click **Submit**.

Your changes take effect typically within seconds.

### To update a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Update the mappings (or click the X to delete a mapping), and then click **Submit**.

Your changes take effect typically within seconds.

### To delete a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. For the mapping you want to delete, click **Delete** next to it.
4. Confirm when prompted.

Your changes take effect typically within seconds.

### Managing Identity Providers in the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations:

#### Identity providers:

- [CreateIdentityProvider](#)
- [ListIdentityProviders](#)

- [GetIdentityProvider](#)
- [UpdateIdentityProvider](#)
- [DeleteIdentityProvider](#): Before you can use this operation, you must first use [DeleteIdpGroupMapping](#) to remove all the group mappings for the identity provider.

### Group mappings:

- [CreateIdpGroupMapping](#): Each group mapping is a separate entity with its own OCID.
- [ListIdpGroupMappings](#)
- [GetIdpGroupMapping](#)
- [UpdateIdpGroupMapping](#)
- [DeleteIdpGroupMapping](#)

## Federating with SAML 2.0 Identity Providers

This topic describes the general steps to federate Oracle Cloud Infrastructure with any identity provider that supports the Security Assertion Markup Language (SAML) 2.0 protocol. If you want specific instructions for Oracle Identity Cloud Service or Microsoft Active Directory, see [Federating with Oracle Identity Cloud Service](#) or [Federating with Microsoft Active Directory](#).



### Tip

Find detailed setup steps for more IdPs in the following white papers:

- [Oracle Cloud Infrastructure Okta Configuration for Federation and Provisioning](#)
- [Federating Oracle Access Manager to Oracle Cloud Infrastructure](#)

### Instructions for Federating

Following is the general process an administrator goes through to set up the identity provider, and below are instructions for each step. It's assumed that the administrator is an Oracle Cloud Infrastructure user with the required credentials and access.



#### Note

Before following the steps in this topic, see [Federating with Identity Providers](#) to ensure that you understand general federation concepts.

1. In the Oracle Cloud Infrastructure Console, get the federation metadata required to establish a trust relationship with the Identity Provider (IdP).
2. In the IdP, configure Oracle Cloud Infrastructure as an application (sometimes called a trusted relying party).
3. In the IdP, assign users and groups to your new Oracle Cloud Infrastructure application.
4. In the IdP, get the required information needed by Oracle Cloud Infrastructure.
5. In Oracle Cloud Infrastructure:
  - a. Add the identity provider to your tenancy and provide information you got from the IdP.
  - b. Map the IdP's groups to IAM groups.
6. In Oracle Cloud Infrastructure, make sure you have IAM policies set up for the groups so you can control users' access to Oracle Cloud Infrastructure resources.
7. Inform your users of the name of your Oracle Cloud Infrastructure tenant and the URL for the Console (for example, <https://console.us-ashburn-1.oraclecloud.com>).

### Step 1: Get information from Oracle Cloud Infrastructure

Summary: Download the federation metadata document.

The federation metadata document is a standard SAML 2.0 document, which provides information about Oracle Cloud Infrastructure you'll need to provide to your IdP. Depending

on your provider's setup requirements, you may need to upload the entire document, or you may be asked to provide only specific metadata values from the document.

1. Sign in to the Oracle Cloud Infrastructure Console as an administrator.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
3. Right-click the **Download this document** link and save the document.

### **Step 2: Set up Oracle Cloud Infrastructure as a trusted application**

Consult your IdP documentation for how to set up a trusted application. Refer to the metadata document you downloaded for required parameters.

### **Step 3: Assign users and groups to the new application.**

Follow your IdP's procedures for adding users and groups to the application you set up for Oracle Cloud Infrastructure.

### **Step 4: Download the IdP's metadata document.**

Your IdP should provide a SAML 2.0 document that contains the information Oracle Cloud Infrastructure needs to complete the federation. See your IdP documentation for instructions on downloading this document.

### **Step 5: Federate the IdP with Oracle Cloud Infrastructure**

**Summary:** Add the identity provider to your tenancy. You can set up the group mappings at the same time, or set them up later.

#### Details:

1. Go to the [Console](#) and sign in with your Oracle Cloud Infrastructure login and password.
2. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.
3. Click **Add Identity Provider**.
4. Enter the following:

- a. **Name:** A unique name for this federation trust. This is the name federated users see when choosing which identity provider to use when signing in to the Console, so consider making this a friendly, intuitive name your users will understand. The name must be unique across all identity providers you add to the tenancy. You cannot change this later.
- b. **Description:** A friendly description.
- c. **Type:** Select **Microsoft Active Directory Federation Service (ADFS) or SAML 2.0 Compliant Identity Provider**.
- d. **XML:** Upload the metadata.xml document that you downloaded from your IdP.
- e. **Encrypt Assertion:** Selecting the check box lets the IAM service know to expect the encryption from the IdP. If you select this check box, you must also set up encryption of the assertion in your IdP. For more information, see [Encrypt Assertion](#). See also your IdP's documentation.
- f. **Force Authentication:** Selected by default. When selected, users are required to provide their credentials to the IdP (re-authenticate) even when they are already signed in to another session.
- g. **Authentication Context Class References:** This field is required for Government Cloud customers. When one or more values are specified, Oracle Cloud Infrastructure (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the IdP must contain an authentication statement with that authentication context class reference. If the SAML response authentication context does not match what is specified here, the Oracle Cloud Infrastructure auth service rejects the SAML response with a 400. Several common authentication context class references are listed in the menu. To use a different context class, select **Custom**, then manually enter the class reference.
- h. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information

about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Continue**.
6. Set up the mappings between the IdP groups and IAM groups in Oracle Cloud Infrastructure. A given IdP group can be mapped to zero, one, or multiple IAM groups, and vice versa. However, each individual mapping is between only a single IdP group and a single IAM group. Changes to group mappings take effect typically within seconds in your home region, but may take several minutes to propagate to all regions.



### Note

If you don't want to set up the group mappings now, you can simply click **Create** and come back to add the mappings later.

To create a group mapping:

- a. Under **Identity Provider Group**, enter the name of the group in your IdP. You must enter the name exactly, including the correct case.  
Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".



### Tip

Requirements for IAM group name: No spaces. Allowed characters: letters, numerals, hyphens, periods, underscores, and plus signs (+). The name cannot be changed later.

- b. Repeat the above sub-steps for each mapping you want to create, and then click **Create**.

The identity provider is now added to your tenancy and appears in the list on the **Federation** page. Click the identity provider to view its details and the group mappings you just set up.

Oracle assigns the identity provider and each group mapping a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

In the future, come to the **Federation** page if you want to edit or add group mappings or delete the identity provider from your tenancy.

### Step 6: Set up IAM policies for the groups

If you haven't already, set up IAM policies to control the access the federated users have to your organization's Oracle Cloud Infrastructure resources. For more information, see [Getting Started with Policies](#) and [Common Policies](#).

### Step 7: Give your federated users the name of the tenant and URL to sign in

The federated users need the URL for the Oracle Cloud Infrastructure Console (for example, <https://console.us-ashburn-1.oraclecloud.com>) and the name of your tenant. They'll be prompted to provide the tenant name when they sign in to the Console.

## Managing Identity Providers in the Console



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To add an identity provider

See [Instructions for Federating](#).

### To delete an identity provider

All the group mappings for the identity provider will also be deleted.

1. Delete the identity provider from your tenancy:
  - a. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
  - b. Click the identity provider to view its details.
  - c. Click **Delete**.
  - d. Confirm when prompted.
2. Follow your IdP's documentation to delete the application from your IdP.

### To add group mappings for an identity provider

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

A list of the identity providers in your tenancy is displayed.

2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. Add at least one mapping:
  - a. Click **+ Add Mapping**.
  - b. Enter the IdP group name exactly in the **Identity Provider Group** text box.
  - c. Choose the IAM group you want to map from the list under **OCI Group**. If you instead want to create a new IAM group, select **New OCI Group** and enter the name of the new group in **New OCI Group Name**. The new group is automatically created in IAM and mapped to the IdP group. It will also automatically be given this description, which you can't change: "Group created during federation".
  - d. Repeat the above sub-steps for each mapping you want to create, and then click **Submit**.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions

### To update a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.

A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. Click **Edit Mapping**.
4. When prompted, provide the client ID and client secret for the Oracle Identity Cloud Service application, and then click **Continue**.
5. Update the mappings (or click the X to delete a mapping), and then click **Submit**.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions

### To delete a group mapping

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider to view its details.
3. For the mapping you want to delete, click **Delete** next to it.
4. Confirm when prompted.

Your changes take effect typically within seconds in your home region. Wait several more minutes for changes to propagate to all regions.

### Managing Identity Providers in the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations:

#### Identity providers:

- [CreateIdentityProvider](#)
- [ListIdentityProviders](#)
- [GetIdentityProvider](#)
- [UpdateIdentityProvider](#)
- [DeleteIdentityProvider](#): Before you can use this operation, you must first use [DeleteIdpGroupMapping](#) to remove all the group mappings for the identity provider.

#### Group mappings:

- [CreateIdpGroupMapping](#): Each group mapping is a separate entity with its own OCID.
- [ListIdpGroupMappings](#)
- [GetIdpGroupMapping](#)
- [UpdateIdpGroupMapping](#)
- [DeleteIdpGroupMapping](#)

## User Provisioning for Federated Users

This topic describes how you can use SCIM to provision federated users in Oracle Cloud Infrastructure. Provisioned federated users can have API keys and other service-specific credentials.

### Overview

[SCIM \(System for Cross-domain Identity Management\)](#) is an IETF standard protocol that enables user provisioning across identity systems. Oracle Cloud Infrastructure hosts a SCIM endpoint for provisioning federated users into Oracle Cloud Infrastructure. If your IdP is Oracle Identity Cloud Service or Okta, you can set up SCIM user provisioning.

After you configure the SCIM integration between your IdP and Oracle Cloud Infrastructure, users that belong to groups mapped to Oracle Cloud Infrastructure groups are automatically provisioned in Oracle Cloud Infrastructure. Provisioned users are assigned a unique OCID, and can have API keys and other service-specific credentials.

The following functionality is supported for provisioned, federated users:

- Provisioned users are assigned a unique OCID
- Provisioned users can have API keys, auth tokens, and other service-specific credentials
- You can list the users in the Console
- Provisioned users can access the **User Settings** page to see and manage these credentials for themselves

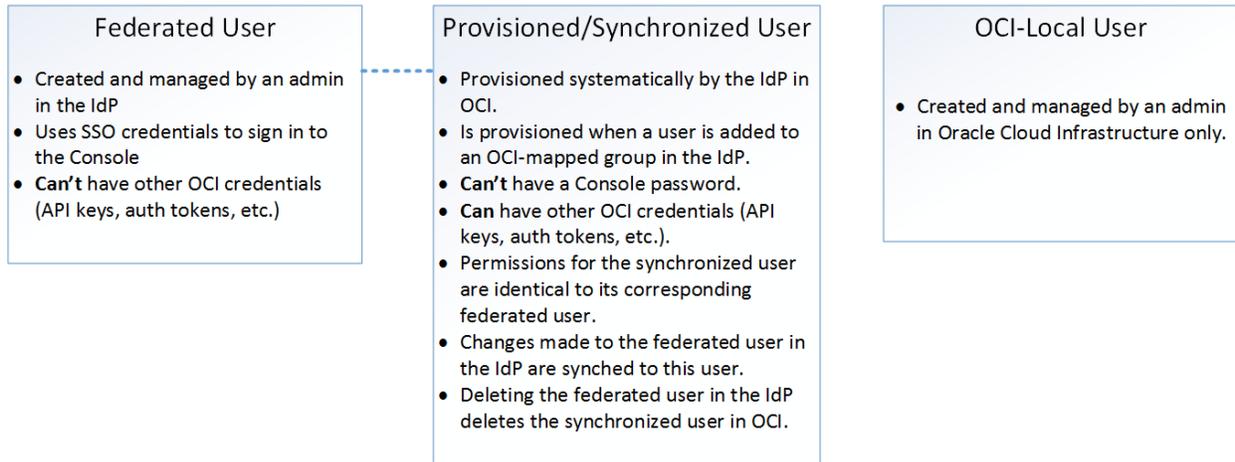
- When you add or remove users to Oracle Cloud Infrastructure-mapped groups in your IdP, the updates are automatically synched with Oracle Cloud Infrastructure

### Understanding User Types

The SCIM configuration introduces the concept of the *provisioned* or *synchronized* user. The following descriptions provide details to help you understand the user types you'll be managing.

- Federated users  
A federated user is created and managed in an identity provider. Federated users can sign in to the Console using a password managed in their identity provider. Federated users are granted access to Oracle Cloud Infrastructure based on their membership in groups that are mapped to Oracle Cloud Infrastructure groups.
- Provisioned (or Synchronized) users  
A synchronized user is systematically provisioned by the identity provider in Oracle Cloud Infrastructure. Synchronized users can have Oracle Cloud Infrastructure credentials, but not Console passwords. When listing users in the Console, you can identify synchronized users using the **User Type** filter.
- Local users  
A local user is a user created and managed in Oracle Cloud Infrastructure's IAM service. Federated tenancies typically would have few, if any, local users. When listing users in the Console, you can identify local users using the **User Type** filter.

The following graphic summarizes the characteristics of the user types:



### Who Should Set Up This Integration?

Set up this integration if your IdP is Oracle Identity Cloud Service or Okta and your federated users need to have the specialized credentials required by some services and features. For example, if you need your federated users to access Oracle Cloud Infrastructure through the SDK or CLI, setting up this integration enables these users to get the API keys needed for this access.

### Prerequisite

Perform this synchronization setup after you have successfully set up a federation between your IdP and Oracle Cloud Infrastructure. See [Supported Identity Providers](#).

### Enabling User Provisioning

#### Instructions for Oracle Identity Cloud Service Federations

If your identity provider is Oracle Identity Cloud Service, you need to perform a one-time upgrade.



### Important

If your tenancy was created December 21, 2018 or later, your tenancy is automatically configured to provision your Oracle Identity Cloud Service users in Oracle Cloud Infrastructure. You do not need to perform the steps in this topic. See [Understanding User Types](#) and [Managing User Capabilities for Federated Users](#) for information on managing your federated users.

## Upgrading Your Oracle Identity Cloud Service Federation

If your federation with Oracle Identity Cloud service was set up before December 21, 2018, perform this one-time upgrade task.

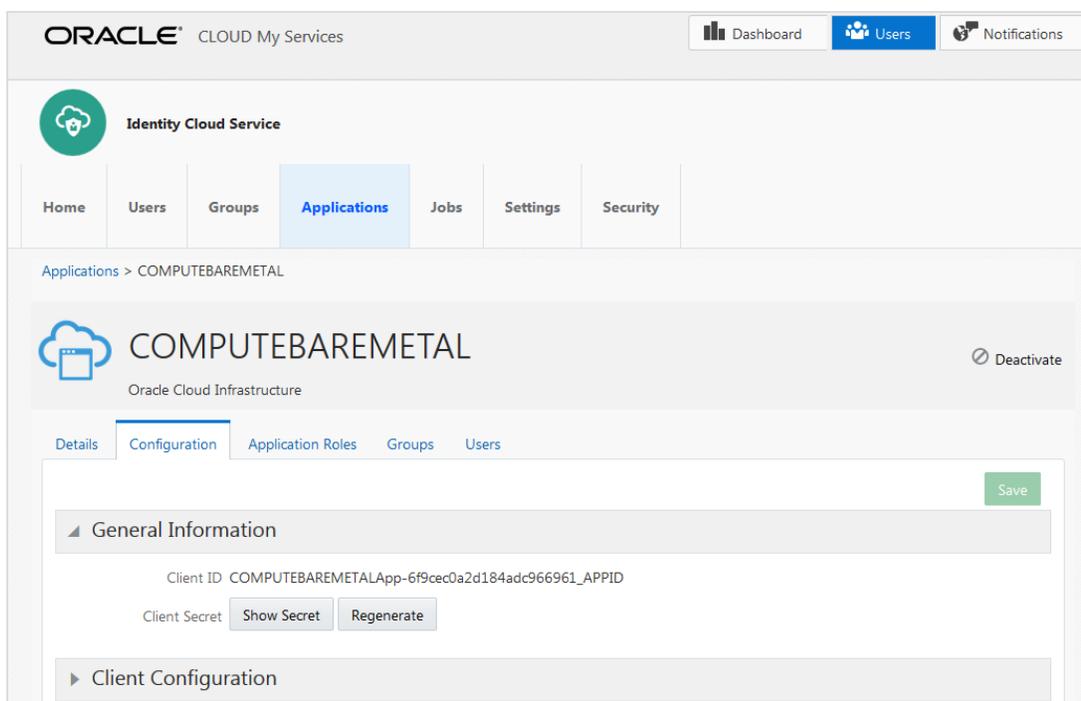
### To upgrade your Oracle Identity Cloud Service federation:

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Federation**.  
A list of the identity providers in your tenancy is displayed.
2. Click your Identity Cloud Service federation to view its details. If your tenancy was auto-federated, it is listed as OracleIdentityCloudService.
3. Click **Edit Mapping**.
4. When prompted, provide the client ID and client secret for the Oracle Identity Cloud Service application, and then click **Continue**.

### Where do I find the client ID and client secret?

The client ID and client secret are stored in Oracle Identity Cloud Service. To get this information:

- a. Sign in to the Oracle Identity Cloud Service console.
- b. In the Identity Cloud Service console, click **Applications**. The list of trusted applications is displayed.
- c. Click COMPUTEBAREMETAL.
- d. Click **Configuration**.
- e. Expand **General Information**. The client ID is displayed. Click **Show Secret** to display the client secret.



Allow several minutes for the changes to take effect.

### Instructions for Okta Federations

If you do not have an existing federation with Okta, follow the instructions in the white paper, [Oracle Cloud Infrastructure Okta Configuration for Federation and Provisioning](#). This paper includes instructions for both setting up your federation and provisioning with SCIM.

If you have an existing federation with Okta with group mappings that you want to maintain, you can add SCIM provisioning as follows:

1. In Okta, delete the existing SAML application you originally set up to federate with Oracle Cloud Infrastructure.
2. Set up a new SAML application in Okta according to the instructions in the white paper, [Oracle Cloud Infrastructure Okta Configuration for Federation and Provisioning](#), with the following exceptions:
  - Skip the steps to **Add Identity Provider** to Oracle Cloud Infrastructure (you already have this resource in Oracle Cloud Infrastructure).
  - Instead, click **Edit Identity Provider** and upload the new metadata.xml document from the new Okta app you created.
  - Then, in Oracle Cloud Infrastructure, ensure that you **Reset Credentials**. Add the new Client ID and Secret to the API integration settings page in Okta (Step 7 in the white paper).

### What to Expect After the Upgrade

When the system has had time to synchronize, you can manage user capabilities for federated users in the Console. Users that belong to a group mapped to a group in Oracle Cloud Infrastructure are listed on the Users page in the Console. Whenever you add new users to mapped groups in Oracle Identity Cloud Service, they will be available in the Console after the system synchronizes.

By default, the following user capabilities are enabled:

- API keys
- auth tokens

- SMTP credentials
- customer secret keys

Notice that you can't enable a local password. The Oracle Cloud Infrastructure console password is still managed only in your IdP.

For more information about user capabilities, see [Managing User Capabilities for Federated Users](#).

### Resetting Credentials

Use the **Reset Credentials** button to reset your SCIM client credentials. You can perform this task periodically as a security measure to rotate your credentials. After you reset these credentials, you'll need to update the SAML app in your identity provider with the new credentials.

**Note:** If your IdP is Oracle Identity Cloud Service, Oracle Cloud Infrastructure automatically resets the credentials with Oracle Identity Cloud Service for you. You don't need to manually reset the configuration.

### Actions You Still Perform in Your Identity Provider

After the integration is set up, continue to perform the following actions in your IdP:

- Create users and assign them to groups.
- Delete users.  
Users that you delete from your IdP are removed from Oracle Cloud Infrastructure when the next syncing cycle completes.
- Query for group membership.
- Manage sign-in passwords for users.

## Managing User Capabilities for Federated Users

This topic describes managing user capabilities for federated users when your tenancy is federated and configured for user provisioning with a [supported identity provider](#).

### About User Capabilities

To access Oracle Cloud Infrastructure, a user must have the required credentials. Users who need to use the Console, must have a password. Users who need access through the API need API keys. Some service features require additional credentials, such as auth tokens, SMTP credentials, and Amazon S3 Compatibility API keys. For a user to get these credentials, the user must be granted the capability to have the credential type.

User capabilities are managed by an Administrator in the user's details. Each user can see their capabilities, but only an Administrator can enable or disable them. The user capabilities available to federated users are:

- API keys
- auth tokens
- SMTP credentials
- customer secret keys

By default, these capabilities are enabled when you provision new users, allowing users to create these credentials for themselves. For information about these user credentials, see [Managing User Credentials](#).



#### **Important**

The capability "Console password" is not available for federated users. Federated users authenticate to the Console through their IdP, where their sign-in passwords are managed.

### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing user capabilities. A user can't enable or disable user capabilities for themselves (except for

Administrators). However, a user can manage their own credentials that have been enabled for them.

### Prerequisites

Management of user capabilities for federated users is supported for Oracle Identity Cloud Service and Okta federations only.

- Oracle Identity Cloud Service federations:  
If your tenancy was created December 21, 2018 or later, your tenancy is automatically configured to manage user capabilities. There are no prerequisites.  
If your tenancy was created before December 21, 2018, you must perform a one-time upgrade. See [Instructions for Oracle Identity Cloud Service Federations](#).
- If your tenancy is federated with Okta, see [User Provisioning for Federated Users](#).

### Viewing Provisioned Federated Users in the Console

After the prerequisites are satisfied, you can view users that you create in your IdP that belong to groups mapped to Oracle Cloud Infrastructure groups. Whenever you add a user to a group mapped to an Oracle Cloud Infrastructure group, the user automatically displays in the Console.

#### To list users in the Console:

Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.

Notice that you can filter the list by user type to include only users that belong to a specified identity provider. **Local Users** are users created in Oracle Cloud Infrastructure's IAM service. The filter list includes all identity providers you have set up.

### Using the Console



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To edit user capabilities

If you're an Administrator, you can edit user capabilities.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user to see its details.
3. Click **Edit User Capabilities**.
4. Select or clear the check box to add or remove a capability.
5. Click **Save**.

### To change a user's description

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user you want to update.  
The user's details are displayed. The description is displayed under the user's login.
3. Click the pencil next to the description.

4. Edit the description and save it. This description is maintained in Oracle Cloud Infrastructure and is not synched back to your identity provider.

### To apply tags to a user

For instructions, see [Resource Tags](#).

### To delete a user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Find the user you want to delete and click the Actions icon (three dots).
3. Click **Delete**.

**Important:** Deleting a user here does not delete the user in your IdP. If you later want the federated user to have a provisioned user in Oracle Cloud Infrastructure, you must remove the user from all OCI-mapped groups in Oracle Identity Cloud Service and re-add the user.

For information about managing user credentials in the Console, see [Managing User Credentials](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage user capabilities:

- [ListUsers](#)
- [GetUser](#)
- [UpdateUser](#): You can update the user capabilities and the user's description.

- [UpdateUserCapabilities](#)
- [DeleteUser](#): This operation deletes the provisioned user in Oracle Cloud Infrastructure, but not the user in the identity provider.

For information about the API operations for managing user credentials, see [Managing User Credentials](#).

The following operations are *not* supported for federated users:

- [ListUserGroupMemberships](#)
- [AddUserToGroup](#)
- [GetUserGroupMembership](#)
- [RemoveUserFromGroup](#)

## Calling Services from an Instance

This topic describes how you can authorize instances to call services in Oracle Cloud Infrastructure.

### Introduction

This procedure describes how you can authorize an instance to make API calls in Oracle Cloud Infrastructure services. After you set up the required resources and policies, an application running on an instance can call Oracle Cloud Infrastructure public services, removing the need to configure user credentials or a configuration file.

### Concepts

#### **DYNAMIC GROUP**

Dynamic groups allow you to group Oracle Cloud Infrastructure instances as principal actors, similar to user groups. You can then create policies to permit instances in these groups to make API calls against Oracle Cloud Infrastructure services. Membership in the group is determined by a set of criteria you define, called *matching rules*.

### **MATCHING RULE**

When you set up a dynamic group, you also define the rules for membership in the group. Resources that match the rule criteria are members of the dynamic group. Matching rules have a specific syntax you follow. See [Writing Matching Rules to Define Dynamic Groups](#).

### **INSTANCE PRINCIPALS**

The IAM service feature that enables instances to be authorized actors (or principals) to perform actions on service resources. Each compute instance has its own identity, and it authenticates using the certificates that are added to it. These certificates are automatically created, assigned to instances and rotated, preventing the need for you to distribute credentials to your hosts and rotate them.

## Security Considerations

Any user who has access to the instance (who can SSH to the instance), automatically inherits the privileges granted to the instance. Before you grant permissions to an instance using this procedure, ensure that you know who can access it, and that they should be authorized with the permissions you are granting to the instance.

## Process Overview

The following steps summarize the process flow for setting up and using instances as principals. The subsequent sections provide more details.

1. Create a [dynamic group](#). In the dynamic group definition, you provide the matching rules to specify which instances you want to allow to make API calls against services.
2. Create a policy granting permissions to the dynamic group to access services in your tenancy (or compartment).
3. A developer in your organization configures the application built using the Oracle Cloud Infrastructure SDK to authenticate using the instance principals provider. The developer deploys the application and the SDK to all the instances that belong to the dynamic group.

4. The deployed SDK makes calls to Oracle Cloud Infrastructure APIs as allowed by the policy (without needing to configure API credentials).
5. For each API call made by an instance, the [Audit service](#) logs the event, recording the OCID of the instance as the value of `principalId` in the event log.

### Steps to Enable Instances to Call Services

Perform these tasks to enable an instance to call services:

[Create a Dynamic Group and Matching Rules](#)

[Write Policies for Dynamic Groups](#)

[Configure the SDK, CLI, or Terraform](#)

#### Creating a Dynamic Group and Matching Rules

See [Managing Dynamic Groups](#).

#### Writing Policies for Dynamic Groups

After you have created a dynamic group, you need to create policies to permit the dynamic groups to access Oracle Cloud Infrastructure services.

Policy for dynamic groups follows the syntax described in [How Policies Work](#). Review that topic to understand basic policy features.

The syntax to permit a dynamic group access to resources in a compartment is:

```
Allow dynamic-group <dynamic_group_name> to <verb> <resource-type> in compartment <compartment_name>
```

The syntax to permit a dynamic group access to a tenancy is:

```
Allow dynamic-group <dynamic_group_name> to <verb> <resource-type> in tenancy
```

Here are a few example policies:

To allow a dynamic group (FrontEnd) to use a load balancer in a specific compartment (ProjectA):

```
Allow dynamic-group FrontEnd to use load-balancers in compartment ProjectA
```

To allow a dynamic group to launch instances in a specific compartment:

```
Allow dynamic-group FrontEnd to manage instance-family in compartment ProjectA
Allow dynamic-group FrontEnd to use volume-family in compartment ProjectA
Allow dynamic-group FrontEnd to use virtual-network-family in compartment ProjectA
```

For more sample policies, see [Common Policies](#).

### Configuring the SDK, CLI, or Terraform

For information about SDKs, see [Software Development Kits and Command Line Interface](#).

#### For the SDK for Java:

In your SDK for Java, create an `InstancePrincipalsAuthenticationDetailsProvider` object. For example:

```
public static void main(String[] args) throws Exception {
 InstancePrincipalsAuthenticationDetailsProvider provider =
 InstancePrincipalsAuthenticationDetailsProvider.builder().build();
 IdentityClient identityClient = new IdentityClient(provider);
 ...
}
```

#### For the Python SDK:

In your Python SDK, create an `oci.auth.signers.InstancePrincipalsSecurityTokenSigner` object. For example:

```
By default this will hit the auth service in the region returned by
http://169.254.169.254/opc/v1/instance/region on the instance.

signer = oci.auth.signers.InstancePrincipalsSecurityTokenSigner()
identity_client = oci.identity.IdentityClient(config={}, signer=signer)
...

```

To refresh the token without waiting, use the following command:

```
signer.refresh_security_token()
```

### Enabling Instance Principal Authorization for the CLI

To enable instance principal authorization from the CLI, you can set the authorization option (`--auth`) for a command. For example:

```
oci os ns get --auth instance_principal
```

Alternatively, you can set the following environment variable:

```
OCI_CLI_AUTH=instance_principal
```

Note that if both are set, the value set for `--auth` takes precedence over the environment variable.

For information about using the CLI, see [Getting Started with the Command Line Interface](#).

### Enabling Instance Principal Authorization for Terraform

To enable instance principal authorization in Terraform, you can set the `auth` attribute to "InstancePrincipal" in the provider definition as shown in the following sample:

```
variable "region" {}

provider "oci" {
 auth = "InstancePrincipal"
 region = "${var.region}"
}
```

Note that when you use instance principal authorization you do not need to include the `tenancy_ocid`, `user_ocid`, `fingerprint`, and `private_key_path` attributes.

## FAQs

How do I query the instance metadata service to query the certificate on the instance?

Use this curl command: `curl http://169.254.169.254/opc/v1/identity/cert.pem`

### How frequently is the certificate rotated on each instance?

The certificate is rotated multiple times each day.

### What happens if I try to use an expired certificate?

You will get a 401-Not Authenticated error.

### Can I change the frequency at which the certificate is rotated?

No. You can't change the frequency at which the certificate is rotated. However, you can change the policy on the dynamic group. If you think an instance has been compromised, you can either change the policy on the dynamic group to revoke permissions for all members of the group, or you can remove the instance from the dynamic group. See [Can I remove an instance from a dynamic group?](#)

### What happens if the certificate is rotated in the middle of a long running operation?

The token expiration is independent of the certificate expiration period. And, it also depends on the application you are interacting with. For example, if Object Storage does not have a multipart PUT operation, then it does not matter how long the operation runs.

### Are the certificates accessible for all users on an instance?

Yes. Ensure that only users who should be granted the access that you have granted to the dynamic group, have access to the instance.

### Are dynamic groups created at the tenancy level?

Yes.

### Can I remove an instance from a dynamic group?

Yes. You can remove it by modifying the matching rule to exclude it. See below for an example.

### Can I exclude specific instances in a compartment from the dynamic group?

Yes. For example, assume you want to exclude two specific instances in a compartment from the dynamic group. Write a matching rule like this:

```
All {instance.compartment.id = '<compartment_ocid>',
 instance.id != '<instance1_to_exclude_ocid>', instance.id != '<instance2_to_exclude_ocid>'}
```

The above rule includes all instances in the compartment except those with the OCIDs specified.

## Managing Users

This topic describes the basics of working with users.



### Important

If your tenancy is federated with Oracle Identity Cloud Service, see [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#) to manage users.

### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing users.

You can create a policy that gives someone power to create new users and credentials, but not control which groups those users are in. See [Let the Help Desk manage users](#).

For the reverse: You can create a policy that gives someone power to determine what groups users are in, but not create or delete users. See [Let group admins manage group membership](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for users or other IAM components, see [Details for IAM](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Working with Users

When creating a user, you must provide a unique, unchangeable *name* for the user. The name must be unique across all users within your tenancy. It will be the user's login to the Console. You might want to use a name that's already in use by your company's own identity system (e.g., Active Directory, LDAP, etc.). You must also provide the user with a *description* (although it can be an empty string), which is a non-unique, changeable description for the user. This could be the user's full name, a nickname, or other descriptive information. Oracle will also assign the user a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).



#### Note

If you delete a user and then create a new user with the same name, they'll be considered different users because they'll have different OCIDs.

Oracle recommends that you supply a password recovery email address for the user. If the user forgets their password, they can request to have a temporary password sent to them using the **Forgot Password** link on the sign-on page. If no email address is present for the user, an administrator must intervene to reset their password.

A new user has no permissions until you place the user in one or more groups, and there's at least one policy that gives that group permission to either the tenancy or a compartment. Exception: each user can manage *their own* credentials they have been enabled to have. An administrator does not need to create a policy to give a user that ability. For more information, see [User Credentials](#).



### **Important**

After creating a new user and putting them in a group, make sure to let them know which compartment(s) they have access to.

You also need to give the new user some credentials so they can access Oracle Cloud Infrastructure. A user can have one or both of the following credentials, depending on the type of access they need: A password for using the Console, and an API signing key for using the API.

## About User Capabilities

To access Oracle Cloud Infrastructure, a user must have the required credentials. Users who need to use the Console, must have a password. Users who need access through the API need API keys. Some service features require additional credentials, such as auth tokens, SMTP credentials, and Amazon S3 Compatibility API keys. For a user to get these credentials, the user must be granted the capability to have the credential type.

User capabilities are managed by an Administrator in the User details. Each user can see their capabilities, but only an Administrator can enable or disable them. The user capabilities are:

- Can use Console password (native users only)
- Can use API keys
- Can use auth tokens

- Can use SMTP credentials
- Can use customer secret keys

By default, all these capabilities are enabled when you create new users, allowing users to create these credentials for themselves. For information about working with user credentials, see [Managing User Credentials](#).

### Enabling Multi-Factor Authentication for a User

See [Managing Multi-Factor Authentication](#) for details.

### Unblocking a User After Unsuccessful Sign-in Attempts

If a user tries 10 times in a row to sign in to the Console unsuccessfully, they will be automatically blocked from further sign-in attempts. An administrator can unblock the user in the Console (see [To unblock a user](#)) or with the [UpdateUserState](#) API operation.

### Deleting a User

You can delete a user, but only if the user is not a member of any groups.

### Limits on Users

For information about the number of users you can have, see [Service Limits](#).

## Using the Console



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To create a user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click **Create User**.
3. Enter the following:
  - **Name:** A unique name or email address for the user (for tips on what value to use, see [Working with Users](#)). The name must be unique across all users in your tenancy. You cannot change this later. The name must meet the following requirements: No spaces. Only Basic Latin letters (ASCII), numerals, hyphens, periods, underscores, +, and @.
  - **Description:** This could be the user's full name, a nickname, or other descriptive information. You can change this later if you want to.
  - **Email:** Enter an email address for the user. This email address is used for password recovery. The email address must be unique in the tenancy.  
If the user forgets their password, they can click **Forgot Password** on the sign on page, and a temporary password will be automatically generated and sent to the email address provided here. The email address can also be updated later by the user or an administrator.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

4. Click **Create**.

Next, you need to give the user permissions by adding them to at least one group. You also need to give the user the credentials they need (see [Managing User Credentials](#)).

### To add a user to a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Locate the user in the list.
3. Click the user.  
Its details are displayed.
4. Click **Groups**.
5. Click **Add User to Group**.
6. Select the group from the drop-down list, and then click **Add**.

Make sure to let the user know which compartment(s) they have access to.

### To remove a user from a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Locate the user in the list.

3. Click the user.  
Its details are displayed.
4. Click **Groups**.
5. Click the Actions icon (three dots), and then click **Remove**.
6. Confirm when prompted.

### To delete a user

Prerequisite: To delete a user, the user must not be in any groups.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. For the user you want to delete, click **Delete**.
3. Confirm when prompted.

### To unblock a user

If you're an administrator, you can use the following procedure to unblock a user who has tried 10 times in a row to sign in to the Console unsuccessfully.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user.  
Its details are displayed, including the current status.
3. Click **Unblock**.
4. Confirm when prompted.

### To change a user's description

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user you want to update.  
The user's details are displayed. The description is displayed under the user's login.
3. Click the pencil next to the description.
4. Edit the description and save it.

### To edit a user's email

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user you want to update.  
The user's details are displayed.
3. Under **User Information**, click the pencil next to **Email**.
4. Enter the email address and click the save icon. The email address must be unique in the tenancy.

### To edit user capabilities

If you're an Administrator, you can edit user capabilities.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user to see its details.
3. Click **Edit User Capabilities**.

4. Select or clear the check box to add or remove a capability.
5. Click **Save**.

### To apply tags to a user

For instructions, see [Resource Tags](#).

For information about managing user credentials in the Console, see [Managing User Credentials](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



### Note

#### *Updates Are Not Immediate Across All Regions*

Your IAM resources reside in your home region. To enforce policy across all regions, the IAM service replicates your resources in each region. Whenever you create or change a policy, user, or group, the changes take effect first in the home region, and then are propagated out to your other regions. It can take several minutes for changes to take effect in all regions. For example, assume you have a group with permissions to launch instances in the tenancy. If you add UserA to this group, UserA will be able to launch instances in your home region within a minute. However, UserA will not be able to launch instances in other regions until the replication process is complete. This process can take up to several minutes. If UserA tries to launch an instance before replication is complete, they will get a not authorized error.

Use these API operations to manage users:

- [CreateUser](#)
- [ListUsers](#)
- [GetUser](#)
- [UpdateUserState](#): Unblocks a user who has tried to sign in 10 times in a row unsuccessfully.
- [UpdateUser](#): You can update the user's description, email, and tags.
- [UpdateUserCapabilities](#)
- [DeleteUser](#)

- [ListUserGroupMemberships](#): Use this operation to get a list of which users are in a group, or which groups a user is in.
- [AddUserToGroup](#): This operation results in a `UserGroupMembership` object with its own OCID.
- [GetUserGroupMembership](#)
- [RemoveUserFromGroup](#): This operation deletes a `UserGroupMembership` object.

For information about the API operations for managing user credentials, see [Managing User Credentials](#).

## Managing Groups

This topic describes the basics of working with groups.



### Important

If your tenancy is federated with Oracle Identity Cloud Service, see [Managing Oracle Identity Cloud Service Users and Groups in the Oracle Cloud Infrastructure Console](#) to manage groups.

## Required IAM Policy

If you're in the Administrators group, then you have the required access for managing groups.

For a policy that only gives someone power to determine what groups users are in, see [Let group admins manage group membership](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Working with Groups

When creating a group, you must provide a unique, unchangeable *name* for the group. The name must be unique across all groups within your tenancy. You must also provide the group with a *description* (although it can be an empty string), which is a non-unique, changeable description for the group. Oracle will also assign the group a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).



#### Note

If you delete a group and then create a new group with the same name, they'll be considered different groups because they'll have different OCIDs.

A group has no permissions until you write at least one policy that gives that group permission to either the tenancy or a compartment. When writing the policy, you can specify the group by using either the unique name or the group's OCID. Per the preceding note, even if you specify the group name in the policy, IAM internally uses the OCID to determine the group. For information about writing policies, see [Managing Policies](#).

You can delete a group, but only if the group is empty.

For information about the number of groups you can have, see [Service Limits](#).

If you're federating with an identity provider, you'll create mappings between the identity provider's groups and your IAM groups. For more information, see [Federating with Identity Providers](#).

### Using the Console



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To create a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.  
A list of the groups in your tenancy is displayed.
2. Click **Create Group**.
3. Enter the following:
  - **Name:** A unique name for the group. The name must be unique across all groups in your tenancy. You cannot change this later.
  - **Description:** A friendly description. You can change this later if you want to.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
4. Click **Create Group**.

Next, you might want to add users to the group, or write a policy for the group. See [To create a policy](#).

### To add a user to a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.  
A list of the groups in your tenancy is displayed.
2. Locate the group in the list.
3. Click the group.  
Its details are displayed
4. Click **Add User to Group**.
5. Select the user from the drop-down list, and then click **Add User**.

### To remove a user from a group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.  
A list of the groups in your tenancy is displayed.
2. Locate the group in the list.
3. Click the group to display its details.  
A list of users in the group is displayed.
4. Locate the user in the list.
5. For the user you want to remove, click **Remove**.
6. Confirm when prompted.

### To delete a group

Prerequisite: To delete a group, it must not have any users in it.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.  
A list of the groups in your tenancy is displayed.

2. Locate the group in the list.
3. For the group you want to delete, click **Delete**.
4. Confirm when prompted.

### To update a group's description

This is available only through the API. If you don't have access to the API and need to update a group's description, contact [Oracle Support](#).

### To apply tags to a group

For instructions, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



### Note

#### *Updates Are Not Immediate Across All Regions*

Your IAM resources reside in your home region. To enforce policy across all regions, the IAM service replicates your resources in each region. Whenever you create or change a policy, user, or group, the changes take effect first in the home region, and then are propagated out to your other regions. It can take several minutes for changes to take effect in all regions. For example, assume you have a group with permissions to launch instances in the tenancy. If you add UserA to this group, UserA will be able to launch instances in your home region within a minute. However, UserA will not be able to launch instances in other regions until the replication process is complete. This process can take up to several minutes. If UserA tries to launch an instance before replication is complete, they will get a not authorized error.

Use these API operations to manage groups:

- [CreateGroup](#)
- [ListGroup](#)s
- [GetGroup](#)
- [UpdateGroup](#): You can update only the group's description.
- [DeleteGroup](#)
- [ListUserGroupMemberships](#): Use to get a list of which users are in a group, or which groups a user is in.

- [AddUserToGroup](#): This operation results in a `UserGroupMembership` object with its own OCID.
- [GetUserGroupMembership](#)
- [RemoveUserFromGroup](#): This operation deletes a `UserGroupMembership` object.

For API operations related to group mappings for identity providers, see [Federating with Identity Providers](#).

## Managing Dynamic Groups

This topic describes how to manage dynamic groups and define the rules to determine a dynamic group's members.

### About Dynamic Groups

Dynamic groups allow you to group Oracle Cloud Infrastructure computer instances as "principal" actors (similar to user groups). You can then create policies to permit instances to [make API calls against Oracle Cloud Infrastructure services](#). When you create a dynamic group, rather than adding members explicitly to the group, you instead define a set of *matching rules* to define the group members. For example, a rule could specify that all instances in a particular compartment are members of the dynamic group. The members can change dynamically as instances are launched and terminated in that compartment.

### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing dynamic groups.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for dynamic groups or other IAM components, see [Details for IAM](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Working with Dynamic Groups

When creating a dynamic group, you must provide a unique, unchangeable *name* for the dynamic group. The name must be unique across all groups within your tenancy. You must also provide the dynamic group with a *description* (although it can be an empty string), which is a non-unique, changeable description for the group. Oracle will also assign the group a unique ID called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).



#### Note

If you delete a dynamic group and then create a new dynamic group with the same name, they'll be considered different groups because they'll have different OCIDs.

A dynamic group has no permissions until you write at least one policy that gives that dynamic group permission to either the tenancy or a compartment. When writing the policy, you can specify the dynamic group by using either the unique name or the dynamic group's OCID. Per the preceding note, even if you specify the dynamic group name in the policy, IAM internally uses the OCID to determine the dynamic group. For information about writing policies, see [Managing Policies](#).

You can delete a dynamic group, but only if the group is empty.

### Updating Dynamic Groups

You can update the matching rules that define the members of a dynamic group. For example, you might change a matching rule that includes all instances in a compartment to exclude a

particular instance. Or, you might update a rule to include a new tag value.



### Important

When you make a change to a matching rule you must allow about one hour for the updated policy to take effect. For example, if you update tags on an instance to either include or exclude that instance from a dynamic group, you must wait for that policy to take effect to include or exclude the instance.

## Limits on Instances in Dynamic Groups

A single compute instance can belong to a maximum of 5 dynamic groups.

## Using the Console



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## To create a dynamic group

1. Open the Console, click **Identity**, and then click **Dynamic Groups**.  
A list of the dynamic groups in your tenancy is displayed.
2. Click **Create Dynamic Group**.

3. Enter the following:
  - **Name:** A unique name for the group. The name must be unique across all groups in your tenancy (dynamic groups and user groups). You can't change this later.
  - **Description:** A friendly description. You can't change this in the Console, but you can change it [Using the API](#).
4. Enter the **Matching Rules**. Resources that meet the rule criteria are members of the group.
  - **Rule 1:** Enter a rule following the guidelines in [Writing Matching Rules to Define Dynamic Groups](#). You can manually enter the rule in the text box or launch the rule builder.
  - Enter additional rules as needed. To add a rule, click **+Additional Rule**.
5. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Create Dynamic Group**.

The matching rule syntax is verified, but the OCIDs are not. Be sure that the OCIDs you enter are correct.

Next, to give the dynamic group permissions, you need to write a policy. See [Writing Policies for Dynamic Groups](#).

### To delete a dynamic group

1. Open the Console, click **Identity**, and then click **Dynamic Groups**.

A list of the dynamic groups in your tenancy is displayed.
2. Locate the dynamic group in the list.
3. For the dynamic group you want to delete, click **Delete**.
4. Confirm when prompted.

### To update a dynamic group's description

This is available only through the API. If you don't have access to the API and need to update a dynamic group's description, contact [Oracle Support](#).

### To update a dynamic group's matching rules

1. Open the Console, click **Identity**, and then click **Dynamic Groups**.  
A list of the dynamic groups in your tenancy is displayed.
2. Click the dynamic group you want to update.  
The dynamic group's details are displayed.
3. Click **Edit All Matching Rules**.
4. Edit the matching rule in the text box; or, you can use the rule builder if the change is [supported by the rule builder](#).

### Writing Matching Rules to Define Dynamic Groups

Matching rules define the resources that belong to the dynamic group. In the Console, you can either enter the rule manually in the provided text box, or you can use the [rule builder](#). The rule builder lets you make selections and entries in a dialog, then writes the rule for you, based on your entries.

You can define the members of the dynamic group based on the following:

- compartment ID - include (or exclude) the instances that reside in that compartment based on compartment OCID
- instance ID - include (or exclude) an instance based on its instance OCID
- tag namespace and tag key - include (or exclude) instances tagged with a specific tag namespace and tag key. All tag values are included. For example, include all instances tagged the with tag namespace `department` and the tag key `operations`.
- tag namespace, tag key, and tag value - include (or exclude) instances tagged with a specific value for the tag namespace and tag key. For example include all instances

tagged with the tag namespace `department` and the tag key `operations` and with the value `'45'`.

A matching rule has the following syntax:

For a single condition:

```
variable =|!= 'value'
```

For multiple conditions:

```
any|all {<condition>,<condition>,...}
```

Supported variables are:

- `instance.compartment.id` - the OCID of the compartment where the instance resides
- `instance.id` - the OCID of the instance
- `tag.<tagnamespace>.<tagkey>.value` - the tag namespace and tag key. For example, `tag.department.operations.value`.
- `tag.<tagnamespace>.<tagkey>.value='<tagvalue>'` - the tag namespace, tag key, and tag value. For example, `tag.department.operations.value='45'`

Here are some examples:

### Include All Instances in a Specific Compartment in the Dynamic Group

To include all instances that are in a specific compartment, add a rule with the following syntax:

```
instance.compartment.id = '<compartment_ocid>'
```

You can add that rule either directly in the text box, or you can use the [rule builder](#).

Example entry in text box:

```
instance.compartment.id = 'ocidv1:compartment:oc1:phx:samplecompartmentocid6q6igvfauxmima74jv'
```

All instances that currently exist or get created in the compartment (identified by the OCID) are members of this group.

### Include All Instances in Any of Two or More Compartments

To include all instances that reside in any of two (or more) compartments, add a rule with the following syntax:

```
Any {instance.compartment.id = '<compartment_ocid>', instance.compartment.id = '<compartment_ocid>'}
```

You can add that rule either directly in the text box, or you can use the [rule builder](#).

Example entry in the text box:

```
Any {instance.compartment.id = 'ocidvl:compartment:oc1:phx:samplecompartmentocid6q6igvfauxmima74jv',
instance.compartment.id = 'ocidvl:compartment:oc1:phx:samplecompartmentocidythksk89eks1soelu2'}
```

Instances that currently exist or get created in either of the specified compartments are members of this group.

### Include All Instances Tagged with a Specific Namespace and Tag Key

To include all instances that are tagged with a specific tag namespace and tag key, add a rule with the following syntax:

```
tag.<tagnamespace>.<tagkey>.value
```

All instances assigned the tagnamespace.tagkey combination are included. Note that the tag value is not evaluated, so all values are included.

**Example:** Assume you have a tag namespace called `department` and a tag key called `operations`. You want to include all instances that are tagged with the namespace and tag key.

Enter the following rule in the text box:

```
tag.department.operations.value
```

All instances that currently exist or get created with the tag namespace and tag key `department.operations` are members of this group.

## Include All Instances In a Specific Compartment with a Specific Tag Namespace, Tag Key, and Tag Value

To include all instances in a specific compartment that are tagged with a specific tag namespace, key, and value, add a rule with the following syntax:

```
All {instance.compartment.id = '<compartment_ocid>',
tag.<tagnamespace>.<tagkey>.value='<tagvalue>'}
```

All instances that are in the identified compartment and that are assigned the tagnamespace.tagkey with the specified tag value are included.

**Example:** Assume you have a tag namespace called `department` and a tag key called `operations`. You want to include all instances that are tagged with the value `45`, that are in a particular compartment.

Enter the following statement in the text box:

```
All
{instance.compartment.id='ocidv1:compartment:oc1:phx:oc1:phx:samplecompartmentocid6q6igvfauxmima74jv',
tag.department.operations.value='45'}
```

## Include Instances in a Specific Compartment Except Those with a Specific Tag

To include all instances in a specific compartment EXCEPT those that are tagged with a specific tag namespace, key, and value, add a rule with the following syntax:

```
All {instance.compartment.id = '<compartment_ocid>',
tag.<tagnamespace>.<tagkey>.value!= '<tagvalue>'}
```

**Example:** Assume you have a tag namespace called `department` and a tag key called `operations`. You want to include all instances in a specific compartment, except those that are tagged with the value `45`.

Enter the following statement in the text box:

```
All
{instance.compartment.id='ocidv1:compartment:oc1:phx:oc1:phx:samplecompartmentocid6q6igvfauxmima74jv',
tag.department.operations.value!='45'}
```

### Using the Rule Builder

The rule builder is a tool available from the Console to help you write matching rules. The rule builder provides menus and text boxes for you to make entries and then writes the rule for you. The rule builder does have some limitations, so you can't use it for all cases.

#### LIMITATIONS OF THE RULE BUILDER

The rule builder does not support the following:

- Exclusion rules - the rule builder lets you select compartment IDs and instance IDs to include only.
- Rules based on tags - the rule builder does not allow you to select tags to include in your rule. To add a rule based on tag values, you need to enter the rule in the Rule text box using the syntax above.

#### LAUNCHING THE RULE BUILDER

When you click **Create Dynamic Group**, the Rule Builder is displayed in the **Create Dynamic Group** dialog.

To create a matching rule using the rule builder

1. Select **Any** or **All** from the menu.  
**Any** includes instances that match any of the statements in the rule.  
**All** includes only instances that match all of the statements in the rule.
2. Select the **Attribute** type for the statement and enter the value:  
**in Compartment ID** includes instances in the compartment you specify.  
**with Instance ID** includes instances with the OCID you specify.
3. Click **+Additional line** to add more statements to this rule.  
When you add multiple statements to a rule, remember that Any includes instances that match any of the statements. If you choose All, instances must match all of the specifications in the statements to be included in the group.

### EXAMPLES USING THE RULE BUILDER

#### Include All Instances in a Specific Compartment in the Dynamic Group

To include all instances that are in a specific compartment, using the rule builder:

- Select ALL.
- **Attribute:** Select in Compartment ID.
- **Value:** Enter `ocidv1:compartment:oc1:yourcompartmentocid`

All instances that currently exist or get created in the compartment (identified by the OCID) are members of this group.

#### Include All Instances in Any of Two or More Compartments

To include all instances that reside in any of two (or more) compartments using the rule builder:

1. Select ANY.
2. Enter:
  - **Attribute:** Select in Compartment ID.
  - **Value:** Enter  
`ocidv1:compartment:oc1:phx:samplecompartmentocid6q6igvfauxmima74jv`
3. Click **+Additional Line**. Enter the following on the second line:
  - **Attribute:** Select in Compartment ID.
  - **Value:** Enter  
`ocidv1:compartment:oc1:phx:samplecompartmentocidythksk89ekslsoelu2`
4. Continue adding additional lines as needed for each compartment you want to include.

Instances that currently exist or get created in any of the specified compartments are members of this group.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage dynamic groups:

- [CreateDynamicGroup](#)
- [ListDynamicGroups](#)
- [GetDynamicGroup](#)
- [UpdateDynamicGroup](#)
- [DeleteDynamicGroup](#)

### Managing Compartments

This topic describes the basics of working with compartments.

#### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing compartments.

For an additional policy related to compartment management, see [Let a compartment admin manage the compartment](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for compartments or other IAM components, see [Details for IAM](#).

#### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource

later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Working with Compartments

When you first start working with Oracle Cloud Infrastructure, you need to think carefully about how you want to use compartments to organize and isolate your cloud resources. Compartments are fundamental to that process. Most resources can be moved between compartments. However, it's important to think through your compartment design for your organization up front, before implementing anything. For more information, see "Setting Up Your Tenancy" in the *Oracle Cloud Infrastructure Getting Started Guide*.

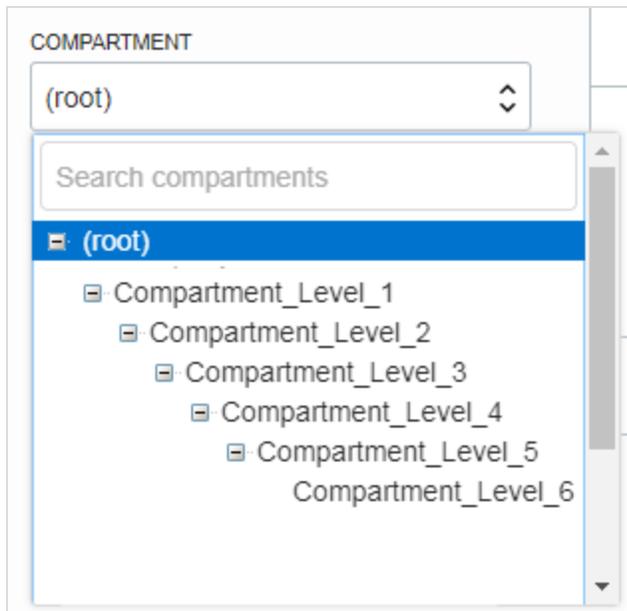
The Console is designed to display your resources by *compartment* within the current region. When you work with your resources in the Console, you must choose which compartment to work in from a list on the page. That list is filtered to show only the compartments in the tenancy that you have permission to access. If you're an administrator, you'll have permission to view all compartments and work with any compartment's resources, but if you're a user with limited access, you probably won't.

Compartments are tenancy-wide, across regions. When you create a compartment, it is available in every region that your tenancy is subscribed to. You can get a cross-region view of your resources in a specific compartment with the compartment explorer. See [Viewing All Resources in a Compartment](#).

### Creating Compartments

When creating a compartment, you must provide a *name* for it (maximum 100 characters, including letters, numbers, periods, hyphens, and underscores) that is unique within its parent compartment. You must also provide a *description*, which is a non-unique, changeable description for the compartment, from 1 through 400 characters. Oracle will also assign the compartment a unique ID called an Oracle Cloud ID. For more information, see [Resource Identifiers](#).

You can create subcompartments in compartments to create hierarchies that are six levels deep.



For information about the number of compartments you can have, see [Service Limits](#).

### Access Control for Compartments

After creating a compartment, you need to write at least one policy for it, otherwise no one can access it (except administrators or users who have permissions set at the tenancy level). When creating a compartment inside another compartment, the compartment inherits access permissions from compartments higher up its hierarchy. For more information, see [Policy Inheritance](#).

When you create an access policy, you need to specify which compartment to attach it to. This controls who can later modify or delete the policy. Depending on how you've designed your compartment hierarchy, you might attach it to the tenancy, a parent, or to the specific compartment itself. For more information, see [Policy Attachment](#).

### **Putting Resources in a Compartment**

To place a new resource in a compartment, you simply specify that compartment when creating the resource (the compartment is one of the required pieces of information to create a resource). If you're working in the Console, you just make sure you're first viewing the compartment where you want to create the resource. Keep in mind that most IAM resources reside in the tenancy (this includes users, groups, compartments, and any policies attached to the tenancy) and can't be created in or managed from a specific compartment.

### **Moving Resources to a Different Compartment**

Most resources can be moved after they are created. There are a few resources that you can't move from one compartment to another.

Some resources have attached resource dependencies and some don't. Not all attached dependencies behave the same way when the parent resource moves.

For some resources, the attached dependencies move with the parent resource to the new compartment. The parent resource moves immediately, but in some cases attached dependencies move asynchronously and are not visible in the new compartment until the move is complete.

For other resources, the attached resource dependencies do not move to the new compartment. You can move these attached resources independently.

After you move the resource to the new compartment, the policies that govern the new compartment apply immediately and affect access to the resource. Depending on the structure of your compartment organization, metering, billing, and alarms can also be affected.

See the service documentation for individual resources to familiarize yourself with the behavior of each resource and its attachments.

### **Viewing Resources in a Compartment**

It's not possible to get a list of all the resources in a compartment by using a single API call. Instead you can list all the resources of a given type in the compartment (e.g., all the instances, all the block storage volumes, etc.).



### Tip

In the Console, the compartment explorer allows you to get a list of resources in a compartment, across regions, with some limitations. For more information, see [Viewing All Resources in a Compartment](#).

### Deleting Compartments

To delete a compartment, it must be empty of all resources. Before you initiate deleting a compartment, be sure that all its resources have been moved, deleted, or terminated, including any policies attached to the compartment.



### Important

Some resource types can't be deleted, therefore, compartments containing these resource types can't be deleted. A resource type that can't be deleted is:

- Data transfer jobs

The delete action is asynchronous and initiates a work request. The state of the compartment changes to Deleting while the work request is executing. It typically takes several minutes for the work request to complete. While it is in the Deleting state it is not displayed on the compartment picker. If the work request fails, the compartment is not deleted and it returns to the Active state.

After a compartment is deleted, its state is updated to Deleted and a random string of characters is appended to its name, for example, CompartmentA might become CompartmentA.qR5hP2BD. Renaming the compartment allows you to reuse the original name for a different compartment. The deleted compartment is displayed on the Compartments page for the number of days specified in your [Audit Retention Period](#) setting (90-365 days). The deleted compartment is removed from the compartment picker. If any policy statements

reference the deleted compartment, the name in the policy statement is updated to the new name.

### Troubleshooting tips for when a compartment fails to delete

If the compartment fails to delete, verify that you have removed all the resources:

- For most resources, you can use the compartment explorer to help you locate them. See [Resources Supported by the Compartment Explorer](#) for the list of supported resources.

#### To view resources in a compartment

Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Compartment Explorer**.

When you open the compartment explorer, the list of all resources that you have permission to view is displayed. The compartment explorer opens with a view of the root compartment. The **Name** and **Description** of the compartment you are viewing are displayed at the top of the page.

To navigate to the compartment you are interested in, use the compartment picker on the left of the Console page.

- Verify that there are no policies in the compartment (policies are not included in Search results).

#### To find policies in a compartment

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.
2. From the compartments list on the left, select the compartment you want to delete.

Policies attached to the compartment are displayed.

- If you can't locate any resources in the compartment, check with your Administrator; you might not have permission to view all resources.



### **Important**

There is a known issue causing deleted compartments to continue to count against your [service limit](#) of compartments. See [Deleted compartments continue to count against service limits](#).

### **Adding Tag Defaults for a Compartment**

Tag defaults let you specify tags to be applied automatically to all resources, at the time of creation, in the current compartment. For more information, see [Managing Tag Defaults](#).

### **Moving a Compartment to a Different Parent Compartment**

You can move a compartment to a different parent compartment within the same tenancy. When you move a compartment, all its contents (subcompartments and resources) are moved with it. Moving a compartment has implications for the contents. These implications are described in the following sections. Ensure that you are aware of these before you move a compartment.

- [Required IAM Policy](#)
- [Restrictions on Moving Compartments](#)
- [Understanding the Policy Implications When You Move a Compartment](#)
- [Understanding Compartment Quota Implications When You Move a Compartment](#)
- [Understanding Tagging Implications When You Move a Compartment](#)

### **Required IAM Policy**

To move a compartment, you must belong to a group that has `manage all-resources` permissions on the lowest shared parent compartment of the current compartment and the destination compartment.

### **Restrictions on Moving Compartments**

- You can't move a compartment to a destination compartment with the same name as the compartment being moved.  
For example, assume compartment A and compartment B are both under the root compartment. Under compartment A is a subcompartment, also called compartment B.

You cannot move the compartment B to the parent compartment B.

- Two compartments within the same parent cannot have the same name. Therefore you can't move a compartment to a destination compartment where a compartment with the same name already exists.

### **Understanding the Policy Implications When You Move a Compartment**

After you move a compartment to a new parent compartment, the access policies of the new parent take effect and the policies of the previous parent no longer apply. Before you move a compartment, ensure that:

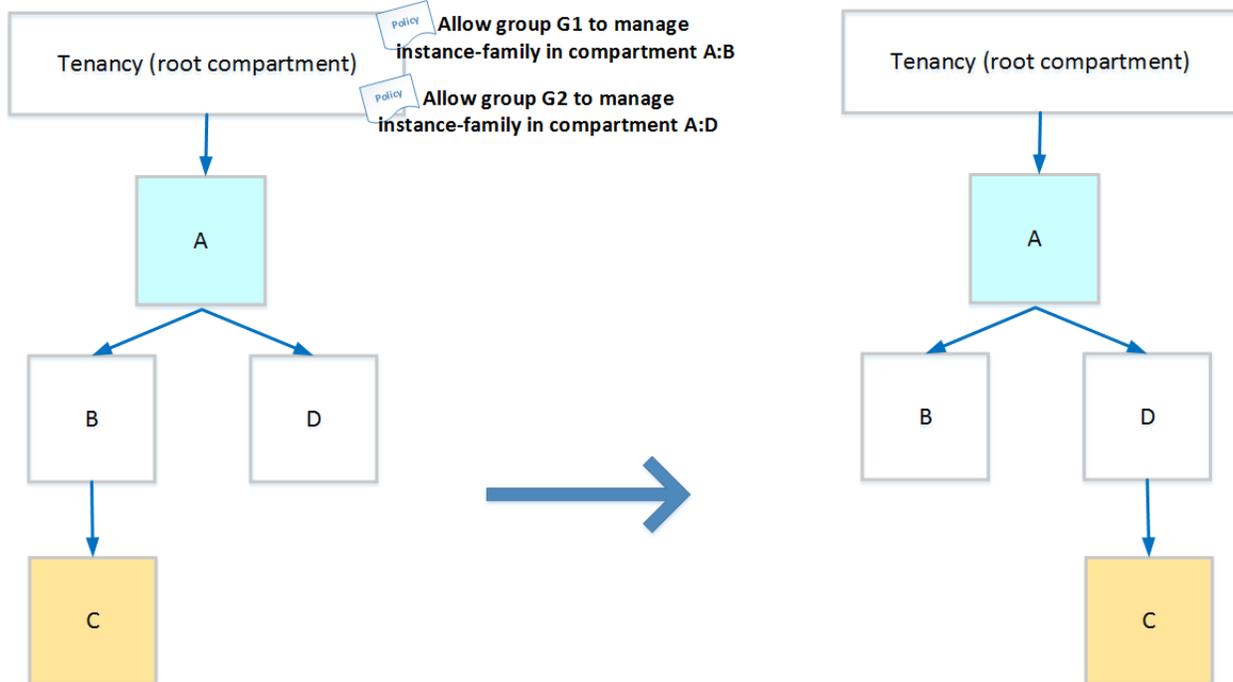
- You are aware of the policies that govern access to the compartment in its current position.
- You are aware of the policies in the new parent compartment that will take effect when you move the compartment.

In some cases, when moving nested compartments with policies that specify the hierarchy, the policies are automatically updated to ensure consistency.

### **Policy Examples**

#### **Groups with Permissions in the Current Compartment Lose Access; Groups with Permissions in the Destination Compartment Gain Access**

The following figure shows a compartment hierarchy in which compartment C, a child of A:B is moved to the hierarchy A:D.



The tenancy has the following policies defined for compartments B and D:

Policy1: Allow group G1 to manage instance-family in compartment A:B

Policy2: Allow group G2 to manage instance-family in compartment A:D

Impact when compartment C is moved from B to D:

Group G1 can no longer manage instance-families in compartment C.

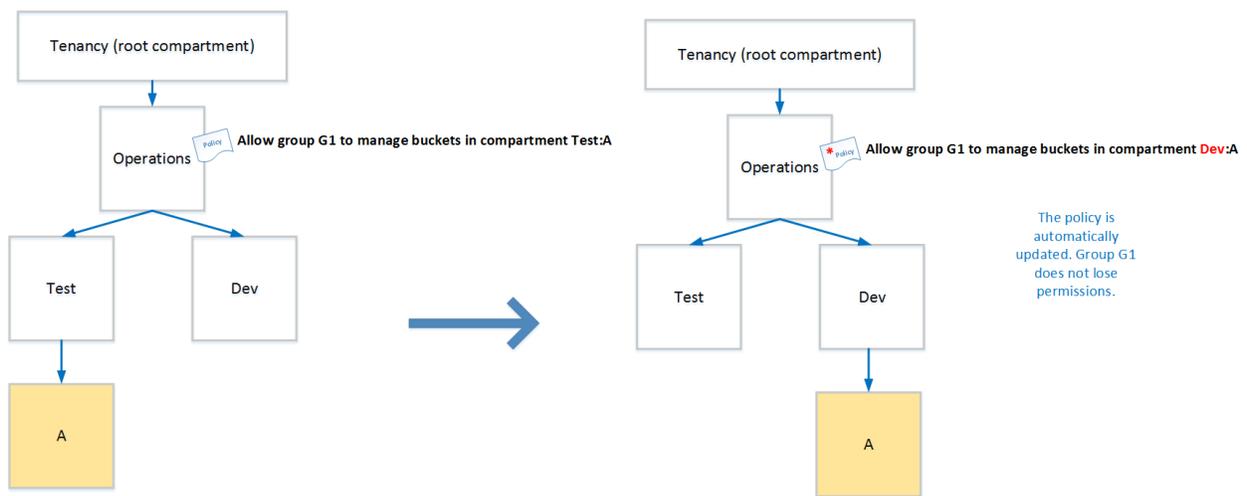
Group G2 can now manage instance-families in compartment C.

Ensure that you are aware not only of what groups lose permissions when you move a compartment, but also what groups will gain permissions.

### Automatic Update of Policies

When you move a compartment, some policies will be automatically updated. Policies that specify the compartment hierarchy down to the compartment being moved will automatically be updated when the policy is attached to a shared ancestor of the current and target parent. Consider the following examples:

#### Example 1: Policy automatically updated



In this example, you move compartment A from Operations:Test to Operations:Dev. The policy that governs compartment A is attached to the shared parent, Operations. When the compartment is moved, the policy statement is automatically updated by the IAM service to specify the new compartment location.

The policy

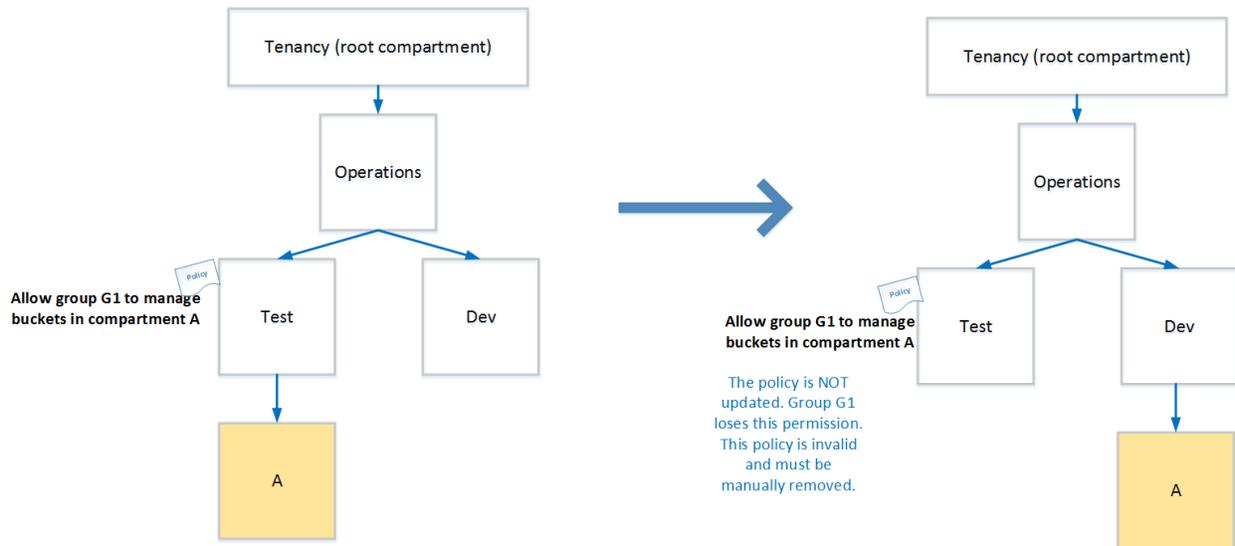
```
Allow group G1 to manage buckets in compartment Test:A
```

is updated to

```
Allow group G1 to manage buckets in compartment Dev:A
```

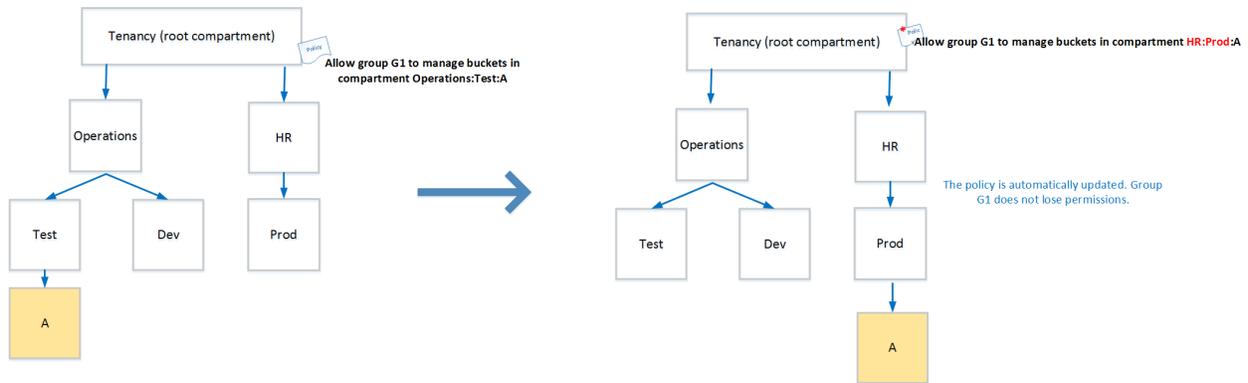
No manual intervention is required to allow group G1 to continue to access compartment A in its location.

### Example 2: Policy not updated



In this example, you move compartment A from Operations:Test to Operations:Dev. However, the policy that governs compartment A here is attached directly to the Test compartment. When the compartment is moved, the policy is not automatically updated. The policy that specifies compartment A is no longer valid and must be manually removed. Group G1 no longer has access to compartment A in its new location under Dev. Unless another existing policy grants access to group G1, you must create a new policy to allow G1 to continue to manage buckets in compartment A.

### Example 3: Policy attached to the tenancy is updated



In this example, you move compartment A from Operations:Test to HR:Prod. The policy that governs compartment A is attached to the tenancy, which is a shared ancestor by the original parent compartment and the new parent compartment. Therefore, when the compartment is moved, the policy statement is automatically updated by the IAM service to specify the new compartment location.

The policy statement:

```
Allow group G1 to manage buckets in compartment Operations:Test:A
```

is updated to

```
Allow group G1 to manage buckets in compartment HR:Prod:A
```

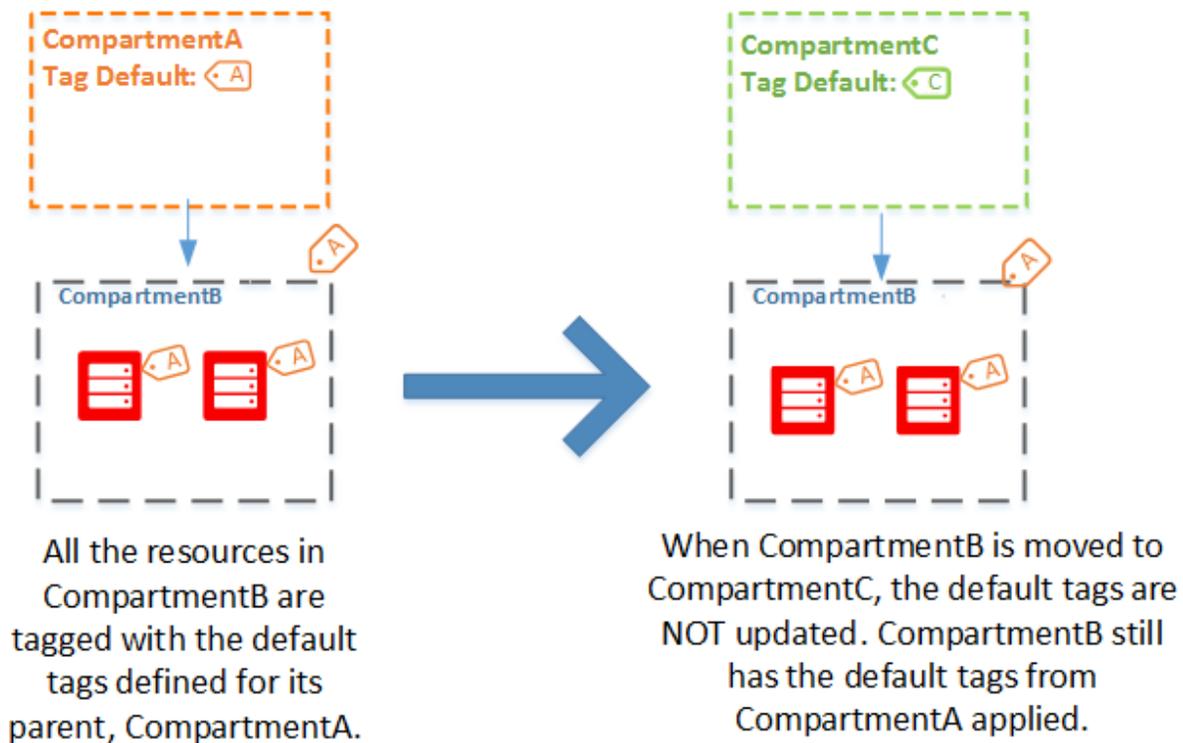
No manual intervention is required to allow group G1 to continue to access compartment A.

### Understanding Compartment Quota Implications When You Move a Compartment

When you move one compartment to another, resource quotas in the destination compartment are not verified and are not enforced. Therefore, if the compartment move results in a quota violation in the destination compartment, the move is not blocked. After the move is complete, the destination compartment will be in an over-quota state. You will not be able to create new resources that are over-quota until you either adjust the quotas for the destination compartment or remove resources to comply with the existing quota. For more information on managing compartment quotas, see [Compartment Quotas](#).

### Understanding Tagging Implications When You Move a Compartment

Tags are not automatically updated after a compartment move. If you have implemented a tagging strategy based on compartment, you must update the tags on the resources after the move. For example, assume CompartmentA has a child compartment, CompartmentB. CompartmentA is set up with tag defaults so that every resource in CompartmentA is tagged with TagA. Therefore CompartmentB and all its resources are tagged with default tag, TagA. When you move CompartmentB to CompartmentC, it will still have the default tags from CompartmentA. If you have set up default tags for CompartmentC, you'll need to add those to the resources in the moved compartment.



Default tags must be manually updated after a compartment is moved

## Using the Console



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## To create a compartment

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments you have access to is displayed.
2. Navigate to the compartment in which you want to create the new compartment:
  - To create the compartment in the tenancy (root compartment) click **Create Compartment**.
  - Otherwise, click through the hierarchy of compartments until you reach the detail page of the compartment in which you want to create the compartment. On the **Compartment Details** page, click **Create Compartment**.
3. Enter the following:
  - **Name:** A unique name for the compartment (maximum 100 characters, including letters, numbers, periods, hyphens, and underscores). The name must be unique across all the compartments in your tenancy. Avoid entering confidential information.
  - **Description:** A friendly description. You can change this later if you want to. Avoid entering confidential information.
  - **Compartment:** The compartment you are in is displayed. To choose another compartment to create this compartment in, select it from the list.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

4. Click **Create Compartment**.

Next, you might want to write a policy for the compartment. See [To create a policy](#).

### To update a compartment's name

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.

A list of the compartments in your tenancy is displayed.

2. For the compartment you want to rename, click the Actions icon (three dots), and then click **Rename Compartment**.



#### Tip

You can't change the name of your root compartment.

3. Enter the new **Name**. The name must be unique across all the compartments in your tenancy. The name can have a maximum of 100 characters, including letters, numbers, periods, hyphens, and underscores. Avoid entering confidential information.
4. Click **Rename Compartment**.

### To update a compartment's description

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments in your tenancy is displayed.
2. For the compartment you want to update, click the Actions icon (three dots), and then click **Edit Compartment Description**.
3. Enter the new description. Avoid entering confidential information.
4. Click **Save**.

### To view the contents of a compartment

1. Open the Console,
2. Open the navigation menu and select the type of resource you want to view. For example, click **Compute** to view all your Compute resources.
3. Choose the compartment from the list on the left side of the page.  
The page updates to show only the resources in that compartment.

Remember that most IAM resources reside in the tenancy (this includes users, groups, and compartments). Policies can reside in either the tenancy (root compartment) or other compartments.

### To move a compartment

To move a compartment, you must belong to a group that has `manage all-resources` permissions on the lowest shared parent compartment of the current compartment and the destination compartment.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.

A list of the compartments in your tenancy is displayed. If the compartment you want to move is not directly beneath the root compartment, click through the hierarchy of compartments to view the wanted compartment.

2. For the compartment you want to move, click the Actions icon (three dots), and then click **Move Compartment**.
3. Select the destination compartment.
4. Confirm that you are aware of the [implications of the move](#).
5. Click **Move Compartment**.

### To move a resource to a different compartment

1. Open the Console.
2. Open the navigation menu and select the type of resource you want to work with. For example, click **Compute** to view all your Compute resources.
3. In the **List Scope** section, select a compartment. Resources in the selected compartment are displayed.
4. Find the resource in the list, click the the Actions icon (three dots), and follow the prompts to move the resource to a new compartment. See the resource documentation for specific steps.

The resource is moved immediately. If attached resource dependencies move with the parent resource, the resource dependencies are moved asynchronously, and do not appear in the new compartment until the move is complete.

### To apply tags to a compartment

For instructions, see [Resource Tags](#).

### To manage tag defaults for a compartment

See [Managing Tag Defaults](#).

### To delete a compartment

You must remove all resources from a compartment before you can delete it.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments in your tenancy is displayed.
2. For the compartment you want to delete, click the Actions icon (three dots), and then click **Delete Compartment**.
3. At the prompt, click **OK**.

After you click **OK**, a work request is submitted to delete the compartment. The compartment state changes to Deleting. If the work request fails, the state returns to Active.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage compartments:

- [CreateCompartment](#)
- [ListCompartments](#)
- [GetCompartment](#): Returns the metadata for the compartment, not its contents.
- [UpdateCompartment](#)
- [DeleteCompartment](#)
- [MoveCompartment](#)
- [GetWorkRequest](#): Gets the work requests spawned by the DeleteCompartment operation.

You can retrieve the contents of a compartment only by resource type. There's no API call that lists *all* resources in the compartment. For example, to list all the instances in a

compartment, call the Core Services API [ListInstances](#) operation and specify the compartment ID as a query parameter.

## Managing Regions

This topic describes the basics of managing your region subscriptions. For more information about regions in Oracle Cloud Infrastructure, see [Regions and Availability Domains](#). For information about Platform Services regions, see [Managing Platform Services Regions](#).

### Required IAM Policy

If you're in the Administrators group, then you have the required access to manage region subscriptions.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for managing regions or other IAM components, see [Details for IAM](#).

### The Home Region

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your *home region*. Your home region is where your IAM resources are defined. When you subscribe to another region, your IAM resources are available in the new region, however, the master definitions reside in your home region and can only be changed there.

Resources that you can create and update only in the home region are:

- Users
- Groups
- Policies
- Compartments
- Dynamic groups
- Federation resources

When you use the API to update your IAM resources, you must use the endpoint for your home region. IAM automatically propagates the updates to all regions in your tenancy.

When you use the Console to update your IAM resources, the Console sends the requests to the home region for you. You don't need to switch to your home region first. IAM then automatically propagates the updates to all regions in your tenancy.

When you subscribe your tenancy to a new region, all the policies from your home region are enforced in the new region. If you want to limit access for groups of users to specific regions, you can write policies to grant access to specific regions only. For an example policy, see [Restrict admin access to a specific region](#).



### Note

#### *IAM Updates Are Not Immediate Across All Regions*

When you create or update an IAM resource, be aware that you need to allow up to several minutes for the changes in your home region to become available in all regions.

## Using the Console to Manage Infrastructure Regions

### To view the list of infrastructure regions

Open the Console, open the **Region** menu, and then click **Manage Regions**.

A list of the regions offered by Oracle Cloud Infrastructure is displayed. Regions that you have not subscribed to provide a button to create a subscription.

### To subscribe to an infrastructure region

1. Open the Console, open the **Region** menu, and then click **Manage Regions**.  
The list of regions available to your tenancy Oracle Cloud Infrastructure is displayed. Your home region is labeled.
2. Locate the region you want to subscribe to and click **Subscribe**.  
Note that it could take several minutes to activate your tenancy in the new region. Remember, your IAM resources are global, so when the subscription becomes active, all your existing policies are enforced in the new region.  
To switch to the new region, use the **Region** menu in the Console. See [Switching Regions](#) for more information.

You cannot unsubscribe from a region.

### Using the API to Work with Infrastructure Regions

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage infrastructure regions:

- [GetTenancy](#)
- [ListRegions](#): Returns a list of regions offered by Oracle Cloud Infrastructure in your selected realm.
- [CreateRegionSubscription](#)
- [ListRegionSubscriptions](#)

You cannot unsubscribe from a region.

### Region FAQs

#### Can an individual user subscribe to a region?

A region subscription is at the tenancy level. An administrator can subscribe the tenancy to a region. All IAM policies are enforced in the new region, so *all* users in the tenancy will have the same access and permissions in the new region.

#### Can I see my existing resources in the new region?

When you select a region in the Console, you are shown a view of the resources in your selected region. Most cloud resources (instances, VCNs, buckets, etc.) exist only in a specific region, so you only see them when you select the region where they were created. The exception is IAM resources: compartments, users, groups, and policies are global across all regions. See also [Working in Multiple Regions](#).

#### How do my service limits apply to the new region?

Service limits can be scoped to the tenant level, the region level, or the availability domain level. When you subscribe to a new region, you get access to the region and its availability domains. Service limits apply accordingly. The [service limits page](#) lists the scope of each resource limit.

#### Can I restrict access to a specific region?

Yes. You can write policies that grant permissions in a specified region only. For an example policy, see [Restrict admin access to a specific region](#).

#### Can I change my home region?

No. Oracle assigns your home region and you can't change it.

## Managing Platform Services Regions

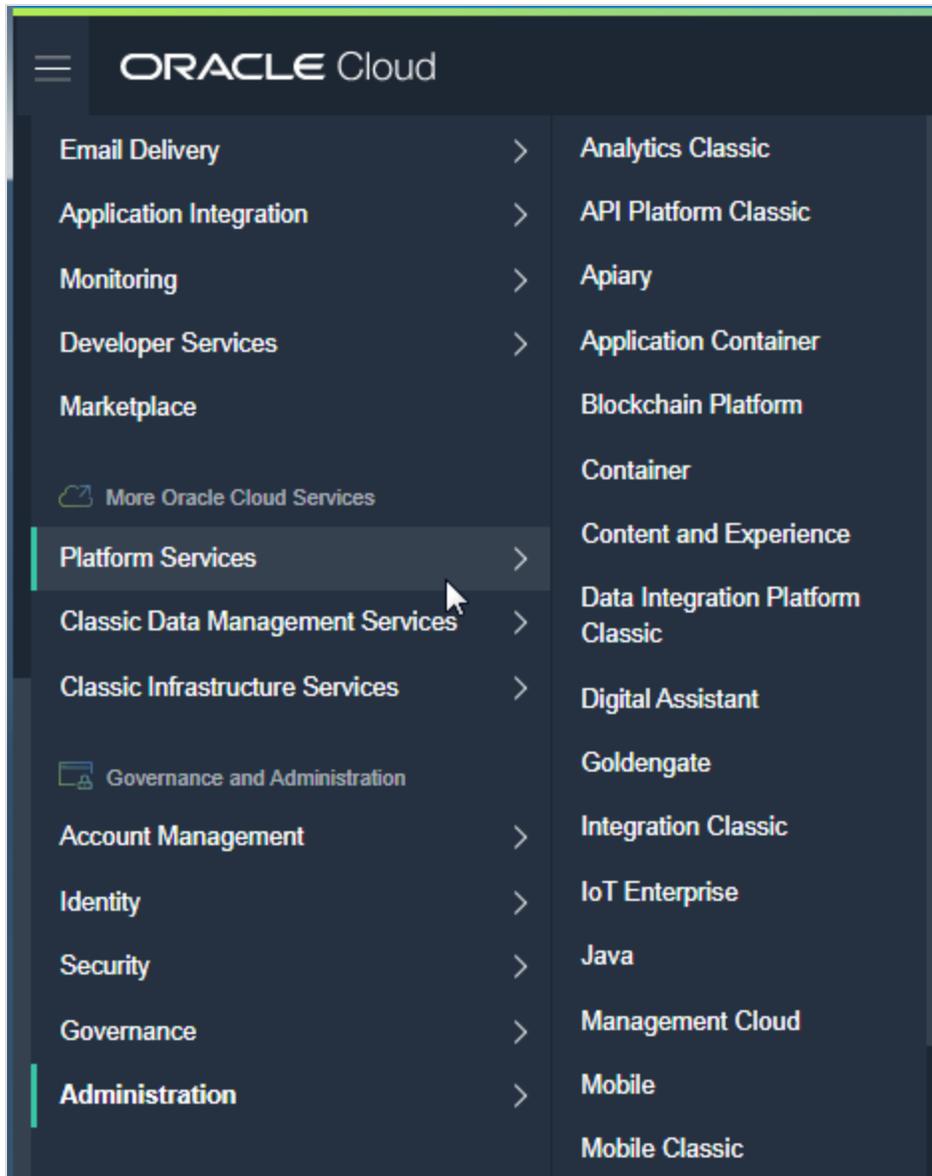
This topic describes how to manage Platform Services region subscriptions.

### About the Platform Services Regions

You can manage Platform Services regions in the Console.

To use Platform Services that are not natively integrated with Oracle Cloud Infrastructure, you need to subscribe to the Platform Services region as well as the [Infrastructure region](#). For example, to create a Platform Service instance in the Germany Central (Frankfurt) region, you need to subscribe your tenancy to both the Infrastructure region: Germany Central (Frankfurt) and to the Platform Service region: Europe and Middle East.

To know which services require the Platform Service region subscription, you can view the list under **More Oracle Services** on the Console navigation menu. A sample of the navigation to these services is shown in the following screenshot:



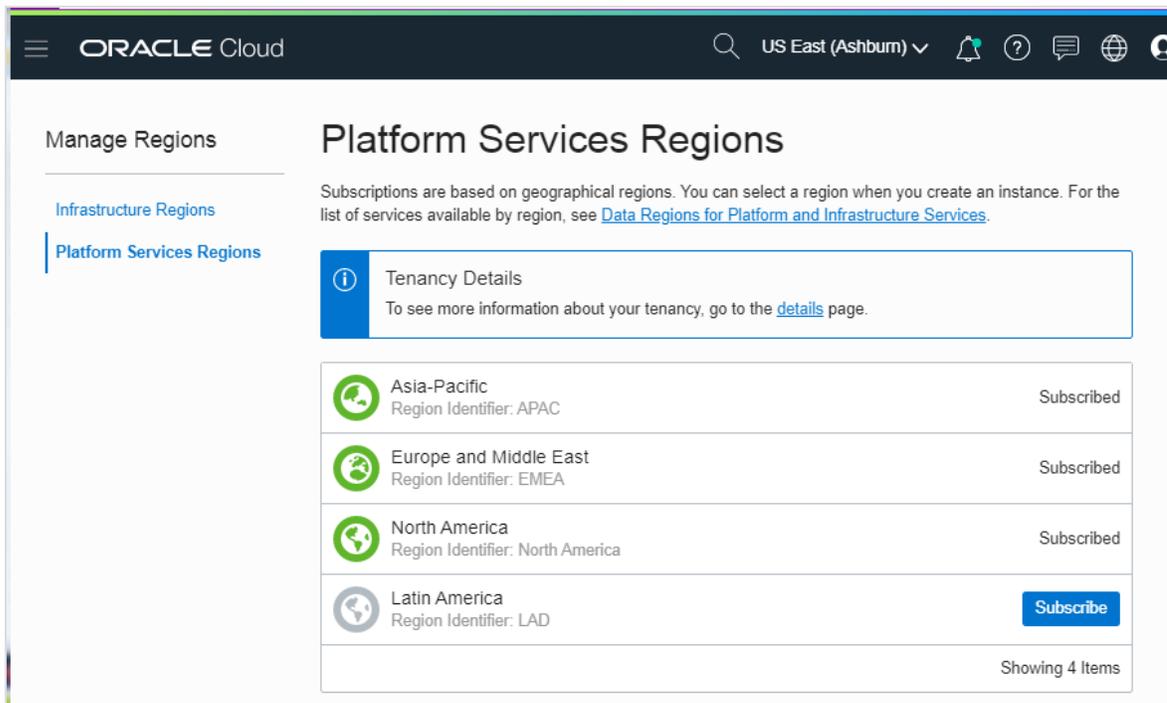
Before you can view these services in the Console or access the **Manage Platform Services Regions** page, your tenancy must have entitlements to use the Platform Services.

## Managing Platform Services Regions

### To view and subscribe to Platform Services regions

1. Open the Console, open the **Region** menu, and then click **Manage Regions**.
2. On the **Manage Regions** page, click **Platform Services Regions**.

The list of geographical regions is displayed. Regions that you have not subscribed to provide a button to create a subscription. A sample of the **Platform Services Regions** page is shown in the following screenshot:



3. To subscribe to a region, locate the region in the list and click **Subscribe**. It might take several minutes to activate your tenancy in the new region.

You cannot unsubscribe from a region.

# Managing the Tenancy

This topic describes options on the tenancy details page in the Console.

## Required IAM Policy

If you're in the Administrators group, then you have the required access to manage the tenancy.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for your tenancy and other IAM components, see [Details for IAM](#).

## Viewing the Tenancy Details Page

To view the tenancy details page:

Open the **Profile** menu () and click **Tenancy: <your\_tenancy\_name>**.

## Details About Your Tenancy

The tenancy details page provides the following information about your tenancy:

### TENANCY OCID

Every Oracle Cloud Infrastructure resource has an Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). You need your tenancy's OCID to use the API. You'll also need it when contacting support.

### HOME REGION

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one of the available regions. This is your *home region*. Your home region is where your IAM resources are defined. For more information about the home region, see [The Home Region](#).

### **NAME**

Your tenancy name. Your tenancy name is typically chosen when you set up your Oracle Cloud account.

### **CSI NUMBER**

Your Customer Service Identifier for Oracle Support.

### **AUDIT RETENTION PERIOD**

The retention period for the Audit service logs. The value of the retention period setting affects all regions and all compartments for this tenancy. You can't set different retention periods for different regions or compartments. For more information about this setting, see [Setting Audit Log Retention Period](#).

### **OBJECT STORAGE DESIGNATED COMPARTMENTS AND NAMESPACE**

The Object Storage service provides API support for both Amazon S3 Compatibility API and Swift API. By default, buckets created using the Amazon S3 Compatibility API or the Swift API are created in the root compartment of the Oracle Cloud Infrastructure tenancy. You can designate a different compartment for the Amazon S3 Compatibility API or Swift API to create buckets in. For more information, see [Designating Compartments for the Amazon S3 Compatibility and Swift APIs](#).

For information about your Object Storage namespace, see [Understanding Object Storage Namespaces](#).

### **TAGS**

Tagging allows you to define keys and values and associate them with resources. You can then use the tags to help you organize and list resources based on your business needs. If you have permissions to manage the tenancy, you also have permissions to apply free-form tags. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#).

### **SERVICE LIMITS**

The limits allotted to your tenancy and usage against these limits. Not all service resources are included in the list shown here on the Console. For more information or to

request an increase, see [Service Limits](#).

### Using the API

Many of the options set on this page are managed through the owning service. For example, the Object Storage settings are managed with the [Object Storage service API](#), and setting the Audit log retention period is handled by the [Audit service API](#).

To get information about your tenancy use the following operation:

- [GetTenancy](#)

To tag a tenancy, use the following operations:

- [GetCompartment](#)
- [UpdateCompartment](#)

In the above operations, use the tenancy OCID for the `compartmentID` parameter.

## Managing Policies

This topic describes the basics of working with policies.

### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing policies.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies to control who else can write policies or manage other IAM components, see [Let a compartment admin manage the compartment](#), and also [Details for IAM](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Working with Policies

If you haven't already, make sure to read [How Policies Work](#) to understand the basics of how policies work.

When creating a policy, you must specify the compartment where it should be *attached*, which is either the tenancy (the root compartment) or another compartment. Where it's attached governs who can later modify or delete it. For more information, see [Policy Attachment](#). When creating the policy in the Console, you attach the policy to the desired compartment by creating the policy *while viewing that compartment*. If you're using the API, you specify the identifier of the desired compartment in the [CreatePolicy](#) request.

Also when creating a policy, you can specify its *version date*. For more information, see [Policy Language Version](#). You can change the version date later if you like.

When creating a policy, you must also provide a unique, non-changeable *name* for it. The name must be unique across all policies in your tenancy. You must also provide a *description* (although it can be an empty string), which is a non-unique, changeable description for the policy. Oracle will also assign the policy a unique ID called an Oracle Cloud ID. For more information, see [Resource Identifiers](#).



#### Note

If you delete a policy and then create a new policy with the same name, they'll be considered different policies because they'll have different OCIDs.

For information about how to write a policy, see [How Policies Work](#) and [Policy Syntax](#).

When you create a policy, make changes to an existing policy, or delete a policy, your changes go into effect typically within 10 seconds.

You can view a list of your policies in the Console or with the API. In the Console, the list is automatically filtered to show only the policies attached to the compartment you're viewing. To determine which policies apply to a particular group, you must view the individual statements inside all your policies. There isn't a way to automatically obtain that information in the Console or API.

For information about the number of policies you can have, see [Service Limits](#).

### Using the Console



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To create a policy

Prerequisite: The group and compartment that you're writing the policy for must already exist.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.  
A list of the policies in the compartment you're viewing is displayed.
2. If you want to attach the policy to a compartment other than the one you're viewing, select the desired compartment from the list on the left. Where the policy is attached controls who can later modify or delete it (see [Policy Attachment](#)).
3. Click **Create Policy**.

4. Enter the following:

- **Name:** A unique name for the policy. The name must be unique across all policies in your tenancy. You cannot change this later.
- **Description:** A friendly description. You can change this later if you want to.
- **Policy Versioning:** Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in format YYYY-MM-DD format. For more information, see [Policy Language Version](#).
- **Statement:** A policy statement. For the correct format to use, see [Policy Basics](#) and also [Policy Syntax](#). If you want to add more than one statement, click **+**.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create**.

The new policy will go into effect typically within 10 seconds.

### To get a list of your policies

Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**. A list of the policies in the compartment you're currently viewing is displayed. If you want to view policies attached to a different compartment, select that compartment from the list on the left. You can't get a single list of all policies; they're always displayed by compartment.

To determine which policies apply to a particular group, you must view the individual statements inside all your policies. There isn't a way to automatically obtain that information in the Console.

### To update the description for an existing policy

This is available only through the API. A workaround is to create a new policy with the new description and delete the old policy.

### To update the statements in an existing policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.  
A list of the policies in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the policy you want to update.  
The policy's details and statements are displayed.
3. Either delete or add new statements (for the required format for statements, see [Policy Basics](#) and [Policy Syntax](#)). If you want to update an existing statement, create a new one with your desired changes and then delete the old one.

Your changes will go into effect typically within 10 seconds.

### To update the version date for an existing policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.  
A list of the policies in the compartment you're currently viewing is displayed. If you don't see the policy you're looking for, make sure you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the policy you want to update.  
The policy's details, version date, and statements are displayed.
3. Click **Update Version Date**.

4. Select **Keep Policy Current** if you'd like the policy to stay current with any future changes to the service's definitions of verbs and resources. Or if you'd prefer to limit access according to the definitions that were current on a specific date, select **Use Version Date** and enter that date in format YYYY-MM-DD format. For more information, see [Policy Language Version](#).
5. Click **Update Version Date**.

Your changes will go into effect typically within 10 seconds.

### To delete a policy



#### Tip

Remember that if you delete a policy and then create a new one with the same name, they'll be considered different policies because they'll have different OCIDs.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.  
A list of the policies in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. For the policy you want to delete, click **Delete**.
3. Confirm when prompted.

Your changes will go into effect typically within 10 seconds.

### To apply tags to a policy

For instructions, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



#### Note

##### *Updates Are Not Immediate Across All Regions*

Your IAM resources reside in your home region. To enforce policy across all regions, the IAM service replicates your resources in each region. Whenever you create or change a policy, user, or group, the changes take effect first in the home region, and then are propagated out to your other regions. It can take several minutes for changes to take effect in all regions. For example, assume you have a group with permissions to launch instances in the tenancy. If you add UserA to this group, UserA will be able to launch instances in your home region within a minute. However, UserA will not be able to launch instances in other regions until the replication process is complete. This process can take up to several minutes. If UserA tries to launch an instance before replication is complete, they will get a not authorized error.

Use these API operations to manage policies:

- [CreatePolicy](#)
- [ListPolicies](#)
- [GetPolicy](#)

- [UpdatePolicy](#)
- [DeletePolicy](#)

## Managing User Credentials

This topic describes the basics of working with Oracle Cloud Infrastructure Identity and Access Management (IAM) user credentials. If you're not already familiar with the available credentials, see [User Credentials](#).

### Working with Console Passwords and API Keys

Each user automatically has the ability to change or reset *their own* Console password, as well as manage *their own* API keys. An administrator does not need to create a policy to give a user those abilities.

To manage credentials for users other than yourself, you must be in the Administrators group or some other group that has permission to work with the tenancy. Having permission to work with a compartment within the tenancy is not sufficient. For more information, see [The Administrators Group and Policy](#).

IAM administrators (or anyone with permission to the tenancy) can use either the Console or the API to manage all aspects of both types of credentials, for themselves and all other users. This includes creating an initial one-time password for a new user, resetting a password, uploading API keys, and deleting API keys.

Users who are not administrators can manage *their own* credentials. In the Console, users can:

- Change or reset their own password.
- Upload an API key in the Console for their own use (and also delete their own API keys).

And with the API, users can:

- Reset their own password with [CreateOrResetUIPassword](#).
- Upload an additional API key to the IAM service for their own use with [UploadApiKey](#) (and also delete their own API keys with [DeleteApiKey](#)). Remember that a user can't use the API to change or delete their own credentials until they themselves upload a key in the Console, or an administrator uploads a key for that user in the Console or the API.

A user can have a maximum of three API keys at a time.

### Working with Auth Tokens



#### Note

"Auth tokens" were previously named "Swift passwords". Any Swift passwords you had created are now listed in the Console as auth tokens. You can continue to use the existing passwords.

Auth tokens are Oracle-generated token strings that you can use to authenticate with third-party APIs that do not support Oracle Cloud Infrastructure's signature-based authentication. Each user created in the IAM service automatically has the ability to create, update, and delete their own auth tokens in the Console or the API. An administrator does not need to create a policy to give a user those abilities. Administrators (or anyone with permission to the tenancy) also have the ability to manage auth tokens for other users.

Note that you cannot change your auth token to a string of your own choice. The token is always an Oracle-generated string.

Auth tokens do not expire. Each user can have up to two auth tokens at a time. To get an auth token in the Console, see [To create an auth token](#).

#### Using an Auth Token with Swift

Swift is the OpenStack object store service. If you already have an existing Swift client, you can use it with the Recovery Manager (RMAN) to back up an Oracle Database System (DB

System) database to Object Storage. You will need to get an auth token to use as your Swift password. When you sign in to your Swift client, you provide the following:

- Your Oracle Cloud Infrastructure Console user login
- Your Swift-specific auth token, provided by Oracle
- Your organization's Oracle tenant name

Any user of a Swift client that integrates with Object Storage needs permission to work with the service. If you're not sure if you have permission, contact your administrator. For information about policies, see [How Policies Work](#). For basic policies that enable use of Object Storage, see [Common Policies](#).

### Working with Customer Secret Keys



#### Note

"Customer Secret keys" were previously named "Amazon S3 Compatibility API keys". Any keys you had created are now listed in the Console as Customer Secret keys. You can continue to use the existing keys.

Object Storage provides an API to enable interoperability with Amazon S3. To use this Amazon S3 Compatibility API, you need to generate the signing key required to authenticate with Amazon S3. This special signing key is an Access Key/Secret Key pair. Oracle provides the Access Key that is associated with your Console user login. You or your administrator generates the Customer Secret key to pair with the Access Key.

Each user created in the IAM service automatically has the ability to create, update, and delete their own Customer Secret keys in the Console or the API. An administrator does not need to create a policy to give a user those abilities. Administrators (or anyone with permission to the tenancy) also have the ability to manage Customer Secret keys for other users.

Any user of the Amazon S3 Compatibility API with Object Storage needs permission to work with the service. If you're not sure if you have permission, contact your administrator. For information about policies, see [How Policies Work](#). For basic policies that enable use of Object Storage, see [Common Policies](#).

Customer Secret keys do not expire. Each user can have up to two Customer Secret keys at a time. To create keys using the Console, see [To create a Customer Secret key](#).

### Working with SMTP Credentials

Simple Mail Transfer Protocol (SMTP) credentials are needed in order to send email through the Email Delivery service. Each user is limited to a maximum of two SMTP credentials. If more than two are required, they must be generated on other existing users or additional users must be created.



#### Note

You cannot change your SMTP username or password to a string of your own choice. The credentials are always Oracle-generated strings.

Each user created in the IAM service automatically has the ability to create and delete their own SMTP credentials in the Console or the API. An administrator does not need to create a policy to give a user those abilities. Administrators (or anyone with permission to the tenancy) also have the ability to manage SMTP credentials for other users.



#### Tip

Although each user can create and delete their own credentials, it is a security best practice to create a new user and generate SMTP credentials on this user rather than generating SMTP credentials on your Console user that already has permissions assigned to it.

SMTP credentials do not expire. Each user can have up to two credentials at a time. To get SMTP credentials in the Console, see [To generate SMTP credentials](#).

For information about using the Email Delivery service, see [Overview of the Email Delivery Service](#).

## Using the Console

### To change your Console password

You're prompted to change your initial one-time password *the first time* you sign in to the Console. The following procedure is for changing your password again later.

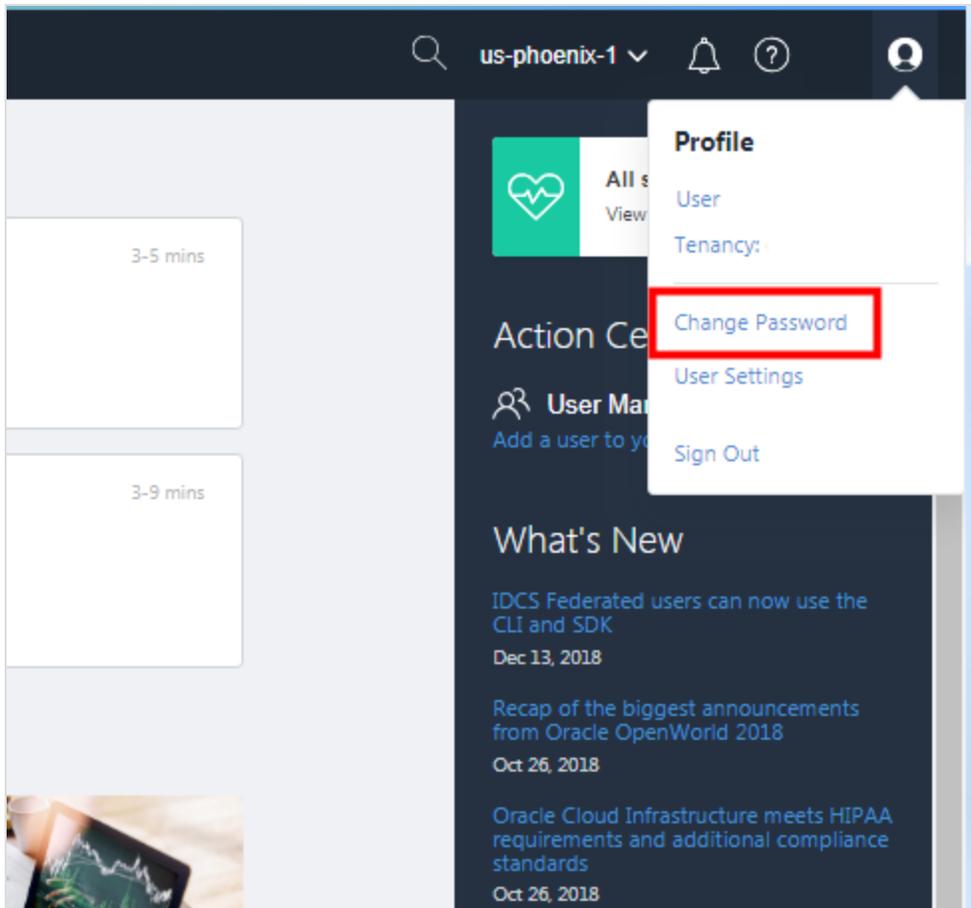


#### Note

*For Federated Users*

If your company uses an identity provider (other than Oracle Identity Cloud Service) to manage user logins and passwords, you can't use the Console to update your password. You do that with your identity provider.

1. Sign in to the Console using the Oracle Cloud Infrastructure Username and Password.
2. After you sign in, go to the top-right corner of the Console, open the **Profile** menu () and then click **Change Password**.



3. Enter the current password.
4. Follow the prompts to enter the new password, and then click **Save New Password**.

### To create or reset another user's Console password

If you're an administrator, you can use the following procedure to create or reset a user's password. The procedure generates a new one-time password that the user must change the next time they sign in to the Console.

1. View the user's details: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. Click **Create/Reset Password**.

The new one-time password is displayed. If you're an administrator performing the task for another user, you need to securely deliver the new password to the user. The user will be prompted to change their password the next time they sign in to the Console. If they don't change it within 7 days, the password will expire and you'll need to create a new one-time password for the user.

### To reset your password if you forgot it

If you have an email address in your user profile, you can use the **Forgot Password** link on the sign on page to have a temporary password sent to you. If you don't have an email address in your user profile, you must ask an administrator to reset your password for you.

### To unblock a user

If you're an administrator, you can unblock a user who has tried 10 times in a row to sign in to the Console unsuccessfully. See [To unblock a user](#).

### To upload an API signing key

The following procedure works for a regular user or an administrator. Administrators can upload an API key for either another user or themselves.

**Important**

The API key must be an **RSA key in PEM format (minimum 2048 bits)**. The PEM format looks something like this:

```
-----BEGIN PUBLIC KEY-----
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAAoTFqF...
...
-----END PUBLIC KEY-----
```

For more information about generating a public PEM key, see [Required Keys and OCIDs](#).

1. View the user's details:
  - If you're uploading an API key for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator uploading an API key for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. Click **Add Public Key**.
3. Paste the key's value into the window and click **Add**.  
The key is added and its fingerprint is displayed (example fingerprint: d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:13).

**Note**

When making API requests, you'll need the key's fingerprint, along with your tenancy's OCID and user OCID. See [Required Keys and OCIDs](#).

### To delete an API signing key

The following procedure works for a regular user or an administrator. Administrators can delete an API key for either another user or themselves.

1. View the user's details:
  - If you're deleting an API key for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator deleting an API key for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. For the API key you want to delete, click **Delete**.
3. Confirm when prompted.

The API key is no longer valid for sending API requests.

### To create an auth token

1. View the user's details:
  - If you're creating an auth token for yourself: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator creating an auth token for another user: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Auth Tokens**.
3. Click **Generate Token**.
4. Enter a description that indicates what this token is for, for example, "Swift password token".
5. Click **Generate Token**.  
The new token string is displayed.

6. Copy the token string immediately, because you can't retrieve it again after closing the dialog box.

If you're an administrator creating an auth token for another user, you need to securely deliver it to the user by providing it verbally, printing it out, or sending it through a secure email service.

### To delete an auth token

The following procedure works for a regular user or an administrator. Administrators can delete an auth token for either another user or themselves.

1. View the user's details:
  - If you're deleting an auth token for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator deleting an auth token for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Auth Tokens**.
3. For the auth token you want to delete, click **Delete**.
4. Confirm when prompted.

The auth token is no longer valid for accessing third-party APIs.

### To create a Customer Secret key

1. View the user's details:
  - If you're creating a Customer Secret key for yourself: Open the **Profile** menu () and click **User Settings**.

- If you're an administrator creating a Customer Secret key for another user: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Customer Secret Keys**.  
A Customer Secret key consists of an Access Key/Secret key pair. Oracle automatically generates the Access Key when you or your administrator generates the Secret Key to create the Customer Secret key.
  3. Click **Generate Secret Key**.
  4. Enter a friendly description for the key and click **Generate Secret Key**.  
The generated **Secret Key** is displayed in the **Generate Secret Key** dialog box. At the same time, Oracle generates the **Access Key** that is paired with the **Secret Key**. The newly generated Customer Secret key is added to the list of **Customer Secret Keys**.
  5. Copy the **Secret Key** immediately, because you can't retrieve the **Secret Key** again after closing the dialog box for security reasons.  
If you're an administrator creating a Secret Key for another user, you need to securely deliver it to the user by providing it verbally, printing it out, or sending it through a secure email service.
  6. Click **Close**.
  7. To show or copy the **Access Key**, click the **Show** or **Copy** action to the left of the **Name** of a particular Customer Secret key.

### To delete a Customer Secret key

The following procedure works for a regular user or an administrator. Administrators can delete a Customer Secret key for either another user or themselves.

1. View the user's details:
  - If you're deleting a Customer Secret key for *yourself*: Open the **Profile** menu () and click **User Settings**.

- If you're an administrator deleting a Customer Secret key for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **Customer Secret Keys**.
  3. For the Customer Secret key you want to delete, click **Delete**.
  4. Confirm when prompted.

The Customer Secret key is no longer available to use with the Amazon S3 Compatibility API.

### To generate SMTP credentials

1. View the user's details:
  - If you're generating SMTP credentials for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator generating SMTP credentials for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. Click **SMTP Credentials**.
3. Click **Generate SMTP Credentials**.
4. Enter a **Description** of the SMTP Credentials in the dialog box.
5. Click **Generate SMTP Credentials**. A user name and password is displayed.
6. Copy the user name and password for your records and click **Close**. Copy the credentials immediately, because you can't retrieve the password again after closing the dialog box for security reasons.

If you're an administrator creating the credential set for another user, you need to securely deliver it to the user by providing it verbally, printing it out, or sending it through a secure email service.

### To delete SMTP credentials

The following procedure works for a regular user or an administrator. Administrators can delete SMTP credentials for either another user or themselves.

1. View the user's details:
  - If you're deleting SMTP credentials for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator deleting SMTP credentials for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. On the left side of the page, click **SMTP Credentials**.
3. For the SMTP credentials you want to delete, click **Delete**.
4. Confirm when prompted.

The SMTP credentials are no longer available to use with the Email Delivery service.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use this API operation to manage Console passwords and access:

- [CreateOrResetUIPassword](#): This generates a new one-time Console password for the user. The next time the user signs in to the Console, they'll be prompted to change the password.
- [UpdateUserState](#): Unblocks a user who has tried to sign in 10 times in a row unsuccessfully.

Use these API operations to manage API signing keys:

- [ListApiKeys](#)
- [UploadApiKey](#)
- [DeleteApiKey](#)

Use these API operations to manage auth tokens:

- [CreateAuthToken](#)
- [UpdateAuthToken](#): You can only update the auth token's description, not change the token string itself.
- [ListAuthTokens](#)
- [DeleteAuthToken](#)

Use these API operations to manage Customer Secret keys:

- [CreateCustomerSecretKey](#)
- [UpdateCustomerSecretKey](#): You can only update the secret key's description, not change the key itself.
- [ListCustomerSecretKeys](#)
- [DeleteCustomerSecretKey](#)

Use these API operations to manage SMTP credentials:

- [CreateSmtCredential](#)
- [UpdateSmtCredential](#): You can only update the description.
- [ListSmtCredentials](#)
- [DeleteSmtCredential](#)

## Managing Authentication Settings

This topic describes how to set password policy rules for local IAM users in your tenancy.

### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing password policy.

To view authentication policy, you must be granted `inspect` access on the `authentication-policies` resource. For example:

```
Allow group GroupA to inspect authentication-policies in tenancy
```

To modify authentication policy, you must be granted the `AUTHENTICATION_POLICY_UPDATE` permission. This permission is included in the `manage` verb. For example:

```
Allow group GroupA to manage authentication-policies in tenancy
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

### Working with Password Policy Rules

A password policy that you set in the IAM service is applicable for all local (or non-federated) users.

When a user is created or when a user changes their password, the IAM service validates the password that is provided against the password policy to ensure that it meets the criteria for the policy. When a user logs in for the first time to change the password, or resets the password at any time, the password policy is evaluated and enforced.

#### **When Do Changes to Password Policy Rules Take Effect**

Changes to password policy rules take effect immediately so that the next time any user changes their password they must create a password that meets the criteria. Existing passwords will continue to work even if they would be invalid under the new rules. Users are not forced to change existing passwords to meet the new criteria. Passwords are evaluated against the rules only at the time they are created or changed.

## About the Password Policy Rules

The following table describes the rules that you can include in your password policy:

Rule	Setting Options	Default IAM Service Setting
Minimum password length	Minimum value is 8 (characters). Maximum value is 100.	12 characters
Special characters	Require passwords to contain at least 1 of the following special characters: !\"#\$%&'()*+,-./:;<=>?@[\\]^_`{ }~	Enforced
Lowercase characters	Require passwords to contain at least 1 lowercase alphabetic character a-z.	Enforced
Uppercase characters	Require passwords to contain at least 1 uppercase alphabetic character A-Z.	Enforced
Numeric characters	Require passwords to contain at least 1 number 0-9.	Enforced

Oracle recommends that you enforce all the password rules.

## Using the Console

### To edit password policy rules

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Authentication Settings**.  
The authentication settings for your tenancy are displayed.

2. Click **Edit**.
3. Enter the following to set the password policy:
  - **Minimum Password Length:** Enter a number to define the minimum number of characters that a user's password must contain. Allowed values are 8 through 100.
4. Select the **Password Rules** you want to enforce:
  - **Must contain at least 1 numeric character:** Select the check box to require at least 1 number (0-9) in the password.
  - **Must contain at least 1 special character:** Select the check box to require at least 1 special character. Special characters are: !\"#\$%&'()\*+,-./:;<=>?@[\\]^\_`{|}~
  - **Must contain at least 1 lowercase character:** Select the check box to require at least 1 lowercase alphabetic character (a-z).
  - **Must contain at least 1 uppercase character:** Select the check box to require at least 1 uppercase alphabetic character (A-Z).
5. Click **Save**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage password rules:

- [GetAuthenticationPolicy](#)
- [UpdateAuthenticationPolicy](#)

## Managing Multi-Factor Authentication

This topic describes how users can manage multi-factor authentication (MFA) in Oracle Cloud Infrastructure.

### Required IAM Policy

Only the user can enable multi-factor authentication (MFA) for their own account. Users can also disable MFA for their own accounts. Members of the Administrators group can disable MFA for other users, but they cannot enable MFA for another user.

### About Multi-Factor Authentication

Multi-factor authentication is a method of authentication that requires the use of more than one factor to verify a user's identity.

With MFA enabled in the IAM service, when a user signs in to Oracle Cloud Infrastructure, they are prompted for their user name and password, which is the first factor (something that they know). The user is then prompted to provide a second verification code from a registered MFA device, which is the second factor (something that they have). The two factors work together, requiring an extra layer of security to verify the user's identity and complete the sign-in process.

In general, MFA may include any two of the following:

- Something that you know, like a password.
- Something that you have, like a device.
- Something that you are, like your fingerprint.

The IAM service supports two-factor authentication using a password (first factor) and a device that can generate a time-based one-time password (TOTP) (second factor).

### General Concepts

Here's a list of the basic concepts you need to be familiar with.

### **MULTI-FACTOR AUTHENTICATION (MFA)**

Multi-factor authentication (MFA) is a method of authentication that requires the use of more than one factor to verify a user's identity. Examples of authentication factors are a password (something you know) and a device (something you have).

### **AUTHENTICATOR APP**

An app you install on your mobile device that can provide software-based secure tokens for identity verification. Examples of authenticator apps are Oracle Mobile Authenticator and Google Authenticator. To enable MFA for the IAM service, you'll need a device with an authenticator app installed. You'll use the app to register your device and then you'll use the same app (on the same device) to generate a time-based one-time passcode every time you sign in.

### **REGISTERED MOBILE DEVICE**

Multi-factor authentication is enabled for a specific user and for a specific device. The procedure to enable MFA for a user includes the registration of the mobile device. This same device must be used to generate the time-based one-time passcode every time the user signs in. If the registered mobile device becomes unavailable, an administrator must disable MFA for the user so that MFA can be re-enabled with a new device.

### **TIME-BASED ONE-TIME PASSWORD (TOTP)**

A TOTP is a password (or passcode) that is generated by an algorithm that computes a one-time password from a shared secret key and the current time, as defined in [RFC 6238](#). The authenticator app on your registered mobile device generates the TOTP that you need to enter every time you sign in to Oracle Cloud Infrastructure.

## Supported Authenticator Apps

The following authenticator apps have been tested with the Oracle Cloud Infrastructure IAM service:

- Oracle Mobile Authenticator
- Google Authenticator

You can find these apps in your mobile device's app store. You must install one of these apps on your mobile device before you can enable MFA.

### Working with MFA

Keep the following in mind when you enable MFA:

- You must install a supported authenticator app on the mobile device you intend to register for MFA.
- Each user must enable MFA for themselves using a device they will have access to every time they sign in. An administrator *cannot* enable MFA for another user.
- To enable MFA, you use your mobile device's authenticator app to scan a QR code that is generated by the IAM service and displayed in the Console. The QR code shares a secret key with the app to enable the app to generate TOTP's that can be verified by the IAM service.
- A user can register only one device to use for MFA.
- After you add your Oracle Cloud Infrastructure account to your authenticator app, the account name displays in the authenticator app as Oracle *<tenancy\_name>* - *<username>*.

### Restricting Access to Only MFA-Verified Users

You can restrict access to resources to only users that have been authenticated through the IAM service's time-based one-time password authentication. You set up this restriction in the policy that allows access to the resource.

To restrict the access granted through a policy to only MFA-verified users, add the following `where` clause to the policy:

```
where request.user.mfaTotpVerified='true'
```

For example, assume your company has this policy in place to allow GroupA to manage instances:

```
allow group GroupA to manage instance-family in tenancy
```

To enhance security, you want to ensure that only users who have been verified through MFA can manage instances. To restrict access to only these users, revise the policy statement as follows:

```
allow group GroupA to manage instance-family in tenancy where request.user.mfaTotpVerified='true'
```

With this policy in place, only the members of GroupA who have successfully signed in by entering both their password and the time-based one-time passcode generated by their registered mobile device, are allowed to access and manage instances. Users who have not enabled MFA and sign in using only their password, will not be allowed access to manage instances.

For information on writing policies, see [Policy Syntax](#).

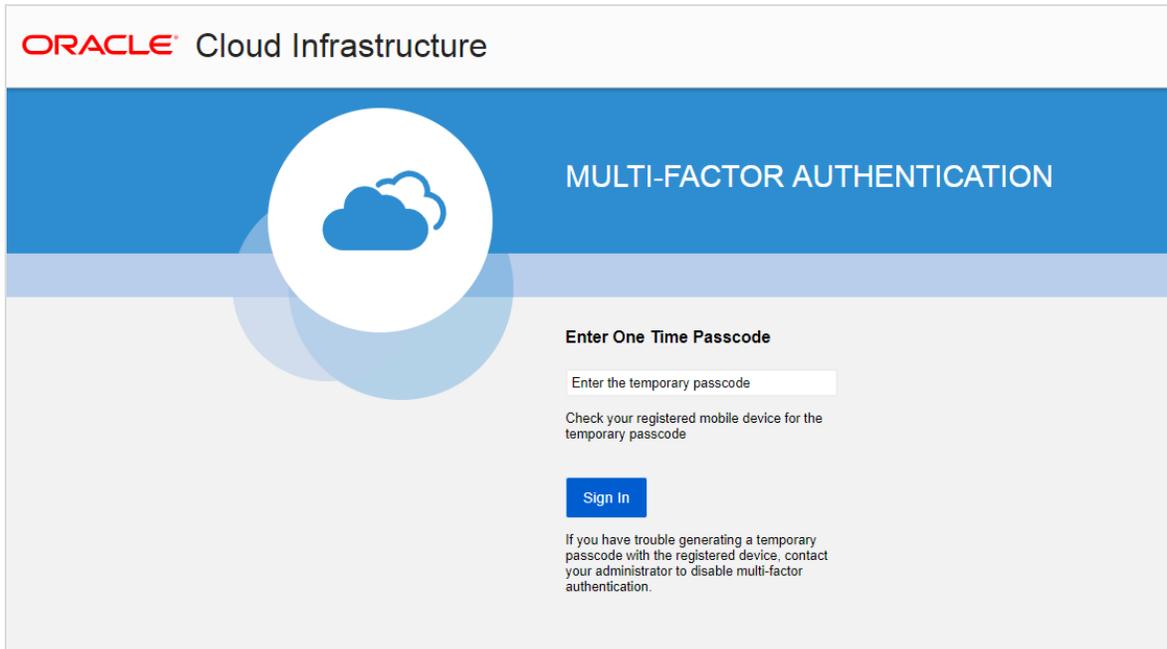
### Sign in Process After Enabling MFA

After you have enabled MFA, use one of the following procedures to sign in to Oracle Cloud Infrastructure:

#### To sign in using the Console

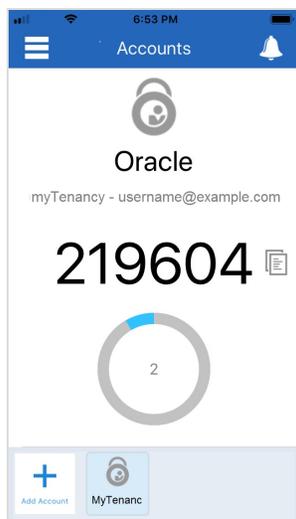
1. Navigate to the Console sign-in page.
2. Enter your Oracle Cloud Infrastructure **User Name** and **Password** and then click **Sign In**.

After your user name and password are authenticated, you have successfully supplied the first factor for authentication. The secondary authentication page displays and prompts you to enter a one-time passcode, as shown in the following screenshot.



The screenshot shows the Oracle Cloud Infrastructure Multi-Factor Authentication page. At the top left, the Oracle logo is followed by the text "Cloud Infrastructure". Below this is a blue header bar with a white circular icon containing a blue cloud. To the right of the icon, the text "MULTI-FACTOR AUTHENTICATION" is displayed in white. Below the header bar, the page has a light gray background. On the left side, there is a large, semi-transparent blue circular graphic. On the right side, the text "Enter One Time Passcode" is displayed in bold. Below this text is a white input field with the placeholder text "Enter the temporary passcode". Below the input field, the text "Check your registered mobile device for the temporary passcode" is displayed. Below this text is a blue button with the text "Sign In". At the bottom of the page, there is a small text block that reads: "If you have trouble generating a temporary passcode with the registered device, contact your administrator to disable multi-factor authentication."

3. Open the authenticator app on your registered mobile device and then open the account for your Oracle Cloud Infrastructure tenancy. The following screenshot shows an example from Oracle Mobile Authenticator.



4. Enter the passcode displayed by your authenticator app (for example, 219604) and then click **Sign In**.

**Important:** The authenticator app generates a new time-based one-time passcode every 30 seconds. You must enter a code while the code is still valid. If you miss the time window for one passcode, you can enter the next one that is generated. Just ensure that you enter the code that is currently displayed by your app.

### To sign in using the command line interface (CLI)

1. To sign in with the CLI, run the following command:

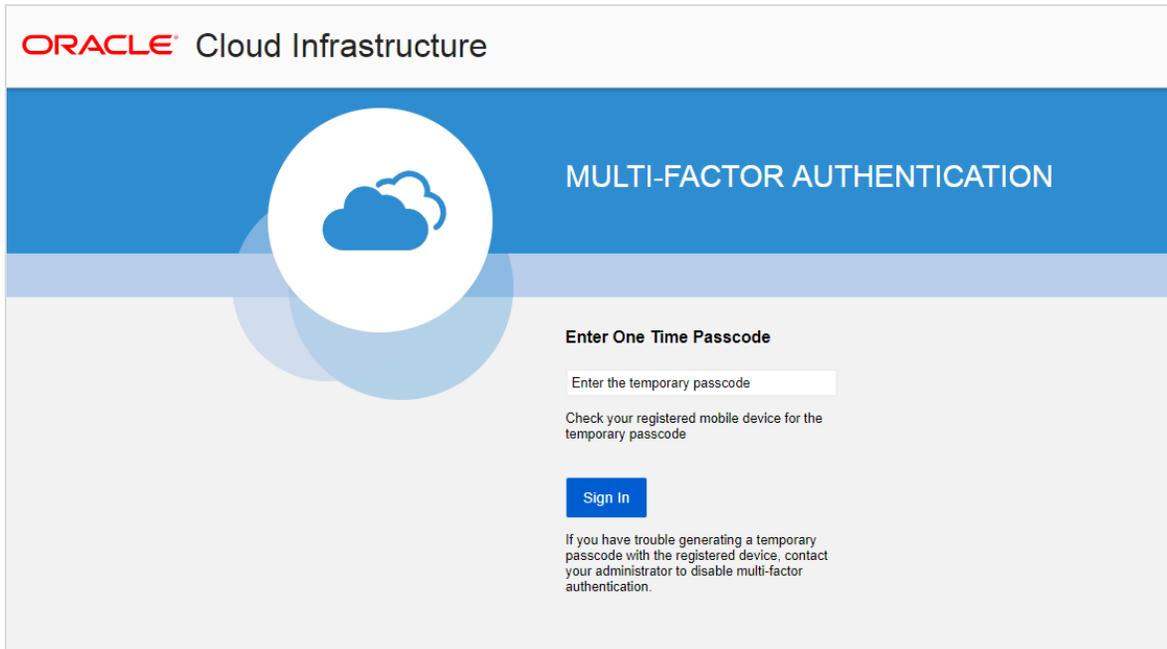
```
oci session authenticate --region US East (Ashburn)
```

A browser window opens, and a prompt instructs you to use the browser to sign in.

```
Please switch to newly opened browser window to log in!
```

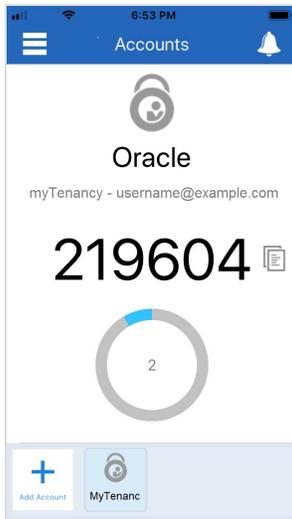
2. In the browser window, enter your Oracle Cloud Infrastructure **User Name** and **Password** and then click **Sign In**.

After your user name and password are authenticated, you have successfully supplied the first factor for authentication. The secondary authentication page displays and prompts you to enter a one-time passcode, as shown in the following screenshot.



The screenshot shows the Oracle Cloud Infrastructure Multi-Factor Authentication page. At the top left, the Oracle logo is followed by the text "Cloud Infrastructure". Below this is a blue header bar with a white circular icon containing a blue cloud. To the right of the icon, the text "MULTI-FACTOR AUTHENTICATION" is displayed in white. Below the header bar, the page has a light gray background. The main heading is "Enter One Time Passcode". Below this heading is a text input field with the placeholder text "Enter the temporary passcode". Underneath the input field is the instruction "Check your registered mobile device for the temporary passcode". A blue "Sign In" button is positioned below the instruction. At the bottom of the page, there is a small block of text: "If you have trouble generating a temporary passcode with the registered device, contact your administrator to disable multi-factor authentication."

3. Open the authenticator app on your registered mobile device and then open the account for your Oracle Cloud Infrastructure tenancy. The following screenshot shows an example from Oracle Mobile Authenticator.



4. Enter the passcode displayed by your authenticator app (for example, 219604) and then click **Sign In**.

**Important:** The authenticator app generates a new time-based one-time passcode every 30 seconds. You must enter a code while the code is still valid. If you miss the time window for one passcode, you can enter the next one that is generated. Just ensure that you enter the code that is currently displayed by your app.

After you authenticate, prompts instruct you to return to the CLI and enter the name of a profile.

5. In the CLI, type a name for the profile.



### Tip

For more information about working with the CLI, see [Quickstart](#) and [Getting Started with the Command Line Interface](#).

### What To Do If You Lose Your Registered Mobile Device

If you lose your registered mobile device, you will not be able to authenticate to Oracle Cloud Infrastructure through the Console. Contact your administrator to disable multi-factor authentication for your account. You can then repeat the process to enable multi-factor authentication with a new mobile device.

### Unblocking a User After Unsuccessful Sign-in Attempts

If a user tries 10 times in a row to sign in to the Console unsuccessfully, they will be automatically blocked from further sign-in attempts. An administrator can unblock the user in the Console (see [To unblock a user](#)) or with the [UpdateUserState](#) API operation.

### Disabling MFA

Each user can disable MFA for themselves. An administrator can also disable MFA for another user.



#### **Warning**

Do not disable MFA unless you are instructed to by your administrator.

### Using the Console

Use the following procedures to manage MFA in the Console.

#### To enable MFA for your user account

**Prerequisite:** You must install a [supported authenticator app](#) on the mobile device you intend to register for MFA.

1. In the upper-right corner of the Console, open the **Profile** menu () and then select **User Settings**. Your user details are displayed.
2. Click **Enable Multi-Factor Authentication**.
3. Scan the QR code displayed in the dialog with your mobile device's authenticator app.  
**Note:** If you close the browser, or if the browser crashes before you can enter the verification code, you must generate a new QR code and scan it again with your app. To generate a new QR code, click the **Enable Multi-Factor Authentication** button again.
4. In the **Verification Code** field, enter the code displayed on your authenticator app.
5. Click **Enable**.

Your mobile device is now registered with the IAM service and your account is enabled for MFA. Every time you sign in, you are prompted for your username and password first. After you provide the correct credentials, you will be prompted for a TOTP code generated by the authenticator app on your registered mobile device. *You must have your registered mobile device available every time you sign in to Oracle Cloud Infrastructure.*

### To disable MFA for your user account

1. In the upper-right corner of the Console, open the **Profile** menu () and then select **User Settings**. Your user details are displayed.
2. Click **Disable Multi-Factor Authentication**.
3. Confirm when prompted.

### To disable MFA for another user

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.  
A list of the users in your tenancy is displayed.
2. Click the user you want to update.  
The user's details are displayed.

3. Click **Disable Multi-Factor Authentication**.
4. Confirm when prompted.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



#### Note

*Updates Are Not Immediate Across All Regions*

Your IAM resources reside in your home region. To enforce policy across all regions, the IAM service replicates your resources in each region. Whenever you create or change a policy, user, or group, the changes take effect first in the home region, and then are propagated out to your other regions. It can take several minutes for changes to take effect in all regions.

Use these API operations to manage multi-factor authentication devices:

- [CreateMfaTotpDevice](#)
- [ListMfaTotpDevices](#)
- [GetMfaTotpDevice](#)
- [DeleteMfaTotpDevice](#)
- [ActivateMfaTotpDevice](#)
- [GenerateTotpSeed](#)

# CHAPTER 19 Key Management

This chapter explains how to create key vaults and encryption keys and how to manage and use them.

## Overview of Key Management

Oracle Cloud Infrastructure Key Management provides you with centralized management of the encryption of your data. You can use Key Management to create master encryption keys and data encryption keys, rotate keys to generate new cryptographic material, enable or disable keys for use in cryptographic operations, assign keys to resources, and use keys for encryption and decryption.

Oracle Cloud Infrastructure Object Storage, Oracle Cloud Infrastructure Block Volume, and Oracle Cloud Infrastructure File Storage integrate with Key Management to support encryption of data in buckets, block or boot volumes, and file systems. Integration with Oracle Cloud Infrastructure Identity and Access Management (IAM) lets you control who and what services can access which keys and what they can do with those keys. Oracle Cloud Infrastructure Audit integration gives you a way to monitor key usage. Audit tracks administrative actions on keys and vaults.

Keys are stored on highly available and durable hardware security modules (HSM) that meet Federal Information Processing Standards (FIPS) 140-2 Security Level 3 security certification. Key Management uses the Advanced Encryption Standard (AES) as its encryption algorithm and its keys are AES symmetric keys.

## Key Management Concepts

The following concepts are integral to understanding Key Management.

### **KEYS**

Keys are logical entities that represent one or more key versions that contain the cryptographic material used to encrypt and decrypt data, protecting the data where it is

stored. When processed as part of an encryption algorithm, a key specifies how to transform plaintext into ciphertext during encryption and how to transform ciphertext into plaintext during decryption. Conceptually, Key Management recognizes two types of encryption keys. You can create master encryption keys using the Console or API. Key Management stores those keys in a key vault. After you have a master encryption key, you can then use the API to generate data encryption keys that the service returns to you. Key Management introduces master encryption keys as an Oracle Cloud Infrastructure resource.

### **VAULTS**

Key vaults are logical entities where Key Management creates and durably stores your keys. Vaults are partitions on a hardware security module that are isolated from one another to ensure the security and integrity of the encryption keys that are stored on them. The type of vault you have determines features and functionality such as degrees of storage isolation, access to management and encryption, scalability, and pricing. At this time, the only type of vault you can create is a virtual private vault. Key Management designates vaults as an Oracle Cloud Infrastructure resource.

### **KEY VERSIONS**

Each master encryption key is automatically assigned a key version. When you rotate a key, Key Management generates a new key version. Periodically rotating keys limits the amount of data encrypted by one key version. Key rotation thereby reduces the risk if a key is ever compromised. A key's unique, Oracle-assigned identifier, called an Oracle Cloud ID (OCID), remains the same across rotations, but the key version enables Key Management to seamlessly rotate keys to meet any compliance requirements you might have. Although you can't use an older key version for encryption after you rotate it, the key version remains available to decrypt any data that it previously encrypted. Key Management removes the need for you to track which key version was used to encrypt what data because the key's ciphertext contains the information that Key Management requires for decryption.

### **HARDWARE SECURITY MODULES**

When you create a master encryption key using the Console or API, Key Management stores the key version within a hardware security module (HSM) to provide a layer of physical security. Any given key version, after it's created, is replicated within the service infrastructure as a measure of protection against hardware failures. Key versions are not otherwise stored anywhere else and cannot be exported from an HSM. Key Management uses HSMs that meet Federal Information Processing Standards (FIPS) 140-2 Security Level 3 security certification. This means that the HSM hardware is tamper-evident, has physical safeguards for tamper-resistance, requires identity-based authentication, and deletes keys from the device when it detects tampering.

### **ENVELOPE ENCRYPTION**

The data encryption key used to encrypt your data is, itself, encrypted with a master encryption key. This concept is known as envelope encryption. Oracle Cloud Infrastructure services do not have access to the plaintext data without interacting with Key Management and without access to the master encryption key that is protected by Oracle Cloud Infrastructure Identity and Access Management (IAM). For decryption purposes, Object Storage, Block Volume, and File Storage store only the encrypted form of the data encryption key.

## Regions and Availability Domains

You can use Key Management in the India West (Mumbai), South Korea Central (Seoul), Australia East (Sydney), Japan East (Tokyo), Canada Southeast (Toronto), Germany Central (Frankfurt), Switzerland North (Zurich), Brazil East (Sao Paulo), UK South (London), US East (Ashburn), and US West (Phoenix) regions. Unlike other Oracle Cloud Infrastructure services, however, Key Management does not have one regional endpoint for all API operations. The service has one regional endpoint for the provisioning service that handles create, update, and list operations for vaults. For create, update, and list operations for keys, service endpoints are distributed across multiple independent clusters.

Because Key Management has public endpoints, you can directly use data encryption keys generated by Key Management for cryptographic operations in your applications. However, if you want to use master encryption keys with a service that has integrated with Key

Management, you can do so only when the service and the key vault that holds the key both exist within the same region.

Key Management maintains copies of encryption keys across all availability domains within a region. This replication makes it possible for Key Management to generate keys even when an availability domain is unavailable.

### Private Access to Key Management

Key Management supports private access from Oracle Cloud Infrastructure resources in a virtual cloud network (VCN) through a service gateway. Setting up and using a service gateway on a VCN lets resources (such as the instances that your encrypted volumes are attached to) access public Oracle Cloud Infrastructure services such as Key Management without exposing them to the public internet. No internet gateway is required and resources can be in a private subnet and use only private IP addresses. For more information, see [Access to Oracle Services: Service Gateway](#).

### Resource Identifiers

Key Management introduces keys and vaults as Oracle Cloud Infrastructure resources. Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the REST API. Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#). Terraform does not currently support Key Management.

To access the Console, you must use a [supported browser](#). You can use the **Console** link at the top of this page to go to the sign-in page. You will be prompted to enter your cloud tenant, your user name, and your password.

For general information about using the API, see [REST APIs](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Limits on Key Management Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

### Managing Keys

This topic describes what you can do with keys and key versions in terms of managing their creation and usage. For information about how you can use keys in cryptographic operations, see [Using Keys](#). For information about what you can do with vaults where you store keys, see [Managing Vaults](#).

Management of keys includes the ability to do the following:

- Create keys
- View key details
- View a list of keys

## CHAPTER 19 Key Management

---

- View a list of key versions for a specific key
- Update a key name
- Manage a key's tags
- Enable keys for use in cryptographic operations
- Rotate keys to generate new cryptographic material
- Disable keys to prevent their usage in cryptographic operations
- Delete keys to permanently prevent their usage in cryptographic operations or assignment to resources
- Assign keys to specific resources
- Remove keys from their assignment to specific resources
- Move a key to a new compartment

### Required IAM Policy



#### Warning

Keys associated with volumes, buckets, and file systems will not work unless you authorize Oracle Cloud Infrastructure Block Volume, Oracle Cloud Infrastructure Object Storage, and Oracle Cloud Infrastructure File Storage to use keys on your behalf. Additionally, you must also authorize users to delegate key usage to these services in the first place. For more information, see [Let a user group delegate key usage in a compartment](#) and [Let Block Volume, Object Storage, and File Storage services encrypt and decrypt volumes, buckets, and file systems](#) in [Common Policies](#).

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK,

CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For typical policies that give access to keys and vaults, see [Let security admins manage vaults and keys](#).

Also, be aware that a policy statement with `inspect vaults` gives the specified group the ability to see *all* information about the vaults. Likewise, a policy statement with `inspect keys` gives the specified group the ability to see *all* information about the keys. For more information, see [Details for the Key Management Service](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring the traffic associated with your master encryption keys, see [Key Management Metrics](#).

### Moving Resources to a Different Compartment

You can move keys from one compartment to another. After you move a key to a new compartment, inherent policies apply immediately and affect access to the key and key versions. Moving a key doesn't affect access to the vault that a key is associated with. You can move a vault from one compartment to another independently of moving any of its keys. For more information, see [Managing Compartments](#).

### Using the Console

#### To create a new key

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment where you want to create a key.
3. From the list of vaults in the compartment, do one of the following:
  - Click the name of the vault where you want to create a key.
  - Create a new vault for the key by following the instructions in [To create a new vault](#), and then click the name of the vault.
4. Click **Keys**, and then click **Create Key**.
5. In the **Create Key** dialog box, choose a compartment from the **Create in Compartment** list. (Keys can exist outside the compartment the vault is in.)
6. Click **Name**, and then enter a name to identify the key. Avoid entering any confidential information in this field.
7. Specify the key length, in bits, by choosing a length from the **Key Shape: Length** list.
8. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
9. When you are finished, click **Create Key**.

#### To view key details

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.

2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you're interested in.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, and then click the name of the key for which you want to see configuration details.
5. The console displays the following information:
  - **OCID:** The unique, Oracle-assigned ID of the key.
  - **Created:** The date and time when you initially created the key.
  - **Compartment:** The unique, Oracle-assigned ID of the compartment that contains the vault that contains the key.
  - **Vault:** The unique, Oracle-assigned ID of the vault that contains the key.
  - **Key Version:** The unique, Oracle-assigned ID of the key version.
  - **Algorithm:** The encryption algorithm used by the key.
  - **Length:** The number of bits in the key length.

### To view a list of keys

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the keys you're interested in.
3. From the list of vaults in the compartment, click the vault name.
4. To see a list of keys in this vault, click **Keys**.

### To view a list of key versions

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.

2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you're interested in.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, click the name of the key for which you want to see a list of key versions, and then click **Versions**.

### To change the name of a key

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to rename.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to rename, and then click the Actions icon (three dots) for that key.
5. In the **Actions** menu, click **Edit Name**.
6. In the **Edit Key Name** dialog box, click **Name**, and then enter a new name. Avoid entering any confidential information in this field.
7. When you are finished, click **Update**.

### To manage a key's tags

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key for which you want to manage tags.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to manage, and then click the key name.

5. On the **Key Details** page, click the **Tags** tab to view or edit existing tags. Or, click **Add Tag(s)** to add new ones.

### To enable a key

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to enable.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to enable, and then select the check box next to the key name.
5. In the **Actions** menu, click **Enable**.

### To rotate a key

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to rotate.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to rotate, and then select the check box next to the key name.
5. In the **Actions** menu, click **Rotate Key**. (You can only rotate keys in an enabled state.)
6. In the **Confirm** dialog box, click **Rotate Key**.

### To disable a key

1. Open the navigation menu. Under the **Governance and Administration** group, go to

**Security** and click **Key Management**.

2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to disable.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to disable, and then click the Actions icon (three dots) for that key.
5. In the **Actions** menu, click **Disable**.

To delete a key



### Warning

When you set a key to the Pending Deletion state, anything encrypted by that key immediately becomes inaccessible. The key also cannot be assigned or unassigned to any resources or otherwise updated. When the key is deleted, all key material and metadata is irreversibly destroyed. Before you delete a key, either assign a new key to resources currently encrypted by the key or preserve your data another way. If you want to restore use of a key before it is permanently deleted, you can cancel its deletion.

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to delete.
3. From the list of vaults in the compartment, click the vault name.

4. Click **Keys**, locate the key you want to delete, and then click the Actions icon (three dots) for that key.
5. In the **Actions** menu, click **Delete Key**.
6. Confirm that you want to delete the key by clicking the box and then typing the key name.
7. Schedule when you want Key Management to delete the key. By default, the service schedules keys for deletion 30 days from the current date and time. You can set a range between 7 days and 30 days.
8. When you're ready, click **Delete Key**. If needed, you can restore use of the key and access to encrypted resources and data by canceling the scheduled deletion.

### To cancel the deletion of a key



#### Tip

You can only cancel the deletion of a key that's in a Pending Deletion state.

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault with the key you want to delete.
3. From the list of vaults in the compartment, click the vault name.
4. Click **Keys**, locate the key you want to delete, and then click the Actions icon (three dots) for that key.
5. In the **Actions** menu, click **Cancel Deletion**.
6. Confirm that you want to cancel the key's deletion by clicking **Cancel Deletion**. Access to the key and any resources or data encrypted by the key are restored when key returns to an Enabled state.

### To assign a key to a new Object Storage bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment where you want to create a bucket that's encrypted with a Key Management master encryption key.
3. Click **Create Bucket**, and then follow the instructions in [To create a bucket](#) in [Managing Buckets](#).

### To assign a key to an existing Object Storage bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the bucket that you want to encrypt with a Key Management master encryption key.
3. From the list of buckets, click the bucket name.
4. Do one of the following:
  - If the bucket already has a key assigned to it, next to **Encryption Key**, click **Edit** to assign a different key.
  - If the bucket does not already have a key assigned to it, next to **Encryption Key**, click **Assign**.
5. Choose the vault compartment, vault, key compartment, and key.
6. When you are finished, click **Assign** or **Update**, as appropriate.

### To assign a key to a new Block Volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment where you want to create a block volume that's encrypted with a Key Management master encryption key.

3. Click **Create Block Volume**, and then follow the instructions in [Creating a Volume](#).

### To assign a key to an existing Block Volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the block volume that you want to encrypt with a Key Management master encryption key.
3. From the list of volumes, click the volume name.
4. If the volume is currently attached to an instance, click **Detach from Instance**. Follow the instructions in the **Detach Block Volume** dialog box as appropriate, click **Continue Detachment**, and then click **OK**.
5. Then, do one of the following:
  - If the volume already has a key assigned to it, next to **Encryption Key**, click **Edit** to assign a different key.
  - If the volume does not already have a key assigned to it, next to **Encryption Key**, click **Assign**.
6. Choose the vault compartment, vault, key compartment, and key.
7. When you are finished, click **Assign** or **Update**, as appropriate.

### To assign a key to a new file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment where you want to create a file system that's encrypted with a Key Management master encryption key.
3. Click **Create File System**, and then follow the instructions in [Creating File Systems](#).

### To create a Compute instance with an encrypted boot volume



#### Note

These instructions assume you have already fulfilled the prerequisites for creating an instance.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment where you want to create an instance with a boot volume that's encrypted with a Key Management master encryption key.
3. Click **Create Instance**, and then follow the instructions under [Using the Console in Launching an Instance](#).

### To assign a key to an existing boot volume



#### Note

To assign a key to an existing boot volume, you must first detach the boot volume from any instance. However, you can only detach a boot volume from an instance when the instance is stopped. For more information, see [Detaching a Boot Volume](#) and [Stopping and Starting an Instance](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the boot volume that you want to encrypt with a Key Management master encryption key.

3. From the list of volumes, click the volume name.
4. Do one of the following:
  - If the volume already has a key assigned to it, next to **Encryption Key**, click **Edit** to assign a different key.
  - If the volume does not already have a key assigned to it, next to **Encryption Key**, click **Assign**.
5. Choose the vault compartment, vault, key compartment, and key.
6. When you are finished, click **Assign** or **Update**, as appropriate.

### To remove a key assignment from a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the bucket from which you want to remove a Key Management key assignment.
3. From the list of buckets, click the bucket name.
4. Next to **Encryption Key**, click **Unassign**.
5. In the **Confirm** dialog box, click **OK** to remove the key assignment from the bucket.

### To remove a key assignment from a Block Volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Block Storage** and click **Block Volumes**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the block volume from which you want to remove a Key Management key assignment.
3. From the list of volumes, click the volume name.
4. Next to **Encryption Key**, click **Unassign**.
5. In the **Confirm** dialog box, click **OK** to remove the key assignment from the volume.

### To remove a key assignment from a boot volume

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Boot Volumes**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the boot volume from which you want to remove a Key Management key assignment.
3. From the list of volumes, click the volume name.
4. Next to **Encryption Key**, click **Unassign**.
5. In the **Confirm** dialog box, click **OK** to remove the key assignment from the volume.

### To change a key assignment for a file system

1. Open the navigation menu. Under **Core Infrastructure**, click **File Storage** and then click **File Systems**.
2. Under **List Scope**, in the **Compartment** list, choose the compartment that contains the file system from which you want to remove or change a Key Management key assignment.
3. From the list of file systems, click the file system name.
4. Next to **Encryption Key**, click **Edit**.
5. If you want to use Oracle-managed keys:
  - In **Encryption Type**, select **Encrypt using Oracle-managed keys**.
6. If you want to assign a different customer-managed key:
  - In **Encryption Type**, select **Encrypt using customer-manged keys**.
  - Choose the vault compartment, vault, key compartment, and key.
7. When you are finished, click **Save Changes**.

### To move a key to a different compartment

1. Open the navigation menu. Under the **Governance and Administration** group, go to

**Security** and click **Key Management**.

2. Under **Table Scope**, in the **Compartment** list, choose the compartment that contains the master encryption key that you want to move.
3. Find the key in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.
6. If there are alarms monitoring the key, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).



#### Tip

Each vault has a unique endpoint for create, update, and list operations for keys. This endpoint is referred to as the control plane URL or management endpoint. Each vault also has a unique endpoint for cryptographic operations. This endpoint is known as the data plane URL or the cryptographic endpoint. When using the CLI for key operations, you must provide the appropriate endpoint for the type of operation. To retrieve a vault's endpoints, see instructions in [To view vault configuration details](#).

## CHAPTER 19 Key Management

---

### To create a new key

Open a command prompt and run `oci kms management key create` to create a new key:

```
oci kms management key create --compartment-id <target_compartment_id> --display-name <key_name> --key-shape <key_encryption_information> --endpoint <control_plane_url>
```

For example, on a MacOS or Linux machine:

```
oci kms management key create --compartment-id ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name key-1 --key-shape '{"algorithm":"AES","length":"16"}' --endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

Or, for example, on a Windows machine:

```
oci kms management key create --compartment-id ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name key-1 --key-shape '{"algorithm":"AES","length":"16"}' --endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```



#### Warning

Avoid entering confidential information in the key name.

### To create a new key with resource tags

Open a command prompt and run `oci kms management key create` with one or both of the `--defined-tags` and `--freeform-tags` options to create a new key with resource tags:

```
oci kms management key create --compartment-id <target_compartment_id> --display-name <key_name> --key-shape <JSON_formatted_key_encryption_information> --defined-tags <JSON_formatted_defined_tag> --freeform-tags <JSON_formatted_freeform_tag> --endpoint <control_plane_url>
```

For example, on a MacOS or Linux machine:

```
oci kms management key create --compartment-id ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name key-1 --key-shape '{"algorithm":"AES","length":"16"}' --defined-tags '{"Operations": {"CostCenter":"42"}}' --
```

## CHAPTER 19 Key Management

---

```
freeform-tags '{"Department":"Finance"}' --endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

Or, for example, on a Windows machine:

```
oci kms management key create --compartment-id
ocid1.compartment.oc1..exampleexample25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name key-1
--key-shape '{"algorithm":"AES","length":"16"}' --defined-tags '{"Operations":
{"CostCenter":"42"}' --freeform-tags '{"Department":"Finance"}' --endpoint
https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```



### Warning

Avoid entering confidential information in the key name.

### To view a key's details

Open a command prompt and run `oci kms management key get` to view a specific key's details:

```
oci kms management key get --key-id <key_OCID> --endpoint <control_plane_url>
```

For example:

```
oci kms management key get --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To view a list of keys

Open a command prompt and run `oci kms management key list` to list keys in a vault:

```
oci kms management key list --compartment-id <target_compartment_id> --endpoint <control_plane_url>
```

For example:

## CHAPTER 19 Key Management

---

```
oci kms management key list --compartment-id
ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --endpoint
https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To view a list of key versions

Open a command prompt and run `oci kms management key-version list` to view a list of key versions for a specific key:

```
oci kms management key-version list --key-id <key_OCID> --endpoint <control_plane_url>
```

For example:

```
oci kms management key-version list --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To change the name of a key

Open a command prompt and run `oci kms management key update` to edit a key's name.



#### Warning

Avoid entering confidential information in the key name.

```
oci kms management key update --key-id <key_OCID> --display-name <new_key_name> --endpoint <control_
plane_url>
```

For example:

```
oci kms management key update --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
display-name key-A --endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To enable a key

Open a command prompt and run `oci kms management key enable` to enable a key:

## CHAPTER 19 Key Management

---

```
oci kms management key enable --key-id <target_key_id> --endpoint <control_plane_url>
```

For example:

```
oci kms management key enable --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To rotate a key

Open a command prompt and run `oci kms management key rotate` to rotate a key:

```
oci kms management key rotate --key-id <target_key_id> --endpoint <control_plane_url>
```

For example:

```
oci kms management key rotate --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To disable a key

Open a command prompt and run `oci kms management key disable` to disable a key:

```
oci kms management key disable --key-id <target_key_id> --endpoint <control_plane_url>
```

For example:

```
oci kms management key disable --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To delete a key



#### Warning

When you set a key to the Pending Deletion state,



anything encrypted by that key immediately becomes inaccessible. The key also cannot be assigned or unassigned to any resources or otherwise updated. When the key is deleted, all key material and metadata is irreversibly destroyed. Before you delete a key, either assign a new key to resources currently encrypted by the key or preserve your data another way. If you want to restore use of a key before it is permanently deleted, you can cancel its deletion.

Open a command prompt and run `oci kms management key schedule-deletion` to schedule a key's deletion:

```
oci kms management key schedule-deletion --key-id <target_key_id> --endpoint <control_plane_url>
```

For example:

```
oci kms management key schedule-deletion --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

By default, the service schedules keys for deletion 30 days from the current date and time. You can set a range between 7 days and 30 days. For example:

```
oci kms management key schedule-deletion --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --time-
of-deletion 2019-06-30T10:00:00Z --endpoint https://exampleaaacu2-management.kms.us-ashburn-
1.oraclecloud.com
```

### To cancel the deletion of a key



#### Tip

You can only cancel the deletion of a key that's in a



Pending Deletion state.

Open a command prompt and run `oci kms management key cancel-deletion` to cancel a key's scheduled deletion:

```
oci kms management key cancel-deletion --key-id <target_key_id> --endpoint <control_plane_url>
```

For example:

```
oci kms management key cancel-deletion --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
endpoint https://exampleaaacu2-management.kms.us-ashburn-1.oraclecloud.com
```

### To move a key to a different compartment

Open a command prompt and run `oci kms management key change-compartment` to move a master encryption key from one compartment to another within the same tenancy:

```
oci kms management key change-compartment --key-id <target_key_id> --compartment-id <new_compartment_id>
```

For example:

```
oci kms management key change-compartment --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
compartment-id ocid1.compartment.oc1..examplelexample25qrlpo4agcmothkbggmuz2zzum45ibploqtabwk3zz
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage keys:

- [CreateKey](#)
- [DisableKey](#)

- [EnableKey](#)
- [GetKey](#)
- [UpdateKey](#)
- [CreateKeyVersion](#)
- [GetKeyVersion](#)
- [ListKeys](#)
- [ListKeyVersions](#)
- [CancelKeyDeletion](#)
- [ScheduleKeyDeletion](#)
- [ChangeKeyCompartment](#)

## Managing Vaults

This topic describes what you can do with vaults. For information about what you can do with keys, see [Managing Keys](#).

Key Management lets you create virtual private vaults in your tenancy. A virtual private vault provides you with a dedicated partition in a hardware security module (HSM), offering a level of storage isolation for encryption keys that's effectively equivalent to a virtual independent HSM. As such, virtual private vaults have dedicated administration and users. They also restrict other tenants from accessing your encryption keys.

Vault management tasks include the following:

- Creating a vault
- Viewing vault configuration details
- Updating the vault name
- Managing vault tags
- Deleting a vault
- Moving a vault to a new compartment

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For typical policies that give access to keys and vaults, see [Let security admins manage vaults and keys](#).

Also, be aware that a policy statement with `inspect vaults` gives the specified group the ability to see *all* information about the vaults. Likewise, a policy statement with `inspect keys` gives the specified group the ability to see *all* information about the keys. For more information, see [Details for the Key Management Service](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Moving Resources to a Different Compartment

You can move vaults from one compartment to another. After you move a vault to a new compartment, inherent policies apply immediately and affect access to the vault. Moving a vault doesn't affect access to any keys that the vault contains. You can move a key from one compartment to another independently of moving the vault it's associated with. For more information, see [Managing Compartments](#).

### Using the Console

#### To view vault configuration details

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault you want to view.
3. From the list of vaults in the compartment, click the name of the vault.
4. The console displays the following information:
  - **OCID:** The unique, Oracle-assigned ID of the vault.
  - **Created:** The date and time when you initially created the vault.
  - **Compartment:** The unique, Oracle-assigned ID of the compartment that contains the vault.
  - **Vault Type:** The type of vault. At this time, you can only have a virtual private vault (VIRTUAL\_PRIVATE).
  - **Control Plane URL:** The service endpoint for [CreateKey](#), [CreateKeyVersion](#), [EnableKey](#), [DisableKey](#), [UpdateKey](#), [ListKeys](#), [ListKeyVersions](#), [GetKey](#), and [GetKeyVersion](#) operations.
  - **Data Plane URL:** The service endpoint for [Encrypt](#), [Decrypt](#), and [GenerateDataEncryptionKey](#) operations.

#### To create a new vault

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment where you want to create the vault.
3. Click **Create Vault**.

4. In the **Create Vault** dialog box, click **Name**, and then enter a display name for the vault. Avoid entering any confidential information in this field.
5. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. When you are finished, click **Create**.

### To change a vault name

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault you want to rename.
3. From the list of vaults in the compartment, click the name of the vault.
4. On the **Vault Details** page, click **Edit Name**.
5. In the **Edit Vault Name** dialog box, click **Name**, and then enter a new display name for the vault. Avoid entering any confidential information in this field.
6. When you are finished, click **Save**.

### To manage a vault's tags

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault for which you want to manage tags.
3. From the list of vaults in the compartment, click the name of the vault.

4. On the **Vault Details** page, click the **Tags** tab to view or editing existing tags. Or, click **Add Tag(s)** to add new ones.

### To delete a vault



#### Note

When you delete a vault, the vault and all its associated keys go into a pending deletion state until the waiting period expires. By default, this is 30 days, but can be set from a minimum of 7 days up to a maximum of 30 days. When a vault is deleted, all its associated keys are also deleted.

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault you want to delete.
3. From the list of vaults in the compartment, click the name of the vault.
4. On the **Vault Details** page, click **Delete**.
5. To confirm that you want to delete the vault, type the name of the vault, and then choose the date and time you want the vault to be deleted.
6. When you are finished, click **Delete Vault**.

### To cancel the deletion of a vault

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.

2. Under **List Scope**, in the **Compartment** list, click the name of the compartment that contains the vault that's in a pending deletion state.
3. From the list of vaults in the compartment, click the name of the vault.
4. On the **Vault Details** page, click **Cancel Deletion**.
5. To confirm that you want to cancel deletion of the vault, click **Cancel Deletion**.

### To move a vault to a different compartment

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Under **Table Scope**, in the **Compartment** list, choose the compartment that contains the vault that you want to move.
3. Find the vault in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To view vault configuration details

Open a command prompt and run `oci kms management vault get` to view the configuration details for a vault:

```
oci kms management vault get --vault-id <target_vault_id>
```

For example:

## CHAPTER 19 Key Management

---

```
oci kms management vault get --vault-id
ocid1.vault.region1.sea.exampleaaaacu2.examplesrcvbtqe5wgrxn2jua3olmeausn5fauxseubwu5my5tf3w3j33edq
```

### To create a new vault

Open a command prompt and run `oci kms management vault create` to create a new vault:

```
oci kms management vault create --compartment-id <target_compartment_id> --display-name <vault_name> -
-vault-type VIRTUAL_PRIVATE
```

For example:

```
oci kms management vault create --compartment-id
ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name
vault-1 --vault-type VIRTUAL_PRIVATE
```



#### Warning

Avoid entering confidential information in the vault name.

### To create a new vault with resource tags

Open a command prompt and run `oci kms management vault create` with one or both of the `--defined-tags` and `--freeform-tags` options to create a new vault with resource tags:

```
oci kms management vault create --compartment-id <target_compartment_id> --display-name <vault_name> -
-vault-type VIRTUAL_PRIVATE --defined-tags <JSON_formatted_defined_tag> --freeform-tags <JSON_
formatted_freeform_tag>
```

For example, on a MacOS or Linux machine:

```
oci kms management vault create --compartment-id
ocid1.compartment.oc1..example1example25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name
```

## CHAPTER 19 Key Management

---

```
vault-1 --vault-type VIRTUAL_PRIVATE --defined-tags '{"Operations": {"CostCenter": "42"}}' --freeform-tags '{"Department": "Finance"}'
```

Or, for example, on a Windows machine:

```
oci kms management vault create --compartment-id
ocid1.compartment.oc1..exampleexample25qrlpo4agcmothkbgqgmuz2zzum45ibploqtabwk3zz --display-name
vault-1 --vault-type VIRTUAL_PRIVATE --defined-tags '{"Operations\": {"CostCenter\":"42\"}}' --
freeform-tags '{"Department\":"Finance\"}'
```



### Warning

Avoid entering confidential information in the vault name.

### To change a vault name

Open a command prompt and run `oci kms management vault update` to change a vault's name:

```
oci kms management vault update --vault-id <target_vault_id>
```

For example:

```
oci kms management vault update --vault-id
ocid1.vault.region1.sea.exampleaaaacu2.examplesrcvbtqe5wgrxn2jua3olmeausn5fauxseubwu5my5tf3w3j33edq --
display-name new-vault-name
```

### To delete a vault

Open a command prompt and run `oci kms management vault schedule-deletion` to delete a vault:

```
oci kms management vault schedule-deletion --vault-id <target_vault_id>
```

For example:

## CHAPTER 19 Key Management

---

```
oci kms management vault schedule-deletion --vault-id
ocid1.vault.region1.sea.exampleaaacu2.examplesrcvbtqe5wgrxn2jua3olmeausn5fauxseubwu5my5tf3w3
```

When you delete a vault, the vault and all its associated keys go into a pending deletion state until the waiting period expires. By default, this is 30 days, but can be set from a minimum of 7 days up to a maximum of 30 days. When a vault is deleted, all its associated keys are also deleted.

### To cancel the deletion of a vault

Open a command prompt and run `oci kms management vault cancel-deletion` to cancel the pending deletion of a vault:

```
oci kms management vault cancel-deletion --vault-id <target_vault_id>
```

For example:

```
oci kms management vault cancel-deletion --vault-id
ocid1.vault.region1.sea.exampleaaacu2.examplesrcvbtqe5wgrxn2jua3olmeausn5fauxseubwu5my5tf3w3
```

### To move a vault to a different compartment

Open a command prompt and run `oci kms management vault change-compartment` to move a vault from one compartment to another within the same tenancy:

```
oci kms management vault change-compartment --vault-id <target_vault_id> --compartment-id <new_
compartment_id>
```

For example:

```
oci kms management vault change-compartment --vault-id
ocid1.vault.region1.sea.exampleaaacu2.examplesrcvbtqe5wgrxn2jua3olmeausn5fauxseubwu5my5tf3w3 --
compartment-id ocid1.compartment.oc1..exampleexample25qrlpo4agcmothkbggmuz2zzum45ibploqtabwk3zz
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage vaults:

- [CancelVaultDeletion](#)
- [CreateVault](#)
- [GetVault](#)
- [ListVaults](#)
- [ScheduleVaultDeletion](#)
- [UpdateVault](#)
- [ChangeVaultCompartment](#)

## Using Keys

This topic describes what you can do with keys in terms of cryptographic operations. For information about managing keys, see [Managing Keys](#). For information about managing the vaults in which you store keys, see [Managing Vaults](#).

Cryptographic operations include the following:

- Encrypting data
- Decrypting data
- Generating data encryption keys

You can use either the command line interface (CLI) or API to perform cryptographic operations.

### Required IAM Policy



#### Warning

Keys associated with volumes and buckets will not work unless you authorize Oracle Cloud Infrastructure Block Volume and Oracle Cloud Infrastructure Object Storage to use keys on your behalf. Additionally, you must also authorize users to delegate key usage to these services in the first place. For more information, see [Let a user group delegate key usage in a compartment](#) and [Let Block Volume and Object Storage services encrypt and decrypt volumes and buckets](#) in [Common Policies](#).

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For typical policies that give access to keys and vaults, see [Let security admins manage vaults and keys](#).

Also, be aware that a policy statement with `inspect vaults` gives the specified group the ability to see *all* information about the vaults. Likewise, a policy statement with `inspect keys` gives the specified group the ability to see *all* information about the keys. For more information, see [Details for the Key Management Service](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring](#)

## CHAPTER 19 Key Management

---

[Overview](#) and [Notifications Overview](#).

For information about monitoring the traffic associated with your master encryption keys, see [Key Management Metrics](#).

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).



#### Tip

Each vault has a unique endpoint for create, update, and list operations for keys. This endpoint is referred to as the control plane URL or management endpoint. Each vault also has a unique endpoint for cryptographic operations. This endpoint is known as the data plane URL or the cryptographic endpoint. When using the CLI for key operations, you must provide the appropriate endpoint for the type of operation. To retrieve a vault's endpoints, see instructions in [To view vault configuration details](#).

### To encrypt data by using your Key Management master encryption key

Open a command prompt and run `oci kms crypto encrypt` to encrypt data:

```
oci kms crypto encrypt --key-id <key_OCID> --plaintext <base64_string> --endpoint <data_plane_url>
```

For example:

```
oci kms crypto encrypt --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux21mqz245gezevsq --
plaintext VGhlIHFlaWNrIGJyb3duIGZveCBqdWlwcYBvdmVyIHROZSBsYXp5IGRvZy4= --endpoint https://exampleaaacu3-
crypto.kms.us-ashburn-1.oraclecloud.com
```

## CHAPTER 19 Key Management

---

Optionally, you can include the `associated-data` option to provide an encryption context that might contain useful, but non-secret, information about the encrypted data. That information is associated with the encrypted data such that the data cannot be decrypted without it, providing an additional layer of protection. Associated data must be properly formatted JSON.

```
oci kms crypto encrypt --key-id
ocidl.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
plaintext VGhlIHFlaWNrIGJyb3duIGZveCBqdWlwcYBvdmVyIHRoZSBsYXp5IGRvZy4= --associated-data '
{"CustomerId":"12345", "Custom Data":"custom data"}' --endpoint https://exampleaaacu3-crypto.kms.us-
ashburn-1.oraclecloud.com
```

### To decrypt data by using your Key Management master encryption key

Open a command prompt and run `oci kms crypto decrypt` to decrypt data:

```
oci kms crypto decrypt --key-id <key_OCID> --plaintext <base64_string> --endpoint <data_plane_url>
```

For example:

```
oci kms crypto decrypt --key-id
ocidl.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
plaintext VGhlIHFlaWNrIGJyb3duIGZveCBqdWlwcYBvdmVyIHRoZSBsYXp5IGRvZy4= --endpoint https://exampleaaacu3-
crypto.kms.us-ashburn-1.oraclecloud.com
```

If the data you want to decrypt had an encryption context associated with it at the time of encryption, the same encryption context is required to decrypt the data. For example, the `--associated-data` in the following sample matches what was provided in the preceding sample command for encrypting data.

```
oci kms crypto decrypt --key-id
ocidl.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --
plaintext VGhlIHFlaWNrIGJyb3duIGZveCBqdWlwcYBvdmVyIHRoZSBsYXp5IGRvZy4= --associated-data '
{"CustomerId":"12345", "Custom Data":"custom data"}' --endpoint https://exampleaaacu3-crypto.kms.us-
ashburn-1.oraclecloud.com
```

### To generate a data encryption key from your Key Management master encryption key

Open a command prompt and run `oci kms crypto generate-data-encryption-key` to

## CHAPTER 19 Key Management

---

generate a data encryption key that you can then use to encrypt and decrypt data:

```
oci kms crypto generate-data-encryption-key --key-id <key_OCID> --key-shape <key_encryption_information> --include-plaintext-key <Boolean_value> --endpoint <data_plane_url>
```

For example:

```
oci kms crypto generate-data-encryption-key --key-id
ocid1.key.region1.sea.exampleaaacu2.examplesmtpsugmoy4m5cvblugmizcoeu2nfc6b3zfaux2lmqz245gezevsq --key-
shape file://path/to/json/file --include-plaintext-key true --endpoint https://exampleaaacu3-
crypto.kms.us-ashburn-1.oraclecloud.com
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to use keys in cryptographic operations:

- [Decrypt](#)
- [Encrypt](#)
- [GenerateDataEncryptionKey](#)

### Key Management Metrics

You can monitor the usage of your Key Management service master encryption keys by using metrics, alarms, and [notifications](#). For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

This topic describes the metrics emitted by the Key Management service in the `oci_kms_keys` namespace.

Resources: master encryption keys.

### Overview of the Key Management Service Metrics

The Key Management service metrics help you measure the success and error count of cryptographic operations. You can use metrics data to diagnose and troubleshoot problems with keys.

To view a default set of metrics charts in the Console, navigate to the key that you're interested in, and then click **Metrics**. You also can use the Monitoring service to create [custom queries](#).

### Prerequisites

**IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_kms\_keys

The metrics listed in the following table are automatically available for any master encryption keys that you create. You do not need to enable monitoring on the resource to get these metrics.

Key Management service metrics for keys include the following dimensions:

**RESOURCEDISPLAYNAME**

The friendly name of the resource to which the metrics apply.

**RESOURCEID**

The OCID of the resource to which the metrics apply.

## CHAPTER 19 Key Management

### RESPONSECODE

The HTTP response code to the cryptographic operation to which the metrics apply.

Metric	Metric Display Name	Unit	Description	Dimensions
EncryptResponseCount	<b>Encrypt Response Count</b>	count	HTTP responses received by the service for Encrypt calls.	resourceDisplayName resourceId responseCode
DecryptResponseCount	<b>Decrypt Response Count</b>	count	HTTP responses received by the service for Decrypt calls.	
GenerateDataEncryptionKeyResponseCount*	<b>GenerateDataEncryptionKey Response Count</b>	count	HTTP responses received by the service for GenerateDataEncryptionKey calls.	

### Using the Console

To view default metric charts for a single master encryption key

1. Open the navigation menu. Under the **Governance and Administration** group, go to **Security** and click **Key Management**.
2. Click a key vault to view the master encryption keys it contains.
3. Click a key name to view its details.
4. Under **Resources**, click **Metrics**.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for multiple master encryption keys

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that contains the master encryption keys that you're interested in.
3. For **Metric Namespace**, select **oci\_kms\_keys**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.



#### Tip

If there are multiple master encryption keys in the compartment, the charts default to show a separate line for each master encryption key. You can instead show a single line aggregated across all master encryption keys in the compartment by selecting the **Aggregate Metric Streams** check box.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

## CHAPTER 19 Key Management

---

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 20 Load Balancing

This chapter explains how to set up a load balancer.

## Overview of Load Balancing

The Oracle Cloud Infrastructure Load Balancing service provides automated traffic distribution from one entry point to multiple servers reachable from your virtual cloud network (VCN). The service offers a load balancer with your choice of a public or private IP address, and provisioned bandwidth.

A load balancer improves resource utilization, facilitates scaling, and helps ensure high availability. You can configure multiple load balancing policies and application-specific health checks to ensure that the load balancer directs traffic only to healthy instances. The load balancer can reduce your maintenance window by draining traffic from an unhealthy application server before you remove it from service for maintenance.

## How Load Balancing Works

The Load Balancing service enables you to create a public or private load balancer within your VCN. A public load balancer has a public IP address that is accessible from the internet. A private load balancer has an IP address from the hosting subnet, which is visible only within your VCN. You can configure multiple listeners for an IP address to load balance transport Layer 4 and Layer 7 (TCP and HTTP) traffic. Both public and private load balancers can route data traffic to any backend server that is reachable from the VCN.

### **Public Load Balancer**

To accept traffic from the internet, you create a public load balancer. The service assigns it a public IP address that serves as the entry point for incoming traffic. You can associate the public IP address with a friendly DNS name through any DNS vendor.

A public load balancer is regional in scope. If your region includes multiple availability domains, a public load balancer requires either a regional subnet (recommended) or two availability domain-specific (AD-specific) subnets, each in a separate availability domain. With a regional subnet, the Load Balancing service creates a primary load balancer and a standby load balancer, each in a different availability domain, to ensure accessibility even during an availability domain outage. If you create a load balancer in two AD-specific subnets, one subnet hosts the primary load balancer and the other hosts a standby load balancer. If the primary load balancer fails, the public IP address switches to the secondary load balancer. The service treats the two load balancers as equivalent and you cannot specify which one is "primary".

Whether you use regional or AD-specific subnets, each load balancer requires one private IP address from its host subnet. The Load Balancing service supplies a floating public IP address to the primary load balancer. The floating public IP address does not come from your backend subnets.

If your region includes only one availability domain, the service requires just one subnet, either regional or AD-specific, to host both the primary and standby load balancers. The primary and standby load balancers each require a private IP address from the host subnet, in addition to the assigned floating public IP address. If there is an availability domain outage, the load balancer has no failover.



### Warning

You cannot specify a [private subnet](#) for your public load balancer.

### Private Load Balancer

To isolate your load balancer from the internet and simplify your security posture, you can create a private load balancer. The Load Balancing service assigns it a private IP address that serves as the entry point for incoming traffic.

When you create a private load balancer, the service requires only one subnet to host both the

primary and standby load balancers. The load balancer can be regional or AD-specific, depending on the scope of the host subnet. The load balancer is accessible only from within the VCN that contains the host subnet, or as further restricted by your security rules.

The assigned floating private IP address is local to the host subnet. The primary and standby load balancers each require an extra private IP address from the host subnet.

If there is an availability domain outage, a private load balancer created in a regional subnet within a multi-AD region provides failover capability. A private load balancer created in an AD-specific subnet, or in a regional subnet within a single availability domain region, has no failover capability in response to an availability domain outage.

### All Load Balancers

Your load balancer has a backend set to route incoming traffic to your Compute instances. The backend set is a logical entity that includes:

- A list of backend servers.
- A load balancing policy.
- A health check policy.
- Optional SSL handling.
- Optional session persistence configuration.

The backend servers (Compute instances) associated with a backend set can exist anywhere, as long as the associated network security groups (NSGs), security lists, and route tables allow the intended traffic flow.

If your VCN uses network security groups (NSGs), you can associate your load balancer with an NSG. An NSG has a set of security rules that controls allowed types of inbound and outbound traffic. The rules apply only to the resources in the group. Contrast NSGs with a security list, where the rules apply to all the resources in any subnet that uses the list. For more information about NSGs, see [Network Security Groups](#).

If you prefer to use security lists for your VCN, the Load Balancing service can suggest appropriate security list rules. You also can configure them yourself through the Networking service. See [Security Lists](#) for more information.

See [Security Rules](#) for detailed information comparing NSGs and security lists.

Oracle recommends that you create your load balancer in a regional subnet.

Oracle recommends that you distribute your backend servers across all availability domains within the region.

To create a minimal system with a functioning load balancer, you must:

- For a public load balancer, create a VCN with an internet gateway and a public regional subnet.



### Warning

You cannot specify a [private subnet](#) for your public load balancer.

- For a private load balancer, create a VCN with at least one private subnet.
- Create at least two Compute instances, each in a separate availability domain.
- Create a load balancer.
- Create a backend set with a health check policy.
- Add backend servers (Compute instances) to the backend set.
- Create a listener, with optional SSL handling.
- Update the load balancer subnet security rules so they allow the intended traffic.



### **Note**

#### *Private IP Address Consumption*

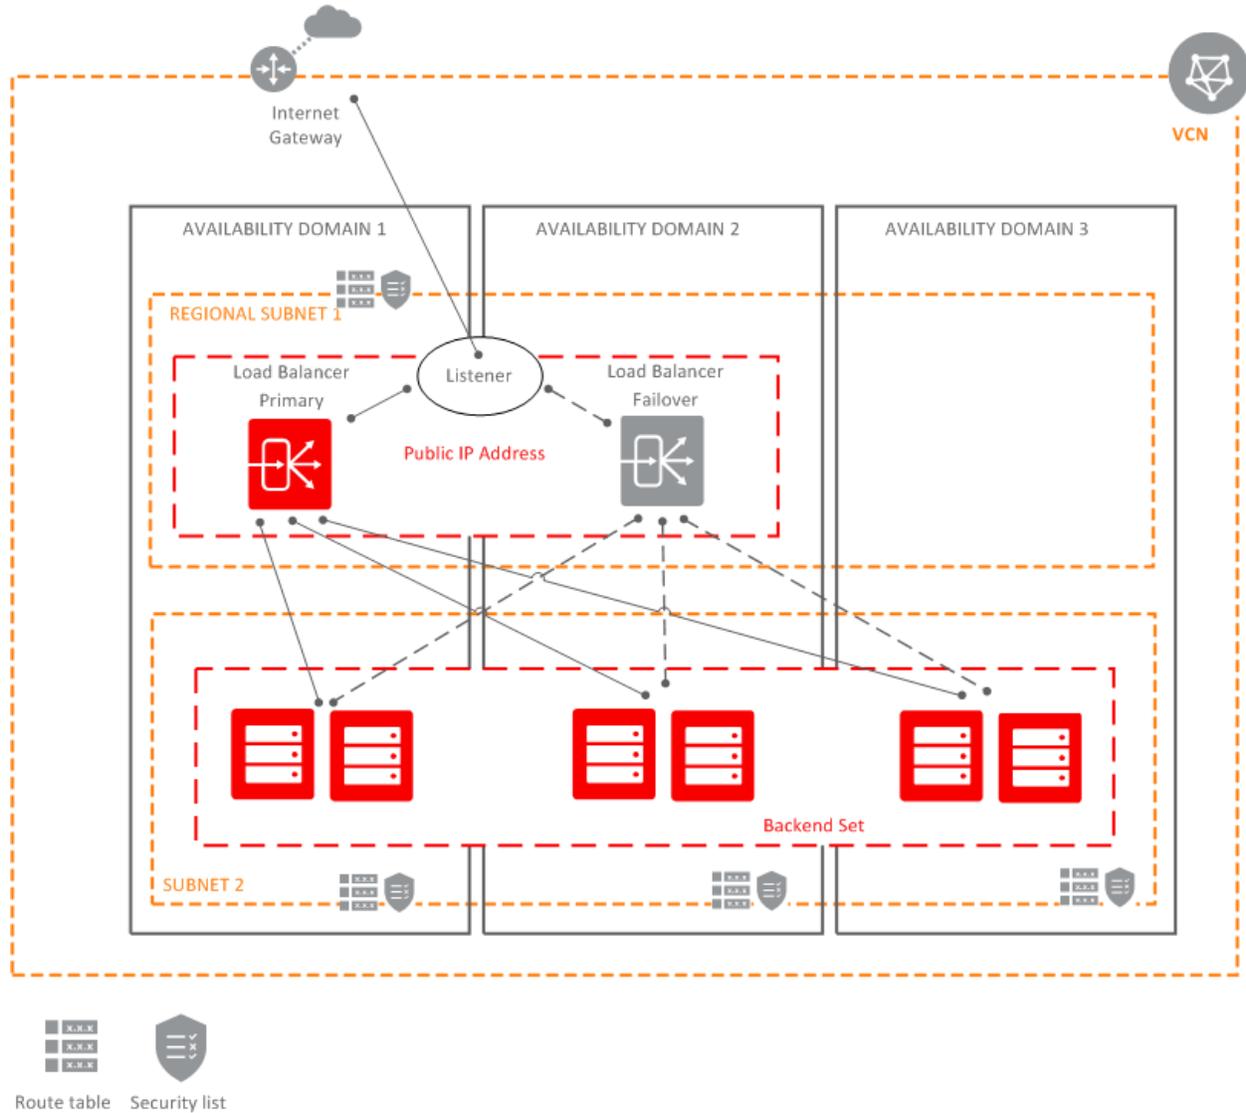
A public load balancer created in one public subnet consumes two private IP addresses from the host subnet.

A public load balancer created in two public subnets consumes two private IP addresses, one from each host subnet.

A private load balancer created in a single subnet consumes three private IP addresses from the host subnet.

## CHAPTER 20 Load Balancing

The following diagram provides a high-level view of a simple public load balancing system configuration. Far more sophisticated and complex configurations are common.



### Load Balancing Concepts

The following concepts are essential to working with Load Balancing.

#### **BACKEND SERVER**

An application server responsible for generating content in reply to the incoming TCP or HTTP traffic. You typically identify application servers with a unique combination of overlay (private) IPv4 address and port, for example, 10.10.10.1:8080 and 10.10.10.2:8080.

For more information, see [Managing Backend Servers](#).

#### **BACKEND SET**

A logical entity defined by a list of backend servers, a load balancing policy, and a health check policy. SSL configuration is optional. The backend set determines how the load balancer directs traffic to the collection of backend servers.

For more information, see [Managing Backend Sets](#).

#### **CERTIFICATES**

If you use HTTPS or SSL for your listener, you must associate an SSL server certificate (X.509) with your load balancer. A certificate enables the load balancer to terminate the connection and decrypt incoming requests before passing them to the backend servers.

For more information, see [Managing SSL Certificates](#).

### HEALTH CHECK

A test to confirm the availability of backend servers. A health check can be a request or a connection attempt. Based on a time interval you specify, the load balancer applies the health check policy to continuously monitor backend servers. If a server fails the health check, the load balancer takes the server temporarily out of rotation. If the server subsequently passes the health check, the load balancer returns it to the rotation.

You configure your health check policy when you [create a backend set](#). You can configure TCP-level or HTTP-level health checks for your backend servers.

- TCP-level health checks attempt to make a TCP connection with the backend servers and validate the response based on the connection status.
- HTTP-level health checks send requests to the backend servers at a specific URI and validate the response based on the status code or entity data (body) returned.

The service provides application-specific health check capabilities to help you increase availability and reduce your application maintenance window.

For more information on health check configuration, see [Editing Health Check Policies](#).

### HEALTH STATUS

An indicator that reports the general health of your load balancers and their components.

For more information, see the **Health Status** section of [Editing Health Check Policies](#).

### LISTENER

A logical entity that checks for incoming traffic on the load balancer's IP address. You configure a listener's protocol and port number, and the optional SSL settings. To handle TCP, HTTP, and HTTPS traffic, you must configure multiple listeners.

Supported protocols include:

- TCP
- HTTP/1.0

- HTTP/1.1

For more information, see [Managing Load Balancer Listeners](#).

### **LOAD BALANCING POLICY**

A load balancing policy tells the load balancer how to distribute incoming traffic to the backend servers. Common load balancer policies include:

- Round robin
- Least connections
- IP hash

For more information, see [How Load Balancing Policies Work](#).

### **PATH ROUTE SET**

A set of path route rules to route traffic to the correct backend set without using multiple listeners or load balancers.

For more information, see [Managing Request Routing](#).

### **REGIONS AND AVAILABILITY DOMAINS**

The Load Balancing service manages application traffic across availability domains within a region. A region is a localized geographic area, and an availability domain is one or more data centers located within a region. A region is composed of several availability domains.

For more information, see [Regions and Availability Domains](#).

### **SESSION PERSISTENCE**

A method to direct all requests originating from a single logical client to a single backend web server.

For more information, see [Session Persistence](#).

### SHAPE

A template that determines the load balancer's total pre-provisioned maximum capacity (bandwidth) for ingress plus egress traffic. Available shapes include 10Mbps, 100 Mbps, 400 Mbps, and 8000 Mbps.

The 10Mbps shape is Always Free eligible. For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).



#### Tip

Pre-provisioned maximum capacity applies to aggregated connections, not to a single client attempting to use the full bandwidth.

### SSL

Secure Sockets Layer (SSL) is a security technology for establishing an encrypted link between a client and a server. You can apply the following SSL configurations to your load balancer:

#### SSL TERMINATION

The load balancer handles incoming SSL traffic and passes the unencrypted request to a backend server.

#### END TO END SSL

The load balancer terminates the SSL connection with an incoming traffic client, and then initiates an SSL connection to a backend server.

#### SSL TUNNELING

If you configure the load balancer's listener for TCP traffic, the load balancer tunnels incoming SSL connections to your application servers.

Load Balancing supports the TLS 1.2 protocol with a default setting of **strong** cipher strength. The default supported ciphers include:

- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-SHA256
- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES256-SHA256
- DHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES128-SHA256

For more information, see [Managing SSL Certificates](#).

### **SUBNET**

A subdivision you define in a VCN, such as 10.0.0.0/24 and 10.0.1.0/24. A subnet can span a region or exist within in a single availability domain. A subnet consists of a contiguous range of IP addresses that do not overlap with other subnets in the VCN. For each subnet, you specify the routing and security rules that apply to it.

For more information on subnets, see [VCNs and Subnets](#) and [Public vs. Private Subnets](#).

### **TAGS**

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### **VIRTUAL HOSTNAME**

A virtual server name applied to a listener to enhance request routing.

For more information, see [Managing Request Routing](#).

### **VIRTUAL CLOUD NETWORK (VCN)**

A private network that you set up in the Oracle data centers, with firewall rules and specific types of communication gateways that you can choose to use. A VCN covers a single, contiguous IPv4 CIDR block of your choice in the [allowed IP address ranges](#).

You need at least one virtual cloud network before you launch a load balancer.

For information about setting up virtual cloud networks, see [Overview of Networking](#).

### **VISIBILITY**

Specifies whether your load balancer is public or private.

#### **PUBLIC**

A public load balancer has a public IP address that clients can access from the internet.

#### **PRIVATE**

A private load balancer has a private IP address from a VCN local subnet. Clients can access the private load balancer using methods and technology that can provide access to a private IP, such as:

- Cross-VCN (via LPG peering)
- From another region (via RPC)
- From on-prem (via FC private peering)

For more information, see [Managing a Load Balancer](#).

### **WORK REQUEST**

An object that reports on the current state of a Load Balancing request.

The Load Balancing service handles requests asynchronously. Each request returns a work request ID (OCID) as the response. You can view the work request item to see the status of the request.

For more information, see [Viewing the State of a Work Request](#).

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring the traffic passing through your load balancer, see [Load Balancing Metrics](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For

example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Limits on Load Balancing Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

Other limits include:

- You cannot dynamically change the load balancer shape to handle more incoming traffic. You can use the API or Console to create a load balancer with the new shape information.
- You cannot convert an AD-specific load balancer to a regional load balancer or the reverse.
- The Load Balancing services supports IPv6 addresses for load balancers in the US Government Cloud only. IPv6 support is [only for the load balancer itself, and not the backend](#).
- The maximum number of concurrent connections is limited when you use stateful security rules for your load balancer subnets. In contrast, there is no theoretical limit on concurrent connections if you use stateless security rules. The practical limitations depend on various factors. The larger your load balancer shape, the greater the connection capacity. Other considerations include system memory, TCP timeout periods, TCP connection state, and so forth.



### Tip

To accommodate high-volume traffic, Oracle strongly recommends that you use [stateless security rules](#) for your load balancer subnets.

- Each load balancer has the following configuration limits:
  - One IP address
  - 16 backend sets
  - 512 backend servers per backend set
  - 1024 backend servers total
  - 16 listeners

## How Load Balancing Policies Work

After you [create a load balancer](#), you can apply policies to control traffic distribution to your backend servers. The Load Balancing service supports three primary policy types:

- [Round Robin](#)
- [Least Connections](#)
- [IP Hash](#)

When processing load or capacity varies among backend servers, you can refine each of these policy types with backend server *weighting*. Weighting affects the proportion of requests directed to each server. For example, a server weighted '3' receives three times the number of connections as a server weighted '1'. You assign weights based on criteria of your choosing, such as each server's traffic-handling capacity.

Load balancer policy decisions apply differently to TCP load balancers, cookie-based session persistent HTTP requests (sticky requests), and non-sticky HTTP requests.

- A TCP load balancer considers policy and weight criteria to direct an initial incoming request to a backend server. All subsequent packets on this connection go to the same endpoint.
- An HTTP load balancer configured to handle cookie-based session persistence forwards requests to the backend server specified by the cookie's session information.
- For non-sticky HTTP requests, the load balancer applies policy and weight criteria to every incoming request and determines an appropriate backend server. Multiple requests from the same client could be directed to different servers.

### Round Robin

Round Robin is the default load balancer policy. This policy distributes incoming traffic sequentially to each server in a backend set list. After each server has received a connection, the load balancer repeats the list in the same order.

Round Robin is a simple load balancing algorithm. It works best when all the backend servers have similar capacity and the processing load required by each request does not vary significantly.

### Least Connections

The Least Connections policy routes incoming non-sticky request traffic to the backend server with the fewest active connections. This policy helps you maintain an equal distribution of active connections with backend servers. As with the round robin policy, you can assign a weight to each backend server and further control traffic distribution.



### Tip

In TCP use cases, a connection can be active but have no current traffic. Such connections do not serve as a good load metric.

## IP Hash

The IP Hash policy uses an incoming request's source IP address as a hashing key to route non-sticky traffic to the same backend server. The load balancer routes requests from the same client to the same backend server as long as that server is available. This policy honors server weight settings when establishing the initial connection.

IP Hash ensures that requests from a particular client are always directed to the same backend server, as long as it is available.

You cannot add a backend server marked as **Backup** to a backend set that uses the IP Hash policy.



### Warning

Multiple clients that connect to the load balancer through a proxy or NAT router appear to have the same IP address. If you apply the IP Hash policy to your backend set, the load balancer routes traffic based on the incoming IP address and sends these proxied client requests to the same backend server. If the proxied client pool is large, the requests could flood a backend server.

# Connection Management

Oracle Cloud Infrastructure load balancers support connection multiplexing. The load balancer can route many incoming requests from multiple clients to the destination backend server through a few (one or multiple) backend connections.

After your load balancer connects a client to a backend server, the connection can be closed due to inactivity. Also, you can configure load balancer listeners to control the maximum idle time allowed during each TCP connection or HTTP request and response pair. Oracle recommends that you do not allow your backend servers to close connections to the load balancer.

## Highlights

Three different timeout settings affect your load balancer's behavior:

- **[Keep-alive setting between the load balancer and backend server](#)**  
The load balancer closes backend server connections that are idle for more than 300 seconds.
- **Keep-alive setting between the load balancer and the client**  
The Load Balancing service sets the keep-alive value to maintain the connection for 10,000 transactions or until it has been idle for 65 seconds, whichever limit occurs first. You cannot change the value of this setting.
- **[Idle timeout](#)**  
You can set the duration of the idle timeout when you [create a listener](#). This setting applies to the time allowed between two successive receive or two successive send network input/output operations during the HTTP request-response phase.

## Keep-Alive Settings

The load balancing service does not honor keep-alive settings from backend servers. The load balancer closes backend server connections that are idle for more than 300 seconds. Oracle recommends that you do not allow your backend servers to close connections to the load

balancer. To prevent possible 502 errors, ensure that your backend servers do not close idle connections in less than 310 seconds.

The Load Balancing service sets the keep-alive value to maintain the connection for 10,000 transactions or until it has been idle for 65 seconds, whichever limit occurs first. You cannot change the value of this setting.

### Connection Configuration

When you [create a TCP or HTTP listener](#), you can specify the maximum idle time in seconds. This setting applies to the time allowed between two successive receive or two successive send network input/output operations during the HTTP request-response phase. If the configured timeout has elapsed with no packets sent or received, the client's connection is closed. For HTTP and WebSocket connections, a send operation does not reset the timer for receive operations and a receive operation does not reset the timer for send operations.



#### Tip

This timeout setting does not apply to idle time between a completed response and a subsequent HTTP request.

The default timeout values are:

- 300 seconds for TCP listeners.
- 60 seconds for HTTP listeners.

Modify the timeout parameter if either the client or the backend server requires more time to transmit data. Some examples include:

- The client sends a database query to the backend server and the database takes over 300 seconds to execute. Therefore, the backend server does not transmit any data within 300 seconds.

- The client uploads data using the HTTP protocol. During the upload, the backend does not transmit any data to the client for more than 60 seconds.
- The client downloads data using the HTTP protocol. After the initial request, it stops transmitting data to the backend server for more than 60 seconds.
- The client starts transmitting data after establishing a WebSocket connection, but the backend server does not transmit data for more than 60 seconds.
- The backend server starts transmitting data after establishing a WebSocket connection, but the client does not transmit data for more than 60 seconds.

The maximum timeout value is 7200 seconds. Contact [My Oracle Support](#) to file a service request if you want to increase this limit for your tenancy. For more information, see [Service Limits](#).

## HTTP "X-" Headers

HTTP requests and responses often include header fields that provide contextual information about the message. [RFC 2616](#) defines a standard set of HTTP header fields. Some non-standard header fields, which begin with `x-`, are common. The Load Balancing service adds or modifies the following `x-` headers when it passes requests to your servers.

### X-Forwarded-For

Provides a list of connection IP addresses.

The load balancer appends the last remote peer address to the `X-Forwarded-For` field from the incoming request. A comma and space precede the appended address. If the client request header does not include an `X-Forwarded-For` field, this value is equal to the `X-Real-IP` value. The original requesting client is the first (left-most) IP address in the list, assuming that the incoming field content is trustworthy. The last address is the last (most recent) peer, that is, the machine from which the load balancer received the request. The format is:

```
X-Forwarded-For: <original_client>, <proxy1>, <proxy2>
```

## CHAPTER 20 Load Balancing

---

Example incoming field:

```
X-Forwarded-For: 202.1.112.187
```

Example field with appended proxy IP address:

```
X-Forwarded-For: 202.1.112.187, 192.168.0.10
```

### X-Forwarded-Host

Identifies the original host and port requested by the client in the `Host` HTTP request header. This header helps you determine the original host, since the hostname or port of the reverse proxy (load balancer) might differ from the original server handling the request.

```
X-Forwarded-Host: www.oracle.com:8080
```

### X-Forwarded-Port

Identifies the listener port number that the client used to connect to the load balancer. For example:

```
X-Forwarded-Port: 443
```

### X-Forwarded-Proto

Identifies the protocol that the client used to connect to the load balancer, either `http` or `https`. For example:

```
X-Forwarded-Proto: https
```

### X-Real-IP

Identifies the client's IP address. For the Load Balancing service, the "client" is the last remote peer.

Your load balancer intercepts traffic between the client and your server. Your server's access logs, therefore, include only the load balancer's IP address. The `X-Real-IP` header provides the client's IP address. For example:

```
X-Real-IP: 192.168.0.10
```

## Session Persistence

Session persistence is a method to direct all requests originating from a single logical client to a single backend web server. Backend servers that use caching to improve performance, or to enable log-in sessions or shopping carts, can benefit from session persistence.

You enable session persistence when you [create a load balancer](#) or when you [create a backend set](#). You can also [edit an existing backend set](#) to enable, disable, or change the session persistence configuration.

## Sticky Cookies

The Load Balancing service offers two mutually exclusive cookie-based configurations for enabling session persistence:

- [Application cookie stickiness](#)
- [Load balancer cookie stickiness](#)



### Note

#### *IP Address-driven Session Persistence*

Some products offer session persistence support without cookies. These products depend on the IP address of the incoming request. ISP proxies and company exit gateways can issue many requests from a single IP address. In this case, a single backend server can be subject to high traffic volumes. Your backend fleet can become overwhelmed, one server at a time, even though effective load balancing is possible.

Another weakness of IP address-driven session persistence is that the originating IP address can change. In this case, session persistence can be lost or the request redirected to the wrong backend server.

### Application Cookie Stickiness

To configure application cookie session persistence, you specify a cookie name and decide whether to [disable fallback](#) for unavailable servers.

The Load Balancing service activates application cookie session persistence (stickiness) when a backend server sends a `Set-Cookie` response header containing a recognized cookie name. The cookie name must match the name specified in the backend set configuration. If the configuration specifies a match-all pattern, '\*', any cookie set by the server activates session persistence. Unless a backend server activates session persistence, the service follows the [load balancing policy](#) specified when you created the load balancer.

Requirements:

- Your load balancer must operate in HTTP mode to support server side, cookie-driven session persistence.

- The client computer must accept cookies for Load Balancing session persistence feature to work.

### HOW IT WORKS

The Load Balancing service calculates a hash of the configured cookie and other request parameters, and sends that value to the client in a cookie. The value stored in the cookie enables the service to route subsequent client requests to the correct backend server. If your backend servers change any of the defined cookies, the service recomputes the cookie's value and resends it to the client.



#### Warning

Oracle recommends that you treat cookie data as an opaque entity. Do not use it in your applications.

The backend server can stop application cookie persistence by deleting the session persistence cookie. If you used the match-all pattern, it must delete all cookies. You can delete cookies by sending a `Set-Cookie` response header with a past expiration date. The Load Balancing service routes subsequent requests using the configured load balancing policy.

### Load Balancer Cookie Stickiness

When you configure load balancer cookie stickiness, the load balancer inserts a cookie into the response. The parameters configured within the cookie enable session stickiness. This method is useful when you have applications and web backend services that cannot generate their own cookies.

To configure load balancer cookie session persistence, you specify:

- *The cookie name.*

If you do not specify a cookie name, the default name is `X-Oracle-BMC-LBS-Route`.



### Note

Ensure that the cookie name used at the backend application servers is different from the cookie name used at the load balancer. To minimize the chance of name collision, Oracle recommends that you use a prefix such as `X-Oracle-OCI-`.

If a backend server and the load balancer both insert cookies with the same name, the client or browser behavior can vary depending on the domain value associated with the cookie. If the name and domain values of the `Set-cookie` header generated by a backend server and the `Set-cookie` header generated by the load balancer are the same, the client or browser treats them as one cookie. The client returns only one of the cookie values in subsequent requests. If both `Set-cookie` names are the same, but the domain names are different, the client or browser treats them as two different cookies.

- *The domain in which the cookie is valid.* The `Set-cookie` header inserted by the load balancer contains a domain attribute with the specified value. This attribute has no default value. If you do not specify a value, the load balancer does not insert the domain attribute into the `Set-cookie` header.



### Note

- [RFC 6265 - HTTP State Management Mechanism](#) describes client and browser behavior when the domain attribute is present or not present in the `Set-cookie` header. If the value of the `Domain` attribute is `example.com` in the `Set-cookie` header, the client includes the same cookie in the `Cookie` header when making HTTP requests to `example.com`, `www.example.com`, and `www.abc.example.com`. If the `Domain` attribute is not present, the client returns the cookie only for the domain to which the original request was made.
- Ensure that this attribute specifies the correct domain value. If the `Domain` attribute in the `Set-cookie` header does not include the domain to which the original request was made, the client or browser might reject the cookie. As specified in RFC 6265, the client accepts a cookie with the `Domain` attribute value `example.com` or `www.example.com` sent from `www.example.com`. It does not accept a cookie with the `Domain` attribute `abc.example.com` or `www.abc.example.com` sent from `www.example.com`.

- *The URI path in which the cookie is valid.* The `Set-cookie` header inserted by the load balancer contains a `Path` attribute with the specified value.

Clients include the cookie in an HTTP request only if the path portion of the `request-uri` matches, or is a subdirectory of, the cookie's `Path` attribute.

The default value is  `'/'`.

- *The amount of time the cookie remains valid.* The `Set-cookie` header inserted by the load balancer contains a `Max-Age` attribute with the specified value. The specified value must be at least one second. There is no default value for this attribute. If you do not specify a value, the load balancer does not include the `Max-Age` attribute in the `Set-cookie` header. Usually, the client or browser retains the cookie until the current session ends, as defined by the client.
- *Whether the `Set-cookie` header should contain the `Secure` attribute.* The `Secure` attribute directs the client or browser to send the cookie only using a secure protocol.



### Tip

If you set this field to true, you cannot associate the corresponding backend set with an HTTP listener.

- *Whether the `Set-cookie` header should contain the `HttpOnly` attribute.* The `HttpOnly` attribute limits the scope of the cookie to HTTP requests. This attribute directs the client or browser to omit the cookie when providing access to cookies through non-HTTP APIs. For example, it restricts the cookie from JavaScript channels.
- *Whether to [disable fallback](#) for unavailable servers.*



### Note

Path route rules take precedence to determine the target backend server. The load balancer verifies that session stickiness is enabled for the backend server and that the cookie configuration is valid for the target. The system ignores invalid cookies.

### Fallback

By default, the Load Balancing service directs traffic from a persistent session client to a different backend server when the original server is unavailable. You can configure the backend set to disable this fallback behavior. When you disable fallback, the load balancer fails the request and returns an HTTP 502 code. The service continues to return an HTTP 502 until the client no longer presents a persistent session cookie.



#### Warning

If fallback is disabled, cookies with a distant future expiration date can cause a client outage.

The Load Balancing service considers a server marked `drain` available for existing persisted sessions. New requests that are not part of an existing persisted session are not sent to that server.

### Managing a Load Balancer

This topic describes how to create or delete a load balancer on your system.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Prerequisites

To implement a working load balancer, you need:

- For a public load balancer in a region with multiple availability domains, you need a VCN with a public regional subnet or at least two public AD-specific subnets. In the latter case, each AD-specific subnet must reside in a separate availability domain. For more information on subnets, See [VCNs and Subnets](#) and [Public vs. Private Subnets](#).



#### **Warning**

You cannot specify a [private subnet](#) for your public load balancer.

- For a public load balancer in a region with only one availability domain, you need a VCN with at least one public subnet.
- For a private load balancer in any region, you need a VCN with at least one private subnet.
- Two or more backend servers (Compute instances) running your applications. For more information on Compute instances, see [Creating an Instance](#).



### Note

#### *Private IP Address Consumption*

A public load balancer created in one public regional subnet consumes two private IP addresses from the host subnet. The primary and secondary load balancers reside within the same subnet. Each load balancer requires a private IP address from that subnet. The Load Balancing service assigns a floating public IP address, which does not come from the host subnet.

A public load balancer created in two public AD-specific subnets consumes two private IP addresses, one from each host subnet. The primary and secondary load balancers reside within different subnets. Each load balancer requires one private IP address from its host subnet. The Load Balancing service assigns a floating public IP address, which does not come from the host subnets.

A private load balancer created in a single subnet consumes three private IP addresses from the host subnet. The primary and secondary load balancers reside within the same subnet. Each load balancer requires a private IP address from that subnet. The floating private IP address also comes from the host subnet.

## Working with Load Balancers

For background information on Oracle Cloud Infrastructure Load Balancing, see [Overview of Load Balancing](#).

For the purposes of access control, you must specify the compartment where you want the load balancer to reside. Consult an administrator in your organization if you're not sure which compartment to use. For information about compartments and access control, see [Managing Compartments](#).

When you create a load balancer within your VCN, you get a public or private IP address, and provisioned total bandwidth. If you need another IP address, you can create another load balancer.

A public load balancer in a region with multiple availability domains requires one public regional subnet or two public AD-specific subnets to host the primary load balancer and a standby. In the latter case, each AD-specific subnet must reside in a separate availability domain. A public load balancer in a region with only one availability domain requires a single public subnet to host the primary load balancer and a standby. For more information on VCNs and subnets, see [Overview of Networking](#). You can associate the public IPv4 address with a DNS name from any vendor. You can use the public IP address as a front end for incoming traffic. The load balancer can route data traffic to any backend server that is reachable from the VCN.

A private load balancer requires only one subnet to host the primary load balancer and a standby. The private IP address is local to the subnet. The load balancer is accessible only from within the VCN that contains the associated subnet, or as further restricted by your security list rules. The load balancer can route data traffic to any backend server that is reachable from the VCN.

The essential components for load balancing include:

- A load balancer with pre-provisioned bandwidth.
- A backend set with a health check policy. See [Managing Backend Sets](#).
- Backend servers for your backend set. See [Managing Backend Servers](#).
- One or more listeners. See [Managing Load Balancer Listeners](#).
- Load balancer subnet security rules to allow the intended traffic. To learn more about these rules, see [Security Rules](#).



### Tip

To accommodate high-volume traffic, Oracle strongly recommends that you use [stateless security rules](#) for your load balancer subnets.

Optionally, you can associate your listeners with SSL server certificate bundles to manage how your system handles SSL traffic. See [Managing SSL Certificates](#).

For information about the number of load balancers you can have, see [Service Limits](#).

### Configuration Changes and Service Disruption

For a running load balancer, some configuration changes lead to service disruptions. The following guidelines help you understand the effect of changes to your load balancer.

- Operations that add, remove, or modify a backend server create no disruptions to the Load Balancing service.
- Operations that edit an existing health check policy create no disruptions to the Load Balancing service.
- Operations that trigger a load balancer reconfiguration can produce a brief service disruption with the possibility of some terminated connections.

### Health Status

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components. You can see health status indicators on the Console *List* and *Details* pages for load balancers, backend sets, and backend servers. You also can use the Load Balancing API to retrieve this information.

For general information about health status indicators, see [Editing Health Check Policies](#).

### Load Balancer Health Summary

The Console list of load balancers provides health status summaries that indicate the overall health of each load balancer. Health status indicators have four levels. The meaning of each level is:

The Console list of load balancers provides health status summaries that indicate the overall health of each load balancer. Health status indicators come in four levels. The meaning of each level is:

- **OK:** All backend sets associated with the load balancer return a status of OK.
- **WARNING:** All the following conditions are true:
  - At least one backend set associated with the load balancer returns a status of WARNING or UNKNOWN.
  - No backend sets return a status of CRITICAL.
  - The load balancer life-cycle state is ACTIVE.
- **CRITICAL:** At least one backend set associated with the load balancer returns a status of CRITICAL.
- **UNKNOWN:** Any one of the following conditions is true:
  - The load balancer life-cycle state is not ACTIVE.
  - No backend sets are defined for the load balancer.
  - All the following conditions are true:
    - More than half of the backend sets associated with the load balancer return a status of UNKNOWN.
    - None of the backend sets return a status of WARNING or CRITICAL.
    - The load balancer life-cycle state is ACTIVE.
  - The system could not retrieve metrics for any reason.

For guidance on detecting and correcting common issues, see [Using Health Status](#).

### Load Balancer Health Details

The load balancer *Details* page provides the same **Overall Health** status indicator found in the list of load balancers. It also includes counters for the **Backend Set Health** status values reported by the load balancer's child backend sets.

The health status counter badges indicate the following:

- The number of child entities reporting the indicated health status level.
- If a counter corresponds to the overall health, the badge has a fill color.
- If a counter has a zero value, the badge has a light gray outline and no fill color.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Using the Console

#### To create a load balancer

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking**

and click **Load Balancers**.

2. Choose a **Compartment** you have permission to work in, and then click **Create Load Balancer**.
3. The workflow to create a load balancer includes three waypoints to configure your load balancer:

### Step 1 - Add Details

Specify the attributes of the load balancer.

- **Load Balancer Name:** Required. Accept the default name or specify a friendly name for the load balancer. It does not have to be unique, but it cannot be changed in the Console. (You can, however, change it with the API.) Avoid entering confidential information.
- **Choose Visibility Type:** Specify whether your load balancer is public or private.
  - **Public:** Choose this option to create a public load balancer. You can use the assigned public IP address as a front end for incoming traffic and to balance that traffic across all backend servers.
  - **Private:** Choose this option to create a private load balancer. You can use the assigned private IP address as a front end for incoming internal VCN traffic and to balance that traffic across all backend servers.

- **Choose the Maximum Total Bandwidth:** Required. Specify a shape to provision the maximum total bandwidth (ingress plus egress) for your load balancer. Available shapes include:
  - Micro - 10 Mbps  
This load balancer shape is Always Free eligible.  
For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).
  - Small - 100 Mbps
  - Medium - 400 Mbps
  - Large - 8000 Mbps



### Tip

After you create a load balancer, you cannot change the provisioned bandwidth. You can create another load balancer with a different maximum bandwidth.

- **Enable IPv6 Address Assignment:** Available only in the US Government Cloud. Specify whether the load balancer supports IPv6 addresses for incoming requests.



### Note

- When you create a load balancer, you can optionally choose to have an IPv4/IPv6 dual-stack configuration. When you choose the IPv6 option, the Load Balancing service assigns both an IPv4 and an IPv6 address to the load balancer. The load balancer receives client traffic sent to the assigned IPv6 address. The load balancer uses only IPv4 addresses to communicate with backend servers. There is no IPv6 communication between the load balancer and the backend servers.
- IPv6 address assignment occurs only at load balancer creation. You cannot assign an IPv6 address to an existing load balancer.
- **Only VCNs in the US Government Cloud currently support IPv6 addressing.** For more information about Oracle Cloud Infrastructure's IPv6 implementation, see [IPv6 Addresses](#).

- **Choose Networking**

If the current compartment contains at least one VCN, the Console provides a drop-down list of VCNs for you to choose from.

- **Virtual Cloud Network in <compartment>**: Required. Specify a VCN for the load balancer.

By default, the Console shows a list of VCNs in the compartment you're currently working in. Click the **Change Compartment** link to select a VCN from a different compartment.

- **Subnet in <compartment>**: Required. Select an available subnet. For a public load balancer, it must be a public subnet.

By default, the Console shows a list of subnets in the compartment you're currently working in. Click the **Change Compartment** link to select a subnet from a different compartment.



### Tip

In addition to *public* or *private*, subnets can be either *regional* or *AD-specific*. Oracle recommends using regional subnets. For more information, see [About Regional Subnets](#).

- **Subnet (2 of 2) in <compartment>**: Required for a public load balancer when you specify an AD-specific subnet for **Subnet**. Select a second public subnet. The second subnet must reside in a separate availability domain from the first subnet.



### Tip

- If you chose to create a private load balancer under **Visibility Type**, the form prompts you to select only one subnet.
- If you're working in a region that includes only one availability domain, a second subnet is not required. The form prompts you to select only one subnet.

If the current compartment contains no virtual cloud networks, the Load Balancing service offers to create a VCN for you.

- **Virtual Cloud Network in <compartment>**: When the current compartment contains no virtual cloud networks, the drop-down list is disabled. The system offers to create a VCN for you.

If you want to use an existing VCN in another compartment, click the **Change Compartment** link and choose that compartment from the drop-down list.

**Virtual Cloud Network Name:** Optional, when the system creates a VCN for you. Specify a friendly name for the new cloud network. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.

If you do not specify a name for the new VCN, the system generates a name for you.

- **Use Network Security Groups to Control Traffic:** Check this box if you want to add your load balancer to a network security group (NSG). For more information about NSGs, see [Network Security Groups](#).

- **Network Security Groups in <compartment>**: Choose an NSG to add your load balancer to.

By default, the Console shows a list of NSGs in the compartment you're currently working in. Click the **Change Compartment** link to select an NSG from a different compartment.

- (Optional) Click **+ Another Network Security Group** to add your load balancer to another NSG.



### Tip

You can change the NSGs that your load balancer belongs to after you create it. On the **Load Balancer Details** page, click the **Edit** link that appears beside the list of associated network security groups.

- **Show Advanced Options:** Click this link to change the host compartment or to create tags for the load balancer.
  - **Management:**
    - **Create in Compartment:** Optionally, you can select a different compartment to host the load balancer.
  - **Tagging:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

### Step 2 - Choose Backends

A load balancer distributes traffic to backend servers within a backend set. A backend set is a logical entity defined by a load balancing policy, a list of backend servers (Compute instances), and a health check policy.

The load balancer creation workflow creates one backend set for your load balancer. Optionally, you can add backend sets and backend servers after you create the load balancer.

- **Specify a Load Balancing Policy:** Required. Choose the load balancer policy for the backend set. The available options are:
  - **Weighted Round Robin:** This policy distributes incoming traffic sequentially to each server in a backend set list.
  - **IP Hash:** This policy ensures that requests from a particular client are always directed to the same backend server.
  - **Least Connections:** This policy routes incoming request traffic to the backend server with the fewest active connections.

For more information on these policies, see [How Load Balancing Policies Work](#).

- **Select Backend Servers:** Optional. Add backend servers to the backend set. Click **Add Backends** to select resources from a list of available Compute instances.



### Important

When you add backend servers, the Load Balancing service automatically creates security list rules for you. If you prefer to create security list rules manually, click **Show Advanced Options** and choose the option to **Manually configure security list rules after the load balancer is created**.

- **Add Backends:** Select (check) the instances you want to include in the load balancer's backend set.

To select instances from a different compartment, use the **Change Compartment** link and choose a compartment from the drop-down list.

After you select the instances you want to add from the current compartment, click **Add Selected Backends**.



### Tip

- You can choose instances from one compartment at a time. After you add instances from one compartment, you can choose **Add More Backends** to add instances from another compartment.
- You cannot add a backend server marked as **Backup** to a backend set that uses the IP Hash policy.

After you add instances to the backend set, they appear in the **Select Backend Servers** table. You can:

- Specify the server **Port** to which the load balancer must direct traffic. The default is port 80.
- Click the Actions icon (three dots) for a server and choose **Delete** to remove it from the backend set.
- **Specify Health Check Policy:** Required. Specify the test parameters that confirm the health of your backend servers.
  - **Protocol:** Required. Specify the protocol to use for health check queries, either HTTP or TCP.



### Important

Configure your health check protocol to [match your application or service](#).

- **Port:** Optional. Specify the backend server port against which to run the health check.



### Tip

You can enter the value '0' to have the health check use the backend server's traffic port.

- **URL Path (URI):** (HTTP only) Required. Specify a URL endpoint against which to run the health check.
- **Interval in ms:** Optional. Specify how frequently to run the health check, in milliseconds. The default is 10000 (10 seconds).
- **Timeout in ms:** Optional. Specify the maximum time in milliseconds to wait for a reply to a health check. A health check is successful only if a reply returns within this timeout period. The default is 3000 (3 seconds).
- **Number of retries:** Optional. Specify the number of retries to attempt before a backend server is considered "unhealthy". This number also applies when recovering a server to the "healthy" state. The default is 3.
- **Status Code:** (HTTP only) Optional. Specify the status code a healthy backend server must return.
- **Response Body Regex:** (HTTP only) Optional. Provide a regular expression for parsing the response body from the backend server.
- **Show Advanced Options:** Click this link to access security list and session persistence options.
  - **Backend Set Name:** Specify a friendly name for the backend set. It must be unique within the load balancer, and it cannot be changed. If you do not specify a name, the Load Balancing service creates one for you.

Valid backend set names include only alphanumeric characters, dashes, and underscores. Backend set names cannot contain spaces. Avoid entering confidential information.

- **Security List:** Choose to manually configure subnet security list rules to allow the intended traffic or allow the system to create security list rules for you. To learn more about these rules, see [Parts of a Security Rule](#).
  - **Manually configure security list rules after the load balancer is created:** When you choose this option, you must configure security list rules after load balancer creation.
  - **Automatically add security list rules:** Default. When you choose this option, the Load Balancing service creates security list rules for you.

The system displays a table for egress rules and a table for ingress rules. Each table lets you choose the security list that applies to the relevant subnet.

You can choose whether to apply the proposed rules for each affected subnet.
- **Session Persistence:** Optional. Specify how the load balancer manages session persistence.



### Important

See [Session Persistence](#) for important information on configuring these settings.

- **Disable Session Persistence:** Choose this option to disable cookie-based session persistence.
- **Enable Application Cookie Persistence:** Choose this option to

enable persistent sessions from a single logical client when the backend application server response includes a `Set-cookie` header with the cookie name you specify.

- **Cookie Name:** The cookie name used to enable session persistence. Specify `*` to match any cookie name. Avoid entering confidential information.
- **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
- **Enable Load Balancer Cookie Persistence:** Choose this option to enable persistent sessions based on a cookie inserted by the load balancer.
  - **Cookie Name:** Specify the name of the cookie used to enable session persistence. If blank, the default cookie name is `x-Oracle-BMC-LBS-Route`.  
Ensure that any cookie names used at the backend application servers are different from the cookie name used at the load balancer. Avoid entering confidential information.
  - **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
  - **Domain Name:** Optional. Specify the domain in which the cookie is valid.  
This attribute has no default value. If you do not specify a value, the load balancer does not insert the domain attribute into the `Set-cookie` header.
  - **Path:** Optional. Specify the path in which the cookie is valid. The default value is `/`.

- **Expiration Period in Seconds:** Optional. Specify the amount of time the cookie remains valid. If blank, the cookie expires at the end of the client session.
- **Attributes**
  - **Secure:** Specify whether the `Set-cookie` header should contain the `Secure` attribute. If selected, the client sends the cookie only using a secure protocol.  
If you enable this setting, you cannot associate the corresponding backend set with an HTTP listener.
  - **HTTP Only:** Specify whether the `Set-cookie` header should contain the `HttpOnly` attribute. If selected, the cookie is limited to HTTP requests. The client omits the cookie when providing access to cookies through non-HTTP APIs such as JavaScript channels.

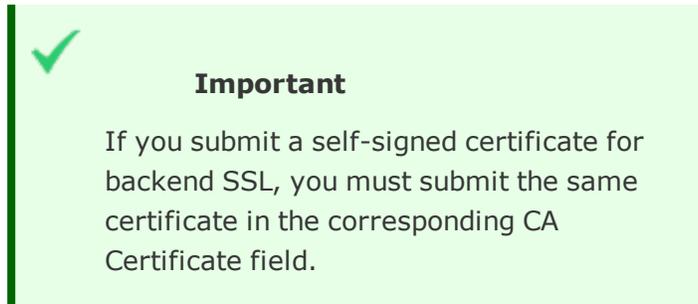
### Step 3 - Configure Listener

- **Listener Name:** Required. Specify a friendly name for the listener. The name must be unique, and cannot be changed. Avoid entering confidential information. If you do not specify a name, the Load Balancing service creates one for you.
- **Specify the type of traffic your listener handles:** Required. Specify the protocol to use.
- **Specify the port your listener monitors for ingress traffic:** Required. Specify the port. Defaults are:
  - 443 for HTTPS
  - 80 for HTTP
  - 22 for TCP
- If you chose the HTTPS protocol, or if you chose the TCP protocol

and selected the Use SSL check box

- **Choose SSL Certificate File:** Required. Drag and drop the certificate file, in PEM format, into the **SSL Certificate** field.

Alternatively, you can choose the **Paste SSL Certificate** option to paste a certificate directly into this field.



- **Specify CA Certificate:** Optional. (Recommended for backend SSL termination configurations.) Select (check) this box if you want to provide a CA certificate. See [Working with SSL Certificates](#) for more information.
  - **Choose CA Certificate File:** Drag and drop the CA certificate file, in PEM format, into the **CA Certificate** field.  
Alternatively, you can choose the **Paste CA Certificate** option to paste a certificate directly into this field.
- **Specify Private Key:** Optional. (Required for SSL termination.) Select (check) this box if you want to provide a private key for the certificate.
  - **Choose Private Key File:** Drag and drop the private key, in PEM format, into the **Private Key** field.  
Alternatively, you can choose the **Paste Private Key** option to paste a private key directly into this field.

- **Enter Private Key Passphrase:** Optional. Specify the private key passphrase.
- **Show Advanced Options:** Click this link to set the idle timeout.
  - **Specify the maximum timeout in seconds:** Optional. Specify the maximum idle time in seconds. This setting applies to the time allowed between two successive receive or two successive send network input/output operations during the HTTP request-response phase.



### Tip

The maximum value is 7200 seconds. For more information, see [Connection Management](#).

4. Click **Create Load Balancer**.

After the system provisions the load balancer, details appear in the load balancer list. To view more details, click the load balancer name.

### To delete a load balancer

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to delete.
3. For the load balancer you want to delete, click the Actions icon (three dots), and then click **Delete**.
4. Confirm when prompted.

### Moving a Load Balancer to a Different Compartment

You can move your load balancer from its current compartment into a different compartment. For information about compartments and access control, see [Managing Compartments](#).

#### To move a load balancer to a different compartment

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. In the **List Scope** section, select a compartment.
3. Find the load balancer in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. In the **Move Resource** dialog box, choose the destination compartment from the list.
5. Click **Move Resource**.

### Managing Tags for a Load Balancer

You can apply tags to your resources, such as load balancers, to help you organize them according to your business needs. You can apply tags at the time you create a load balancer, or you can update the load balancer later with the wanted tags. For general information about applying tags, see [Resource Tags](#).

#### To manage tags for a load balancer

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to tag, and then click the load balancer's name.
3. Click the **Tags** tab to view or edit existing tags, or click **Apply Tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring the traffic passing through your load balancer, see [Load Balancing Metrics](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage load balancers:

- [ChangeLoadBalancerCompartment](#)
- [CreateLoadBalancer](#)
- [DeleteLoadBalancer](#)
- [GetLoadBalancer](#)
- [GetLoadBalancerHealth](#)
- [ListLoadBalancers](#)
- [ListLoadBalancerHealths](#)
- [UpdateLoadBalancer](#): You can update the load balancer's display name.
- [UpdateNetworkSecurityGroups](#)

### Managing Backend Sets

This topic describes how to create and delete backend sets for use with a load balancer. For information about managing load balancers, see [Managing a Load Balancer](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

## Working with Backend Sets

A backend set is a logical entity defined by a load balancing policy, a health check policy, and a list of backend servers. To create a backend set, you must specify a load balancing policy and health check script, and then add a list of backend servers (Compute instances). SSL and session persistence configuration is optional. A backend set must be associated with one or more listeners for the load balancer to work.

You cannot delete a backend set used by an active listener.

Changing the load balancing policy of a backend set temporarily interrupts traffic and can drop active connections.

For background information on the Oracle Cloud Infrastructure Load Balancing, see [Overview of Load Balancing](#).

### Health Status

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components. You can see health status indicators on the Console *List* and *Details* pages for load balancers, backend sets, and backend servers. You also can use the Load Balancing API to retrieve this information.

For general information about health status indicators, see [Editing Health Check Policies](#).

### Backend Set Health Summary

The Console list of a load balancer's backend sets provides health status summaries that indicate the overall health of each backend set. Health status indicators have four levels. The meaning of each level is:

- **OK:** All backend servers in the backend set return a status of OK.
- **WARNING:** Both of the following conditions are true:
  - Half or more of the backend set's backend servers return a status of OK.
  - At least one backend server returns a status of WARNING, CRITICAL, or UNKNOWN.
- **CRITICAL:** Fewer than half of the backend set's backend servers return a status of OK.
- **UNKNOWN:** At least one of the following conditions is true:
  - More than half of the backend set's backend servers return a status of UNKNOWN.
  - The system could not retrieve metrics for any reason.
  - The backend set does not have a listener attached.

For guidance on detecting and correcting common issues, see [Using Health Status](#).

### Backend Set Health Details

The backend set *Details* page provides the same **Overall Health** status indicator found in the load balancer's list of backend sets. It also includes counters for the **Backend Health** status values reported by the backend set's child backend servers.

The health status counter badges indicate the following:

- The number of child entities reporting the indicated health status level.
- If a counter corresponds to the overall health, the badge has a fill color.
- If a counter has a zero value, the badge has a light gray outline and no fill color.

### Using the Console

#### To create a backend set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets** (if necessary), and then click **Create Backend Set**.
4. In the **Create Backend Set** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the backend set. It must be unique within the load balancer, and it cannot be changed.  
Valid backend set names include only alphanumeric characters, dashes, and underscores. Backend set names cannot contain spaces. Avoid entering confidential information.

- **Traffic Distribution Policy:** Required. Choose the load balancer policy for the backend set. The available options are:
  - **IP Hash**
  - **Least Connections**
  - **Weighted Round Robin**

For more information on these policies, see [How Load Balancing Policies Work](#).



### Tip

You cannot add a backend server marked as **Backup** to a backend set that uses the IP Hash policy.

- **Use SSL:** Optional. Check this box to associate an SSL certificate bundle with the backend set.

If there are no certificate bundles attached to the load balancer, this option is disabled.

  - **Certificate Name:** Required. Select the certificate bundle to use. You can choose any certificate bundle that is attached to the current load balancer. See [Managing SSL Certificates](#) for more information.
  - **Verify Peer Certificate:** Optional. Select this option to enable peer certificate verification.
  - **Verify Depth:** Optional. Specify the maximum depth for certificate chain verification.
- **Session Persistence:** Optional. Specify how the load balancer manages session persistence.



### Important

See [Session Persistence](#) for important information on configuring these settings.

- **Disable Session Persistence:** Choose this option to disable cookie-based session persistence.
- **Enable Application Cookie Persistence:** Choose this option to enable persistent sessions from a single logical client when the response from a backend application server includes a `Set-cookie` header with the cookie name you specify.
  - **Cookie Name:** The cookie name used to enable session persistence. Specify `*` to match any cookie name. Avoid entering confidential information.
  - **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
- **Enable Load Balancer Cookie Persistence:** Choose this option to enable persistent sessions based on a cookie inserted by the load balancer.
  - **Cookie Name:** Specify the name of the cookie used to enable session persistence. If blank, the default cookie name is `X-Oracle-BMC-LBS-Route`.  
Ensure that any cookie names used at the backend application servers are different from the cookie name used at the load balancer. Avoid entering confidential information.
  - **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
  - **Domain Name:** Optional. Specify the domain in which the cookie is valid.

This attribute has no default value. If you do not specify a value, the load balancer does not insert the domain attribute into the `Set-cookie` header.

- **Path:** Optional. Specify the path in which the cookie is valid. The default value is `/`.
- **Expiration Period in Seconds:** Optional. Specify the amount of time the cookie remains valid. If blank, the cookie expires at the end of the client session.
- **Attributes**
  - **Secure:** Specify whether the `Set-cookie` header should contain the `Secure` attribute. If selected, the client sends the cookie only using a secure protocol.  
If you enable this setting, you cannot associate the corresponding backend set with an HTTP listener.
  - **HTTP Only:** Specify whether the `Set-cookie` header should contain the `HttpOnly` attribute. If selected, the cookie is limited to HTTP requests. The client omits the cookie when providing access to cookies through non-HTTP APIs such as JavaScript channels.
- **Health Check:** Required. Specify the test parameters to confirm the health of backend servers.
  - **Protocol:** Required. Specify the protocol to use, either HTTP or TCP.



### Important

Configure your health check protocol to [match your application or service](#).

- **Port:** Optional. Specify the backend server port against which to run the health check.



### Tip

You can enter the value '0' to have the health check use the backend server's traffic port.

- **URL Path (URI):** (HTTP only) Required. Specify a URL endpoint against which to run the health check.
- **Interval in ms:** Optional. Specify how frequently to run the health check, in milliseconds. The default is 10000 (10 seconds).
- **Timeout in ms:** Optional. Specify the maximum time in milliseconds to wait for a reply to a health check. A health check is successful only if a reply returns within this timeout period. The default is 3000 (3 seconds).
- **Number of retries:** Optional. Specify the number of retries to attempt before a backend server is considered "unhealthy". This number also applies when recovering a server to the "healthy" state. The default is '3'.
- **Status Code:** (HTTP only) Optional. Specify the status code a healthy backend server must return.
- **Response Body Regex:** (HTTP only) Optional. Provide a regular expression for parsing the response body from the backend server.

5. Click **Create**.

After your backend set is provisioned, you must specify backend servers for the set. See [Managing Backend Servers](#) for more information.

### To edit a backend set



#### Warning

Updating the backend set temporarily interrupts traffic and can drop active connections.

When you edit a backed set, you can choose a new load balancing policy and modify the SSL configuration.

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set you want to edit.
4. Click **Edit Backend Set**.
5. Make the configuration changes you need, and then click **Submit**.

If you want to modify the backend set's health check policy, see [Editing Health Check Policies](#).

If you want to add or remove backend servers from the backend set, see [Managing Backend Servers](#).

### To delete a backend set



#### Tip

You cannot delete a backend set used by an active listener. First, remove any backend sets you want to



delete from the associated listeners.

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets**.
4. For the backend set you want to delete, click the Actions icon (three dots), and then click **Delete**.
5. Confirm when prompted.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage load balancer backend sets:

- [CreateBackendSet](#)
- [DeleteBackendSet](#)
- [GetBackendSet](#)
- [GetBackendSetHealth](#)
- [ListBackendSets](#)
- [UpdateBackendSet](#)

### Managing Backend Servers

This topic describes how to manage backend servers for use with a load balancer. For information about managing load balancers, see [Managing a Load Balancer](#).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Working with Backend Servers

When you implement a load balancer, you must specify the backend servers (Compute instances) to include in each backend set. The load balancer routes incoming traffic to these backend servers based on the policies you specified for the backend set. You can use the Console to add and remove backend servers in a backend set.

To route traffic to a backend server, the Load Balancing service requires the IP address of the compute instance and the relevant application port. If the backend server resides within the same VCN as the load balancer, Oracle recommends that you specify the compute instance's private IP address. If the backend server resides within a different VCN, you must specify the public IP address of the compute instance. You also must ensure that the VCN's security rules allow Internet traffic.



### Warning

When you add backend servers to a backend set, you specify either the instance OCID or an IP address for the server to add. An instance with multiple VNICs attached can have multiple IP addresses pointing to it.

- If you identify a backend server by OCID, Load Balancing uses the primary VNIC's primary private IP address.
- If you identify the backend servers to add to a backend set by their IP addresses, it is possible to point to the same instance more than once.

To enable backend traffic, your backend server subnets must have appropriate ingress and egress security rules. When you add backend servers to a backend set, you can specify the applicable network security groups (NSGs). If you prefer to use security lists for your VCN, the Load Balancing service Console can suggest security list rules for you. You also can configure them yourself through the Networking service. See [Security Lists](#) for more information.



### Tip

To accommodate high-volume traffic, Oracle strongly recommends that you use [stateless security rules](#) for your load balancer subnets.

You can add and remove backend servers without disrupting traffic.

### Health Status

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components. You can see health status indicators on the Console *List* and *Details* pages for load balancers, backend sets, and backend servers. You also can use the Load Balancing API to retrieve this information.

For general information about health status indicators, see [Editing Health Check Policies](#).

#### Backend Server Health Summary

The Console list of a backend set's backend servers provides health status summaries that indicate the overall health of each backend server. The primary and standby load balancers both provide health check results that contribute to the health status. Health status indicators have four levels. The meaning of each level is:

- **OK:** The primary and standby load balancer health checks both return a status of OK.
- **WARNING:** One health check returned a status of OK and one did not.
- **CRITICAL:** Neither health check returned a status of OK.
- **UNKNOWN:** One or both health checks returned a status of UNKNOWN or the system was unable to retrieve metrics.

To view the health status details for a specific backend server, click its **IP Address**.

For guidance on detecting and correcting common issues, see [Using Health Status](#).

#### Backend Server Health Details

The *Details* page for a backend set provides the same **Overall Health** status indicator found in the backend set's list of backend servers. It also reports the following data for the two health checks performed against each backend server:

### **IP ADDRESS**

The IP address of the health check status report provider, which is a Compute instance managed by the Load Balancing service. This identifier helps you differentiate same-subnet load balancers that report health check status.

The Load Balancing service ensures high availability by providing one primary and one standby load balancer. To diagnose a backend server issue, you must know the source of the health check report. For example, a misconfigured security rule might cause one load balancer instance to report that a backend server is healthy. The other load balancer instance might return an unhealthy status. In this case, one of the two load balancer instances cannot communicate with the backend server. Reconfigure the security rules to restore the backend server's health status.

### **STATUS**

The status returned by the health check. Possible values include:

- **OK**  
The backend server's response satisfied the health check policy requirements.
- **INVALID\_STATUS\_CODE**  
The HTTP response status code did not match the expected status code specified by the health policy.
- **TIMED\_OUT**  
The backend server did not respond within the timeout interval specified by the health policy.
- **REGEX\_MISMATCH**  
The backend server response did not satisfy the regular expression specified by the health policy.
- **CONNECT\_FAILED**  
The health check server could not connect to the backend server.

- **IO\_ERROR**  
An input or output communication error occurred while reading or writing a response or request to the backend server.
- **OFFLINE**  
The backend server is set to **offline**, so health checks are not run.
- **UNKNOWN**  
Health check status is not available.

### LAST CHECKED

The date and time of the most recent health check.

Health status is updated every three minutes. No finer granularity is available.

## Using the Console

### To add one or more servers to a backend set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set to which you want to add one or more backend servers.



#### Tip

If the load balancer has no backend sets, you must [create one](#) before you can specify a backend server.

4. In the **Resources** menu, click **Backends**, and then click **Add Backends**.



### Tip

You cannot add a backend server marked as **Backup** to a backend set that uses the IP Hash policy.

5. **Choose how to add backend servers:** Specify how you want to add backend servers to the backend set:
  - **Compute Instances:** Choose this option to select from a list of available Compute instances.
    - **Instances in <compartment>:** Select (check) the instances you want to include in the backend set.

To select instances from a different compartment, use the **Change Compartment** link and choose a compartment from the drop-down list.



### Tip

You can choose instances from one compartment at a time. After you add instances from one compartment, you must repeat the **Add Backends** process to add instances from another compartment.

Once you select an instance to add to the backend set, you can specify:

- **Port:** Required. The backend server port to which the load balancer must direct traffic.
- **Weight:** The load balancing weight assigned to the server. For more information, see [How Load Balancing Policies Work](#).
- Choose to manually configure subnet security list rules that allow the

intended traffic or let the Load Balancing service create security list rules for you. To learn more about these rules, see [Parts of a Security Rule](#).

- **Manually configure security list rules after the load balancer is created:** When you choose this option, you must [create your own rules](#) after adding the backend servers.
- **Automatically add security list rules:** When you choose this option, the Load Balancing service creates security list rules for you. The system displays a table for egress rules and a table for ingress rules. Each table lets you choose the security list that applies to the relevant subnet. You can then choose whether to apply the proposed rules for each affected subnet.
- **IP Addresses:** Choose this option to enter the IP addresses of the backend servers (Compute instances) to add.
  - **IP Address:** Required. Specify the IP address of a backend server you want to add to the backend set.
  - **Port:** Required. Specify the server port to which the load balancer must direct traffic.
  - **Weight:** Required. Specify the load balancing weight to apply to this server. For more information, see [How Load Balancing Policies Work](#).

You can click the plus + icon to add another server to the list or click the x icon to remove a list item.

6. Click **Add**.

### To edit backend server settings

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.

3. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set that includes the backend servers you want to edit.
4. In the **Resources** menu, click **Backends**. A list of servers in the backend set appears.
5. Select (check) the row corresponding to the backend server you want to edit.
6. Choose an action from the **Actions** button drop-down list. The available actions include:
  - a. **Edit:** Opens a single dialog box in which you can edit the port, weight, drain, offline, and backup settings.
  - b. **Edit Port:** Opens a dialog box in which you can change the application port setting.
  - c. **Edit Weight:** Opens a dialog box in which you can change the load balancing weight.
  - d. **Edit Drain State:** Opens a dialog box in which you can change the drain state. If you set the server's drain status to **true**, the load balancer stops forwarding new TCP connections and new non-sticky HTTP requests to this backend server. This setting allows an administrator to take the server out of rotation for maintenance purposes.
  - e. **Edit Offline State:** Opens a dialog box in which you can change the offline status. If you set the server's offline status to **true**, the load balance forwards no ingress traffic to this backend server.
  - f. **Edit Backup State:** Opens a dialog box in which you can change the backup status. If you set the server's backup status to **true**, the load balancer forwards ingress traffic to this backend server only when all other backend servers not marked as backup fail the [health check policy](#). This configuration is useful for handling disaster recovery scenarios.



### Warning

Backend servers marked as **Backup** are not compatible with a load balancer that uses the IP Hash policy.

- g. **Delete:** Removes the server from the backend set.



### Tip

You can select multiple servers to apply the same action to each one.

- 7. Click **Save Changes**.

### To remove a server from a backend set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set from which you want to remove a server.
4. In the **Resources** menu, click **Backends**. A list of servers in the backend set appears.
5. Select (check) the row corresponding to the backend server you want to edit.
6. Choose the **Delete** action from the **Actions** button drop-down list.
7. Confirm when prompted.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage the backend servers in a backend set:

- [CreateBackend](#)
- [DeleteBackend](#)
- [GetBackend](#)
- [GetBackendHealth](#)
- [ListBackends](#)
- [UpdateBackend](#)

### Managing Load Balancer Listeners

This topic is part of the setup and maintenance of a load balancer. For more Load Balancing information about managing load balancers, see [Managing a Load Balancer](#).



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have

permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Working with Listeners

A listener is a logical entity that checks for incoming traffic on the load balancer's IP address.

To handle TCP, HTTP, and HTTPS traffic, you must configure at least one listener per traffic type.

When you create a listener, you must ensure that your VCN's [security rules](#) allow the listener to accept traffic.



#### Tip

To accommodate high-volume traffic, Oracle strongly recommends that you use [stateless security rules](#) for your load balancer subnets.

You can have one SSL certificate bundle per listener. You can configure two listeners, one each for ports 443 and 8443, and associate SSL certificate bundles with each listener. For more information about SSL certificates for load balancers, see [Managing SSL Certificates](#).

### Using the Console

#### To create a listener

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Listeners**, and then click **Create Listener**.
4. In the **Create Listener** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the listener. The name must be unique, and cannot be changed. Avoid entering confidential information.
  - **Hostname:** Optional. Select up to 16 [virtual hostnames](#) for this listener.



#### Important

To apply a virtual hostname to a listener, the name must be part of the load balancer's configuration. If the load balancer has no associated hostnames, you can [create one](#) on the **Hostnames** page.

- **Protocol:** Required. Specify the protocol to use, either HTTP or TCP.
- **Port:** Required. Specify the port on which to listen for incoming traffic.
- **Use SSL:** Optional. Check this box to associate an SSL certificate bundle with the listener. The following settings are required to enable SSL handling. See [Managing SSL Certificates](#) for more information.
  - **Certificate Name:** Required. The friendly name of the SSL certificate bundle to use.

- **Verify Peer Certificate:** Optional. Select this option to enable peer certificate verification.
- **Verify Depth:** Optional. Specify the maximum depth for certificate chain verification.
- **Backend Set:** Required. Specify the default backend set to which the listener routes traffic.
- **Idle Timeout in Seconds:** Optional. Specify the maximum idle time in seconds. This setting applies to the time allowed between two successive receive or two successive send network input/output operations during the HTTP request-response phase.



### Tip

The maximum value is 7200 seconds. For more information, see [Connection Management](#).

- **Path Route Set:** Optional. Specify the name of the set of path-based routing rules that applies to this listener's traffic.



### Important

- To apply a [path route set](#) to a listener, the set must be part of the load balancer's configuration.
- To remove a path route set from an existing listener, choose **None** as the **Path Route Set** option. The path route set remains available for use by other listeners on this load balancer.

5. Click **Create**.

When you create a listener, you must also [update your VCN's security rules](#) to allow traffic to that listener.

### To edit a listener

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Listeners**.
4. For the listener you want to edit, click the Actions icon (three dots), and then click **Edit Listener**.
5. Make the configuration changes you need, and then click **Save Changes**.

### To delete a listener

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Listeners**.
4. For the listener you want to delete, click **Delete**.
5. Confirm when prompted.

### To enable a listener to accept traffic

To enable a listener to accept traffic, you must update your VCN's security rules:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.

The list of VCNs in the current compartment appears.

2. Click the name of the VCN containing your load balancer, and then click **Security Groups** or **Security Lists**.

A list of the security groups or lists in the cloud network appears.

3. Click the name of the NSG or security list that applies to your load balancer.

4. Add or edit the existing rules to allow access from the appropriate resources.

An NSG's security rules appear on the **Network Security Group Details** page. From there you can add, edit, or remove rules.

The **Security List Details** page provides access to separate tables in which you can add or edit **Ingress Rules** or **Egress Rules**.

For details on rule configuration, see [Security Rules](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage listeners:

- [CreateListener](#)
- [DeleteListener](#)
- [UpdateListener](#)

### Managing Request Routing

This topic describes how to manage your load balancer's request routing. For information about managing load balancers, see [Managing a Load Balancer](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

## Routing Incoming Requests

The Load Balancing service enables you to route incoming requests to various backend sets. You can:

- Assign [virtual hostnames](#) to a listener.
- Create [path route rules](#).
- [Combine](#) these techniques.

### Virtual Hostnames

You can assign virtual hostnames to any [listener you create](#) for your load balancer. Each hostname can correspond to an application served from your backend. Some advantages of virtual hostnames include:

- A single associated IP address. Multiple hostnames, backed by DNS entries, can point to the same load balancer IP address.
- A single load balancer. You do not need a separate load balancer for each application.
- A single load balancer shape. Running multiple applications behind a single load balancer helps you manage aggregate bandwidth demands and optimize utilization.
- Simpler backend set management. Managing a set of backend servers under a single resource simplifies network configuration and administration.

You can define exact virtual hostnames, such as "app.example.com", or you can use wildcard names. Wildcard names include an asterisk (\*) in place of the first or last part of the name. When searching for a virtual hostname, the service chooses the first matching variant in the following priority order:

1. Exact name match (no asterisk), such as `app.example.com`.
2. Longest wildcard name that begins with an asterisk, such as `*.example.com`.



#### Tip

Prefix wildcard names might require a wildcard certificate for HTTPS sites.

3. Longest wildcard name that ends with an asterisk, such as `app.example.*`.



### Tip

Suffix wildcard names might require a multi-domain Subject Alternative Name (SAN) certificate for HTTPS sites.

You do not need to specify the matching pattern to apply. The pattern is inherent in the asterisk position, that is, starting, ending, or none.

The following considerations apply to virtual hostnames:

- You cannot use regular expressions.
- To apply virtual hostnames to a listener, you first [create one or more virtual hostnames](#) associated with a load balancer.
- Virtual hostname selection priority is not related to the listener's configuration order.
- You can apply a maximum of 16 virtual hostnames to a listener.
- You can associate a maximum of 16 virtual hostnames with a load balancer.



### Tip

The virtual hostnames feature supports HTTP and HTTPS listeners only, but does not support TCP listeners.



### Note

#### *Default Listener*

If a listener has no virtual hostname specified, that listener is the default for the assigned port.

If all listeners on a port have virtual hostnames, the first virtual hostname configured for that port serves as the default listener.

### Path Route Rules

Some applications have multiple endpoints or content types, each distinguished by a unique URI path. For example, `/admin/`, `/data/`, `/video/`, or `/cgi/`. You can use path route rules to route traffic to the correct backend set without using multiple listeners or load balancers.

A *path route* is a string that the Load Balancing service matches against an incoming URI to determine the appropriate destination backend set.

- You cannot use asterisks in path route strings.
- You cannot use regular expressions.
- Path route string matching is case-insensitive.



### Important

Browsers often add an ending slash to the path in a request. If you specify a path such as `/admin`, you might want to configure the path both with and without the trailing slash. For example, `/admin` and `/admin/`.

A *path route rule* consists of a path route string and a pattern match type.

- Pattern match types include:

- **EXACT\_MATCH**

Looks for a path string that exactly matches the incoming URI path.

Applies case-insensitive regex:

```
^<path_string>$
```

- **FORCE\_LONGEST\_PREFIX\_MATCH**

Looks for the path string with the best, longest match of the beginning portion of the incoming URI path.

Applies case-insensitive regex:

```
<path_string>.*
```

- **PREFIX\_MATCH**

Looks for a path string that matches the beginning portion of the incoming URI path.

Applies case-insensitive regex:

```
^<path_string>.*
```

- **SUFFIX\_MATCH**

Looks for a path string that matches the ending portion of the incoming URI path.

Applies case-insensitive regex:

```
.*<path_string>$
```

- Path route rules apply only to HTTP and HTTPS requests and have no effect on TCP requests.

A *path route set* includes all path route rules that define the data routing for a particular listener.

- You can specify up to 20 path route rules per path route set.
- You can have one path route set per listener. The [maximum number of listeners](#) limits the number of path route sets you can specify for a load balancer.

### RULE PRIORITY

The system applies the following priorities, based on match type, to the path route rules within a set:

- For one path route rule that specifies the EXACT\_MATCH type, there is no cascade of priorities. The listener looks for an exact match only.
- For two path route rules, one that specifies the EXACT\_MATCH type and one that specifies any other match type, the exact match rule is evaluated first. If no match is found, then the system looks for the second match type.
- For multiple path route rules specifying various match types, the system applies the following priority cascade:
  1. EXACT\_MATCH
  2. FORCE\_LONGEST\_PREFIX\_MATCH
  3. PREFIX\_MATCH or SUFFIX\_MATCH
- The order of the rules within the path route set does not matter for EXACT\_MATCH and FORCE\_LONGEST\_PREFIX\_MATCH. The system applies the priority cascade no matter where these match types appear in the path route set.
- If matching cascades down to prefix or suffix matching, the order of the rules within the path route set DOES matter. The system chooses the first prefix or suffix rule that matches the incoming URI path.

### Virtual Hostname and Path Route Rules Combinations

Virtual hostnames and path route rules route requests to backend sets. Listeners with a virtual hostname receive priority over the default (no hostname) listener. The following example shows the results of a simple routing interaction.

The example system includes three listeners and one path route set:

#### Listener 1

- Virtual hostname: *none*
- Default backend set: *A*

- Path route set: `PathRouteSet1`

### Listener 2

- Virtual hostname: `captive.com`
- Default backend set: `B`
- Path route set: `PathRouteSet1`

### Listener 3

- Virtual hostname: `wild.com`
- Default backend set: `C`
- Path route set: `PathRouteSet1`

### Path Route Set

- Path route set name: `PathRouteSet1`
  - Exact match on path string `/tame/` routes to backend set `B`.
  - Exact match on path string `/feral/` routes to backend set `C`.

The example configuration routes incoming URLs as follows:

`http://animals.com/` is routed to backend set `A`

- Virtual hostname `animals.com` matches **Listener 1**.
- Path `/` is not an `EXACT_MATCH` for any path route string in `PathRouteSet1`.

`http://animals.com/tame/` is routed to backend set `B`

- Virtual hostname `animals.com` matches **Listener 1**.
- Path `/tame/` is an `EXACT_MATCH` for path route string `/tame/` in `PathRouteSet1`.

`http://animals.com/feral/` is routed to backend set C

- Virtual hostname `animals.com` matches **Listener 1**.
- Path `/feral/` is an EXACT\_MATCH for path route string `/feral/` in `PathRouteSet1`.

`http://captive.com/` is routed to backend set B

- Virtual hostname `captive.com` matches **Listener 2**.
- Path `/` is not an EXACT\_MATCH for any path route string in `PathRouteSet1`.

`http://captive.com/tame/` is routed to backend set B

- Virtual hostname `captive.com` matches **Listener 2**.
- Path `/tame/` is an EXACT\_MATCH for path route string `/tame/` in `PathRouteSet1`.

`http://captive.com/feral/` is routed to backend set C

- Virtual hostname `captive.com` matches **Listener 2**.
- Path `/feral/` is an EXACT\_MATCH for path route string `/feral/` in `PathRouteSet1`.

`http://wild.com/` is routed to backend set C

- Virtual hostname `wild.com` matches **Listener 3**.
- Path `/` is not an EXACT\_MATCH for any path route string in `PathRouteSet1`.

`http://wild.com/tame/` is routed to backend set B

- Virtual hostname `wild.com` matches **Listener 3**.
- Path `/tame/` is an EXACT\_MATCH for path route string `/tame/` in `PathRouteSet1`.

`http://wild.com/feral/` is routed to backend set C

- Virtual hostname `wild.com` matches **Listener 3**.
- Path `/feral/` is an EXACT\_MATCH for path route string `/feral/` in `PathRouteSet1`.

### Using the Console

You can specify virtual hostnames and path route sets when you [create or update a listener](#).

#### Creating Virtual Hostnames

To apply virtual hostnames to a listener, you first create one or more virtual hostnames. The virtual hostnames become a part of the load balancer's configuration. You then specify one or more virtual hostnames to use when you create or update a listener for the load balancer.

#### To create a virtual hostname

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Hostnames**, and then click **Create Hostname**.
4. In the **Create Hostname** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the hostname. The name must be unique, and cannot be changed. Avoid entering confidential information.
  - **Hostname:** Required. Specify the virtual hostname. See [Virtual Hostnames](#) for a description of valid hostname construction and behavior.
5. Click **Create**. The **Work Request Submitted** dialog box opens.
6. To close the dialog box, click **Close**. To open the **Work Requests** page and view the status of the work request, click **View All Work Requests**.

After you create a virtual hostname, the name becomes available for use with the associated load balance. To apply the hostname, [create or update a listener](#).

### To update a virtual hostname

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Hostnames**.
4. For the hostname you want to edit, click the Actions icon (three dots), and then click **Edit**.
5. In the **Edit Hostname** dialog box, enter your updates to the **Hostname** field. You cannot edit the **Name** field of an existing virtual hostname.
6. Click **Update**. The **Work Request Submitted** dialog box opens.
7. To close the dialog box, click **Close**. To open the **Work Requests** page and view the status of the work request, click **View All Work Requests**.

### To delete a virtual hostname

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Hostnames**.
4. For the hostname you want to edit, click the Actions icon (three dots), and then click **Delete**. The **Work Request Submitted** dialog box opens.
5. To close the dialog box, click **Close**. To open the **Work Requests** page and view the status of the work request, click **View All Work Requests**.

### Creating Path Route Sets

To apply path route rules to a listener, you first create a path route set that contains the rules. The path route set becomes a part of the load balancer's configuration. You then specify the path route set to use when you create or update a listener for the load balancer. To remove a path route set from a listener, [edit the listener](#) and choose **None** as the **Path Route Set** option.

#### To create a path route set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Path Route Sets**, and then click **Create Path Route Set**.
4. In the **Create Path Route Set** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the path route set. The name must be unique, and cannot be changed.  
The path route set name cannot begin with a period and cannot contain the characters `;`, `?`, `#`, `%`, `/`, `\`, `[`, or `]`.  
Avoid entering confidential information.
  - **Path Route Rules**
    - **Order:** Optional. If you have multiple path route rules, you can click the up or down arrows to move the corresponding rule.



### Tip

The [order of the rules within the path route set](#) does not matter in most cases. However, if matching cascades down to prefix or suffix matching, the system chooses the first prefix or suffix rule that matches the incoming URI path.

- **Match Style:** Required. The [type of matching](#) to apply to incoming URIs.
  - **URL String:** Required. The path string to match against the incoming URI path, for example `/admin/`.
  - **Backend Set Name:** Required. The name of the target backend set for requests where the incoming URI matches the specified path.
5. (Optional) Click **+ Additional Rule** to create another path route rule or click the red box to delete an existing rule. You can have up to 20 path route rules in a set.
  6. Click **Create**.

After you create a path route set, the set becomes available for use with the associated load balance. [Create or update a listener](#) to apply the path route set.

### To update a path route set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Path Route Sets**.
4. Click the name of the path route set you want to update, and then click **Edit Path Route Rules**.

5. In the **Edit Path Route Rules** dialog box, edit the following as needed for each rule you want to change:
  - **Order:** Optional. If you have multiple path route rules, you can click the up or down arrows to move the corresponding rule.



### Tip

The [order of the rules within the path route set](#) does not matter in most cases. However, if matching cascades down to prefix or suffix matching, the system chooses the first prefix or suffix rule that matches the incoming URI path.

- **Match Style:** Required. The [type of matching](#) to apply to incoming URIs.
  - **URL String:** Required. The path string to match against the incoming URI path, for example `/admin/`.
  - **Backend Set Name:** Required. The name of the target backend set for requests where the incoming URI matches the specified path.
6. (Optional) Click **+ Additional Rule** to create another path route rule or click the red box to delete an existing rule. You can have up to 20 path route rules in a set.
  7. Click **Save Changes**.

### To update a single path route rule

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Path Route Sets**, and then click the name of the path route set you want to update.

4. For the path route rule you want to edit, click the Actions icon (three dots), and then click **Edit Path Route**.
5. In the **Edit Path Route Rule** dialog box, edit the following as needed for each rule you want to change:
  - **Order:** Optional. If you have multiple path route rules, you can click the up or down arrows to move the corresponding rule.



### Tip

The [order of the rules within the path route set](#) does not matter in most cases. However, if matching cascades down to prefix or suffix matching, the system chooses the first prefix or suffix rule that matches the incoming URI path.

- **Match Style:** Required. The [type of matching](#) to apply to incoming URIs.
  - **URL String:** Required. The path string to match against the incoming URI path, for example `/admin/`.
  - **Backend Set Name:** Required. The name of the target backend set for requests where the incoming URI matches the specified path.
6. Click **Save Changes**.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage request routing:

- [CreateListener](#)
- [CreatePathRouteSet](#)

- [DeleteListener](#)
- [DeletePathRouteSet](#)
- [GetPathRouteSet](#)
- [ListPathRouteSets](#)
- [UpdateListener](#)
- [UpdatePathRouteSet](#)

## Managing Rule Sets

This topic describes how you can create rule sets composed of actions to apply to traffic at an HTTP listener.

For more information about managing load balancer listeners, see [Managing Load Balancer Listeners](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Working with Rule Sets

A rule set is a named set of rules associated with a load balancer and applied to one or more listeners on that load balancer. Rules are objects that represent actions applied to traffic at a load balancer listener.

You can include the following types of rules in a rule set:

- [Access control rules](#), which restrict access to application resources based on the source of the request.
- [Access method rules](#), which specify the permitted HTTP methods.
- [URL redirect rules](#), which route incoming HTTP requests to a different destination URL.
- [Request and response header rules](#), which add, alter, or remove HTTP request or response headers.

Rule sets apply only to HTTP listeners.

You can apply an existing rule set when you edit a listener. You can apply the same rule set to multiple listeners on the same load balancer.

Rule sets are not shared between load balancers. To use the same set of rules on another load balancer, you must create a new, identical rule set under that load balancer.

You can have up to 20 rules in a rule set. You can associate a maximum of 50 rules with a load balancer.

### Access Control Rules

Access control rules permit access to application resources based on user-specified IP address or address range match conditions. If you do not specify any access control rules, the

default rule is to allow all traffic. If you add access control rules, the load balancer denies any traffic that does not match the rules.

The service accepts only classless inter-domain routing (CIDR) format (`x.x.x.x/y` or `x::x::x/y`) strings for the match condition.

Specify `0.0.0.0/0` or `::/0` to match all incoming traffic.



### Note

Only US Government Cloud regions currently permit IPv6 values.

### Access Method Rules

Access method rules specify the HTTP methods allowed at the associated listener. The load balancer does not forward a disallowed request to the backend servers and returns a `405 Method Not Allowed` response with a list of the allowed methods. You can associate only one list of allowed methods with a given listener.

By default, you can specify only the standard HTTP methods defined in the [HTTP Method Registry](#). The list of HTTP methods is extensible. If you need to configure custom HTTP methods, contact [My Oracle Support](#) to remove the restriction from your tenancy. Your backend application must be able to handle the specified methods.

### Default HTTP Methods

ACL

BASELINE-CONTROL

BIND

CHECKIN

CHECKOUT

## CHAPTER 20 Load Balancing

---

CONNECT

COPY

DELETE

GET

HEAD

LABEL

LINK

LOCK

MERGE

MKACTIVITY

MKCALENDAR

MKCOL

MKREDIRECTREF

MKWORKSPACE

MOVE

OPTIONS

ORDERPATCH

PATCH

POST

PRI

PROPFIND

PROPPATCH

PUT

REBIND

REPORT

SEARCH

TRACE

UNBIND

UNCHECKOUT

UNLINK

UNLOCK

UPDATE

UPDATEREDIRECTREF

VERSION-CONTROL

### URL Redirect rules

URL redirect rules specify how to route incoming HTTP requests to a different destination URL. URL redirect rules apply only to HTTP listeners. You configure each redirect rule for a particular listener and a designated path. A listener can have only one redirect rule for a given incoming URL path.

When you create a URL redirect rule, you specify the path string and match condition the service uses to evaluate an incoming URL for redirection. You also define the redirect URL and response code.

#### **Incoming path string evaluation**

You specify the path string, or pattern, to evaluate in the incoming URL. For example:

```
/video
```

You also specify the match condition to apply when evaluating the incoming URL for redirection. The available match types are:

- **FORCE\_LONGEST\_PREFIX\_MATCH**  
The system looks for a redirect rule path string with the best, longest match of the beginning portion of the incoming URL path.
- **EXACT\_MATCH**  
The incoming URL path must exactly and completely match the specified path string.
- **PREFIX\_MATCH**  
The beginning portion of the incoming URL path must exactly match the specified path string.
- **SUFFIX\_MATCH**  
The ending portion of the incoming URL path must exactly match the specified path string.

### Redirection URL construction

You define the redirect URL applied to the original request. URL redirect rules recognize the following URL components:

```
<protocol>://<host>:<port>/<path>?<query>
```

You can specify a literal string or provide a token for any component. Tokens extract values from the incoming HTTP request URL. Tokens are case-sensitive. For example, `{host}` is a valid token, but `{HOST}` is not.

- **Protocol**  
The HTTP protocol to use in the redirect URL. Valid values are HTTP and HTTPS.  
The `{protocol}` token extracts the protocol from the incoming HTTP request URL. It is the only valid token for this property.
- **Host**  
The valid domain name or IP address to use in the redirect URL.

The `{host}` token extracts the host from the incoming HTTP request URL. All URL Redirect tokens are valid for this property. You can use any token more than once. Curly braces `{ }` are valid in this property only to surround tokens.

- **Port**

The communication port to use in the redirect URL. Valid values include integers from 1 to 65535.

The `{port}` token extracts the port from the incoming HTTP request URL. It is the only valid token for this property.

- **Path**

The HTTP URL path to use in the redirect URL. To omit the path from the redirect URL, set this value to an empty string.

The `{path}` token extracts the path string from the incoming HTTP request URL. All URL Redirect tokens are valid for this property. You can use any token more than once.

If the path string does not begin with the `{path}` token, it must begin with a forward slash `/`.

- **Query**

The query string to use in the redirect URL. To omit all incoming query parameters from the redirect URL, set this value to an empty string.

The `{query}` token extracts the query string from the incoming HTTP request URL. All URL Redirect tokens are valid for this property. You can use any token more than once.

If the query string does not begin with the `{query}` token, it must begin with a question mark `?`.

You can specify multiple query parameters as a single string. Separate each query parameter with an ampersand `&`.

If the specified query string results in a redirect URL ending with `?` or `&`, the last character is truncated. For example, if the incoming URL is

`http://host.com:8080/documents` and the query property value is `?lang=en&{query}`, the redirect URL is `http://host.com:8080/documents?lang=en`. The system truncates the final ampersand `&` because the incoming URL included no value to replace the `{query}` token.



### Warning

Failure to specify a value for at least one URL component field can result in a redirect loop.

### Manual redirect URL construction

The Console provides text entry fields for each URL component. Alternatively, you can manually specify the full redirect URL.

You can retain the literal characters of a token when you specify values for the path and query properties of the redirect URL. Use a backslash `\` as the escape character for the `\`, `{`, and `}` characters. For example, if the incoming HTTP request URL is `/video`, the path property value `/example{path}123\{path\}` appears in the constructed redirect URL as `/example/video123{path}`.

Some path and query string examples:

- **`/example/video/123`** appears as `/example/video/123` in the redirect URL.
- **`/example{path}`** appears as `/example/video/123` in the redirect URL when `/video/123` is the path in the incoming HTTP request URL.
- **`{path}/123`** appears as `/example/video/123` in the redirect URL when `/example/video` is the path in the incoming HTTP request URL.
- **`{path}123`** appears as `/example/video123` in the redirect URL when `/example/video` is the path in the incoming HTTP request URL.
- **`/ {host} /123`** appears as `/example.com/123` in the redirect URL when `example.com` is the hostname in the incoming HTTP request URL.
- **`/ {host} / {port}`** appears as `/example.com/123` in the redirect URL when `example.com` is the hostname and `123` is the port in the incoming HTTP request URL.
- **`/ {query}`** appears as `/lang=en` in the redirect URL when the query is `lang=en` in the incoming HTTP request URL.

- **lang=en&time\_zone=PST** appears as `lang=en&time_zone=PST` in the redirect URL.
- **{query}** appears as `lang=en&time_zone=PST` in the redirect URL when `lang=en&time_zone=PST` is the query string in the incoming HTTP request. If the incoming HTTP request has no query parameters, the `{query}` token renders as an empty string.
- **lang=en&{query}&time\_zone=PST** appears as `lang=en&country=us&time_zone=PST` in the redirect URL when `country=us` is the query string in the incoming HTTP request. If the incoming HTTP request has no query parameters, this value renders as `lang=en&time_zone=PST`.
- **protocol={protocol}&hostname={host}** appears as `protocol=http&hostname=example.com` in the redirect URL when the protocol is `http` and the hostname is `example.com` in the incoming HTTP request.
- **port={port}&hostname={host}** appears as `port=8080&hostname=example.com` in the redirect URL when the port is `8080` and the hostname is `example.com` in the incoming HTTP request URL.

### Response code

You can specify the HTTP status code to return when the incoming request is redirected. Valid response codes for redirection from the standard HTTP specification are:

- 301 Moved Permanently
- 302 Found
- 303 See Other
- 307 Temporary Redirect
- 308 Permanent Redirect

The default value is 302 Found.

### Request and Response Header Rules

Request and response header rules add, alter, or remove HTTP request or response headers. These rules can help you pass metadata to your backend servers to do things like:

- Identify which listener sent a request.
- Notify a backend server about SSL termination.

Examples of how rule sets can help you enhance site security include:

- Adding headers to prevent external domains from iframing your site.
- Removing debug headers, such as "Server", sent by backend servers. This action helps you hide the implementation details of your backend.
- Adding the "strict-transport-security" header, with a proper value, to responses. This header helps guarantee that access to your site is HTTPS only.
- Adding the "x-xss-protection" header with a proper value. This header helps you enforce the cross-site scripting (XSS) protection built into modern browsers.
- Adding the "x-content-type" header with a proper value. This header helps you prevent attacks based on content type shifting.

### **Example: Notify WebLogic that the load balancer terminated SSL**

You can configure your load balancer to perform SSL termination. Often, your backend applications require notification of this action. For example, HTTPS [WebLogic](#) e-commerce online transaction processing looks for the `WL-Proxy-SSL` header to confirm that a request came in over SSL. You can use rule sets to add this header at the load balancer listener.



#### **Tip**

For security reasons, WebLogic ignores this header unless you select (check) the **WebLogic Plugin Enabled** check box in WebLogic's Administration Console.

1. Follow the instructions to [create a rule set](#) and:
  - a. Choose the **Add Request Header** option from the **Action** drop-down list.
  - b. Enter `WL-Proxy-SSL` as the **Header** name.
  - c. Set the header **Value**:
    - If your load balancer is configured to perform SSL termination, set this value to "true".
    - If the SSL termination point is in the web server where the plug-in operates, set this value to "false".
2. [Create a listener](#), or [edit an existing listener](#), and add the new rule set.

### Using the Console

To apply a rule set to a listener, you first create the rule set that contains the rules. The rule set becomes a part of the load balancer's configuration. You can specify the rule set to use when you create or update a listener for the load balancer.

### To create a rule set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Rule Sets**, and then click **Create Rule Set**.
4. In the **Create Rule Set** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the rule set. The name must be unique, and cannot be changed. Avoid entering confidential information.
  - **Specify Access Control Rules:** Optional. Select (check) this box to add access control rules.

- **IP Address CIDR:** Enter the IP address CIDR block from which access is allowed.
- **+ Another Access Control Rule:** Optional. Click this button to enter another IP address CIDR or click the corresponding **X** to remove an existing entry.
- **Specify Access Method Rules:** Optional. Select (check) this box to add access method rules.
  - **Allowed Methods:** From the drop-down list, select the HTTP methods to allow. You can select multiple methods. Click the label's **X** to remove an existing method.
- **Specify URL Redirect Rules:** Optional. Select (check) this box to add URL redirect rules.
  - **Source Path:** Specify the incoming path string that triggers the redirect rule. For example, `/video`.
  - **Match Type:** Choose the match condition to apply when evaluating an incoming path string. The available match types are:
    - **FORCE\_LONGEST\_PREFIX\_MATCH**  
The system looks for a redirect rule path string with the best, longest match of the beginning portion of the incoming URL path.
    - **EXACT\_MATCH**  
The incoming URL path must exactly and completely match the specified path string.
    - **PREFIX\_MATCH**  
The beginning portion of the incoming URL path must exactly match the specified path string.
    - **SUFFIX\_MATCH**  
The ending portion of the incoming URL path must exactly match the specified path string.

- **Redirect to:** Specify a value for at least one URL component field. Component fields you do not modify retain the values of the incoming URL. Optionally, click the **Switch to full URL** link to enter the redirect URL manually.



### Warning

Failure to specify a value for at least one URL component field can result in a redirect loop.

- **Protocol:** Specify the HTTP protocol to use in the redirect URL. Valid values are:
  - {protocol}
  - HTTPS
  - HTTP
- **Host:** Specify a valid domain name (hostname) or IP address for the redirect URL. All redirect URL tokens are valid for this property.
- **Port:** Specify the communication port to use in the redirect URL. Valid values include integers from 1 to 65535.
- **Path:** The HTTP URL path to use in the redirect URL. All redirect URL tokens are valid for this property.



### Important

If the path string does not begin with the `{path}` token, it must begin with the forward slash character `/`.

- **Query:** Specify the query string to use in the redirect URL. All redirect URL tokens are valid for this property.



### Important

If the query string does not begin with the `{query}` token, it must begin with the question mark `?` character.

- **Response Code:** Specify the HTTP status code to return when the incoming request is redirected. The default response code is **302 Found**.

Valid response codes for redirection from the standard HTTP specification are:

- 301 Moved Permanently
- 302 Found
- 303 See Other
- 307 Temporary Redirect
- 308 Permanent Redirect

- **+ Another URL Redirect Rule** Optional. Click this button to create another rule or click the corresponding **X** to delete an existing rule.
- **Specify Request Header Rules:** Optional. Select (check) this box to add request header rules.
  - **Order:** Optional. If you have multiple rules, you can click the up or down arrows to move the corresponding rule.
  - **Action:** Select the action that the rule applies. Available actions include:
    - **Add Request Header**

Adds the specified header and value to the incoming request. If the specified header is already present, the system replaces it. If more than one header with the same name is present, the system removes all of them and adds one header corresponding to the specified header and value.
    - **Extend Request Header**

Adds the specified prefix or suffix to the incoming request. You must provide a prefix value, a suffix value, or both when you choose this action. The system does not support this rule for headers with multiple values.
    - **Remove Request Header**

Removes the specified header. If the same header appears more than once in the request, the load balancer removes all occurrences of the specified header.



### Note

These rules apply only to HTTP or HTTP2 headers.

- **Header:** A header name that conforms to RFC 7230.



### Warning

The system does not distinguish between underscore and dash characters in headers. That is, it treats `example_header_name` and `example-header-name` as identical. Oracle recommends that you do not rely on underscore or dash characters to uniquely distinguish header names.

- **Value:** (Add rules only.) A header value that conforms to RFC 7230.
- **Prefix:** (Extend rules only.) A character string to add to the beginning of the existing header name. The resulting header must conform to RFC 7230.
- **Suffix:** (Extend rules only.) A character string to add to the end of the existing header name. The resulting header must conform to RFC 7230.
- **+ Another Request Header Rule** Optional. Click this button to create another rule or click the corresponding **X** to delete an existing rule.
- **Specify Response Header Rules:** Optional. Select (check) this box to add response header rules.
  - **Order:** Optional. If you have multiple rules, you can click the up or down arrows to move the corresponding rule.
  - **Action:** Select the action that the rule applies. Available actions include:
    - **Add Response Header**  
Adds the specified header and value to the outgoing response.  
If the specified header is already present, the system replaces it.

If more than one header with the same name is present, the system removes all of them and adds one header corresponding to the specified header and value.

- **Extend Response Header**

Adds the specified prefix or suffix to the incoming request.

You must provide a prefix value, a suffix value, or both when you choose this action.

The system does not support this rule for headers with multiple values.

- **Remove Response Header**

Removes the specified header.

If the same header appears more than once in the response, the load balancer removes all occurrences of the specified header.



**Note**

These rules apply only to HTTP or HTTP2 headers.

- **Header:** A header name that conforms to RFC 7230.



### Warning

The system does not distinguish between underscore and dash characters in headers. That is, it treats `example_header_name` and `example-header-name` as identical. Oracle recommends that you do not rely on underscore or dash characters to uniquely distinguish header names.

- **Value:** (Add rules only.) A header value that conforms to RFC 7230.
- **Prefix:** (Extend rules only.) A character string to add to the beginning of the existing header name. The resulting header must conform to RFC 7230.
- **Suffix:** (Extend rules only.) A character string to add to the end of the existing header name. The resulting header must conform to RFC 7230.
- **+ Another Response Header Rule** Optional. Click this button to create another rule or click the corresponding **X** to delete an existing rule.

5. Click **Create**.

After you create a rule set, the set becomes available for use with the associated load balancer. [Update a listener](#) to apply the rule set.

### To update a rule set

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer associated with the rule set you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Rule Sets**.
4. Click the name of the rule set you want to edit, and then click **Edit Rules**.

5. In the **Edit Rule Set** dialog box, add or remove rules as needed.  
You cannot edit the **Name** field of an existing rule set.
6. Click **Save Changes**.

### To remove a rule set from a listener

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer associated with the rule set you want to delete, and then click the load balancer's name.
3. In the **Resources** menu, click **Listeners**.
4. For the listener you want to edit, click the Actions icon (three dots), and then click **Edit Listener**.
5. In the **Rule Sets** section of the dialog box, click the corresponding **X** to remove an existing rule set.



#### Tip

This action removes the rule set from the current listener, but the rule set remains available for application to other listeners on the load balancer.

6. Click **Save Changes**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage rule sets:

- [CreateRuleSet](#)
- [DeleteRuleSet](#)
- [GetRuleSet](#)
- [ListRuleSets](#)
- [UpdateRuleSet](#)

## Managing SSL Certificates

This topic is part of the setup and maintenance of a load balancer. For more information about managing load balancers, see [Managing a Load Balancer](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Working with SSL Certificates

To use SSL with your load balancer, you must [add one or more certificate bundles](#) to your system. The certificate bundle you upload includes the public certificate, the corresponding private key, and any associated Certificate Authority (CA) certificates. For the easiest workflow, upload the certificate bundles you want to use *before* you create the listeners or backend sets you want to associate them with.

Load balancers commonly use single domain certificates. However, load balancers with listeners that include [request routing configuration](#) might require a subject alternative name (SAN) certificate (also called multi-domain certificate) or a wildcard certificate. The Load Balancing service supports each of these certificate types.



### Important

- The Load Balancing service does not generate SSL certificates. It can only import an existing certificate that you already own. The certificate can be one issued by a vendor, such as Verisign or GoDaddy. You can also use a self-signed certificate that you generate with an open source tool, such as [OpenSSL](#) or [Let's Encrypt](#). Refer to the corresponding tool's documentation for instructions on how to generate a self-signed certificate.
- If you submit a self-signed certificate for backend SSL, you must submit the same certificate in the corresponding CA Certificate field.

Oracle Cloud Infrastructure accepts x.509 type certificates in PEM format only. The following is an example PEM encoded certificate:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
```

### Converting to PEM format

If you receive your certificates and keys in formats other than PEM, you must convert them before you can upload them to the system. You can use OpenSSL to convert certificates and keys to PEM format. The following example commands provide guidance.

#### CERTIFICATE OR CERTIFICATE CHAIN FROM DER TO PEM

```
openssl x509 -inform DER -in <certificate_name>.der -outform PEM -out <certificate_name>.pem
```

#### PRIVATE KEY FROM DER TO PEM

```
openssl rsa -inform DER -in <private_key_name>.der -outform PEM -out <private_key_name>.pem
```

## CHAPTER 20 Load Balancing

---

### CERTIFICATE BUNDLE FROM PKCS#12 (PFX) TO PEM

```
openssl pkcs12 -in <certificate_bundle_name>.p12 -out <certificate_bundle_name>.pem -nodes
```

### CERTIFICATE BUNDLE FROM PKCS#7 TO PEM

```
openssl pkcs7 -in <certificate_bundle_name>.p7b -print_certs -out <certificate_bundle_name>.pem
```

### Uploading Certificate Chains

If you have multiple certificates that form a single certification chain, you must include all relevant certificates in one file before you upload them to the system. The following example of a certificate chain file includes four certificates:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
```

### Submitting Private Keys



#### Tip

Oracle recommends a minimum length of 2048 bits for your RSA private key.

## CHAPTER 20 Load Balancing

---

If your private key submission returns an error, the three most common reasons are:

- You provided an incorrect passphrase.
- Your private key is malformed.
- The system does not recognize the encryption method used for your key.

### PRIVATE KEY CONSISTENCY

If you receive an error related to the private key, you can use OpenSSL to check its consistency:

```
openssl rsa -check -in <private_key>.pem
```

This command verifies that the key is intact, the passphrase is correct, and the file contains a valid RSA private key.

### DECRYPTING A PRIVATE KEY

If the system does not recognize the encryption technology used for your private key, decrypt the key. Upload the unencrypted version of the key with your certificate bundle. You can use OpenSSL to decrypt a private key:

```
openssl rsa -in <private_key>.pem -out <decrypted_private_key>.pem
```

### Updating an Expiring Certificate

To ensure consistent service, you must update (rotate) expiring certificates:

1. Update your client or backend server to work with a new certificate bundle.

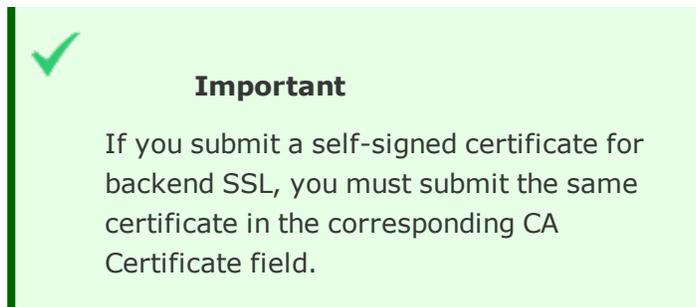


#### Note

The steps to update your client or backend server are unique to your system.

2. Upload the new SSL certificate bundle to the load balancer

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. Click the load balancer you want to configure.
- d. In the **Resources** menu, click **Certificates**, and then click **Add Certificate**.
- e. In the **Add Certificate** dialog box, enter the following:
  - **Certificate Name:** Required. Specify a friendly name for the certificate bundle. It must be unique within the load balancer, and it cannot be changed in the Console. (It can be changed using the API.) Avoid entering confidential information.
  - **Choose SSL Certificate File:** Required. Drag and drop the certificate file, in PEM format, into the **SSL Certificate** field.  
Alternatively, you can choose the **Paste SSL Certificate** option to paste a certificate directly into this field.



- **Specify CA Certificate:** Optional. (Recommended for backend SSL termination configurations.) Select (check) this box if you want to provide a CA certificate.
  - **Choose CA Certificate File:** Drag and drop the CA certificate file, in PEM format, into the **CA Certificate** field.

Alternatively, you can choose the **Paste CA Certificate** option to paste a certificate directly into this field.

- **Specify Private Key:** Optional. (Required for SSL termination.) Select (check) this box if you want to provide a private key for the certificate.

- **Choose Private Key File:** Drag and drop the private key, in PEM format, into the **Private Key** field.

Alternatively, you can choose the **Paste Private Key** option to paste a private key directly into this field.

- **Enter Private Key Passphrase:** Optional. Specify the private key passphrase.

- f. Click **Add Certificate**.

3. Edit listeners or backend sets (as needed) so they use the new certificate bundle

#### EDITING A LISTENER:

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. In the **Resources** menu, click **Listeners**.
- d. For the listener you want to edit, click the Actions icon (three dots), and then click **Edit Listener**.
- e. In the **Certificate Name** drop-down list, choose the new certificate bundle.
- f. Click **Submit**.

### EDITING A BACKEND SET:



#### Warning

Updating the backend set temporarily interrupts traffic and can drop active connections.

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set you want to edit.
- d. Click **Edit Backend Set**.
- e. In the **Edit Backend Set** dialog box, select (check) **Use SSL**.
- f. In the **Certificate Name** drop-down list, choose the new certificate bundle.
- g. Click **Save Changes**.

#### 4. (Optional) Remove the expiring SSL certificate bundle



#### Important

You cannot delete an SSL certificate bundle that is associated with a listener or backend set. Remove the bundle from any additional listeners or backend sets before deleting.

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.

- b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. Click the load balancer you want to configure.
- d. In the **Resources** menu, click **Certificates**.
- e. For the certificate you want to delete, click the Actions icon (three dots), and then click **Delete**.
- f. Confirm when prompted.

### Configuring SSL Handling

With Oracle Cloud Infrastructure Load Balancing, you can:

- Terminate SSL at the load balancer. This configuration is *frontend SSL*. Your load balancer can accept encrypted traffic from a client. There is no encryption of traffic between the load balancer and the backend servers.
- Implement SSL between the load balancer and your backend servers. This configuration is *backend SSL*. Your load balancer does not accept encrypted traffic from client servers. Traffic between the load balancer and the backend servers is encrypted.
- Implement end to end SSL. Your load balancer can accept SSL encrypted traffic from clients and encrypts traffic to the backend servers.

#### Terminating SSL at the Load Balancer

To terminate SSL at the load balancer, you must [create a listener](#) at a port such as 443, and then associate an uploaded certificate bundle with the listener.

#### Implementing Backend SSL

To implement SSL between the load balancer and your backend servers, you must [associate an uploaded certificate bundle with the backend set](#).



### Tip

- If you want to have more than one backend server in the backend set, sign your backend servers with an intermediate CA certificate. The intermediate CA certificate must be included as part of the certificate bundle.
- Your backend services must be able to accept and terminate SSL.

### Implementing End to End SSL

To implement end to end SSL, you must associate uploaded certificate bundles with both the [listener](#) and the [backend set](#).

### Using the Console

#### To upload an SSL certificate bundle to your load balancing system

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. Click the load balancer you want to configure.
4. In the **Resources** menu, click **Certificates**, and then click **Add Certificate**.
5. In the **Add Certificate** dialog box, enter the following:
  - **Certificate Name:** Required. Specify a friendly name for the certificate bundle. It must be unique within the load balancer, and it cannot be changed in the Console. (It can be changed using the API.) Avoid entering confidential

information.

- **Choose SSL Certificate File:** Required. Drag and drop the certificate file, in PEM format, into the **SSL Certificate** field.

Alternatively, you can choose the **Paste SSL Certificate** option to paste a certificate directly into this field.



### Important

If you submit a self-signed certificate for backend SSL, you must submit the same certificate in the corresponding CA Certificate field.

- **Specify CA Certificate:** Optional. (Recommended for backend SSL termination configurations.) Select (check) this box if you want to provide a CA certificate.
    - **Choose CA Certificate File:** Drag and drop the CA certificate file, in PEM format, into the **CA Certificate** field.  
Alternatively, you can choose the **Paste CA Certificate** option to paste a certificate directly into this field.
  - **Specify Private Key:** Optional. (Required for SSL termination.) Select (check) this box if you want to provide a private key for the certificate.
    - **Choose Private Key File:** Drag and drop the private key, in PEM format, into the **Private Key** field.  
Alternatively, you can choose the **Paste Private Key** option to paste a private key directly into this field.
    - **Enter Private Key Passphrase:** Optional. Specify the private key passphrase.
6. Click **Add Certificate**.

To delete an SSL certificate bundle from your load balancing system



### Important

You cannot delete an SSL certificate bundle that is associated with a listener or backend set. Remove the bundle from any listeners or backend sets before deleting.

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. Click the load balancer you want to configure.
4. In the **Resources** menu, click **Certificates**.
5. For the certificate you want to delete, click the Actions icon (three dots), and then click **Delete**.
6. Confirm when prompted.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage load balancer certificates:

- [CreateCertificate](#)
- [DeleteCertificate](#)
- [ListCertificates](#)

# Editing Health Check Policies

This topic describes how to modify health check policies for a backend set.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

## Working with Health Check Policies

A health check is a test to confirm the availability of backend servers. A health check can be a request or a connection attempt. Based on a time interval you specify, the load balancer applies the health check policy to continuously monitor backend servers. If a server fails the health check, the load balancer takes the server temporarily out of rotation. If the server subsequently passes the health check, the load balancer returns it to the rotation.

You configure your health check policy when you [create a backend set](#). You can configure TCP-level or HTTP-level health checks for your backend servers.

- TCP-level health checks attempt to make a TCP connection with the backend servers and validate the response based on the connection status.
- HTTP-level health checks send requests to the backend servers at a specific URI and validate the response based on the status code or entity data (body) returned.

The service provides application-specific health check capabilities to help you increase availability and reduce your application maintenance window.



### **Important**

*Configure your health check protocol to match your application or service.*

If you run an HTTP service, be sure to configure an HTTP-level health check. If you run a TCP-level health check against an HTTP service, you might not get an accurate response. The TCP handshake can succeed and indicate that the service is up even when the HTTP service is incorrectly configured or having other issues. Although the health check appears good, customers might experience transaction failures. For example:

- The backend HTTP service has issues and returns 5XX messages. An HTTP health check catches the message and marks the service as down. In this case, a TCP health check handshake succeeds and marks the service as healthy, even though the HTTP service is not usable.
- The backend HTTP service responds with 4XX messages because of authorization issues or no configured content. A TCP health check does not catch these errors.

## Health Status

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components. You can see health status indicators on the Console *List* and *Details* pages for load balancers, backend

sets, and backend servers. You also can use the Load Balancing API to retrieve this information.

Health status indicators have four levels. The general meaning of each level is:

**OK (GREEN)**

No attention required.

The resource is functioning as expected.

**WARNING (YELLOW)**

Some reporting entities require attention.

The resource is not functioning at peak efficiency or the resource is incomplete and requires further work.

**CRITICAL (RED)**

Some or all reporting entities require immediate attention.

The resource is not functioning or unexpected failure is imminent.

**UNKNOWN (GRAY)**

Health status cannot be determined.

The resource is not responding or is in transition and might resolve to another status over time.

The precise meaning of each level differs among the following components:

- [Load balancers](#)
- [Backend sets](#)
- [Backend servers](#)

### Using Health Status

At the highest level, load balancer health reflects the health of its components. The health status indicators provide information you might need to drill down and investigate an existing

issue. Some common issues that the health status indicators can help you detect and correct include:

**A HEALTH CHECK IS MISCONFIGURED.**

In this case, all the backend servers for one or more of the affected listeners report as unhealthy. If your investigation finds that the backend servers do not have problems, then a backend set probably includes a misconfigured health check.

**A LISTENER IS MISCONFIGURED.**

All the backend server health status indicators report **OK**, but the load balancer does not pass traffic on a listener.

The listener might be configured to:

- Listen on the wrong port.
- Use the wrong protocol.
- Use the wrong policy.

If your investigation shows that the listener is not at fault, check the security list configuration.

**A SECURITY RULE IS MISCONFIGURED.**

Health status indicators help you diagnose two cases of misconfigured security rules:

- All entity health status indicators report **OK**, but traffic does not flow (as with misconfigured listeners). If the listener is not at fault, check the security rule configuration.
- All entity health statuses report as unhealthy. You have checked your health check configuration and your services run properly on your backend servers.

In this case, your security rules might not include the IP range for the source of the health check requests. You can find the health check source IP on the [Details](#) page for each backend server. You can also use the API to find the IP in the `sourceIpAddress` field of the [HealthCheckResult](#) object.



### Note

#### Source IP

The source IP for health check requests comes from a Compute instance managed by the Load Balancing service.

#### ONE OR MORE OF THE BACKEND SERVERS REPORTS AS UNHEALTHY.

A backend server might be unhealthy or the health check might be misconfigured. To see the corresponding error code, check the **status** field on the backend server's [Details](#) page. You can also use the API to find the error code in the `healthCheckStatus` field of the [HealthCheckResult](#) object.

#### OTHER CASES IN WHICH HEALTH STATUS MIGHT PROVE HELPFUL INCLUDE:

- VCN [network security groups](#) or [security lists](#) block traffic.
- Compute instances have misconfigured route tables.

### Health Status Limitations

Health status is updated every three minutes. No finer granularity is available.

Health status does not provide historical health data.

### Using the Console

You create your health check tests when you [create a backend set](#).

### To edit an existing health check policy

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.

2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set you want to modify.
4. Click **Update Health Check**.
5. In the **Health Check** section, specify the test parameters to confirm the health of backend servers.



### Tip

All parameters are required when updating an existing health check policy.

- **Protocol:** Required. Specify the protocol to use, either HTTP or TCP.



### Important

Configure your health check protocol to [match your application or service](#).

- **Port:** Required. Specify the backend server port against which to run the health check.



### Tip

You can enter the value '0' to have the health check use the backend server's traffic port.

- **URL Path (URI):** (HTTP only) Required. Specify a URL endpoint against which to run the health check.
- **Interval in ms:** Required. Specify how frequently to run the health check, in milliseconds. Default is 10000 (10 seconds).
- **Timeout in ms:** Required. Specify the maximum time in milliseconds to wait for a reply to a health check. A health check is successful only if a reply returns within this timeout period. Default is 3000 (3 seconds).
- **Number of retries:** Required. Specify the number of retries to attempt before a backend server is considered "unhealthy". This number also applies when recovering a server to the "healthy" state. Default is 3.
- **Status Code:** (HTTP only) Required. Specify the status code a healthy backend server must return.
- **Response Body Regex:** (HTTP only) Optional. Provide a regular expression for parsing the response body from the backend server. The system treats a blank entry here as the value ".\*".



### Tip

Health checks require all fields to match. Your status code and response body both must match, as specified.

6. Click **Save**.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use this API operation to edit a backend set's health check policy:

[UpdateBackendSet](#)

Use these API operations to retrieve health status information:

- [GetBackendHealth](#)
- [GetBackendSetHealth](#)
- [GetLoadBalancerHealth](#)
- [ListLoadBalancerHealths](#)

## Viewing the State of a Work Request

This topic describes how to view the state of work requests associated with a given load balancer.



### Note

The Load Balancing service does not use the common [Work Requests API](#) to support work request operations. Instead, Load Balancing work requests are supported by the Load Balancing API. See [Using the Console to View Work Requests](#) for information on viewing work requests for other services.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: For a typical policy that gives access to load balancers and their components, see [Let network admins manage load balancers](#).

Also, be aware that a policy statement with `inspect load-balancers` gives the specified group the ability to see *all* information about the load balancers. For more information, see [Details for Load Balancing](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Monitoring Work Requests

Many of the Oracle Cloud Infrastructure Load Balancing service requests do not take effect immediately. In these cases, the request spawns an asynchronous workflow for fulfillment. To provide visibility for in-progress workflows, the Load Balancing service creates a work request object. Because some operations depend on the completion of other operations, you must monitor each operation's work request and confirm it has succeeded before proceeding to the next operation. For example, if you want to create a backend set and add a backend server to the new set, you first must create the backend set. After that operation completes, you can add the backend server. If you try to add a backend server before the backend set creation completes, the system cannot ensure that the request to add the server succeeds. You can monitor the request to add a backend set to determine when that workflow is complete, and then add the backend server.

The work request states are:

#### **ACCEPTED**

The request is in the work request queue to be processed.

#### **IN PROGRESS**

A work request record exists for the specified request, but there is no associated WORK\_COMPLETED record.

#### **SUCCEEDED**

A work request record exists for this request and an associated WORK\_COMPLETED record has the state SUCCEEDED.

### **FAILED**

A work request record exists for this request and an associated WORK\_COMPLETED record has the state FAILED.

## Using the Console

The Oracle Cloud Infrastructure Console consumes the REST API and is subject to the same considerations as any Oracle Cloud Infrastructure client. You can view the state of a load balancing work request in the Console:

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to review, and then click the load balancer's name.
3. In the **Resources** menu, click **Work Requests**. The status of all work requests appears on the page.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these operations to monitor the state of work requests:

- [ListWorkRequests](#)
- [GetWorkRequest](#)

## Load Balancing Metrics

You can monitor the health, capacity, and performance of your load balancers by using metrics, alarms, and [notifications](#).

This topic describes the metrics emitted by the Load Balancing service in the `oci_lbaas` metric namespace.

Resources: Load balancers, listeners, and backend sets.

### Overview of the Load Balancing Service Metrics

Your load balancer acts as an intermediary for data traffic between clients and your application servers. Clients send requests to your load balancer and the load balancer distributes the requests to your backend servers according to rules you establish. See the diagram in [Overview of Load Balancing](#) for a high-level view of a simple public load balancing system configuration.

The Load Balancing service metrics help you measure the number and type of connections, and quantity of data managed by your load balancer. You can use metrics data to diagnose and troubleshoot load balancer and client issues. The metrics also help you analyze the HTTP responses returned by the servers in your backend set.

To view a default set of metrics charts in the Console, navigate to the load balancer or backend set you're interested in, and then click **Metrics**. You also can use the Monitoring service to create [custom queries](#).

### Prerequisites

- **IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- The metrics listed on this page are automatically available for any load balancer, listener, and backend set you create. You do not need to enable monitoring on the resource to get these metrics.

### Available Metrics: oci\_lbaas

Load Balancing service metrics include the following dimensions:

#### **AVAILABILITYDOMAIN**

The availability domain in which the load balancer resides.

#### **BACKENDSETNAME**

The name of the backend set to which the metrics apply.

#### **LBCOMPONENT**

The load balancer component to which the metrics apply.

Valid metrics for the Load Balancing service vary among the three **lbComponent** dimension values:

- **Backendset**
- **Listener**
- **Loadbalancer**

The tables on this page describe which data is valid for each of these dimension values. If you choose a metric that does not apply to the specified dimension value, the metric returns no data.

#### **LBHOSTID**

A unique ID that represents the current load balancer host. This ID is subject to change.

#### **LISTENERNAME**

The name of the listener to which the metrics apply.

#### **REGION**

The region in which the load balancer resides.

#### **RESOURCEID**

The OCID of the resource to which the metrics apply.

**Metrics for the lbComponent Dimension Value "Backendset"**

Metric	Metric Display Name	Unit	Description	Dimensions
ActiveConnections	<b>Active Connections</b>	count	The number of active connections from the load balancer to all backend servers.	availabilityDomain backendSetName lbComponent lbHostId region resourceId
BackendServers	<b>Backend Servers</b>	count	The number of backend servers in the backend set.	
BackendTimeouts	<b>Backend Timeouts</b>	count	The number of timeouts across all backend servers.	
BytesReceived	<b>Bytes Received</b>	bytes	The number of bytes received across all backend servers.	

## CHAPTER 20 Load Balancing

---

<b>Metric</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Description</b>	<b>Dimensions</b>
BytesSent	<b>Bytes Sent</b>	bytes	The number of bytes sent across all backend servers.	
ClosedConnections	<b>Closed Connections</b>	count	The number of connections closed between the load balancer and backend servers.	
HttpRequests	<b>Inbound Requests</b>	count	The number of incoming client requests to the backend set.	
HttpResponses	<b>Responses</b>	count	The number of HTTP responses across all backend servers.	

Metric	Metric Display Name	Unit	Description	Dimensions
HttpResponses200	<b>HTTP 200 Responses</b>	count	The number of HTTP 200 responses received from backend servers.	
HttpResponses2xx	<b>HTTP 2xx Responses</b>	count	The number of HTTP 2xx responses received from backend servers.	
HttpResponses3xx	<b>HTTP 3xx Responses</b>	count	The number of HTTP 3xx responses received from backend servers.	
HttpResponses4xx	<b>HTTP 4xx Responses</b>	count	The number of HTTP 4xx responses received from backend servers.	

Metric	Metric Display Name	Unit	Description	Dimensions
HttpResponses502	<b>HTTP 502 Responses</b>	count	The number of HTTP 502 responses received from backend servers.	
HttpResponses504	<b>HTTP 504 Responses</b>	count	The number of HTTP 504 responses received from backend servers.	
HttpResponses5xx	<b>HTTP 5xx Responses</b>	count	The number of HTTP 5xx responses received from backend servers.	
InvalidHeaderResponses	<b>Invalid Header Responses</b>	count	The number of invalid header responses across all backend servers.	

Metric	Metric Display Name	Unit	Description	Dimensions
KeepAliveConnections	<b>Keep-alive Connections</b>	count	The number of keep-alive connections.	
ResponseTimeFirstByte	<b>Average Response Time (TCP only)</b>	ms	Average time to the first byte of response from backend servers. TCP only.	
ResponseTimeHttpHeader	<b>Average Response Time (HTTP only)</b>	ms	Average response time of backend servers. HTTP only.	
UnhealthyBackendServers	<b>Unhealthy Backend Servers</b>	count	The number of unhealthy backend servers in the backend set.	

**Metrics for the lbComponent Dimension Value "Loadbalancer"**

Metric	Metric Display Name	Unit	Description	Dimensions
AcceptedConnections	<b>Accepted Connections</b>	count	The number of connections accepted by the load balancer.	availabilityDomain lbComponent lbHostId region
AcceptedSSLHandshake	<b>Accepted SSL Handshakes</b>	count	The number of accepted SSL handshakes.	resourceId
ActiveConnections	<b>Active Connections</b>	count	The number of active connections from clients to the load balancer.	
ActiveSSLConnections	<b>Active SSL Connections</b>	count	The number of active SSL connections.	
BytesReceived	<b>Bytes Received</b>	bytes	The number of bytes received by the load balancer.	

Metric	Metric Display Name	Unit	Description	Dimensions
BytesSent	<b>Bytes Sent</b>	bytes	The number of bytes sent by the load balancer.	
FailedSSLClientCertVerify	<b>Failed Client SSL Cert Verifications</b>	count	The number of failed client SSL certificate verifications.	
FailedSSLHandshake	<b>Failed SSL Handshakes</b>	count	The number of failed SSL handshakes.	
HandledConnections	<b>Handled Connections</b>	count	The number of connections handled by the load balancer.	
HttpRequests	<b>Inbound Requests</b>	count	The number of incoming client requests to the load balancer.	

**Metrics for the lbComponent Dimension Value "Listener"**

Metric	Metric Display Name	Unit	Description	Dimensions
HttpResponses200	<b>HTTP 200 Responses</b>	count	The number of HTTP 200 responses received from backend sets.	availabilityDomain lbComponent lbHostId listenerName region resourceId
HttpResponses2xx	<b>HTTP 2xx Responses</b>	count	The number of HTTP 2xx responses received from backend sets.	
HttpResponses3xx	<b>HTTP 3xx Responses</b>	count	The number of HTTP 3xx responses received from backend sets.	
HttpResponses4xx	<b>HTTP 4xx Responses</b>	count	The number of HTTP 4xx responses received from backend sets.	
HttpResponses502	<b>HTTP 502 Responses</b>	count	The number of HTTP 502 responses received from backend sets.	
HttpResponses504	<b>HTTP 504 Responses</b>	count	The number of HTTP 504 responses received from backend sets.	
HttpResponses5xx	<b>HTTP 5xx Responses</b>	count	The number of HTTP 5xx responses received from backend sets.	
HttpResponses	<b>Responses</b>	count	The number of incoming responses received from backend sets.	

### Using the Console

#### To view default metric charts for a single load balancer

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Choose the **Compartment** that contains the load balancer you want to view, and then click the load balancer's name.
3. In the **Resources** menu, click **Metrics** (if necessary).  
The **Metrics** page displays a default set of charts for the current load balancer.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

#### To view default metric charts for multiple load balancers

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Metric Namespace**, select **oci\_lbaas**.  
The **Service Metrics** page displays a default set of charts for the selected metric namespace. For more information about the emitted metrics, see the foregoing table.  
You can also use the Monitoringservice to create [custom queries](#).

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 21 Marketplace

This chapter explains how to work with listings in the Oracle Cloud Infrastructure Marketplace.

## Overview of Marketplace

Oracle Cloud Infrastructure Marketplace is an online store that offers solutions specifically for customers of Oracle Cloud Infrastructure. In the Oracle Cloud Infrastructure Marketplace catalog, you can find listings for two types of solutions from Oracle and trusted partners: images and stacks. These listing types include different categories of applications. Also, some listings are free and others belong to certain pricing models.

Images are templates of virtual hard drives that determine the operating system and software to run on an instance. You can deploy image listings on an Oracle Cloud Infrastructure Compute instance. Marketplace also offers stack listings. Stacks represent definitions of groups of Oracle Cloud Infrastructure resources that you can act on as a group. Each stack has a configuration consisting of one or more declarative configuration files. With an image or a stack, you have a customized, more streamlined way of getting started with a publisher's software.



### Note

Marketplace is not available in Oracle Cloud Infrastructure Government Cloud realms.

## Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the REST API. Instructions for the Console and API are included in topics throughout this guide. However, you cannot specifically access Marketplace using the Command Line Interface (CLI). Marketplace does not have CLI support. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Working with Listings

This topic describes how to work with listings in the Oracle Cloud Infrastructure Marketplace. You can do the following:

- Filter listings to find what you want
- View a listing to learn about the product that it offers
- Launch an instance from an image listing
- Launch stack resources from a stack listing

By default, Oracle Cloud Infrastructure Marketplace displays all listings in its catalog. Image listings have a **Launch Instance** button. Stack listings have a **Launch Stack** button.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators, the following policies provide access to Marketplace. These policies specify different target groups as examples, but you can specify the same target group (for example, MarketplaceUsers) when applying the policies in your own tenancy.

- The policy [Let users list and subscribe to images from the Partner Image catalog](#) gives the specified group the ability to list, read, and use Marketplace image listings.
- The policy [Let users launch Compute instances](#) gives the specified group general access to managing instances and images, along with the required level of access to attach existing block volumes to the instances. Use this policy in conjunction with the preceding policy for users who need to launch instances from image listings. Use this policy in conjunction with the following policy for users who need to launch stacks from stack listings.
- The policies described in [Managing Stacks and Jobs](#) grant access to stacks and jobs in the tenancy. Use the appropriate policy statements to give a group the ability to list, read, and use Marketplace stack listings. (Users do not need permission to run destroy jobs to launch a stack from a Marketplace listing, but they do need permissions to run plan jobs and apply jobs.)

- If you need to write more restrictive policies for Marketplace, see [Details for the Core Services](#) and [Details for Resource Manager](#), as needed.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Using the Console

#### To filter listings

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Marketplace**.
2. Under **Filters**, do one or more of the following:
  - To display listings of a certain deployment type, click **Type**, and then click either **Image** or **Stack**.
  - To display listings from a specific publisher, click **Publisher**, and then click a publisher name.
  - To display listings from a particular product category, click **Category**, and then click a category name.
  - To display listings according to price, click **Price**, and then click a pricing model.

You can combine multiple filters to further narrow down listings. You can also clear filters to expand the list of listings that you see.

#### To view a listing's details

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Marketplace**.
2. Click the listing that you're interested in.
3. Marketplace displays the listing overview by default. To view other details, do the following:

- To view information about the publisher, click **Provider**.
- To view other listings from the same publisher, click **More Apps**.
- To view instructions for using the instance that you create from the listing, click **Usage Information**.

### To launch an instance based on an image



#### Tip

When you create an instance, several other resources are involved (for example, an image, a cloud network, or a subnet). Those other resources can be in the same compartment with the instance or in other compartments. You must have the required level of access to each of the compartments involved in order to launch the instance. This is also true when you attach a volume to an instance; they don't have to be in the same compartment, but if they're not, you need the required level of access to each of the compartments.

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Marketplace**.
2. Click the listing that you're interested in.
3. Review the **Usage Instructions** tab to ensure you understand what you will need to deploy and to access the instance after you launch it. Both Linux and Windows instances require a cloud network to launch the instance into. For more information, see [Overview of Networking](#). Depending on the type of instance, to access it, you might need an SSH key pair or a security list that enables Remote Desktop Protocol. For more information, see [Managing Key Pairs on Linux Instances](#) and [To enable RDP access](#).

4. Under **Version**, click the package version of the image that you want to install. By default, the menu displays the latest version.
5. Under **Compartment**, click the name of the compartment where you want to launch the instance.
6. Select the check box to accept the terms of use, and then click **Launch Instance**.
7. To finish launching the instance, follow the instructions in [Creating an Instance](#).

The information you need to connect to an instance after you create it might be in the **Usage Information** or the **Related Documents** sections of the listing.

### To launch a stack



#### Tip

When you create a stack, potentially many other resources are involved (for example, an instance, a cloud network, or a subnet), aside from the stacks and jobs resources. You must have the required access to all involved resources to create a stack. Those other resources can be in the same compartment with the instance or in other compartments. You must have the required level of access to each of the compartments involved in order to launch the instance. This is also true when you attach a volume to an instance; they don't have to be in the same compartment, but if they're not, you need the required level of access to each of the compartments.

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Marketplace**.

2. Click the listing that you're interested in.
3. Review the **Usage Instructions** tab and ensure you understand what you will need to deploy and to access the instance after the stack finishes deployment.
4. Under **Version**, click the package version of the stack that you want to install. By default, the menu displays the latest version.
5. Under **Compartment**, click the name of the compartment where you want to launch the instance.
6. Select the check box to accept the terms of use, and then click **Launch Stack**.
7. On the **Stack Information** page, configure the following:
  - **Name**. Optionally, provide a name by which you can refer to the stack after it's deployed.
  - **Description**. Optionally, provide a description of the stack. For example, you can specify the name of the application that will run on the instance after the stack is deployed.
  - **Create in Compartment**. This is the compartment where the stack will be created in the tenancy. (Stacks are attached to a specific region. However, where necessary, the resources on a given stack can be deployed across multiple regions.)
  - **Tags**. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

When you are ready, click **Next**.

8. On the **Configure Variables** page, verify that the values for variables extracted from the Terraform configuration file are as you want them. Some variables might be required, but don't have a default value and must be configured before you can proceed. These vary from listing to listing, but often include the following: availability domain

and compartment. Optionally, you can change default values, such as any display names automatically given to resources, to help differentiate them. For some stacks, you can customize additional variables by selecting the **Additional Customization** or **WLS Instance Advanced Configuration** check box. The variables in these sections otherwise use default values. When you are ready, click **Next**.

9. On the **Review** page, confirm that variables have been configured properly. (Marketplace does not display variables that have default values or variables that you didn't change.) Then, click **Create**.

Resource Manager runs the plan job and the apply job to create stack resources accordingly. The information you need to connect to the instance created as part of the stack can appear in the **Application Information** tab or in the **Usage Information** or **Related Documents** sections of the listing.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to work with listings:

- [GetListing](#)
- [GetPackage](#)
- [ListCategories](#)
- [ListListings](#)
- [ListPackages](#)
- [ListPublishers](#)

# Viewing Terms of Use Agreements for Deployed Applications

Before you can launch an image or a stack from a Marketplace listing, you must first read and accept all terms of use agreements associated with the package version that you choose. This topic describes how to see what software terms of use agreements you have accepted through Oracle Cloud Infrastructure Marketplace. Your organization might need or want to review the specific terms of use associated with a particular package version after you deploy one or more Marketplace solutions.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators, the following policies provide access to Marketplace. These policies specify different target groups as examples, but you can specify the same target group (for example, `MarketplaceUsers`) when applying the policies in your own tenancy.

- The policy [Let users list and subscribe to images from the Partner Image catalog](#) gives the specified group the ability to list, read, and use Marketplace image listings.
- The policy [Let users launch Compute instances](#) gives the specified group general access to managing instances and images, along with the required level of access to attach existing block volumes to the instances. Use this policy in conjunction with the preceding policy for users who need to launch instances from image listings. Use this policy in conjunction with the following policy for users who need to launch stacks from stack listings.
- The policies described in [Managing Stacks and Jobs](#) grant access to stacks and jobs in the tenancy. Use the appropriate policy statements to give a group the ability to list, read, and use Marketplace stack listings. (Users do not need permission to run destroy

## CHAPTER 21 Marketplace

---

jobs to launch a stack from a Marketplace listing, but they do need permissions to run plan jobs and apply jobs.)

- If you need to write more restrictive policies for Marketplace, see [Details for the Core Services](#) and [Details for Resource Manager](#), as needed.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Using the Console

To view the accepted terms of use agreements for a given compartment

1. Open the navigation menu. Under the **Solutions and Platform** group, go to **Marketplace**.
2. Click **Deployed Applications**.
3. Under **Scope**, click **Compartment**, and then click the name of the compartment where you deployed a solution you now want to view the terms of use agreements for.
4. In the list, find the listing that you're interested in, and then click the Actions menu (⋮). (Marketplace independently lists package versions with their distinct terms of use agreements. Also, Marketplace tracks what agreements you accept regardless of whether you actually complete the software deployment. Therefore, the list might show the names of solutions that you didn't actually finish deploying in your tenancy.)
5. To review any of the terms of use agreements you accepted for the specified listing and package version, click its name.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to work with a listing's terms of use agreements:

## CHAPTER 21 Marketplace

---

- [CreateAcceptedAgreement](#)
- [DeleteAcceptedAgreement](#)
- [GetAcceptedAgreement](#)
- [GetAgreement](#)
- [ListAcceptedAgreements](#)
- [ListAgreements](#)
- [UpdateAcceptedAgreement](#)

# CHAPTER 22 Monitoring

This chapter explains how to actively and passively monitor performance and usage metrics for your resources.

## Monitoring Overview

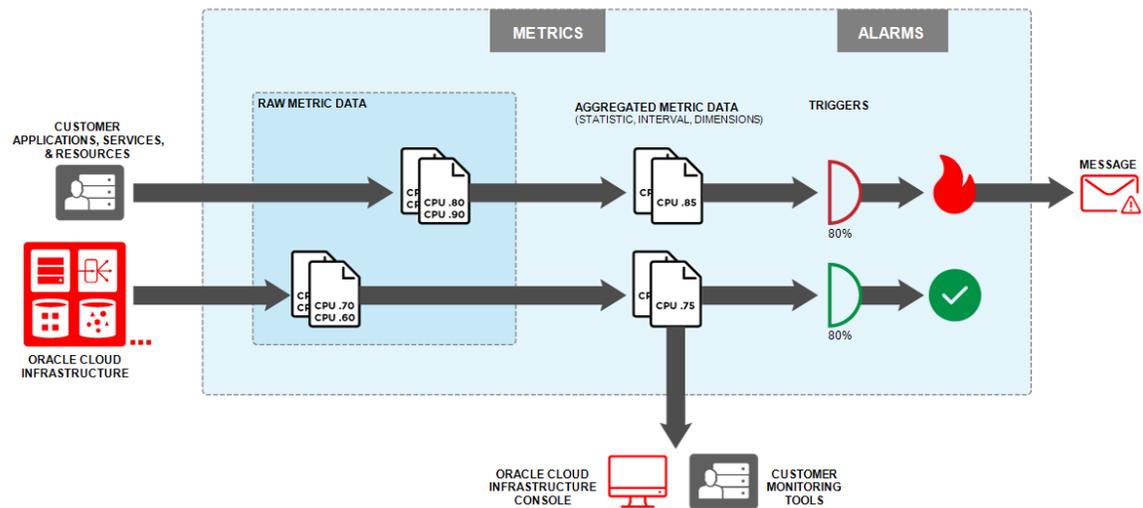
The Oracle Cloud Infrastructure Monitoring service enables you to actively and passively monitor your cloud resources using the Metrics and Alarms features.



### Note

Monitoring is not available in Oracle Cloud Infrastructure Government Cloudrealms.

## Oracle Cloud Infrastructure Monitoring



### How Monitoring Works

The Monitoring service uses metrics to monitor resources and alarms to notify you when these metrics meet alarm-specified triggers.

Metrics are emitted to the Monitoring service by resources as raw data points, or timestamp-value pairs, along with dimensions and metadata. For example, the Compute service (metric namespace "oci\_computeagent") posts this data for monitoring-enabled Compute instances. The posted data includes all oci\_computeagent metrics, such as CpuUtilization. Metric data posted to the Monitoring service is only presented to you or consumed by the Oracle Cloud Infrastructure features that you enable to use metric data.

When you query a metric, the Monitoring service returns aggregated data according to the specified parameters. You can specify a range (such as the last 24 hours), statistic, and interval. The Console displays one monitoring chart per metric for selected resources. The

aggregated data in each chart reflects your selected statistic and interval. API requests can optionally filter by dimension and specify a resolution. API responses include the metric name along with its source compartment and metric namespace. You can feed the aggregated data into a visualization or graphing library.

Metric and alarm data is accessible via the Console, CLI, and API. For retention periods, see [Storage Limits](#).

The Alarms feature of the Monitoring service publishes alarm messages to configured destinations managed by the Notifications service. Each destination is a topic with a set of subscribers. For more information about the Notifications service, see [Notifications Overview](#).

### Message types

The message type indicates the reason that the message was sent.

- **OK\_TO\_FIRING:** The alarm changed from `OK` status to `FIRING` status.
- **FIRING\_TO\_OK:** The alarm changed from `FIRING` status to `OK` status.
- **REPEAT:** The alarm is maintaining a `FIRING` status and repeat notifications are configured.
- **RESET:** The alarm is not detecting the metric firing; the metric is no longer being emitted. The resource that was emitting the metric might have been moved or terminated.



#### **Important**

When a `RESET` status change occurs, determine the health of the resource.

## Message format and examples

Alarm message format:

Parameter	Description
<b>dedupekey</b> Required	<b>string</b> Unique identifier that can be used for de-duplication.
<b>title</b> Required	<b>string</b> The alarm's configured display name.
<b>body</b>	<b>string</b> The alarm's configured message body.
<b>type</b> Required	<b>string</b> The reason for sending the notification message. Valid values: See <a href="#">Message types</a> .
<b>severity</b> Required	<b>string</b> The highest severity level of the listed alarms. Valid values: CRITICAL, ERROR, WARNING, and INFO
<b>timestampEpochMillis</b> Required	<b>long</b> The time when the alarm was triggered, in milliseconds since epoch time.
<b>alarmMetadata</b> Required	<b>array of objects</b> List of alarms related to this notification message.
<b>version</b> Required	<b>int</b> The version of the alarm message format.

**alarmMetadata** format:

Parameter	Description
<b>id</b> Required	<b>string</b> The alarm OCID.
<b>status</b> Required	<b>string</b> The alarm state. Valid values: OK, FIRING
<b>severity</b> Required	<b>string</b> The alarm severity level. Valid values: CRITICAL, ERROR, WARNING, INFO
<b>query</b> Required	<b>string</b> The alarm's configured query. <pre>CpuUtilization[1m]{availabilityDomain="cumS:PHX-AD-1"}.absent()</pre>
<b>totalMetricsFiring</b> Required	<b>int</b> The number of <a href="#">metric streams</a> represented in this notification message.
<b>dimensions</b>	<b>array of objects</b> List of dimension key-value pairs that identify each <a href="#">metric stream</a> . The list is limited to a hundred entries. Empty for an alarm with a status of OK.

Example message "High CPU Utilization" for an alarm that is continuing to be in the FIRING state. In this example, the message includes two metric streams: one for "myinstance1" and another for "myinstance2."

```
{
 "dedupeKey": "exampleuniqueID",
```

## CHAPTER 22 Monitoring

---

```
"title": "High CPU Utilization",
"body": "Follow runbook at http://example.com/runbooks",
"type": "REPEAT",
"severity": "CRITICAL",
"timestampEpochMillis": 1542406320000,
"alarmMetaData": [
 {
 "id": "ocidl.alarm.oc1.iad.exampleuniqueID",
 "status": "FIRING",
 "severity": "CRITICAL",
 "query": "CpuUtilization[1m].mean() > 0",
 "totalMetricsFiring": 2,
 "dimensions": [
 {
 "instancePoolId": "Default",
 "resourceDisplayName": "myinstance1",
 "faultDomain": "FAULT-DOMAIN-1",
 "resourceId": "ocidl.instance.oc1.iad.exampleuniqueID",
 "imageId": "ocidl.image.oc1.iad.exampleuniqueID",
 "availabilityDomain": "szYB:US-ASHBURN-AD-1",
 "shape": "VM.Standard2.1",
 "region": "us-ashburn-1"
 },
 {
 "instancePoolId": "Default",
 "resourceDisplayName": "myinstance2",
 "faultDomain": "FAULT-DOMAIN-3",
 "resourceId": "ocidl.instance.oc1.iad.exampleuniqueID",
 "imageId": "ocidl.image.oc1.iad.exampleuniqueID",
 "availabilityDomain": "szYB:US-ASHBURN-AD-1",
 "shape": "VM.Standard2.1",
 "region": "us-ashburn-1"
 }
]
 }
],
"version": 1.0
}
```

### Metrics Feature Overview

The Metrics feature relays metric data about the health, capacity, and performance of your cloud resources. A metric is a measurement related to health, capacity, or performance of a given resource. Resources, services, and applications emit metrics to the Monitoring service. Common metrics reflect data related to:

- Availability and latency
- Application uptime and downtime
- Completed transactions
- Failed and successful operations
- Key performance indicators (KPIs), such as sales and engagement quantifiers

By querying Monitoring for this data, you can understand how well the systems and processes are working to achieve the service levels you commit to your customers. For example, you can monitor the CPU utilization and disk reads of your Compute instances. You can then use this data to determine when to launch more instances to handle increased load, troubleshoot issues with your instance, or better understand system behavior.

### Example Metric: Failure Rate

For application health, one of the common KPIs is failure rate, for which a common definition is the number of failed transactions divided by total transactions. This KPI is usually delivered through application monitoring and management software.

As a developer, you can capture this KPI from your applications using [custom metrics](#). Simply record observations every time an application transaction takes place and then post that data to the Monitoring service. In this case, set up metrics to capture failed transactions, successful transactions, and transaction latency (time spent per completed transaction).

### Alarms Feature Overview

The Alarms feature of the Monitoring service works with the Notifications service to notify you when metrics meet alarm-specified triggers. When configured, repeat notifications remind

you of a continued firing state at the configured repeat interval. You are also notified when an alarm transitions back to the OK state, or when an alarm is reset.

You can search for alarms using Search-supported attributes. For more information about Search, see [Overview of Search](#).

### Search-Supported Attributes for Alarms

For attribute descriptions, see [Alarm Reference](#).

- id
- displayName
- compartmentId
- metricCompartmentId
- namespace
- query
- severity
- destinations
- suppression
- isEnabled
- lifecycleState
- timeCreated
- timeUpdated
- tags

### Monitoring Concepts

The following concepts are essential to working with Monitoring.

### AGGREGATED DATA

The result of applying a *statistic* and *interval* to a selection of raw *data points* for a given *metric*. For example, you can apply the *statistic* `max` and *interval* `1h` (one hour) to the last 24 hours of raw *data points* for the *metric* `CpuUtilization`. Aggregated data is displayed in default metric charts in the Console. You can also build metric queries for specific sets of aggregated data. For instructions, see [Viewing Default Metric Charts](#) and [Building Metric Queries](#).

### ALARM

The *alarm query* to evaluate and the *notification destination* to use when the alarm is in the firing state, along with other alarm properties. For instructions on managing alarms, see [Managing Alarms](#).

### ALARM QUERY

The Monitoring Query Language (MQL) expression to evaluate for the *alarm*. An alarm query must specify a *metric*, *statistic*, *interval*, and a *trigger rule* (threshold or absence). The Alarms feature of the Monitoring service interprets results for each returned time series as a Boolean value, where zero represents false and a non-zero value represents true. A true value means that the *trigger rule* condition has been met. For more information, see [Building Metric Queries](#) and the query attribute description in the [Alarm](#) API reference.

### DATA POINT

A timestamp-value pair for the specified *metric*. Example: 2018-05-10T22:19:00Z, 10.4

A data point is either raw or aggregated. Raw data points are posted by the *metric namespace* to the Monitoring service using the [PostMetricData](#) operation. The *frequency* of the data points posted varies by *metric namespace*. For example, your custom namespace might send data points for a given *metric* at a 20-second *frequency*.

Aggregated data points are the result of applying a *statistic* and *interval* to raw data points. The *interval* of the aggregated data points is determined by the [SummarizeMetricsData](#) request. For example, a request specifying the *statistic* `sum` and

*interval* 1h (one hour) returns a `sum` value for each hour of available raw data points for the given *metric*.

### DIMENSION

A qualifier provided in a *metric definition*. Example: Resource identifier (`resourceId`), provided in the definitions of `oci_computeagent` metrics. Use dimensions to filter or group metric data. Example dimension name-value pair for filtering by availability domain: `availabilityDomain = "VeBZ:PHX-AD-1"`

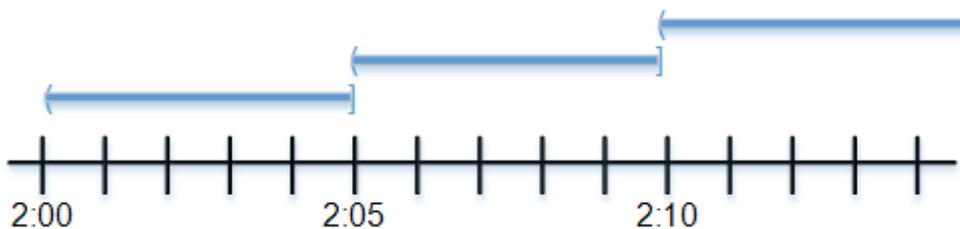
### FREQUENCY

The time period between each posted raw *data point* for a given *metric*. (Raw data points are posted by the *metric namespace* to the Monitoring service.) While frequency varies by metric, default service metrics typically have a frequency of 60 seconds (that is, one data point posted per minute). See also [resolution](#).

### INTERVAL

The time window used to convert the given set of raw *data points*.

The timestamp of the aggregated data point corresponds to the end of the time window during which raw data points are assessed. For example, for a five-minute interval, the timestamp "2:05" corresponds to the five-minute time window from 2:00:n to 2:05:00.



The following example query specifies a 5-minute interval. `CpuUtilization[5m].max()`  
For supported values, see [Monitoring Query Language \(MQL\) Reference](#).

See also *resolution*.

### MESSAGE

The content that the Alarms feature of the Monitoring service publishes to topics in the *alarm's* configured *notification destinations*. A message is sent when the *alarm* transitions to another state, such as from "OK" to "FIRING." For more information about messages, see [How Monitoring Works](#).

### METADATA

A reference provided in a *metric definition*. Example: unit (bytes), provided in the definition of the `oci_computeagent metricDiskBytesRead`. Use metadata to determine additional information about a given metric. For metric definitions, see [Supported Services](#).

### METRIC

A measurement related to health, capacity, or performance of a given resource. Example: The `oci_computeagent metricCpuUtilization`, which measures usage of a Compute instance. For metric definitions, see [Supported Services](#).



#### Note

Metric resources do not have OCIDs.

### METRIC DEFINITION

A set of references, qualifiers, and other information provided by a *metric namespace* for a given *metric*. For example, the `oci_computeagent metricDiskBytesRead` is defined by *dimensions* (such as resource identifier) and *metadata* (specifying bytes for unit) as well as identification of its *metric namespace* (`oci_computeagent`). Each posted set of *data points* carries this information. Use the [ListMetricData](#) API operation to get metric definitions. For metric definitions, see [Supported Services](#).

### METRIC NAMESPACE

Indicator of the resource, service, or application that emits the *metric*. Provided in the *metric definition*. For example, the `CpuUtilization` *metric definition* emitted by the

OracleCloudAgent software on Compute instances lists the *metric namespace* `oci_computeagent` as the source of the `CpuUtilization` *metric*. For metric definitions, see [Supported Services](#).

### **METRIC STREAM**

An individual set of *aggregated data* for a *metric*. A stream can be either specific to a single resource or aggregated across all resources in the compartment. Within a [metric chart](#) in the Console, each metric stream is represented as a line. By default, metric streams are resource-specific, so the chart displays a line for each resource. If you choose to [aggregate all metric streams](#), then the chart displays one line for all resources.

### **NOTIFICATION DESTINATION**

Protocol and other details for sending *messages* when the *alarm* transitions to another state, such as from "OK" to "FIRING." The details and setup may vary by destination service. For the Notifications service, each destination includes a topic and subscription protocol (such as PagerDuty). For more information about messages, topics, and subscriptions, see [Notifications Overview](#).

### **ORACLECLOUDAGENT SOFTWARE**

Software that allows a Compute instance to post raw *data points* to the Monitoring service. Automatically installed with the latest versions of supported images. See [Enabling Monitoring for Compute Instances](#).

### **QUERY**

The Monitoring Query Language (MQL) expression to evaluate for returning *aggregated data*. The query must specify a *metric*, *statistic*, and *interval*. For more information, see [Building Metric Queries](#).

### **RESOLUTION**

The period between time windows, or the regularity at which time windows shift. For example, use a resolution of `1m` to retrieve aggregations every minute.

**Note**

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

**Maximum time range returned for a query**

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

<b>Interval</b>	<b>Default resolution (metric queries)</b>	<b>Maximum time range returned</b>
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days

**See examples of returned data**

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No



resolution or end time is specified, so the resolution defaults to the interval value of `1m`, and the end time defaults to the current time (`2019-01-08T10:00:00.789Z`). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`).

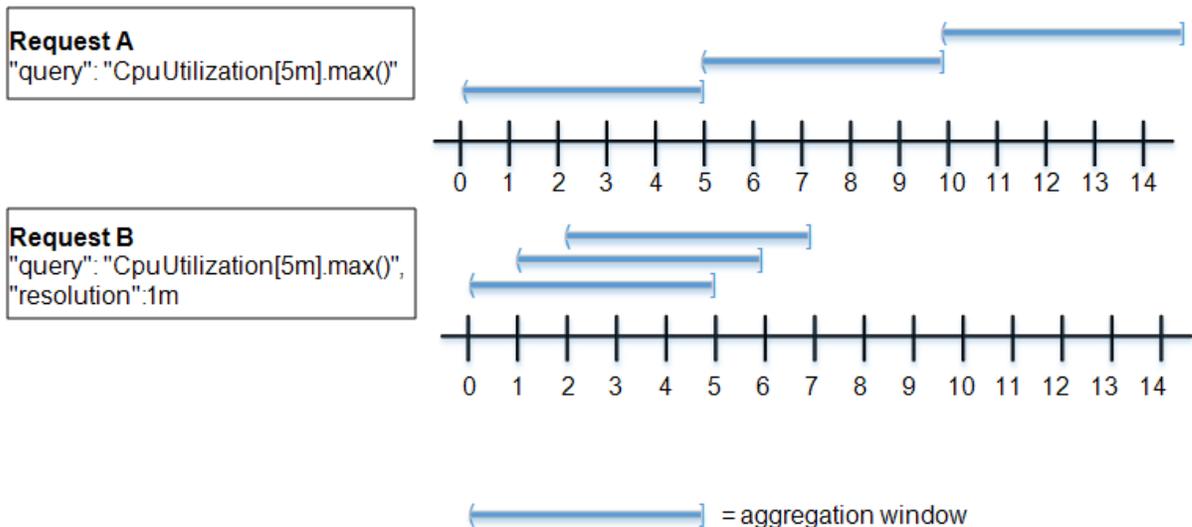
Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (`2019-01-06T10:00:00.789Z`), the earliest data point possible within this seven-day period would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

As shown in the following illustration, *resolution* controls the start time of each aggregation window relative to the previous window while *interval* controls the length of the windows. Both requests apply the statistic `max` to the data within each five-minute

window (from the interval), resulting in a single aggregated *data point* representing the highest `CPUUtilization` counter for that window. Only the resolution value differs. This resolution changes the regularity at which the aggregation windows shift, or the start times of successive aggregation windows. Request A does not specify a resolution and thus uses the default value equal to the interval (5 minutes). This request's five-minute aggregation windows are thus taken from the sets of data points emitted from 0:n to 5:00, 5:n to 10:00, and so forth. Request B specifies a 1-minute resolution, so its five-minute aggregation windows are taken from the set of data points emitted every minute from 0:n to 5:00, 1:n to 6:00, and so forth.



**RESOURCE GROUP**

A custom string provided with a custom metric that can be used as a filter or to aggregate results. The resource group must exist in the definition of the posted metric. Only one resource group can be applied per metric.

## CHAPTER 22 Monitoring

---

### STATISTIC

The aggregation function applied to the given set of raw *data points*. For supported statistics, see [Monitoring Query Language \(MQL\) Reference](#).

### SUPPRESSION

A configuration to avoid publishing *messages* during the specified time range. Useful for suspending alarm notifications during system maintenance. Each suppression applies to a single alarm. In the Console, you can apply one definition of a suppression to multiple alarms. The result is an individual suppression for each alarm. For instructions on suppressing alarms, see [To suppress alarms](#).

### TRIGGER RULE

The condition that must be met for the alarm to be in the firing state. A trigger rule can be based on a threshold or absence of a metric.

## Availability

Monitoring is currently available in the following regions:

Region Name	Region Location	Region Key
India West (Mumbai)	Asia-Pacific: Mumbai, India	BOM
South Korea Central (Seoul)	Asia-Pacific: Seoul, South Korea	ICN
Australia East (Sydney)	Asia-Pacific: Sydney, Australia	SYD
Japan East (Tokyo)	Asia-Pacific: Tokyo, Japan	NRT
Canada Southeast (Toronto)	Canada: Toronto	YYZ
Germany Central (Frankfurt)	Europe: Frankfurt, Germany	FRA
Switzerland North (Zurich)	Europe: Zurich, Switzerland	ZRH

Region Name	Region Location	Region Key
Brazil East (Sao Paulo)	South America: Sao Paulo	GRU
UK South (London)	United Kingdom: London	LHR
US East (Ashburn)	United States: Ashburn, VA	IAD
US West (Phoenix)	United States: Phoenix, AZ	PHX

### Supported Services

The following services have resources or components that can emit metrics to Monitoring:

- Block Storage - see [Block Volume Metrics](#)
- Compute - see these topics:
  - [Compute Instance Metrics](#)
  - [Infrastructure Health Metrics](#)
- Database - see [Database Metrics](#)
- Events - see [Events Metrics](#)
- Health Checks - see [Health Checks Metrics](#)
- Load Balancing - see [Load Balancing Metrics](#)
- Key Management - see [Key Management Metrics](#)
- Networking - see these topics:
  - [VNIC Metrics](#)
  - [FastConnect Metrics](#)
  - [VPN Connect Metrics](#)
- Notifications - see [Notifications Metrics](#)
- Object Storage - see [Object Storage Metrics](#)
- Oracle Functions - see [Function Metrics](#)

- Streaming - see [Streaming Metrics](#)
- WAF - see [WAF Metrics](#)

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).



#### Note

Metric resources do not have OCIDs.

### Ways to Access Monitoring

You can access the Monitoring service using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

**Console:** To access Monitoring using the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring**.

**API:** To access Monitoring through APIs, use [Monitoring API](#) for metrics and alarms and [Notifications API](#) for notifications (used with alarms).

### Moving Alarms to a Different Compartment

You can [move alarms](#) from one compartment to another. When you move an alarm to a new compartment, its associated metrics remain where they are. After you move the alarm to the new compartment, inherent policies apply immediately and affect access to the alarm through

the Console. For more information on moving resources to other compartments, see [Moving Resources to a Different Compartment](#).



### Important

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Monitoring resources, see [Details for Monitoring](#).

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Administrators: For common policies that give groups access to metrics, see [Let users view metric definitions in a compartment](#) and [Restrict user access to a specific metric namespace](#). For a common alarms policy, see [Let users view alarms](#). To authorize resources, such as instances, to make API calls, add the resources to a [dynamic group](#). Use the dynamic group's matching rules to add the resources, and then create a policy that allows that dynamic group

access to metrics. See [Let instances make API calls to access monitoring metrics in the tenancy](#).

### Limits on Monitoring

See [Monitoring Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

Other limits include the following.

#### Storage Limits

Item	Time range stored
Metric definitions	14 days
Alarm history entries	90 days

#### Returned Data Limits (Metrics)

When you [query metrics](#) and [view metric charts](#), the returned data is subject to certain limits. Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

#### Troubleshooting Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

## Viewing Default Metric Charts

This topic describes how to view metric charts for selected resources or a single resource and create alarms based on queries used for charts. Charts are available using the Console.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Prerequisites

- IAM policies: Viewing metric charts is part of monitoring. To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- Metrics exist in Monitoring: The resources that you want to monitor must emit metrics to the Monitoring service.
- Compute instances: To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).

### Working with Default Metric Charts

For background information on metrics in Oracle Cloud Infrastructure, see [Metrics Feature Overview](#). For default metrics by service, see [Supported Services](#).

Default metric charts use predefined service queries. You can select resources of interest and update the interval, statistic, and time range.



#### Note

Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power). Units correspond to the selected metric and do not change by statistic.

### Using the Console

#### To view default metric charts for all resources

Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.

The **Service Metrics** page displays the default charts for all resources in the first accessible **Compartment** and **Metric Namespace**. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

#### Don't see all expected resources or metrics?

- Try a [different time range](#).
- Make sure the correct **Compartment** is selected.

On the **Service Metrics** page, metric namespaces are shown only when associated resources exist in the selected compartment. For example, the `oci_autonomous_database` namespace is shown only when Autonomous Databases exist in the selected compartment.

- Confirm that the missing resources are emitting metrics. See [Enabling Monitoring for Compute Instances](#).
- Review limits information. Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

### To investigate missing resources or metrics

- Try a [different time range](#).
- Make sure the correct **Compartment** is selected.

On the **Service Metrics** page, metric namespaces are shown only when associated resources exist in the selected compartment. For example, the `oci_autonomous_database` namespace is shown only when Autonomous Databases exist in the selected compartment.

- Confirm that the missing resources are emitting metrics. See [Enabling Monitoring for Compute Instances](#).
- Review limits information. Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

### To filter results

Filter results to limit the data plotted on the metric chart. For example, filter results to a resource or region of interest.

Filtering of default metric charts is done through selected dimensions; available dimensions vary by metric.

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. To the right of **Dimensions**, click **Add**.
3. In the **Edit dimensions** dialog box, select a **Dimension Name** and **Dimension Value**.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
4. Click **Done**.  
The default charts show the filtered results of your query.

### To select different resources

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. To select resources on a different compartment, select the **Compartment**. The default charts update to show results for the selected compartment.
3. To select a specific resource within the selected compartment, filter results by a resource-specific dimension, such as **resourceDisplayName**:
  - a. To the right of **Dimensions**, click **Add**.
  - b. For **Dimension Name**, select **resourceDisplayName** or other resource-specific dimension.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- c. For **Dimension Value**, select the value corresponding to the resource you want.
- d. Click **Done**.

The default charts update to show filtered results.

### To aggregate data from all metric streams

Aggregate all metric streams to view the average. For example, aggregate all metric streams

for CPU Utilization to view the average data across all resources. By default, a chart represents each metric stream with a line, which results in multiple lines per chart. When you aggregate metric streams, a chart represents all metric streams with a single line, which results in just one line per chart.

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. Select **Aggregate Metric Streams**.

### To change the time range

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

### Maximum time range returned for a query

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

Interval	Default resolution (metric queries)	Maximum time range returned
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days

### See examples of returned data

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on

January 8th. No resolution or end time is specified, so the resolution defaults to the interval value of 1m, and the end time defaults to the current time (2019-01-08T10:00:00.789Z). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (2019-01-06T10:00:00.789Z), the earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. To select a period of time, such as **Last hour**, click **Start Time** or **End Time**.
3. To enter a time value, click in **Start Time** or **End Time** and then type a value.

### To change a chart interval or statistic

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

### Maximum time range returned for a query

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

Interval	Default resolution (metric queries)	Maximum time range returned
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days

### See examples of returned data

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No resolution or end time is specified, so the resolution defaults to the interval value of 1m, and the end time defaults to the current time (2019-01-08T10:00:00.789Z). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (2019-01-06T10:00:00.789Z), the earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. At the top of the chart you want, select an **Interval** or **Statistic**.  
For supported values, see [Monitoring Query Language \(MQL\) Reference](#).

### To go back to the default charts

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. On the upper right, click **Reset charts**.

### To view chart details

Chart details include the query as a Monitoring Query Language (MQL) expression and the names and OCIDs of represented resources.

1. View the default metric charts: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. Click the chart you want.
3. To view a list of resources represented in the chart, click the arrow to the left of the query displayed under the chart.

You can copy the OCID for a resource by clicking **Copy** to the right of the resource OCID.

### To share a chart



#### Note

The person you share the chart with must have the required IAM policies for access to metrics.

On the upper right of the chart you want, go to **Options**, and then click **Copy Chart URL**.

### To view a query in Metrics Explorer

On the upper right of the chart you want, go to **Options**, and then click **View Query in**

### Metrics Explorer.

#### To copy a query (MQL expression)

On the upper right of the chart you want, go to **Options**, and then click **Copy Query (MQL)**.

#### To view default metric charts for a single resource

On the page for the resource of interest, under **Resources**, click **Metrics**.

For example, to view metric data for a Compute instance:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance you're interested in.
3. On the instance detail page, under **Resources**, click **Metrics**.

A chart is shown for each metric. For a list of metrics related to Compute instances, see [Compute Instance Metrics](#).

The Console displays the last hour of metric data for the selected resource. A chart is shown for each metric emitted by the selected resource.

For a list of metrics emitted by your resource, see [Supported Services](#).

#### To create an alarm from a chart query

Follow the instructions for the page on which the query appears: Service Metrics or Metrics Explorer.

### Service Metrics page

#### To create an alarm from a chart query (Service Metrics)

1. View the **Service Metrics** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. At the top of the chart you're interested in, go to Options, and then select **Create an Alarm on this Query**.
3. On the **Create Alarm** page, under **Define alarm**, add the trigger, and fill in or update other alarm settings as needed:

### Alarm settings

#### Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

#### Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.

- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified. Example: **CpuUtilization**

- **Interval:** The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT**- The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.

- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the

storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

4. Under **Notifications**, select or create at least one notification destination:

### Notifications settings

#### . **Destinations:**

- **Destination Service:** The provider of the destination to use for notifications.

Available options:

- **[Notifications Service](#)**.
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.
- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.
- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com or
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

### Slack subscription



#### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The *<webhook-token>* portion of the URL contains two slashes (/). Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
- **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
  - **Notification Interval:** The period of time to wait before resending the notification.
  - **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
5. If you want to disable the new alarm, clear **Enable This Alarm?**
6. Click **Save alarm**.  
The new alarm is listed on the **Alarm Definitions** page.  
For more information about alarms, see [Alarms Feature Overview](#).

### Metrics Explorer page

#### To create an alarm from a chart query (Metrics Explorer)

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Click **Create Alarm**.
4. On the **Create Alarm** page, under **Define alarm**, add the trigger, and fill in or update other alarm settings as needed:

#### Alarm settings

##### Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

#### Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified. Example: **CpuUtilization**

- **Interval:** The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT**- The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.

- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the

storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

5. Under **Notifications**, select or create at least one notification destination:

### Notifications settings

#### . **Destinations:**

- **Destination Service:** The provider of the destination to use for notifications.

Available options:

- **Notifications Service.**
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.
- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.
- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com or
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

### Slack subscription



#### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The *<webhook-token>* portion of the URL contains two slashes (/). Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
- **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
  - **Notification Interval:** The period of time to wait before resending the notification.
  - **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
6. If you want to disable the new alarm, clear **Enable This Alarm?**
7. Click **Save alarm**.  
The new alarm is listed on the **Alarm Definitions** page.  
For more information about alarms, see [Alarms Feature Overview](#).

### resource page

Examples of resource pages are Compute instance detail pages and Block Volume volume detail pages. Alarms are available from these pages for resources that emit metrics.

### To create an alarm from a chart query (resource page)

1. To view charts: On the resource page, under **Resources**, click **Metrics**.
2. At the top of the chart you're interested in, go to **Options**, and then select **Create an Alarm on this Query**.
3. On the **Create Alarm** page, under **Define alarm**, add the trigger, and fill in or update other alarm settings as needed:

### Alarm settings

#### Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name**: User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

#### Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.

Protocol	Rendering of the title
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.

- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Metric Name**: The name of the metric. Only one metric can be specified. Example: **CpuUtilization**
- **Interval**: The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic**: The aggregation function.

### Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the

storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

4. Under **Notifications**, select or create at least one notification destination:

### Notifications settings

#### . **Destinations:**

- **Destination Service:** The provider of the destination to use for notifications.

Available options:

- **[Notifications Service](#)**.
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.
- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.
- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com or
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

### Slack subscription



#### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The *<webhook-token>* portion of the URL contains two slashes (/). Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
- **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
  - **Notification Interval:** The period of time to wait before resending the notification.
  - **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
5. If you want to disable the new alarm, clear **Enable This Alarm?**
  6. Click **Save alarm**.

The new alarm is listed on the **Alarm Definitions** page.  
For more information about alarms, see [Alarms Feature Overview](#).

## Building Metric Queries

This topic describes how to query metrics for resources of interest, create alarms from a given query, and share Console charts.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Prerequisites

- IAM policies: Querying metrics is part of monitoring. To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- Metrics exist in Monitoring: The resources that you want to monitor must emit metrics to the Monitoring service.
- Compute instances: To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).

### Working with Metric Queries

This section shows MQL syntax of metric and alarm queries.

Use metric queries to actively and passively monitor your cloud resources. Actively monitor with metric queries that you generate spontaneously, on demand. In the Console, update a chart to show data from multiple queries. Store queries you want to reuse. Passively monitor with alarms that add a condition, or trigger rule, to a metric query.

Metric query syntax (boldface elements are required):

```
metric[interval] {dimensionname=dimensionvalue}.groupingfunction.statistic
```

Threshold Alarm query syntax (boldface elements are required):

```
metric[interval] {dimensionname=dimensionvalue}.groupingfunction.statistic
alarmoperator alarmvalue
```

For supported parameter values, see [Monitoring Query Language \(MQL\) Reference](#).

### Example queries

#### Simple metric query

Maximum CPU Utilization at a one-minute interval.

Number of lines displayed in the metric chart (Console): 1 per resource.

```
CpuUtilization[1m].max()
```

#### Filtered metric query

Maximum CPU Utilization at a one-minute interval, filtered to a single resource.

Number of lines displayed in the metric chart (Console): 1.

```
CpuUtilization[1m]{resourceId="ocid1.instance.oc1.phx.exampleuniqueID"}.max()
```

#### Aggregated metric query

All IopsRead at a one-minute interval, filtered to a compartment, aggregated for the maximum.

Number of lines displayed in the metric chart (Console): 1.

```
IopsRead[1m]
{compartmentId="ocid1.compartment.oc1.phx..exampleuniqueID"}.grouping().max()
```

### Group-aggregated metric query

Aggregated average of CPU Utilization by availability domain and pool ID, filtered to Compute instances that use the specified shape.

Number of lines displayed in the metric chart (Console): 1 per pool and 1 per availability domain.

```
CPUUtilization[1m]{shape="VM.Standard2.8"}.groupBy
(availabilityDomain,poolId).mean()
```

### Alarm query (threshold)

Triggered when the 90th percentile of CPU Utilization, aggregated by pool ID, and filtered to the specified availability domain, exceeds 85.

Number of lines displayed in the metric chart (Console): 1 per pool.

```
CpuUtilization[1m]{availabilityDomain="VeBZ:PHX-AD-1"}.groupBy
(poolId).percentile(0.9) > 85
```

### Grouped count of resources (alarm query or metric query)



#### Note

Nested alarm queries are not currently supported in the Console. Use the API to create alarms with nested queries.

An example of a nested query is a grouped count of hosts with up time greater than zero, where the alarm query to identify these hosts is defined within parentheses:

```
(metric[1h].groupBy(host).min() > 0).grouping().count()
```

You can use such a query to either define an alarm or query a metric.

For background information on metrics in Oracle Cloud Infrastructure, see [Metrics Feature Overview](#).

## Using the Console

### To create a query

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.

The **Metrics Explorer** page displays an empty chart with fields to build a query.

2. Fill in the fields for a new query.
  - **Compartment:** The compartment containing the resources that you want to monitor. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified. Metric selections depend on the selected compartment and metric namespace.  
Example: **CpuUtilization**
  - **Interval:** The aggregation window.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour

**Note**

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

**Maximum time range returned for a query**

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

<b>Interval</b>	<b>Default resolution (metric queries)</b>	<b>Maximum time range returned</b>
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days



### See examples of returned data

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No resolution or end time is specified, so the resolution defaults to the interval value of 1m, and the end time defaults to the current time (2019-01-08T10:00:00.789Z). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (2019-01-06T10:00:00.789Z), the earliest data point possible within this seven-day period would be 10:00 on January 1st (2019-01-01T10:00:00.789Z). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Aggregate Metric Streams:** Aggregates all results to plot a single aggregated average for all metric streams. This average is plotted as a single line on the metric chart. This operation is helpful when you want to plot a metric as one line for all resources.

3. Click **Update Chart**.

The chart shows the results of your new query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power). Units correspond to the selected metric and do not change by statistic.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

4. To customize the y-axis label or range, type the label you want into **Y-Axis Label** or type the minimum and maximum values you want into **Y-Axis Min Value** and **Y-Axis Max Value**.

Only numeric characters are allowed for custom ranges. Custom labels and ranges are not persisted in shared queries (MQL).

5. To view the query as a Monitoring Query Language (MQL) expression, click **Advanced Mode**.

**Advanced Mode** is located on the right, under the chart.

Use Advanced Mode to edit your query using MQL syntax to [aggregate results by group](#). The MQL syntax also supports additional parameter values. For more information about query parameters in Basic Mode and Advanced Mode, see [Monitoring Query Language \(MQL\) Reference](#).

6. To create another query, click **Add Query** below the chart.

## To change the time range

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

## Maximum time range returned for a query

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

Interval	Default resolution (metric queries)	Maximum time range returned
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days

## See examples of returned data

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No resolution or end time is specified, so the resolution defaults to the interval value of `1m`, and the end time defaults to the current time (`2019-01-08T10:00:00.789Z`). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (`2019-01-06T10:00:00.789Z`), the earliest data point possible within this seven-day period

would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. To select a period of time, such as **Last hour**, click **Start Time** or **End Time**.
3. To enter a time value, click in **Start Time** or **End Time** and then type a value.

### To filter results

Filter results to limit the data plotted on the metric chart. For example, filter results to a resource or pool of interest.

Filtering is done through selected dimensions; available dimensions vary by metric. You can also filter by [resource groups](#) when provided with the metric.

### To filter by dimensions

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Under **Metric dimensions**, select a **Dimension Name** and **Dimension Value**.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
4. To add a dimension name-value pair, click **+ Additional dimension**.
  5. Click **Update Chart**.  
The chart shows the filtered results of your query.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

6. To view the query as a Monitoring Query Language (MQL) expression, click **Advanced Mode**.

The dimension name-value fragment appears after the metric-interval fragment.

In the following example query, the dimension name-value fragment is

`{resourceId="ocidl.instance.oc1.phx.exampleuniqueID"}`, which filters results by the specified resource identifier.

```
CpuUtilization[1m]
```

```
{resourceId="ocidl.instance.oc1.phx.exampleuniqueID"}.max()
```

The MQL syntax supports more parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

### To filter by a resource group

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Select the **Resource Group** you want to use as a filter.
4. Click **Update Chart**.

The chart shows the filtered results of your query.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

### To aggregate all results

Aggregate all results to plot a single aggregated average for all metric streams. This average is plotted as a single line on the metric chart. This operation is helpful when you want to plot a metric as one line for all resources.

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Click **Aggregate Metric Streams**.
4. Click **Update Chart**.

The chart shows the results of your query.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

5. To view the query as a Monitoring Query Language (MQL) expression, click **Advanced Mode**.

The `grouping()` function appears before the statistic. For example, the following query returns the maximum (`max()`) IopsRead metric data at a one-minute interval, filtered to a compartment, with all results aggregated.

```
IopsRead[1m]{compartmentID = "<compartment_OCID>".grouping().max()
```

Edit your query in MQL to [aggregate results by group](#). The MQL syntax also supports more parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

### To aggregate results by group



#### Note

Aggregating query results by group requires the `groupBy()` function, which is available in **Advanced Mode** only.

Aggregate query results by group to plot group-specific aggregated averages. Each group's average is plotted as a single line on the metric chart. This operation is helpful when you want to identify trends by group rather than individual resource.

For example, the following query returns the average (`mean()`) CPU Utilization metric data at a one-minute interval. Results are filtered to the specified shape and grouped by availability domain. If you have Compute instances of this shape sending metrics across three availability domains, then three lines are plotted on the metric chart.

```
CPUtilization[1m]{shape="VM.Standard1.1"}.groupBy(availabilityDomain).mean()
```

You can also aggregate by [resource groups](#) when provided with the metric (`groupBy(resourceGroup)`).

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Select **Advanced Mode** below the chart on the right.

4. In the **Query Code Editor** box, insert the `groupBy({dimension})` function between the metric-interval fragment and the statistic, where `{dimension}` is the name of a dimension provided in the definition of the indicated metric.

For example, insert the following fragment to group by availability domain, assuming that the dimension is available for the selected metric.

```
groupBy(availabilityDomain)
```

The MQL syntax supports more parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

5. Click **Update Chart**.

The chart is updated to show a single line for each grouped result.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

### To edit a query using MQL syntax

Edit your query using MQL syntax to [aggregate results by group](#) or for more parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit Query** icon.

3. Click **Advanced Mode**.

The query is displayed as a Monitoring Query Language (MQL) expression.

4. In the **Query Code Editor** box, edit the query as needed.

5. Click **Update Chart**.

The chart is updated.

### Troubleshooting Errors and Query Limits

If you see an error that the query has exceeded the maximum number of metric streams, then update the query to evaluate a number of metric streams that is within the limit. For example, you can reduce the metric streams by specifying dimensions. You can continue to evaluate all metric streams that were in the original query by spreading the metric streams across multiple queries (or alarms).

Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

### To create an alarm from a query

Create an alarm to passively monitor for a condition in results from metric queries. Creating an alarm from a query involves adding a trigger rule to the query and setting up notifications.

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit Query** icon.
3. Click **Create Alarm**.
4. On the **Create Alarm** page, under **Define alarm**, add the trigger, and fill in or update other alarm settings as needed:

## Alarm settings

### Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

### Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are

not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified.  
Example: **CpuUtilization**

- **Interval:** The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.

- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

5. Under **Notifications**, select or create at least one notification destination:

### Notifications settings

#### . **Destinations:**

- **Destination Service:** The provider of the destination to use for notifications.

Available options:

- [Notifications Service](#).

- **Compartment:** The compartment storing the topic to be used for

notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.

- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.
- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com or
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

## Slack subscription



### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The `<webhook-token>` portion of the URL contains two slashes (/).

Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
  - **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
  - **Notification Interval:** The period of time to wait before resending the notification.
  - **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
6. If you want to disable the new alarm, clear **Enable This Alarm?**
  7. Click **Save alarm**.

The new alarm is listed on the **Alarm Definitions** page.  
For more information about alarms, see [Alarms Feature Overview](#).

### To hide a query from the chart

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. Click the **Toggle query on chart** icon for the query that you want to hide.

### To share a query

1. View the **Metrics Explorer** page: Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.
2. If necessary, open the query for editing: Click the **Edit query** icon.
3. Click **Advanced Mode**.
4. In the **Query Code Editor** box, copy the query.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use this API operation to find metric names and dimensions (view metric definitions):

[ListMetrics](#)

Use this API operation to query metrics by name (and optionally filter by dimension):

[SummarizeMetricsData](#)

### Publishing Custom Metrics

This topic describes how to publish your own custom metrics to the Monitoring service.

You can publish your own metrics to Monitoring using the API. You can view charts of your published metrics using the Console , query metrics using the API, and set up alarms using the Console or API.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Prerequisites

IAM policies: To publish custom metrics, you must be given the required type of access in a policy written by an administrator. This requirement applies whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a message that you don't have permission or are unauthorized, check with your administrator. You may not have the

required type of access in the current compartment. Administrators: For a related common policy, see [Let users publish custom metrics](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



#### Note

When defining your custom metrics, note the following:

- Ensure that your custom metrics do not exceed limits. For example, note the valid range of dimensions and maximum number of streams for custom metrics. See [PostMetricData](#).
- Define your metrics with aggregation in mind. While custom metrics can be posted as frequently as every second (minimum frequency of one second), the minimum aggregation interval is one minute.
- Define your metrics with return limits in mind. Limits information for returned data includes the 100,000 data point maximum and [time range maximums \(determined by resolution, which relates to interval\)](#). See [MetricData Reference](#).

Use this API operation to publish custom metrics:

[PostMetricData](#)



### Note

Oracle recommends the following:

- Send batched requests to maximize metric streams per request. A batched request contains multiple metrics or metric namespaces. Note limits. See [PostMetricData](#).
- Publish metrics only when relevant contexts require monitoring; that is, when data points need to be collected. If you want to publish metrics during inactive periods when no observations exist, then you can manually create "0" values for publishing.

### Example of a batched request

This example shows a single request containing data points for metrics across two metric namespaces.

```
[
 {
 "namespace": "myFirstNamespace",
 "compartmentId": "ocid1.compartment.oc1..exampleuniqueID",
 "resourceGroup": "myFirstResourceGroup",
 "name": "successRate",
 "dimensions": {
 "resourceId": "ocid1.exampleresource.region1.phx.exampleuniqueID",
 "appName": "myAppA"
 },
 "metadata": {
 "unit": "percent",
 "displayName": "MyAppA Success Rate"
 },
 "datapoints": [
```

## CHAPTER 22 Monitoring

```
{
 "timestamp": "2019-03-10T22:19:20Z",
 "value": 83.0
},
{
 "timestamp": "2019-03-10T22:19:40Z",
 "value": 90.1
}
],
{
 "namespace": "myFirstNamespace",
 "compartmentId": "ocid1.compartment.oc1..exampleuniqueID",
 "resourceGroup": "mySecondResourceGroup",
 "name": "successRate",
 "dimensions": {
 "resourceId": "ocid1.exampleresource.region1.phx.differentuniqId",
 "appName": "myAppA"
 },
 "metadata": {
 "unit": "percent",
 "displayName": "MyAppA Success Rate"
 },
 "datapoints": [
 {
 "timestamp": "2019-03-10T22:19:10Z",
 "value": 100.0
 },
 {
 "timestamp": "2019-03-10T22:19:30Z",
 "value": 100.0
 }
]
},
{
 "namespace": "mySecondNamespace",
 "compartmentId": "ocid1.compartment.oc1..exampleuniqueID",
 "name": "deliveryRate",
 "dimensions": {
 "resourceId": "ocid1.exampleresource.region1.phx.exampleuniqueID",
 "appName": "myAppB"
 }
}
```

```
 },
 "metadata": {
 "unit": "bytes",
 "displayName": "MyAppB Delivery Rate"
 },
 "datapoints": [
 {
 "timestamp": "2019-03-10T22:19:00Z",
 "value": 87.0,
 "count": 60
 },
 {
 "timestamp": "2019-03-10T22:19:00Z",
 "value": 96.0,
 "count": 30
 }
]
 }
}
```

You can access your published custom metrics the same way you access any other metrics stored by the Monitoring service. [View charts from queries](#) using the Console, [query metrics](#) using the CLI or API, and [set up alarms](#) using the Console, CLI, or API.

## Managing Alarms

This topic describes how to create, update, suppress, and delete alarms, as well as how to retrieve alarm history. See also [Best Practices for your Alarms](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Prerequisites

- **IAM policies:** Managing alarms is part of monitoring. To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).
- **Metrics exist in Monitoring:** The resources that you want to monitor must emit metrics to the Monitoring service.
- **Compute instances:** To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### Using the Console

#### To see all firing alarms

Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Status**.

You can [suppress alarms during a given time range](#). You can also [disable](#) and [delete alarms](#).

### To create an alarm

This section includes steps to create example alarms as well as any kind of alarm.

### To create an example threshold alarm

This procedure walks through creation of an [example threshold alarm](#) to detect Compute instances operating at non-optimal thresholds. A *threshold alarm* is an alarm that checks for metric values outside a given range or value. The procedure uses options as displayed in Basic Mode.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click **Create alarm**.
3. On the **Create Alarm** page, under **Define alarm**, fill in or update the alarm settings:
  - **Alarm Name:** Non-Optimal Alarm
  - **Alarm Severity:** Warning
  - **Alarm Body:** Non-optimal utilization detected. An application or process may be consuming more CPU than usual.
  - **Metric description:**
    - **Compartment:** (select your compartment )
    - **Metric Namespace:** `oci_computeagent`
    - **Metric Name:** `CpuUtilization`
    - **Interval:** `1m`
    - **Statistic:** `Count`
  - **Trigger rule:**
    - **Operator:** `between`
    - **Value:** `60`

- **Value:** 80
  - **Trigger Delay Minutes:** 10
4. Set up an email notification under **Notifications, Destinations:**
    - **Destination Service:** [Notifications Service](#)
    - **Compartment:** (select your compartment )
    - **Topic:** Click **Create a topic**
      - **Topic Name:** Operations Team
      - **Topic Description:** Resource Monitoring Channel
      - **Subscription Protocol:** **Email**
      - **Email Addresses:** (type an email address for the operations team here)
  5. Repeat notifications every day:
    - **Repeat Notification?:** (select this option)
    - **Notification Interval:** 24 hours
  6. Click **Save alarm.**

### To create an example absence alarm

This procedure walks through creation of an [example absence alarm](#) to detect resources that may be down or unreachable. An *absence alarm* is an alarm that checks for absent metrics (using the absent operator). The procedure uses options as displayed in Basic Mode.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click **Create alarm**.
3. On the **Create Alarm** page, under **Define alarm**, fill in or update the alarm settings:
  - **Alarm Name:** Up/Down Resource Alarm
  - **Alarm Severity:** Critical

- **Alarm Body:** Resource may be down. Please investigate. Move workloads to another available resource.
  - **Metric description:**
    - **Compartment:** (select your compartment )
    - **Metric Namespace:** `oci_computeagent`
    - **Metric Name:** `CpuUtilization`
    - **Interval:** `1m`
    - **Statistic:** `Count`
  - **Trigger rule:**
    - **Operator:** `absent`
    - **Trigger Delay Minutes:** `5`
4. Set up an email notification under **Notifications, Destinations:**
- **Destination Service:** [Notifications Service](#)
  - **Compartment:** (select your compartment )
  - **Topic:** Click **Create a topic**
    - **Topic Name:** Operations Team
    - **Topic Description:** Resource Up/Down Channel
    - **Subscription Protocol:** `Email`
    - **Email Addresses:** (type an email address for the operations team here)



### Note

To add a notification (subscription) for another protocol, such as PagerDuty, create a copy of this alarm and choose the corresponding protocol. For more information about subscription protocols, see [To create a subscription](#).

5. Repeat notifications every minute:
  - **Repeat Notification?:** (select this option)
  - **Notification Interval:** 1 minute
6. Click **Save alarm**.

### To create an alarm (any kind)

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click **Create alarm**.



### Note

You can also create an alarm from a predefined query on the **Service Metrics** page. Expand **Options** and click **Create an Alarm on this Query**. For more information about service metrics, see [Viewing Default Metric Charts](#).

3. On the **Create Alarm** page, under **Define alarm**, fill in or update the alarm settings:

**Note**

To toggle between Basic Mode and Advanced Mode, click **Switch to Advanced Mode** or **Switch to Basic Mode** (to the right of **Define Alarm**).

### Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

### Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition.

Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."

- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group (optional):** The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified. Example: **CpuUtilization**

- **Interval:** The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.

- **MIN** - The lowest value observed during the specified time period.
- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

4. Set up notifications: Under **Notifications**, fill in the fields.

### . Destinations:

- **Destination Service:** The provider of the destination to use for notifications.  
Available options:
  - [Notifications Service](#).
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.

- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.
- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com OR
```

`https://user@domain.com`. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see [RFC-2617](#). Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

## Slack subscription



### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The `<webhook-token>` portion of the URL contains two slashes (/).

Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
  - **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
  - **Notification Interval:** The period of time to wait before resending the notification.
  - **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
5. If you want to disable the new alarm, clear **Enable This Alarm?**
  6. Click **Save alarm**.

The new alarm is listed on the **Alarm Definitions** page.

### To disable or enable an alarm

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click the alarm that you want to disable or enable.
3. On the alarm detail page, select or clear **Alarm is Enabled**.



#### Note

You can also disable and enable alarms when [creating](#) or [editing an alarm](#).

### To move an alarm to a different compartment

Associated metrics remain in their current compartments. For more information, see [Moving Alarms to a Different Compartment](#).



#### Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Monitoring resources, see [Details for Monitoring](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. In the **List Scope** section, select a compartment.
3. Click the alarm that you want to move.
4. On the alarm detail page, click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

### To update an alarm

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click the alarm that you want to update.
3. Go to **Actions** on the right, and then click **Edit Alarm**.
4. On the **Edit Alarm** page, under **Define alarm**, update alarm settings as needed:

## Basic Mode (default)

By default, this page uses **Basic Mode**, which separates the metric from its dimensions and its trigger rule.

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

- **Metric description:** The metric to evaluate for the alarm condition.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.
  - **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
  - **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
  - **Metric Name:** The name of the metric. Only one metric can be specified.  
Example: **CpuUtilization**

- **Interval:** The aggregation window, or the frequency at which data points are aggregated.

### Interval values

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour



#### Note

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

- **Statistic:** The aggregation function.

### Statistic values

- **COUNT** - The number of observations received in the specified time period.
- **MAX** - The highest value observed during the specified time period.
- **MEAN** - The value of Sum divided by Count during the specified time period.
- **MIN** - The lowest value observed during the specified time period.

- **P50** - The value of the 50th percentile.
- **P90** - The value of the 90th percentile.
- **P95** - The value of the 95th percentile.
- **P99** - The value of the 99th percentile.
- **P99.5** - The value of the 99.5th percentile.
- **RATE** - The per-interval average rate of change.
- **SUM** - All values added together.

- **Metric dimensions:** Optional filters to narrow the metric data evaluated.

### Dimension fields

- **Dimension Name:** A qualifier specified in the metric definition. For example, the dimension `resourceId` is specified in the metric definition for `CpuUtilization`.



#### Note

Long lists of dimensions are trimmed.

- To view dimensions by name, type one or more characters in the box. A refreshed (trimmed) list shows matching dimension names.
- To retrieve all dimensions for a given metric, use the following API operation: [ListMetrics](#)

- **Dimension Value:** The value you want to use for the specified dimension. For example, the resource identifier for your instance of interest.
  - **+ Additional dimension:** Adds another name-value pair for a dimension.
- **Trigger rule:** The condition that must be satisfied for the alarm to be in the firing state. The condition can specify a threshold, such as 90% for CPU Utilization, or an absence.

- **Operator:** The operator used in the condition threshold.

### Operator values

- **greater than**
  - **greater than or equal to**
  - **equal to**
  - **less than**
  - **less than or equal to**
  - **between** (inclusive of specified values)
  - **outside** (inclusive of specified values)
  - **absent**
- **Value:** The value to use for the condition threshold.
  - **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

### Advanced Mode

Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression. Edit your query using MQL syntax to [aggregate results by group](#) or for additional parameter values. See [Monitoring Query Language \(MQL\) Reference](#).

- **Alarm Name:** User-friendly name for the new alarm. This name is sent as the title for notifications related to this alarm.

## Rendering of the title by protocol

Protocol	Rendering of the title
Email	Subject line of the email message.
HTTPS (Custom URL)	Not rendered.
PagerDuty	Title field of the published message.
Slack	Not rendered.

- **Alarm Severity:** The perceived type of response required when the alarm is in the firing state.
- **Alarm Body:** The human-readable content of the notification delivered. Oracle recommends providing guidance to operators for resolving the alarm condition. Consider adding links to standard runbook practices. Example: "High CPU usage alert. Follow runbook instructions for resolution."
- **Tags (optional):** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- **Metric description, dimensions, and trigger rule:** The metric to evaluate for the alarm condition, including dimensions and the trigger rule.
  - **Compartment:** The compartment containing the resources that emit the metrics evaluated by the alarm. The selected compartment is also the storage location of the alarm. By default, the first accessible compartment is selected.

- **Metric Namespace:** The service or application emitting metrics for the resources that you want to monitor.
- **Resource Group** (optional): The group that the metric belongs to. A [resource group](#) is a custom string provided with a custom metric. Not applicable to service metrics.
- **Query Code Editor** box: The alarm query as a Monitoring Query Language (MQL) expression.

Example alarm query:

```
CpuUtilization[1m]{availabilityDomain=AD1}.groupBy(poolId).percentile(0.9) > 85
```

For query syntax and examples, see [Working with Metric Queries](#).

- **Trigger Delay Minutes:** The number of minutes that the condition must be maintained before the alarm is in firing state.

The chart below the **Define alarm** section dynamically displays the last six hours of emitted metrics according to currently selected fields for the query. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

5. Under **Notifications**, update settings as needed:

### . **Destinations:**

- **Destination Service:** The provider of the destination to use for notifications.  
Available options:
  - [Notifications Service](#).
- **Compartment:** The compartment storing the topic to be used for notifications. Can be a different compartment from the alarm and metric. By default, the first accessible compartment is selected.
- **Topic:** The [topic](#) to use for notifications. Each topic supports a [subscription](#) protocol, such as PagerDuty.

- **Create a topic:** Sets up a [topic](#) and [subscription](#) protocol in the selected compartment, using the specified destination service.
  - **Topic Name:** User-friendly name for the new topic. Example: "Operations Team " for a topic used to notify operations staff of firing alarms.
  - **Topic Description:** Description of the new topic.
  - **Subscription Protocol:** Medium of communication to use for the new topic. Configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Subscription Protocol:** Select **Email**.
- **Subscription Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in

```
https://user:password@domain.com OR
```

```
https://user@domain.com. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see RFC-2617.
```

Query parameters are not allowed in URLs.

- **Subscription Protocol:** Select **HTTPS (Custom URL)**.
- **Subscription URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Subscription Protocol:** Select **PagerDuty**.
- **Subscription URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

### Slack subscription



#### Note

See the following [known issue](#) for up-to-date information about creating



### Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The *<webhook-token>* portion of the URL contains two slashes (/).

Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Subscription Protocol:** Select **Slack**.
  - **Subscription URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.
- **+ Additional destination service:** Adds another destination service and [topic](#) to use for notifications.



### Note

Each alarm is limited to one destination per supported destination service.

- **Repeat Notification?:** While the alarm is in the firing state, resends notifications at the specified interval.
- **Notification Interval:** The period of time to wait before resending the notification.

- **Suppress Notifications:** Sets up a suppression time window during which to suspend evaluations and notifications. Useful for avoiding alarm notifications during system maintenance periods.
    - **Suppression Description**
    - **Start Time**
    - **End Time**
6. Select or clear **Enable this alarm?**
  7. Click **Save alarm**.

The updated alarm settings are listed on the **Alarm Definitions** page.

### To update an alarm after moving a resource

This section shows how to update the metric compartment of an alarm after you move a resource that is emitting metrics monitored by the alarm. For example, if you move a block volume to another compartment, then the alarm must be updated if you want to continue monitoring metrics from the moved block volume.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. Click the alarm that you want to update.
3. Go to **Actions** on the right, and then click **Edit Alarm**.
4. Update the metric compartment: On the **Edit Alarm** page, under **Metric description** (or **Metric description, dimensions, and trigger rule** for Advanced mode), change the **Compartment** to the compartment where the resource has been moved.

The chart below the **Define alarm** section dynamically updates according to the selected compartment, displaying the last six hours of emitted metrics. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

If the chart is not showing the expected data, then the old compartment might be specified in the query (MQL), as in the following example:

```
IopsRead[1m]{compartmentId="ocidl.compartment.oc1.phx..oldcompartmentexampleuniqueID"}.grouping().max()
```

5. If the old compartment is specified in the query, then update the query to reference the new compartment:
  - a. Click **Advanced Mode** or **Switch to Advanced Mode** to view the alarm query as a Monitoring Query Language (MQL) expression.
  - b. In **Query Code Editor**, update the query to reference the new compartment.

### View example

Original query:

```
IopsRead[1m]
{compartmentId="ocidl.compartment.oc1.phx..oldcompartmentexampleuniqueID"}.grouping().max()
()
```

Updated query:

```
Read[1m]
{compartmentId="ocidl.compartment.oc1.phx..newcompartmentexampleuniqueID"}.grouping().max()
()
```

For more information about query syntax and more examples, see [Working with Metric Queries](#).

The chart below the **Define alarm** section dynamically updates according to the updated query, displaying the last six hours of emitted metrics. Very small or large values are indicated by International System of Units (SI units), such as M for mega (10 to the sixth power).

If the chart is not showing the expected data, then confirm that every compartment reference (**Compartment, Query Code Editor**) points to the new compartment.

6. Click **Save alarm**.  
The alarm now monitors metrics from the new compartment.

### To suppress alarms



#### Important

Only one suppression can be configured per alarm. Any existing suppression for the alarm is overwritten when you apply a new suppression.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. On the **Alarm Definitions** page, select the check boxes for the alarms you want to suppress.



#### Note

You can also suppress alarms from the Alarm Status page or when [creating](#) or [editing an alarm](#).

3. Go to **Actions** and select **Add Suppressions**.
4. In the **Suppress alarms** dialog box, select a **Start Time** and **End Time** and then optionally fill in a **Suppression Description**.
5. Click **Apply suppressions**.  
A suppression is created for each selected alarm. The updated alarm settings are listed on the **Alarm Definitions** page.

### To delete alarms

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.

2. On the **Alarm Definitions** page, select the check boxes for the alarms you want to delete.



### Note

You can also delete an alarm from its detail page.

3. Go to **Actions** and select **Delete Alarms**.  
The deleted alarms are removed from the compartment and are no longer displayed on the **Alarm definitions** page.

### To view alarm history

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Alarm Definitions**.
2. On the **Alarm Definitions** page, click the alarm that you want to view history for.  
The alarm detail page displays a chart showing data for the indicated time range and a list of timestamped transitions, such as Firing to OK.  
Alarm history is retained for 30 days.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage alarms:

- [ListAlarms](#)
- [GetAlarm](#)
- [CreateAlarm](#)
- [ChangeAlarmCompartment](#)

- [UpdateAlarm](#)
- [DeleteAlarm](#)
- [ListAlarmsStatus](#)
- [RemoveAlarmSuppression](#)
- [GetAlarmHistory](#)

## Best Practices for your Alarms

This topic describes best practices for working with your alarms.

### Create a Set of Alarms for Each Metric

For each metric emitted by your resources, [create alarms](#) that define the following resource behaviors:

- At risk. The resource is at risk of becoming inoperable, as indicated by metric values.
- Non-optimal. The resource is performing at non-optimal levels, as indicated by metric values.
- Resource is up or down. The resource is either not reachable or not operating.

The following examples use the CpuUtilization metric emitted by [the oci\\_computeagent metric namespace](#). This metric monitors the utilization of the Compute instance and the activity level of any services and applications running on the instance. CpuUtilization is a key performance metric for a cloud service because it indicates CPU usage for the Compute instance and it can be used to investigate performance issues. To learn more about CPU usage, see the following URL: [https://en.wikipedia.org/wiki/CPU\\_time](https://en.wikipedia.org/wiki/CPU_time).

#### **At-Risk Example**

A typical at-risk threshold for the CpuUtilization metric is any value greater than 80 percent. A Compute instance breaching this threshold is at risk of becoming inoperable. Often the cause of this behavior is one or more applications consuming a high percentage of the CPU.

In this example, you decide to notify the operations team immediately, setting the severity of the alarm as “Critical” because repair is required to bring the instances back to optimal operational levels. You configure alarm notifications to the responsible team by both PagerDuty and email, requesting an investigation and appropriate fixes before the instances go into an inoperable state. You set repeat notifications every minute. When someone responds to the alarm notifications, you temporarily stop notifications using the best practice of suppressing the alarm. Once metrics return to optimal values, you remove the suppression.

### **Non-Optimal Example**

A typical non-optimal threshold for the CpuUtilization metric is from 60 to 80 percent. When the metric values for a Compute instance are within this range, the instance is above the optimal operational range.

In this example, you decide to notify the appropriate individual or team that an application or process is consuming more CPU than usual. You configure a [threshold alarm](#) to notify the appropriate contacts, setting the severity of the alarm as “Warning,” as no immediate actions are required to investigate and reduce the CPU. You set notification to email only, directed to the appropriate developer or team, with repeat notifications every 24 hours to reduce email notification noise.

### **Resource is Up or Down Example**

A typical indicator of resource availability is a five-minute absence of the CpuUtilization metric. A Compute instance breaching this threshold is either not reachable or not operating. The resource may have stopped responding, or it might have become unavailable because of connectivity issues.

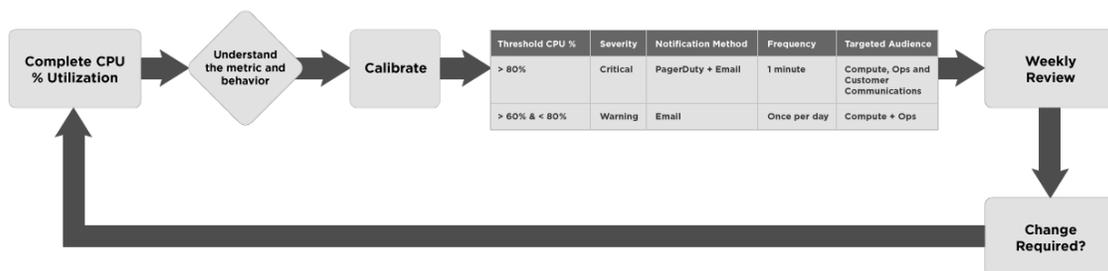
In this example, you decide to notify the operations team immediately, setting the severity of your [absence alarm](#) as “Critical” because repair is required to bring the instances online. You configure alarm notifications to the responsible team by both PagerDuty and email, requesting an investigation and a move of the workloads to another available resource. You set repeat notifications every minute. When someone responds to the alarm notifications, you temporarily stop notifications using the best practice of [suppressing the alarm](#). When the CpuUtilization metric is once again detected from the resource, you remove the suppression.

## Suppress Alarms During Investigations

Once a team member responds to an alarm, suppress notifications during the effort to investigate or mitigate the issue. Temporarily stopping notifications helps to avoid distractions during the investigation and mitigation. Remove the suppression when the issue has been resolved. For instructions, see [To suppress alarms](#).

## Routinely Tune Your Alarms

On a regular basis, such as weekly, review your alarms to ensure optimal configuration. Calibrate each alarm's threshold, severity, and notification details, including method, frequency, and targeted audience.



Optimal alarm configuration addresses the following factors:

- Criticality of the resource.
- Appropriate resource behavior. Assess behavior singly and within the context of the service ecosystem. Review metric value fluctuations for a given period of time and then adjust thresholds as needed.
- Acceptable notification noise. Assess the notification method (for example, email or PagerDuty), the appropriate recipients, and the frequency of repeated notifications.

For instructions, see [To update an alarm](#).

## Monitoring Query Language (MQL) Reference

This topic describes the components that appear in Monitoring Query Language (MQL) expressions, the order that they appear in, and valid values.

MQL syntax governs expressions for querying metrics that are published to the Monitoring service. In the Console, MQL expressions appear in **Advanced Mode**. If you don't need to [aggregate results by group](#) or to use other advanced query functionality, then you can create simpler versions of metric queries using **Basic Mode** in the Console.

### Components in an MQL Expression

An MQL expression includes the following components:

- metric
- interval
- dimensions, as one or more name-value pairs (optional)
- grouping function (optional)
- statistic
- comparison operation (optional). Useful for defining [alarms](#).

The query components appear in the following order (boldface components are required):

**metric**[**interval**] {dimensionname="dimensionvalue"}.groupingfunction.**statistic**

Comparison operation queries used for alarms can take the following formats (boldface components are required):

- **metric**[**interval**] {dimensionname="dimensionvalue"}.groupingfunction.**statistic** (where the **statistic** is absent ())
- **metric**[**interval**] {dimensionname="dimensionvalue"}.groupingfunction.**statistic operator value**

- `metric[interval]`  
`{dimensionname="dimensionvalue"}.groupingfunction.statistic operator`  
`(value1, value2)`

You can nest alarm queries and metric queries.



### Note

Nested alarm queries are not currently supported in the Console. Use the API to create alarms with nested queries.

When nesting queries, the query components appear in the following order (boldface components are required):

```
(metric[interval] {dimensionname="dimensionvalue"}.groupingfunction.statistic).groupingfunction.statistic
```

An example of a nested query is a grouped count of hosts with up time greater than zero, where the alarm query to identify these hosts is defined within parentheses:

```
(metric[1h].groupBy(host).min() > 0).grouping().count()
```

## Metric Query Component

The *metric* component of the query appears before the interval.

```
metric[interval]{dimensionname="dimensionvalue"}.groupingfunction.statistic
```

Valid values for *metric* depend on the resource. An example of a metric is `CpuUtilization`, sent by Compute instances. For a list of supported resources with links to their metric references, see [Supported Services](#). You can also use the [ListMetrics](#) operation to find metrics sent by a particular service, such as the Compute service. This operation returns metric definitions.

### Interval Query Component

The *interval* component of the query appears between the metric and statistic (before the optional dimension name-value pair and grouping function).

```
metric[interval]{dimensionname="dimensionvalue"}.groupingfunction.statistic
```

The Monitoring Query Language (MQL) syntax (**Advanced Mode** in the Console) supports the following range of values for *interval*:

1m-60m (also, 1h)

The **Interval** option in the Console (**Basic Mode**) supports the following range of values:

- **1m** - 1 minute
- **5m** - 5 minutes
- **1h** - 1 hour

**Note**

For metric queries, the interval you select drives the default resolution of the request, which determines the maximum time range of data returned.

**Maximum time range returned for a query**

The maximum time range returned for a metric query depends on the resolution. By default, for metric queries, the resolution is the same as the query interval. The maximum time range is calculated using the current time, regardless of any specified end time. Following are the maximum time ranges returned for each interval selection available in the Console.

<b>Interval</b>	<b>Default resolution (metric queries)</b>	<b>Maximum time range returned</b>
<b>1h</b>	1 hour	90 days
<b>5m</b>	5 minutes	30 days
<b>1m</b>	1 minute	7 days

**See examples of returned data**

Example 1: One-minute interval and resolution up to the current time, sent at 10:00 on January 8th. No resolution or end time is specified, so the resolution



defaults to the interval value of `1m`, and the end time defaults to the current time (`2019-01-08T10:00:00.789Z`). This request returns a maximum of 7 days of metric data points. The earliest data point possible within this seven-day period would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`).

Example 2: Five-minute interval with one-minute resolution up to two days ago, sent at 10:00 on January 8th. Because the resolution drives the maximum time range, a maximum of 7 days of metric data points is returned. While the end time specified was 10:00 on January 6th (`2019-01-06T10:00:00.789Z`), the earliest data point possible within this seven-day period would be 10:00 on January 1st (`2019-01-01T10:00:00.789Z`). Therefore, only 5 days of metric data points can be returned in this example.

For more information about the resolution parameter as used in metric queries, see [SummarizeMetricsData](#).

For alarm queries, the specified interval has no effect on the resolution of the request. The only valid value of the resolution for an alarm query request is `1m`. For more information about the resolution parameter as used in alarm queries, see [Alarm](#).

## Dimension Query Component

The `dimensionname="dimensionvalue"` component of the query appears between the interval and statistic (before the optional grouping function).

```
metric[interval]{dimensionname="dimensionvalue"}.groupingfunction.statistic
```

Surround the dimension value with double quotes. Example dimension name-value pair for filtering by availability domain: `availabilityDomain = "VeBZ:PHX-AD-1"`

You can specify multiple dimension name-value pairs. Place each pair within the brackets and separate the pairs with commas.

Valid values for *dimensionname* depend on the *metric*. An example of a dimension name is `resourceDisplayName`, included with the `CpuUtilization` metric sent by Compute instances. For a list of supported resources with links to their metric references, including dimensions, see [Supported Services](#). You can also use the [ListMetrics](#) operation to find metrics (and their dimensions) sent by a particular application or service, such as the Compute service.

## Grouping Function Query Component

The *groupingfunction* component of the query appears between the interval and statistic (after the optional dimension name-value pair).

```
metric[interval]{dimensionname="dimensionvalue"}.groupingfunction.statistic
```

Valid grouping functions are as follows.

Grouping function (MQL expression; Advanced Mode in the Console)	Grouping function option (Basic Mode in the Console)	Description
<code>groupBy()</code>	(not available)	Aggregates query results by group (dimension or <a href="#">resource group</a> ).  For example, <code>groupBy(availabilityDomain)</code> groups results by availability domain so that results from each availability domain are together.
<code>grouping()</code>	<b>Aggregate Metric Streams</b>	Aggregates all query results.

## Statistic Query Component

The *statistic* component of the query appears after the interval and optional dimension name-value pair and grouping function.

```
metric[interval]{dimensionname="dimensionvalue"}.groupingfunction.statistic
```

Valid statistics are as follows.

Statistic (MQL expression; Advanced Mode in the Console)	Statistic option (Basic Mode in the Console)	Description
<code>absent()</code>	(see <a href="#">absent</a> )	Sets the trigger condition as absence of the specified metric. (Applies to comparison operation queries only. Useful for defining <a href="#">alarms</a> .)
<code>avg()</code>	(not available)	Returns the value of Sum divided by Count during the specified time period. Identical to <code>mean()</code> .
<code>count()</code>	<b>COUNT</b>	Returns the number of observations received in the specified time period.
<code>increment()</code>	(not available)	Returns the per-interval change.
<code>max()</code>	<b>MAX</b>	Returns the highest value observed during the specified time period.
<code>mean()</code>	<b>MEAN</b>	Returns the value of Sum divided by Count during the specified time period.
<code>min()</code>	<b>MIN</b>	Returns the lowest value observed during the specified time period.

Statistic (MQL expression; Advanced Mode in the Console)	Statistic option (Basic Mode in the Console)	Description
<code>percentile()</code>	<b>P50</b> <b>P90</b> <b>P95</b> <b>P99</b> <b>P99.9</b>	Returns the estimated value of the specified percentile. Valid values are greater than 0.0 and less than 1.0.  For example, <code>percentile(0.8)</code> returns the value of the 80th percentile.
<code>rate()</code>	<b>RATE</b>	Returns the per-interval average rate of change. The unit is per-second.
<code>sum()</code>	<b>SUM</b>	Returns all values added together.

## Operator and Value Query Component

The *operator value* component of the query appears after the statistic in threshold alarm queries. Either one or two values are needed, depending on the operator:

- `metric[interval]`  
`{dimensionname="dimensionvalue"}.groupingfunction.statistic operator value`
- `metric[interval]`  
`{dimensionname="dimensionvalue"}.groupingfunction.statistic operator (value1, value2)`

## CHAPTER 22 Monitoring

---

Valid operators are as follows.

<b>Operator (MQL expression; Advanced Mode in the Console)</b>	<b>Operator option (Basic Mode in the Console)</b>	<b>Number of values</b>
>	<b>greater than</b>	1
>=	<b>greater than or equal to</b>	1
==	<b>equal to</b>	1
!= (not equal to)	(not available)	1
<	<b>less than</b>	1
<=	<b>less than or equal to</b>	1
in (inclusive of specified values)	<b>between</b> (inclusive of specified values)	2
not in (inclusive of specified values)	<b>outside</b> (inclusive of specified values)	2
Not applicable. See <a href="#">absent()</a> .	<b>absent</b>	0

# CHAPTER 23 Networking

This chapter explains how to set up cloud networks.

## Overview of Networking

When you work with Oracle Cloud Infrastructure, one of the first steps is to set up a virtual cloud network (VCN) for your cloud resources. This topic gives you an overview of Oracle Cloud Infrastructure Networking components and typical scenarios for using a VCN.

## Networking Components

The Networking service uses virtual versions of traditional network components you might already be familiar with:

### **VIRTUAL CLOUD NETWORK (VCN)**

A virtual, private network that you set up in Oracle data centers. It closely resembles a traditional network, with firewall rules and specific types of communication gateways that you can choose to use. A VCN resides in a single Oracle Cloud Infrastructure region and covers a single, contiguous IPv4 CIDR block of your choice. See [Allowed VCN Size and Address Ranges](#). The terms *virtual cloud network*, *VCN*, and *cloud network* are used interchangeably in this documentation. For more information, see [VCNs and Subnets](#).

### **SUBNETS**

Subdivisions you define in a VCN (for example, 10.0.0.0/24 and 10.0.1.0/24). Subnets contain virtual network interface cards (VNICs), which attach to instances. Each subnet consists of a contiguous range of IP addresses that do not overlap with other subnets in the VCN. You can designate a subnet to exist either in a single availability domain or across an entire region (regional subnets are recommended). Subnets act as a unit of configuration within the VCN: All VNICs in a given subnet use the same route table, security lists, and DHCP options (see the definitions that follow). You can designate a subnet as either public or private when you create it. Private means VNICs in the subnet

can't have public IP addresses. Public means VNICs in the subnet can have public IP addresses at your discretion. See [Access to the Internet](#).

### **VNIC**

A virtual network interface card (VNIC), which attaches to an instance and resides in a subnet to enable a connection to the subnet's VCN. The VNIC determines how the instance connects with endpoints inside and outside the VCN. Each instance has a primary VNIC that's created during instance launch and cannot be removed. You can add secondary VNICs to an existing instance (in the same availability domain as the primary VNIC), and remove them as you like. Each secondary VNIC can be in a subnet in the same VCN as the primary VNIC, or in a different subnet that is either in the same VCN or a different one. However, all the VNICs must be in the same availability domain as the instance. For more information, see [Virtual Network Interface Cards \(VNICs\)](#).

### **PRIVATE IP**

A private IPv4 address and related information for addressing an instance (for example, a hostname for DNS). Each VNIC has a primary private IP, and you can add and remove secondary private IPs. The primary private IP address on an instance doesn't change during the instance's lifetime and cannot be removed from the instance. For more information, see [Private IP Addresses](#).

### **PUBLIC IP**

A public IPv4 address and related information. You can optionally assign a public IP to your instances or other resources that have a private IP. Public IPs can be either *ephemeral* or *reserved*. For more information, see [Public IP Addresses](#).

### **IPv6**

An IPv6 address and related information. IPv6 is currently supported only in the Government Cloud. For more information, see [IPv6 Addresses](#).

### **DYNAMIC ROUTING GATEWAY (DRG)**

An optional virtual router that you can add to your VCN. It provides a path for *private* network traffic between your VCN and on-premises network. You can use it with other

Networking components and a router in your on-premises network to establish a connection by way of IPsec VPN or Oracle Cloud Infrastructure FastConnect. It can also provide a path for private network traffic between your VCN and another VCN in a different region. For more information, see [Access to Your On-Premises Network](#), [Dynamic Routing Gateways \(DRGs\)](#), and [Remote VCN Peering \(Across Regions\)](#).

### **INTERNET GATEWAY**

Another optional virtual router that you can add to your VCN for direct internet access. For more information, see [Access to the Internet](#) and also [Scenario A: Public Subnet](#).

### **NETWORK ADDRESS TRANSLATION (NAT) GATEWAY**

Another optional virtual router that you can add to your VCN. It gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections. For more information, see [Access to the Internet](#) and also [NAT Gateway](#).

### **SERVICE GATEWAY**

Another optional virtual router that you can add to your VCN. It provides a path for *private* network traffic between your VCN and [supported services in the Oracle Services Network](#) (examples: Oracle Cloud Infrastructure Object Storage and Autonomous Database). For example, DB Systems in a private subnet in your VCN can back up data to Object Storage without needing public IP addresses or access to the internet. For more information, see [Access to Oracle Services: Service Gateway](#).

### **LOCAL PEERING GATEWAY (LPG)**

Another optional virtual router that you can add to your VCN. It lets you peer one VCN with another VCN in the same region. *Peering* means the VCNs communicate using private IP addresses, without the traffic traversing the internet or routing through your on-premises network. A given VCN must have a separate LPG for each peering it establishes. For more information, see [Local VCN Peering \(Within Region\)](#).

### REMOTE PEERING CONNECTION (RPC)

A component that you can add to a DRG. It lets you peer one VCN with another VCN in a *different* region. For more information, see [Remote VCN Peering \(Across Regions\)](#).

### ROUTE TABLES

Virtual route tables for your VCN. They have rules to route traffic from subnets to destinations outside the VCN by way of gateways or specially configured instances. Your VCN comes with an empty default route table, and you can add custom route tables of your own. For more information, see [Route Tables](#).

### SECURITY RULES

Virtual firewall rules for your VCN. They are ingress and egress rules that specify the types of traffic (protocol and port) allowed in and out of the instances. You can choose whether a given rule is stateful or stateless. For example, you can allow incoming SSH traffic from anywhere to a set of instances by setting up a stateful ingress rule with source CIDR 0.0.0.0/0, and destination TCP port 22. To implement security rules, you can use *network security groups* or *security lists*. A network security group consists of a set of security rules that apply only to the resources in that group. Contrast this with a security list, where the rules apply to all the resources in any subnet that uses the list. Your VCN comes with a default security list with default security rules. For more information, see [Security Rules](#).

### DHCP OPTIONS

Configuration information that is automatically provided to the instances when they boot up. For more information, see [DHCP Options](#).

## Allowed VCN Size and Address Ranges

A VCN covers a single, contiguous IPv4 CIDR block of your choice. The allowable VCN size range is /16 to /30. Example: 10.0.0.0/16. The Networking service reserves the first two IP addresses and the last one in each subnet's CIDR. After you've created a VCN or subnet, you can't change its size, so it's important to think about the size of VCN and subnets you need before creating them.

For your VCN, Oracle recommends using one of the private IP address ranges specified in [RFC 1918](#) (10.0.0.0/8, 172.16/12, and 192.168/16). However, you can use a publicly routable range. Regardless, this documentation uses the term *private IP address* when referring to IP addresses in your VCN's CIDR.

The VCN's CIDR must not overlap with your on-premises network or [another VCN you peer with](#). The subnets in a given VCN must not overlap with each other. For reference, here's a [CIDR calculator](#).

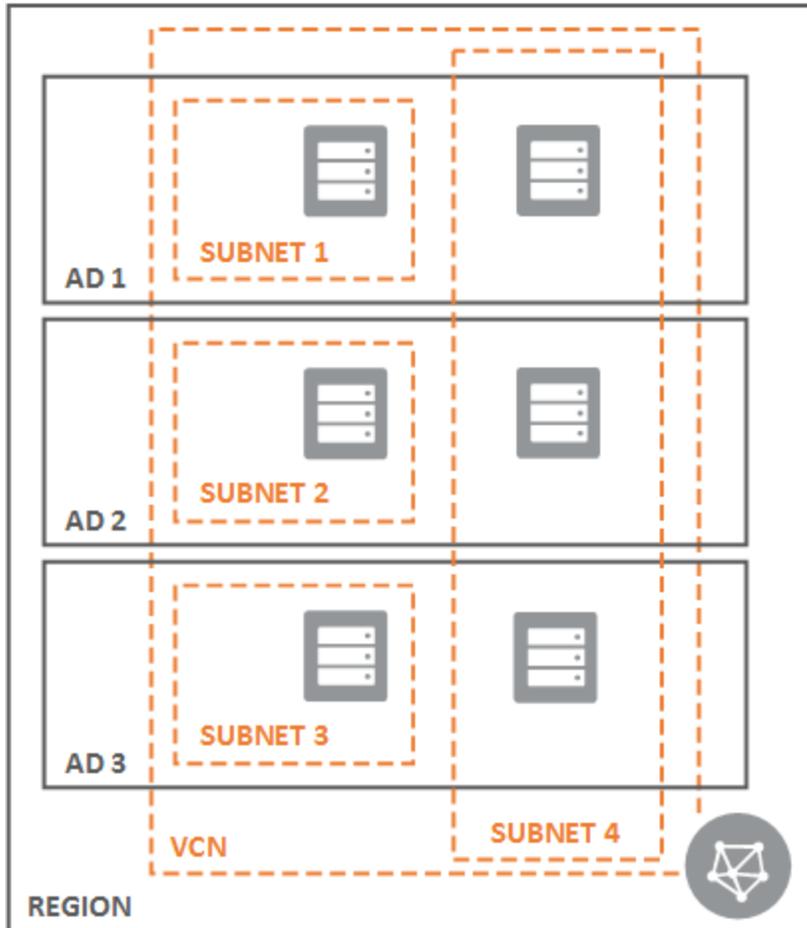
IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

### Availability Domains and Your VCN

Your VCN resides in a single Oracle Cloud Infrastructure region. A region can have multiple availability domains to provide isolation and redundancy. For more information, see [Regions and Availability Domains](#).

Originally subnets were designed to cover only one availability domain (AD) in a region. They were all *AD-specific*, which means the subnet's resources were required to reside in a particular availability domain. Now subnets can be either *AD-specific* or *regional*. You choose the type when you create the subnet. Both types of subnets can co-exist in the same VCN. In the following diagram, subnets 1-3 are AD-specific, and subnet 4 is regional.

Subnets can be regional or specific to an AD



Aside from the removal of the AD constraint, regional subnets behave the same as AD-specific subnets. **Oracle recommends using regional subnets** because they're more flexible. They make it easier to efficiently divide your VCN into subnets while also designing for availability domain failure.

When you create a resource such as a Compute instance, you choose which availability domain the resource will be in. From a virtual networking standpoint, you must also choose

which VCN and subnet the instance will be in. You can either choose a regional subnet, or choose an AD-specific subnet that matches the AD you chose for the instance.

### Default Components that Come With Your VCN

Your VCN automatically comes with these *default* components:

- Default [route table](#), with no route rules
- Default [security list](#), with default security rules
- Default [set of DHCP options](#), with default values

You can't delete these default components. However, you can change their contents (for example, the rules in the default security list). And you can create your own custom versions of each kind of component in your VCN. There are limits to how many you can create and the maximum number of rules. For more information, see [Service Limits](#).

Each subnet always has these components associated with it:

- One route table
- One or more security lists (for the maximum number, see [Service Limits](#))
- One set of DHCP options

During subnet creation, you can choose which route table, security list, and set of DHCP options the subnet uses. If you don't specify a particular component, the subnet automatically uses the VCN's default component. You can [change which components the subnet uses](#) at any time.



#### Tip

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

### Connectivity Choices

You can control whether subnets are public or private, and whether instances get public IP addresses. You can set up your VCN to have access to the internet if you like. You can also privately connect your VCN to public Oracle Cloud Infrastructure services such as Object Storage, to your on-premises network, or to another VCN.

#### **Public vs. Private Subnets**

When you create a subnet, by default it's considered public, which means instances in that subnet are allowed to have public IP addresses. Whoever launches the instance chooses whether it will have a public IP address. You can override that behavior when creating the subnet and request that it be private, which means instances launched in the subnet are prohibited from having public IP addresses. Network administrators can therefore ensure that instances in the subnet have no internet access, even if the VCN has a working internet gateway, and security rules and firewall rules allow the traffic.

#### **How IP Addresses Are Assigned**

Each instance has a primary VNIC that's created during instance launch and cannot be removed. You can add [secondary VNICs](#) to an existing instance (in the same availability domain as the primary VNIC) and remove them as you like.

Every VNIC has a private IP address from the associated subnet's CIDR. You can choose the particular IP address (during instance launch or secondary VNIC creation), or Oracle can choose it for you. The private IP address does not change during the lifetime of the instance and cannot be removed. You can also add [secondary private IPs](#) to a VNIC.

If the VNIC is in a public subnet, then each private IP on that VNIC can have a [public IP](#) assigned to it at your discretion. Oracle chooses the particular IP address. There are two types of public IPs: *ephemeral* and *reserved*. An ephemeral public IP exists only for the lifetime of the private IP it's assigned to. In contrast, a reserved public IP exists as long as you want it to. You maintain a pool of reserved public IPs and allocate them to your instances at your discretion. You can move them from resource to resource in a region as you need to.

### Access to the Internet

There are two optional gateways (virtual routers) that you can add to your VCN depending on the type of internet access you need:

- **Internet gateway:** For resources with public IP addresses that need to be reached from the internet (example: a web server) or need to initiate connections to the internet.
- **NAT gateway:** For resources without public IP addresses that need to initiate connections to the internet (example: for software updates) but need to be protected from inbound connections from the internet.

Just having an internet gateway alone does not expose the instances in the VCN's subnets directly to the internet. The following requirements must also be met:

- The [internet gateway](#) must be enabled (by default, the internet gateway is enabled upon creation).
- The subnet must be [public](#).
- The subnet must have a [route rule](#) that directs traffic to the internet gateway.
- The subnet must have [security list rules](#) that allow the traffic (and each instance's firewall must allow the traffic).
- The instance must have a [public IP address](#).



### Tip

To access public services such as Object Storage from your VCN without the traffic going over the internet, use a [service gateway](#).

Also, be aware that when an internet gateway receives traffic from your VCN destined for a public IP address *that is part of Oracle Cloud Infrastructure* (such as Object Storage), the internet gateway routes the traffic to the destination without sending the traffic over the internet.

You can also give a subnet *indirect* access to the internet by setting up an internet proxy in your on-premises network and then connecting that network to your VCN by way of a DRG. For more information, see [Access to Your On-Premises Network](#).

### Access to Public Oracle Cloud Infrastructure Services

You can use a service gateway with your VCN to enable private access to public Oracle Cloud Infrastructure services such as Object Storage. For example, DB Systems in a private subnet in your VCN can back up data to Object Storage without needing public IP addresses or access to the internet. No internet gateway or NAT is required. For more information, see [Access to Oracle Services: Service Gateway](#).

### Access to Your On-Premises Network

There are two ways to connect your on-premises network to Oracle Cloud Infrastructure:

- [VPN Connect](#): Offers multiple IPsec tunnels between your existing network's edge and your VCN, by way of a [DRG](#) that you create and attach to your VCN.
- [Oracle Cloud Infrastructure FastConnect](#): Offers a private connection between your existing network's edge and Oracle Cloud Infrastructure. Traffic does not traverse the internet. Both private peering and public peering are supported. That means your on-

premises hosts can access private IPv4 addresses in your VCN as well as regional public IPv4 addresses in Oracle Cloud Infrastructure (for example, Object Storage or public [load balancers](#) in your VCN).

You can use one or both types of the preceding connections. If you use both, you can use them simultaneously, or in a redundant configuration. These connections come to your VCN by way of a single DRG that you create and attach to your VCN. Without that DRG attachment and a route rule for the DRG, traffic does not flow between your VCN and on-premises network. At any time, you can detach the DRG from your VCN but maintain all the remaining components that form the rest of the connection. You could then reattach the DRG again, or attach it to another VCN.

### **Access to Another VCN**

You can connect your VCN to another VCN over a private connection that doesn't require the traffic to traverse the internet. In general, this type of connection is referred to as *VCN peering*. Each VCN must have specific components to enable peering. The VCNs must also have specific IAM policies, route rules, and security rules that permit the connection to be made and the desired network traffic to flow over the connection. For more information, see [Access to Other VCNs: Peering](#).

### **Connection to Oracle Cloud Infrastructure Classic**

You can set up a connection between your Oracle Cloud Infrastructure environment and Oracle Cloud Infrastructure Classic environment. This connection can facilitate hybrid deployments between the two environments, or migration from Oracle Cloud Infrastructure Classic to Oracle Cloud Infrastructure. For more information, see [Access to Oracle Cloud Infrastructure Classic](#).

### **Connection to Microsoft Azure**

Oracle and Microsoft have created a cross-cloud connection between Oracle Cloud Infrastructure and Microsoft Azure in certain regions. This connection lets you set up cross-cloud workloads without the traffic between the clouds going over the internet. For more information, see [Access to Microsoft Azure](#).

### Connection to Other Clouds with Libreswan

You can connect your VCN to another cloud provider by using an IPsec VPN with a [Libreswan](#) VM as the customer-premises equipment (CPE). For more information, see [Access to Other Clouds with Libreswan](#).

### Networking Scenarios

This documentation includes a few basic networking scenarios to help you understand the Networking service and generally how the components work together. See these topics:

- [Scenario A: Public Subnet](#)
- [Scenario B: Private Subnet with a VPN](#)
- [Scenario C: Public and Private Subnets with a VPN](#)

### Transit Routing

Scenarios A–C show your on-premises network connected to a VCN by way of [FastConnect](#) or [VPN Connect](#), and accessing only the resources in that VCN.

The following advanced routing scenarios give your on-premises network additional access beyond the resources in the connected VCN. Traffic travels from your on-premises network to the VCN, and then *transits through* the VCN to its destination. See these topics:

- [Transit Routing: Access to Multiple VCNs in the Same Region](#): Your on-premises network has access to *multiple* VCNs in the same region over a single FastConnect private virtual circuit or VPN Connect. The VCNs are in a hub-and-spoke layout, with the on-premises network connected to the VCN that acts as the hub. The spoke VCNs are peered with the hub VCN.
- [Transit Routing: Private Access to Oracle Services](#): Your on-premises network has *private access* to Oracle services in the [Oracle Services Network](#) by way of the connected VCN and the VCN's service gateway. The traffic does not go over the internet.

### Regions and Availability Domains

Your VCN resides in a single Oracle Cloud Infrastructure region. Each subnet resides in a single availability domain (AD). Availability domains are designed to provide isolation and redundancy in your VCN, as illustrated in Scenario B and C earlier. For example, you could set up your primary set of subnets in a single AD, and then set up a duplicate set of subnets in a secondary AD. The two ADs are isolated from each other in the Oracle data centers, so if one fails, you can easily switch over to the other AD. For more information, see [Regions and Availability Domains](#).

### Public IP Address Ranges

For a list of Oracle Cloud Infrastructure public IP ranges, see [IP Address Ranges](#).

### IP Addresses Reserved for Use by Oracle

Certain IP addresses are reserved for Oracle Cloud Infrastructure use and may not be used in your address numbering scheme.

#### **169.254.0.0/16**

These addresses are used for iSCSI connections to the boot and block volumes, instance metadata, and other services.

#### **Three IP Addresses in Each Subnet**

These addresses consist of:

- The first IP address in the CIDR (the network address)
- The last IP address in the CIDR (the broadcast address)
- The first host address in the CIDR (the subnet default gateway address)

For example, in a subnet with CIDR 192.168.0.0/24, these addresses are reserved:

- 192.168.0.0 (the network address)
- 192.168.0.255 (the broadcast address)
- 192.168.0.1 (the subnet default gateway address)

The remaining addresses in the CIDR (192.168.0.2 to 192.168.0.254) are available for use.

### Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

### Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Limits on Your Networking Components

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

### Scenario A: Public Subnet

This topic explains how to set up Scenario A, which consists of a virtual cloud network (VCN) and a regional public subnet. There are public servers in separate availability domains for redundancy. The VCN is directly connected to the internet by way of an internet gateway. The gateway is also used for connectivity to your on-premises network. Any resource in the on-premises network that needs to communicate with resources in the VCN must have a public IP address and access to the internet.

The subnet uses the [default security list](#), which has default rules that are designed to make it easy to get started with Oracle Cloud Infrastructure. The rules enable typical required access (for example, inbound SSH connections and any type of outbound connections). Remember

that security list rules only *allow* traffic. Any traffic not explicitly covered by a security list rule is implicitly denied.

In this scenario, you add additional rules to the default security list. You could instead create a custom security list for those rules. You would then set up the subnet to use both the default security list and the custom security list.

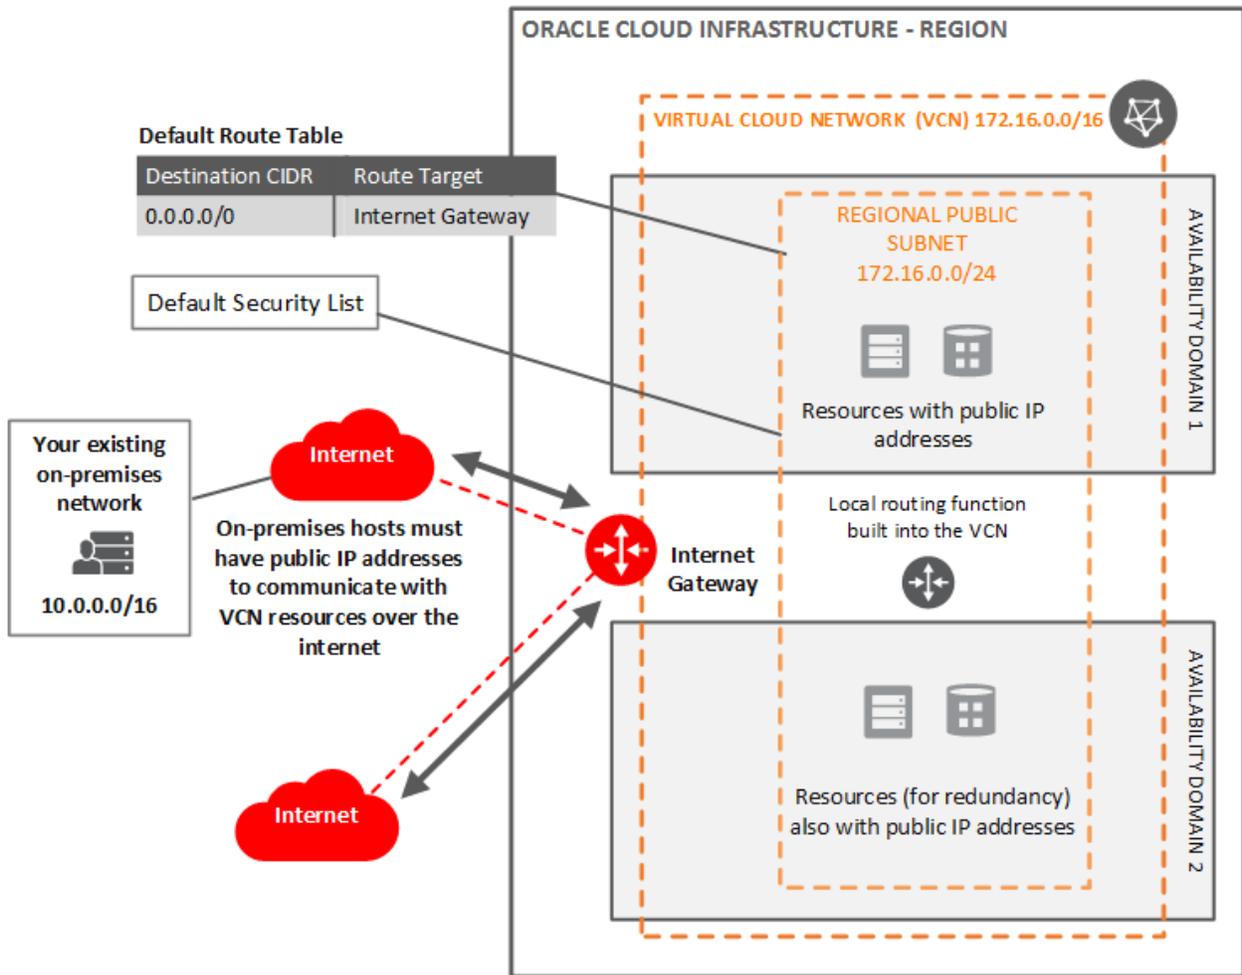


### Tip

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

The subnet uses the default route table, which starts out with no rules when the VCN is created. In this scenario, the table has only a single rule for the internet gateway.

See the following figure.



### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're a member of the Administrators group, you already have the required access to execute Scenario A. Otherwise, you need access to Networking, and you need the ability to launch instances. See [IAM Policies for Networking](#).

### Setting Up Scenario A in the Console

Setup is easy in the Console.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

#### Task 1: Create the VCN

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Virtual Cloud Network**.
4. Enter the following:
  - **Name:** A friendly name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** Leave as is.

- **Create Virtual Cloud Network Only:** Make sure this radio button is selected (the default).
  - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the Government Cloud. For more information, see [IPv6 Addresses](#).
  - **CIDR Block:** A single, contiguous CIDR block for the VCN. For example: 172.16.0.0/16. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
  - **Use DNS Hostnames in this VCN:** Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (`<VCN DNS label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create Virtual Cloud Network**.
- The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose.

### Task 2: Create the regional public subnet

1. While still viewing the VCN, click **Create Subnet**.
2. Enter the following:
  - **Name:** A friendly name for the subnet (for example, Regional Public Subnet). It doesn't have to be unique, and it cannot be changed later in the Console (but you

can change it with the API). Avoid entering confidential information.

- **Regional or Availability Domain-Specific:** Select **Regional** (recommended), which means the subnet spans all availability domains in the region. Later when you launch an instance, you can create it in any availability domain in the region. For more information, see [About Regional Subnets](#).
  - **CIDR Block:** A single, contiguous CIDR block within the VCN's CIDR block. For example: 172.16.0.0/24. You *cannot* change this value later. For reference, here's a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
  - **Route Table:** Select the default route table.
  - **Private or public subnet:** Select **Public Subnet**, which means instances in the subnet can optionally have public IP addresses. For more information, see [Access to the Internet](#).
  - **Use DNS Hostnames in this Subnet:** This option is available only if you provided a DNS label for the VCN during creation. The option is required for assignment of DNS hostnames to hosts in the subnet, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the subnet, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the subnet (`<subnet_dns_label>.<VCN_dns_label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).
  - **DHCP Options:** Select the default set of DHCP options.
  - **Security Lists:** Make sure the default security list is selected (the default).
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
3. Click **Create Subnet**.
- The subnet is then created and displayed on the **Subnets** page.

### Task 3: Create the internet gateway

1. Under **Resources**, click **Internet Gateways**.
2. Click **Create Internet Gateway**.
3. Enter the following:
  - **Name:** A friendly name for the internet gateway. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** Leave as is.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
4. Click **Create Internet Gateway**.

Your internet gateway is created and displayed on the **Internet Gateways** page. It's already enabled, but you must add a route rule that allows traffic to flow to the gateway.

### Task 4: Update the default route table to use the internet gateway

The default route table starts out with no rules. Here you add a rule that routes all traffic destined for addresses outside the VCN to the internet gateway. The existence of this rule also enables inbound connections to come from the internet to the subnet, through the internet gateway. You use security list rules to control the *types of traffic* that are allowed in and out of the instances in the subnet (see the next task).

No route rule is required in order to route traffic within the VCN itself.

1. Under **Resources**, click **Route Tables**.
2. Click the default route table to view its details.
3. Click **Add Route Rule**.
4. Enter the following:

- **Target Type:** Internet Gateway
  - **Destination CIDR block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table goes to the target specified in this rule)
  - **Compartment:** The compartment where the internet gateway is located.
  - **Target:** The internet gateway you created.
5. Click **Add Route Rule**.

The default route table now has a rule for the internet gateway. Because the subnet was set up to use the default route table, the resources in the subnet can now use the internet gateway. The next step is to specify the types of traffic you want to allow in and out of the instances you later create in the subnet.

### Task 5: Update the default security list

Earlier you set up the subnet to use the VCN's [default security list](#). Now you add security list rules that allow the types of connections that the instances in the VCN will need.

For example: This is a public subnet with an internet gateway, so the instances you launch might need to receive inbound HTTPS connections from the internet (if they're web servers). Here's how to add another rule to the default security list to enable that traffic:

1. Under **Resources**, click **Security Lists**.
2. Click the default security list to view its details. By default, you land on the **Ingress Rules** page.
3. Click **Add Ingress Rule**.
4. To enable inbound connections for HTTPS (TCP port 443), enter the following:
  - **Stateless:** Unselected (this is a [stateful rule](#))
  - **Source Type:** CIDR
  - **Source CIDR:** 0.0.0.0/0

- **IP Protocol:** TCP
  - **Source Port Range:** All
  - **Destination Port Range:** 443
5. Click **Add Ingress Rule**.



### Important

#### *Security List Rule for Windows Instances*

If you're going to launch Windows instances, you need to add a security list rule to enable Remote Desktop Protocol (RDP) access. Specifically, you need a stateful ingress rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port. For more information, see [Security Lists](#).

For a production VCN, you typically set up one or more *custom* security lists for each subnet. If you like, you can edit the subnet to [use different security lists](#). If you choose not to use the default security list, do so only after carefully assessing which of its default rules you want to duplicate in your custom security list. For example: the [default ICMP rules in the default security list](#) are important for receiving connectivity messages.

### Task 6: Create instances in separate availability domains

Your next step is to create one or more instances in the subnet. The scenario's diagram shows instances in two different availability domains. When you create the instance, you choose the AD, which VCN and subnet to use, and several other characteristics.

Each instance automatically gets a private IP address. When you create an instance in a *public subnet*, you choose whether the instance gets a public IP address. With this network setup in Scenario A, you *must* give each instance a public IP address, or else you can't access them

through the internet gateway. The default (for a public subnet) is for the instance to get a public IP address.

After creating an instance in this scenario, you can connect to it over the internet with SSH or RDP from your on-premises network or other location on the internet. For more information and instructions, see [Launching an Instance](#).

### Setting Up Scenario A with the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations:

1. [CreateVcn](#): Make sure to include a DNS label for the VCN if you want the instances to have hostnames (see [DNS in Your Virtual Cloud Network](#)).
2. [CreateSubnet](#): Create one regional public subnet. Include a DNS label for the subnet if you want the instances to have hostnames. Use the default route table, default security list, and default set of DHCP options.
3. [CreateInternetGateway](#)
4. [UpdateRouteTable](#): To enable communication with the internet gateway, update the default route table to include a route rule with destination = 0.0.0.0/0, and destination target = the internet gateway. This rule routes all traffic destined for addresses outside the VCN to the internet gateway. No route rule is required in order to route traffic within the VCN itself.
5. [UpdateSecurityList](#): To allow specific types of connections to and from the instances in the subnet.



### Important

#### *Security List Rule for Windows Instances*

If you're going to launch Windows instances, you need to add a security list rule to enable Remote Desktop Protocol (RDP) access. Specifically, you need a stateful ingress rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port. For more information, see [Security Lists](#).

Your next step is to create one or more instances in the subnet. The scenario's diagram shows instances in two different availability domains. When you create the instance, you choose the AD, which VCN and subnet to use, and several other characteristics.

Each instance automatically gets a private IP address. When you create an instance in a *public subnet*, you choose whether the instance gets a public IP address. With this network setup in Scenario A, you *must* give each instance a public IP address, or else you can't access them through the internet gateway. The default (for a public subnet) is for the instance to get a public IP address.

After creating an instance in this scenario, you can connect to it over the internet with SSH or RDP from your on-premises network or other location on the internet. For more information and instructions, see [Launching an Instance](#).

## Scenario B: Private Subnet with a VPN

This topic explains how to set up Scenario B, which consists of a virtual cloud network (VCN) with a regional private subnet. There are servers in separate availability domains for redundancy. The VCN has a dynamic routing gateway (DRG) and IPSec VPN for connectivity to your on-premises network. The VCN has no direct connection to the internet. Any connection to the internet would need to come indirectly by way of the on-premises network.

The subnet uses the [default security list](#), which has default rules that are designed to make it easy to get started with Oracle Cloud Infrastructure. The rules enable typical required access

(for example, inbound SSH connections and any type of outbound connections). Remember that security list rules only *allow* traffic. Any traffic not explicitly covered by a security list rule is denied.

In this scenario, you add additional rules to the default security list. You could instead create a custom security list for those rules. You would then set up the subnet to use both the default security list and the custom security list.

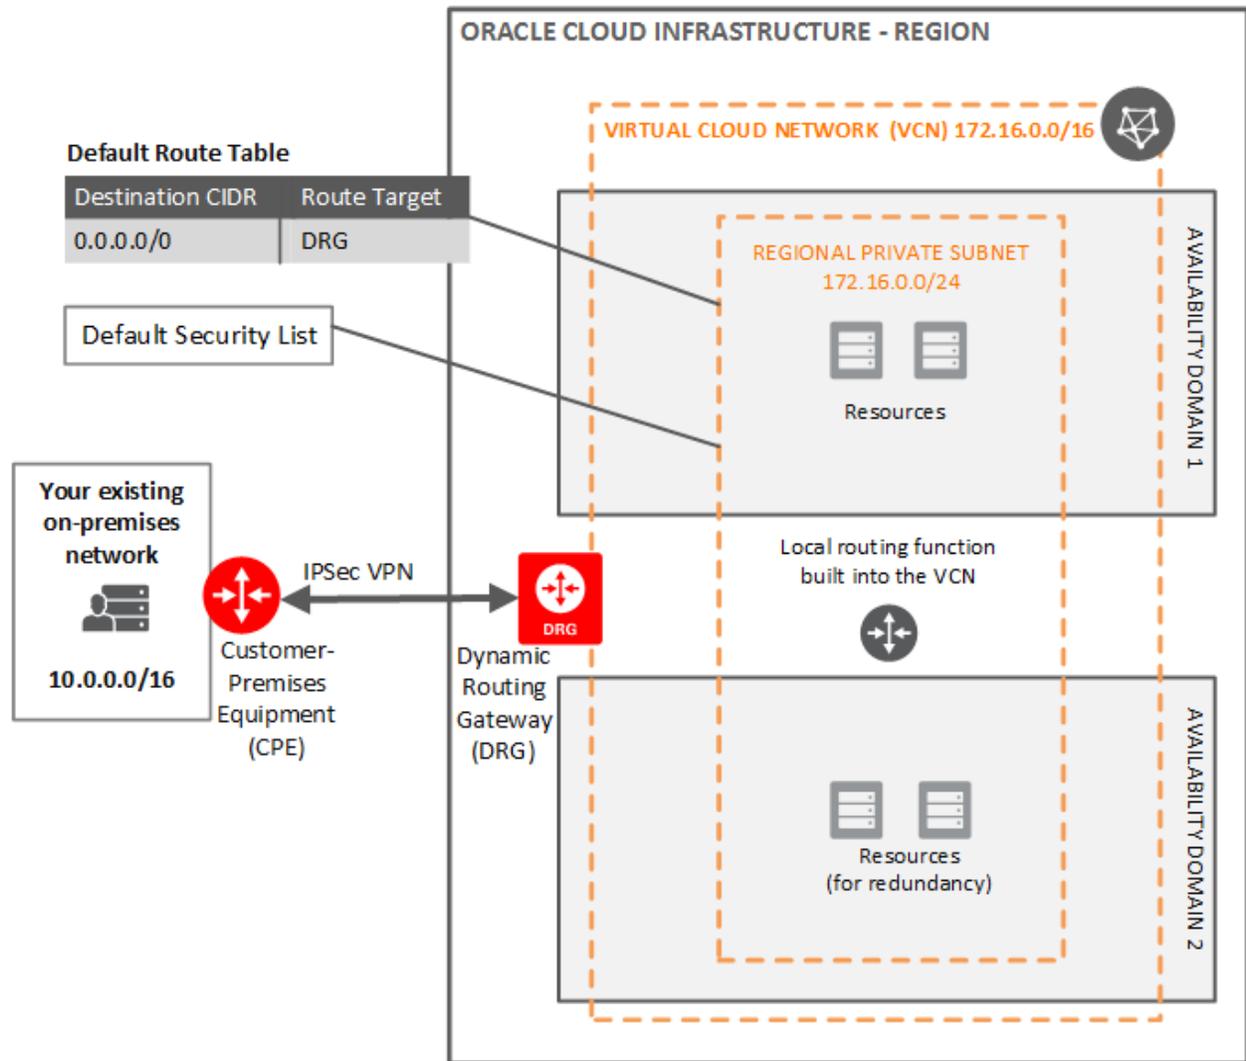


### Tip

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

The subnet uses the default route table, which starts out with no rules when the VCN is created. In this scenario, the table has only a single rule for the DRG. No route rule is required in order to route traffic within the VCN itself.

See the following figure.





### Tip

The scenario uses an IPsec VPN for connectivity. However, you could instead use [Oracle Cloud Infrastructure FastConnect](#).

## Prerequisites

To set up the VPN in this scenario, you need to get the following information from a network administrator:

- Public IP address of the customer-premises equipment (CPE) at your end of the VPN
- Static routes for your on-premises network (this scenario uses static routing for the VPN tunnels, but you could instead use [BGP dynamic routing](#))

You will provide Oracle this information and in return receive the information your network administrator must have to configure the CPE at your end of the VPN.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're a member of the Administrators group, you already have the required access to execute Scenario B. Otherwise, you need access to Networking, and you need the ability to launch instances. See [IAM Policies for Networking](#).

### Setting Up Scenario B

Setup is easy in the Console. Alternatively, you can use the Oracle Cloud Infrastructure API, which lets you execute the individual operations yourself.



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.



#### **Important**

Most of this process involves working with the Console or API (whichever you choose) for a short period to set up the desired Networking components. But there's also a critical step that requires a network administrator in your organization to take information you receive from setting up the components and use it to configure the CPE at your end of the VPN. Therefore you can't complete this process in one short session. You'll need to break for an unknown period of time while the network administrator completes the configuration and then return afterward to confirm communication with your instances over the VPN.

### Using the Console

#### Task 1: Set up the VCN and subnet

1. Create the VCN:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
  - c. Click **Create Virtual Cloud Network**.
  - d. Enter the following:
    - **Name:** A friendly name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave as is.
    - **Create Virtual Cloud Network Only:** Make sure this radio button is selected (the default).
    - **CIDR Block:** A single, contiguous CIDR block for the VCN. For example: 172.16.0.0/16. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
    - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
    - **Use DNS Hostnames in this VCN:** Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will

generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (`<VCN DNS label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).

- **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- e. Click **Create Virtual Cloud Network**.

The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose.

2. Create the regional private subnet:

- a. While still viewing the VCN, click **Create Subnet**.
- b. Enter the following:
  - **Name:** A friendly name for the subnet (for example, *Regional Private Subnet*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Regional or Availability Domain-Specific:** Select **Regional** (recommended), which means the subnet spans all availability domains in the region. Later when you launch an instance, you can create it any availability domain in the region. For more information, see [About Regional Subnets](#).
  - **CIDR Block:** A single, contiguous CIDR block within the VCN's CIDR block. For example: 172.16.0.0/24. You *cannot* change this value later. For reference, here's a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
  - **Route Table:** Select the default route table.

- **Private or public subnet:** Select **Private Subnet**, which means instances in the subnet cannot have public IP addresses. For more information, see [Access to the Internet](#).
  - **Use DNS Hostnames in this Subnet:** This option is available only if you provided a DNS label for the VCN during creation. The option is required for assignment of DNS hostnames to hosts in the subnet, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the subnet, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the subnet (`<subnet_dns_label>.<VCN_dns_label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).
  - **DHCP Options:** Select the default set of DHCP options.
  - **Security Lists:** Select the default security list.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- c. Click **Create Subnet**.
- The subnet is then created and displayed on the **Subnets** page.
3. Update the [default security list](#) to include rules to allow the types of connections that your instances in the VCN will need:
- a. While still on the page displaying your VCN's subnets, click **Security Lists**, and then click the default security list.
  - b. Under **Resources**, click either **Ingress Rules** or **Egress Rules** depending on the type of rule you want to work with. You can add one rule at a time by clicking either **Add Ingress Rule** or **Add Egress Rule**.
  - c. Add your desired rules. Here are suggested ones to add to the default ones already in the default security list:

### Example: Ingress HTTP access

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 80

### Example: Ingress HTTPS access

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443

### Example: Ingress SQL\*Net access for Oracle databases

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0

- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521

Example: Ingress RDP access required for Windows instances

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 3389



### Tip

For additional security, you could modify all the stateful ingress rules to allow traffic only from within your VCN and your on-premises network. You would need to create separate rules for each, one with the VCN's CIDR as the source, and one with the on-premises network's CIDR as the source.

For a production VCN, you typically set up one or more *custom* security lists for each subnet. You can edit the subnet to [use different security lists](#) if you like. If you choose not to use the default security list, do so only after carefully assessing which of its default rules you want to duplicate in your custom security list. For example: the [default ICMP rules in the default security list](#) are important for receiving connectivity messages.

### Task 2: Create instances in separate availability domains

You can now create one or more instances in the subnet (see [Launching an Instance](#)). The scenario's diagram shows instances in two different availability domains. When you create the instance, you choose the AD, which VCN and subnet to use, and several other characteristics.

However, you can't yet communicate with the instances because there's no gateway connecting the VCN to your on-premises network. The next procedure walks you through creating an IPSec VPN connection to enable that communication.

### Task 3: Add an IPSec VPN to your VCN

1. Create a customer-premises equipment (CPE) object:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.
  - b. Click **Create Customer-Premises Equipment**.
  - c. Enter the following:
    - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
    - **Name:** A friendly name for the customer-premises equipment object. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **IP Address:** The public IP address of the CPE at your end of the VPN (see [Prerequisites](#)).
  - d. Click **Create**.

The CPE object will be in the "Provisioning" state for a short period.

2. Create a dynamic routing gateway (DRG):
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.

- b. Click **Create Dynamic Routing Gateway**.
- c. For **Create in Compartment**: Leave the default value (the compartment you're currently working in).
- d. Enter a friendly name for the DRG. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
- e. Click **Create**.

The DRG will be in the "Provisioning" state for a short period. Make sure it is done being provisioned before continuing.

3. Attach the DRG to your VCN:
  - a. Click the DRG that you just created.
  - b. Under **Resources**, click **Virtual Cloud Networks**.
  - c. Click **Attach to Virtual Cloud Network**.
  - d. Select the VCN. Ignore the section for advanced options, which is only for an advanced routing scenario called *transit routing*, which is not relevant here.
  - e. Click **Attach**.

The attachment will be in the "Attaching" state for a short period before it's ready.

4. Update the default route table (which has no rules yet):
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click your VCN.
  - c. Under **Resources**, click **Route Tables**, and then click the default route table.
  - d. Click **Add Route Rule**.
  - e. Enter the following:
    - **Target Type**: Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the

target yourself.

- **Destination CIDR Block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule).

f. Click **Add Route Rule**.

The VCN's default route table now directs outbound traffic to the DRG and ultimately to your on-premises network.

5. Create an IPSec Connection:

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.
- b. Click **Create IPSec Connection**.
- c. Enter the following:
  - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - **Name:** Enter a friendly name for the IPSec connection. It doesn't have to be unique. Avoid entering confidential information.
  - **Customer-Premises Equipment Compartment:** Leave as is (the VCN's compartment).
  - **Customer-Premises Equipment:** Select the CPE object you created earlier.
  - **Dynamic Routing Gateway Compartment:** Leave as is (the VCN's compartment).
  - **Dynamic Routing Gateway:** Select the DRG that you created earlier.
  - **Static Route CIDR:** Enter at least one static route CIDR (see [Prerequisites](#)). If you need to add another, click **Add Static Route**. You can enter up to 10 static routes, and you can [change the static routes](#) later if you like.

- d. Click **Show Advanced Options** and optionally provide the following items:
- **CPE IKE Identifier:** Oracle defaults to using the public IP address of the CPE. But if your [CPE is behind a NAT device](#), you might need to enter a different value. You can either enter the new value here, or [change the value](#) later.
  - **Tunnel 1** and **Tunnel 2:** Leave as is. Later if you want to use BGP dynamic routing instead of static routing for the VPN tunnels, see [Changing from Static Routing to BGP Dynamic Routing](#).
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- e. Click **Create IPSec Connection**.
- The IPSec connection is created and displayed on the page. It will be in the Provisioning state for a short period.
- The displayed tunnel information includes the IP address of the VPN headend and the tunnel's IPSec status (possible values are Up, Down, and Down for Maintenance). At this point, the status is Down. To view the tunnel's shared secret, click the Actions icon (three dots), and then click **View Shared Secret**.
- f. Copy the Oracle VPN IP address and shared secret for each of the tunnels to an email or other location so you can deliver it to the network engineer who will configure the on-premises router.
- For more information, see [CPE Configuration](#). You can view this tunnel information here in the Console at any time.

You have now created all the components required for the IPSec VPN. But your network administrator must configure the CPE before network traffic can flow between your on-premises network and VCN.

### Task 4: Configure your CPE

These instructions are for the network administrator.

1. Make sure you have the tunnel configuration information that Oracle provided during IPsec VPN setup. See [Task 3: Add an IPsec VPN to your VCN](#).
2. Configure your CPE according to the information in [CPE Configuration](#).

If there are already instances in the subnet, you can confirm the IPsec connection is up and running by connecting to the instances from your on-premises network.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations:

1. [CreateVcn](#): Make sure to include a DNS label for the VCN if you want the instances to have hostnames (see [DNS in Your Virtual Cloud Network](#)).
2. [CreateSubnet](#): Create one regional private subnet. Include a DNS label for the subnet if you want the instances to have hostnames. Use the default route table, default security list, and default set of DHCP options.
3. [CreateDrg](#): This creates a new dynamic routing gateway (DRG)
4. [CreateDrgAttachment](#): This attaches the DRG to the VCN.
5. [CreateCpe](#): Here you'll provide the public IP address of the CPE at your end of the VPN (see [Prerequisites](#)).
6. [CreateIPsecConnection](#): Here you'll provide the static routes for your on-premises network (see [Prerequisites](#)). In return, you'll receive the configuration information that your network administrator needs in order to configure your CPE. If you need that information later, you can get it with [GetIPsecConnectionDeviceConfig](#). For more information about the configuration, see [CPE Configuration](#).
7. [UpdateRouteTable](#): To enable communication via the VPN, update the default route table to include this route: a route rule with destination = 0.0.0.0/0, and destination target = the DRG you created earlier.

8. First call [GetSecurityList](#) to get the default security list, and then call [UpdateSecurityList](#) to add rules for the types of connections that your instances in the VCN will need. Be aware that `UpdateSecurityList` overwrites the entire set of rules. Here are some suggested rules to add:
  - Stateful ingress: Source type=CIDR, source CIDR=0.0.0.0/0, protocol=TCP, source port = all, destination port=80 (for HTTP).
  - Stateful ingress: Source type=CIDR, source CIDR=0.0.0.0/0, protocol=TCP, source port = all, destination port=443 (for HTTPS).
  - Stateful ingress: Source type=CIDR, source CIDR=0.0.0.0/0, protocol=TCP, source port = all, destination port=1521 (for SQL\*Net access to Oracle databases).
  - Stateful ingress: Source type=CIDR, source CIDR=0.0.0.0/0, protocol=TCP, source port=all, destination port=3389 (for RDP; required only if using Windows instances).



### Tip

For additional security, you could modify all the stateful ingress rules to allow traffic only from within your VCN and your on-premises network. You would need to create separate rules for each, one with the VCN's CIDR as the source, and one with the on-premises network's CIDR as the source.

9. [LaunchInstance](#): Create one or more instances in the subnet. The scenario's diagram shows instances in two different availability domains. When you create the instance, you choose the AD, which VCN and subnet to use, and several other characteristics. For more information, see [Launching an Instance](#).



**Important**

Although you can create instances in the subnet, you won't be able to communicate with them from your on-premises network until your network administrator configures your CPE (see [CPE Configuration](#)). After that, your IPSec connection should be up and running. You can confirm its status by using [GetIPSecConnectionDeviceStatus](#). You can also confirm the IPSec connection is up by connecting to the instances from your on-premises network.

## Scenario C: Public and Private Subnets with a VPN

This topic explains how to set up Scenario C, which is a simple example of a multi-tier setup. It consists of a virtual cloud network (VCN) with a regional public subnet to hold public servers (such as web servers), and a regional private subnet to hold private servers (such as database servers). There are servers in separate availability domains for redundancy.

The VCN has a dynamic routing gateway (DRG) and IPSec VPN for connectivity to your on-premises network. Instances in the public subnet have direct access to the internet by way of an internet gateway. Instances in the private subnet can initiate connections to the internet by way of a NAT gateway (for example, to get software updates), but cannot receive inbound connections from the internet through that gateway.

Each subnet uses the [default security list](#), which has default rules that are designed to make it easy to get started with Oracle Cloud Infrastructure. The rules enable typical required access (for example, inbound SSH connections and any type of outbound connections). Remember that security list rules only *allow* traffic. Any traffic not explicitly covered by a security list rule is denied.

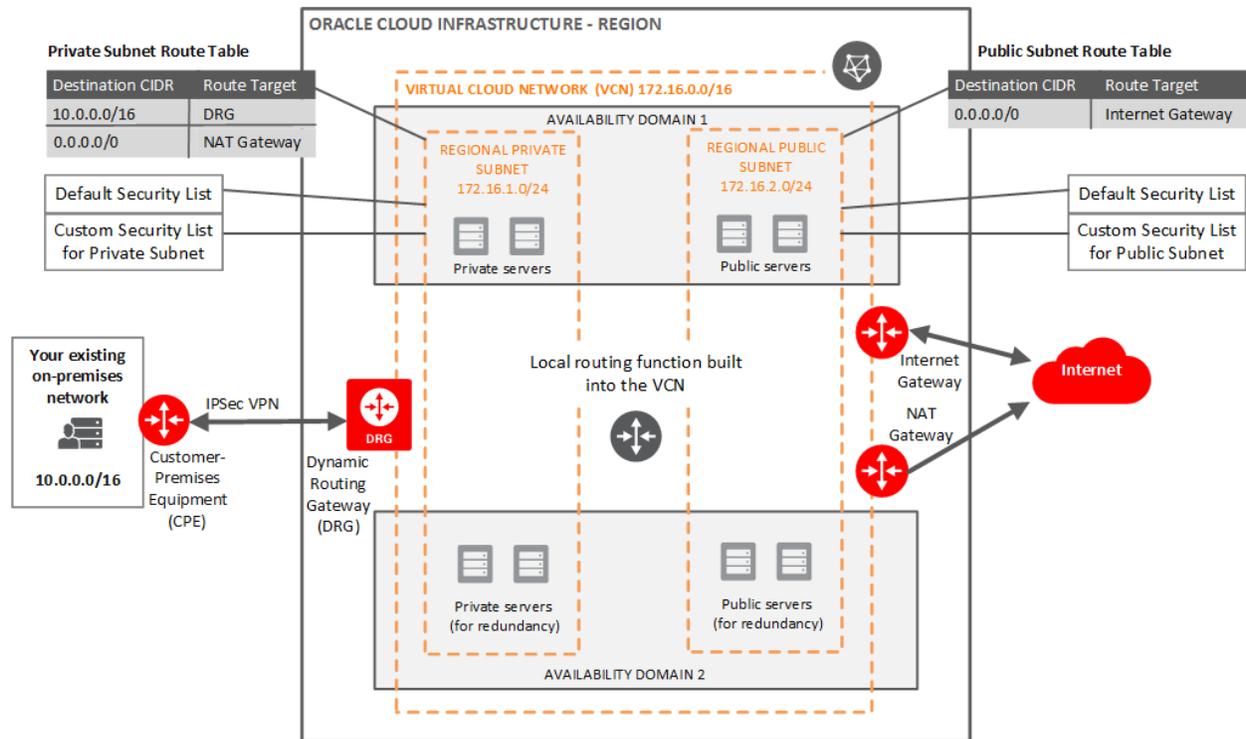


**Tip**

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

Each subnet also has its own custom security list and custom route table with rules specific to the needs of the subnet's instances. In this scenario, the VCN's default route table (which is always empty to start with) is not used.

See the following figure.





### Tip

The scenario uses an IPsec VPN for connectivity. However, you could instead use [Oracle Cloud Infrastructure FastConnect](#).

## Prerequisites

To set up the VPN in this scenario, you need to get the following information from a network administrator:

- Public IP address of the customer-premises equipment (CPE) at your end of the VPN
- Static routes for your on-premises network (this scenario uses static routing for the VPN tunnels, but you could instead use [BGP dynamic routing](#))

You will provide Oracle this information and in return receive the information your network administrator needs in order to configure the on-premises router at your end of the VPN.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're a member of the Administrators group, you already have the required access to execute Scenario C. Otherwise, you need access to Networking, and you need the ability to launch instances. See [IAM Policies for Networking](#).

### Setting Up Scenario C

Setup is easy in the Console. Alternatively, you can use the Oracle Cloud Infrastructure API, which lets you execute the individual operations yourself.



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.



#### **Important**

Most of this process involves working with the Console or API (whichever you choose) for a short period to set up the desired Networking components. But there's also a critical step that requires a network administrator in your organization to take information you receive from setting up the components and use it to configure the on-premises router at your end of the VPN. Therefore you can't complete this process in one short session. You'll need to break for an unknown period of time while the network administrator completes the configuration and then return afterward to confirm communication with your instances over the VPN.

### Using the Console

#### Task 1: Set up the VCN and subnets

1. Create the VCN:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
  - c. Click **Create Virtual Cloud Network**.
  - d. Enter the following:
    - **Name:** A friendly name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave as is.
    - **Create Virtual Cloud Network Only:** Make sure this radio button is selected (the default).
    - **CIDR Block:** A single, contiguous CIDR block for the VCN. For example: 172.16.0.0/16. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
    - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
    - **Use DNS Hostnames in this VCN:** Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the

corresponding **DNS Domain Name** for the VCN (<*VCN DNS label*>.oraclevcn.com). For more information, see [DNS in Your Virtual Cloud Network](#).

- **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- e. Click **Create Virtual Cloud Network**.  
The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose.
2. Create an internet gateway for your VCN:
    - a. Under **Resources**, click **Internet Gateways**.
    - b. Click **Create Internet Gateway**.
    - c. Enter the following:
      - **Name:** A friendly name for the internet gateway. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
      - **Create in Compartment:** Leave as is.
      - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
    - d. Click **Create Internet Gateway**.  
The internet gateway is then created and listed on the page.
  3. Create a NAT gateway for your VCN:
    - a. Under **Resources**, click **NAT Gateways**.
    - b. Click **Create NAT Gateway**.
    - c. Enter the following:
      - **Name:** A friendly name for the NAT gateway. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.

- **Create in Compartment:** Leave as is.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- d. Click **Create NAT Gateway**.
- The NAT gateway is then created and listed on the page.
4. Create the custom route table for the public subnet (which you will create later):
- a. Under **Resources**, click **Route Tables**.
  - b. Click **Create Route Table**.
  - c. Enter the following:
    - **Name:** A friendly name for the route table (for example, *Public Subnet Route Table*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
    - Click **+ Additional Route Rule** and enter the following:
      - **Target Type:** Internet Gateway.
      - **Destination CIDR Block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule).
      - **Compartment:** Leave as is.
      - **Target:** The internet gateway you just created.
  - d. **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
  - e. Click **Create Route Table**.
- The route table is then created and listed on the page.
5. Create the custom route table for the private subnet (which you will create later):

- a. Click **Create Route Table**.
  - b. Enter the following:
    - **Name:** A friendly name for the route table (for example, *Private Subnet Route Table*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
    - Click **+ Additional Route Rule** and enter the following:
      - **Target Type:** NAT Gateway.
      - **Destination CIDR Block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule).
      - **Compartment:** Leave as is.
      - **Target:** The NAT gateway you just created.
  - c. **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
  - d. Click **Create Route Table**.

The route table is then created and listed on the page. Later on after you've set up the IPSec VPN, you will update the Private Subnet Route Table so it routes traffic from the private subnet to the on-premises network by way of the DRG.
6. Update the [default security list](#) to include rules to allow the types of connections that your instances in the VCN will need:
- a. Under **Resources**, click **Security Lists**.
  - b. Click the default security list to view its details. By default, you land on the **Ingress Rules** page.
  - c. Edit each of the existing stateful ingress rules so that the **Source CIDR** is the CIDR for your on-premises network (10.0.0.0/16 in this example) and not

0.0.0.0/0. To edit an existing rule, click the Actions icon (three dots) for the rule, and then click **Edit**.

- d. If you plan to launch Windows instances, add a rule to enable RDP access:

### Example: Ingress RDP access required for Windows instances

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** Your on-premises network (10.0.0.0/16 in this example)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 3389

7. Create a custom security list for the public subnet:
  - a. Return to the **Security Lists** page for the VCN.
  - b. Click **Create Security List**.
  - c. Enter the following:
    - **Name:** Enter a friendly name for the list (for example, *Public Subnet Security List*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - d. Add the following ingress rules:

### Example: Ingress HTTP access

- **Type:** Ingress

- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 80

### Example: Ingress HTTPS access

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443

e. Add the following egress rule:

### Example: Egress SQL\*Net access to Oracle databases

- **Type:** Egress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Destination Type:** CIDR
- **Destination CIDR:** CIDR for the private subnet (172.16.1.0/24 in this example)
- **IP Protocol:** TCP

- **Source Port Range:** All
  - **Destination Port Range:** 1521
- f. Click **Create Security List**.  
The custom security list for the public subnet is then created and listed on the page.
8. Create a custom security list for the *private* subnet:
- a. Click **Create Security List**.
  - b. Enter the following:
    - **Name:** Enter a friendly name for the list (for example, *Private Subnet Security List*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
    - **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - c. Add the following ingress rules:

### Example: Ingress SQL\*Net access from clients in the public subnet

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** CIDR for the public subnet (172.16.2.0/24 in this example)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521

Example: Ingress SQL\*Net access from clients in the private subnet

- **Type:** Ingress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Source Type:** CIDR
- **Source CIDR:** CIDR for the private subnet (172.16.2.1/24 in this example)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521

d. Add the following egress rules:

Example: Egress SQL\*Net access to instances in the private subnet

- **Type:** Egress
- **Stateless:** Unselected (this is a [stateful rule](#))
- **Destination Type:** CIDR
- **Destination CIDR:** CIDR for the private subnet (172.16.1.0/24 in this example)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521

e. Click **Create Security List**.

The custom security list for the private subnet is then created and listed on the page.

9. Create the subnets in the VCN:

- a. Under **Resources**, click **Subnets**.
- b. Click **Create Subnet**.

c. Enter the following:

- **Name:** A friendly name for the regional public subnet (for example, *Regional Private Subnet*). It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
- **Regional or Availability Domain-Specific:** Select **Regional** (recommended), which means the subnet spans all availability domains in the region. Later when you launch an instance, you can create it any availability domain in the region. For more information, see [About Regional Subnets](#).
- **CIDR Block:** A single, contiguous CIDR block within the VCN's CIDR block. For example: 172.16.1.0/24. You *cannot* change this value later. For reference, here's a [CIDR calculator](#).
- **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
- **Route Table:** Select the Private Subnet Route Table you created earlier.
- **Private or public subnet:** Select **Private Subnet**, which means VNICs in the subnet are not allowed to have public IP addresses. For more information, see [Access to the Internet](#).
- **Use DNS Hostnames in this Subnet:** This option is available only if you provided a DNS label for the VCN during creation. The option is required for assignment of DNS hostnames to hosts in the subnet, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the subnet, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the subnet (`<subnet_dns_label>.<vcn_dns_label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).

- **DHCP Options:** Select the default set of DHCP options.
  - **Security Lists:** Select two security lists: Both the default security list and the Private Subnet Security List you created earlier.
- d. Click **Create Subnet**.  
The private subnet is then created and displayed on the **Subnets** page.
- e. Repeat the preceding steps a-d to create the *regional public subnet*. Instead use a name such as *Regional Public Subnet*, select **Public Subnet** instead of **Private Subnet**, use the Public Subnet Route Table, and use both the default security list and Public Subnet Security List you created earlier.

### Task 2: Create instances in separate availability domains

You can now create one or more instances in the subnet (see [Launching an Instance](#)). The scenario's diagram shows instances in two different availability domains. When you create the instance, you choose the AD, which VCN and subnet to use, and several other characteristics.

For each instance in the public subnet, make sure to assign the instance a public IP address. Otherwise, you won't be able to reach the instance from your on-premises network.

You can't yet reach the instances in the private subnet because there's no gateway connecting the VCN to your on-premises network. The next procedure walks you through creating an IPsec VPN connection to enable that communication.

### Task 3: Add an IPsec VPN to your VCN

1. Create a customer-premises equipment object:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.
  - b. Click **Create Customer-Premises Equipment**.
  - c. Enter the following:

- **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - **Name:** A friendly name for the customer-premises equipment object. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **IP Address:** The IP address of the on-premises router at your end of the VPN (see [Prerequisites](#)).
- d. Click **Create**.

The CPE object is in the "Provisioning" state for a short period.

2. Create a dynamic routing gateway (DRG):
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
  - b. Click **Create Dynamic Routing Gateway**.
  - c. For **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - d. Enter a friendly name for the DRG. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - e. Click **Create**.

The DRG will be in the "Provisioning" state for a short period. Make sure it is done being provisioned before continuing.

3. Attach the DRG to your VCN:
  - a. Click the DRG that you just created.
  - b. Under **Resources**, click **Virtual Cloud Networks**.
  - c. Click **Attach to Virtual Cloud Network**.

- d. Select the VCN. Ignore the section for advanced options, which is only for an advanced routing scenario called *transit routing*, which is not relevant here.
- e. Click **Attach**.

The attachment will be in the "Attaching" state for a short period before it's ready.

4. Update the private subnet's route table (which already has one rule for the NAT gateway):
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click your VCN.
  - c. Click **Route Tables**, and then click the Private Subnet Route Table you created earlier.
  - d. Click **Add Route Rule**.
  - e. Enter the following:
    - **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
    - **Destination CIDR Block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule).
  - f. Click **Add Route Rule**.

The table is updated to route any traffic destined for your on-premises network to the DRG. The original rule for 0.0.0.0/0 routes any remaining traffic leaving the subnet to the NAT gateway.
5. Create an IPSec Connection:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.
  - b. Click **Create IPSec Connection**.

- c. Enter the following:
- **Create in Compartment:** Leave the default value (the compartment you're currently working in).
  - **Name:** Enter a friendly name for the IPSec connection. It doesn't have to be unique. Avoid entering confidential information.
  - **Customer-Premises Equipment Compartment:** Leave as is (the VCN's compartment).
  - **Customer-Premises Equipment:** Select the CPE object you created earlier.
  - **Dynamic Routing Gateway Compartment:** Leave as is (the VCN's compartment).
  - **Dynamic Routing Gateway:** Select the DRG that you created earlier.
  - **Static Route CIDR:** Enter at least one static route CIDR (see [Prerequisites](#)). If you need to add another, click **Add Static Route**. You can enter up to 10 static routes, and you can [change the static routes](#) later if you like.
- d. Click **Show Advanced Options** and optionally provide the following items:
- **CPE IKE Identifier:** Oracle defaults to using the public IP address of the CPE. But if your [CPE is behind a NAT device](#), you might need to enter a different value. You can either enter the new value here, or [change the value](#) later.
  - **Tunnel 1** and **Tunnel 2:** Leave as is. Later if you want to use BGP dynamic routing instead of static routing for the VPN tunnels, see [Changing from Static Routing to BGP Dynamic Routing](#).
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
- e. Click **Create IPSec Connection**.

The IPSec connection is created and displayed on the page. It will be in the Provisioning state for a short period.

The displayed tunnel information includes the IP address of the VPN headend and the tunnel's IPSec status (possible values are Up, Down, and Down for Maintenance). At this point, the status is Down. To view the tunnel's shared secret, click the Actions icon (three dots), and then click **View Shared Secret**.

- f. Copy the Oracle VPN IP address and shared secret for each of the tunnels to an email or other location so you can deliver it to the network engineer who will configure the on-premises router.

For more information, see [CPE Configuration](#). You can view this tunnel information here in the Console at any time.

You have now created all the components required for the IPSec VPN. But your network administrator must configure the on-premises router before network traffic can flow between your on-premises network and VCN.

### Task 4: Configure your on-premises router (CPE)

These instructions are for the network administrator.

1. Make sure you have the tunnel configuration information that Oracle provided during VPN setup. See [Task 3: Add an IPSec VPN to your VCN](#).
2. Configure your on-premises router according to the information in [CPE Configuration](#).

If there are already instances in one of the subnets, you can confirm the IPSec connection is up and running by connecting to the instances from your on-premises network. To connect to instances in the public subnet, you must connect to the instance's public IP address.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations:

1. [CreateVcn](#): Make sure to include a DNS label if you want the VCN to use the VCN Resolver (see [DNS in Your Virtual Cloud Network](#)).
2. [CreateInternetGateway](#)
3. [CreateNatGateway](#)
4. [CreateRouteTable](#): Call it to create the Public Subnet Route Table. To enable communication by way of the internet gateway, add a route rule with destination = 0.0.0.0/0, and destination target = the internet gateway you created earlier.
5. [CreateRouteTable](#): Call it again to create the Private Subnet Route Table. To enable communication by way of the NAT gateway, add a route rule with destination = 0.0.0.0/0, and destination target = the NAT gateway you created earlier.
6. First call [GetSecurityList](#) to get the default security list, and then call [UpdateSecurityList](#):
  - Change the existing stateful ingress rules to use your on-premises network's CIDR as the source CIDR, instead of 0.0.0.0/0.
  - If you plan to launch Windows instances, add this stateful ingress rule: Source type = CIDR, source CIDR = your on-premises network on TCP, source port = all, destination port = 3389 (for RDP).
7. [CreateSecurityList](#): Call it to create the Public Subnet Security List with these rules:
  - Stateful ingress: Source type = CIDR, source 0.0.0.0/0 on TCP, source port = all, destination port = 80 (HTTP)
  - Stateful ingress: Source type = CIDR, source 0.0.0.0/0 on TCP, source port = all, destination port = 443 (HTTPS)
  - Stateful egress: Destination type = CIDR, destination CIDR blocks of private subnets on TCP, source port = all, destination port = 1521 (for Oracle databases)
8. [CreateSecurityList](#): Call it again to create the Private Subnet Security List with these rules:
  - Stateful ingress: Source type = CIDR, source CIDR blocks of public subnets on TCP, source port = all, destination port = 1521 (for Oracle databases)

- Stateful ingress: Source type = CIDR, source CIDR blocks of private subnets on TCP, source port = all, destination port = 1521 (for Oracle databases)
  - Stateful egress: Destination type = CIDR, destination CIDR blocks of private subnets on TCP, source port = all, destination port = 1521 (for Oracle databases)
9. [CreateSubnet](#): Call it to create regional public subnet. Include a DNS label for the subnet if you want the [VCN Resolver](#) to resolve hostnames for VNICs in the subnet. Use the Public Subnet Route Table you created earlier. Use both the default security list and the Public Subnet Security List that you created earlier. Use the default set of DHCP options.
  10. [CreateSubnet](#): Call it again to create regional private subnet. Include a DNS label for the subnet if you want the [VCN Resolver](#) to resolve hostnames for VNICs in the subnet. Use the Private Subnet Route Table you created earlier. Use both the default security list and the Private Subnet Security List that you created earlier. Use the default set of DHCP options.
  11. [CreateDrg](#): This creates a new dynamic routing gateway (DRG).
  12. [CreateDrgAttachment](#): This attaches the DRG to the VCN.
  13. [CreateCpe](#): Here you provide the IP address of the router at your end of the VPN (see [Prerequisites](#)).
  14. [CreateIPSecConnection](#): Here you provide the static routes for your on-premises network (see [Prerequisites](#)). In return, you receive the configuration information your network administrator needs in order to configure your router. If you need that information later, you can get it with [GetIPSecConnectionDeviceConfig](#). For more information about the configuration, see [CPE Configuration](#).
  15. First call [GetRouteTable](#) to get the Private Subnet Route Table. Then call [UpdateRouteTable](#) to add a route rule with destination = the on-premises network CIDR (10.0.0.0/16 in this example), and destination target = the DRG you created earlier.
  16. [LaunchInstance](#): Launch at least one instance in each subnet. By default, the instances in the public subnets are assigned public IP addresses. For more information, see [Launching an Instance](#).

You can now communicate from your on-premises network with the instances in the public subnet over the internet gateway.



### Important

Although you can launch instances into the private subnets, you can't communicate with them from your on-premises network until your network administrator configures your on-premises router (see [CPE Configuration](#)). After that, your IPSec connection should be up and running. You can confirm its status by using [GetIPSecConnectionDeviceStatus](#). You can also confirm the IPSec connection is up by connecting to the instances from your on-premises network.

## Transit Routing: Access to Multiple VCNs in the Same Region

*Transit routing* refers to a network setup in which your on-premises network uses a connected virtual cloud network (VCN) to reach Oracle resources or services beyond that VCN. You connect the on-premises network to the VCN with [FastConnect](#) or [VPN Connect](#), and then configure the VCN routing so that traffic *transits through the VCN* to its destination beyond the VCN.

There are two primary transit routing scenarios:

- **Access to multiple VCNs in the same region:** The scenario covered in this topic. This scenario enables communication between your on-premises network and multiple VCNs in the same region over a single FastConnect private virtual circuit or VPN Connect.
- **Private access to Oracle services:** This scenario gives your on-premises network *private access* to Oracle services, so that your on-premises hosts can use their private

IP addresses and the traffic does not go over the internet. See [Transit Routing: Private Access to Oracle Services](#).

### Highlights

- You can use a single FastConnect or IPSec VPN to connect your on-premises network with *multiple* VCNs in the same region, in a *hub-and-spoke* layout.
- The VCNs must be in the same region but can be in different tenancies. For accurate routing, the CIDR blocks of the various subnets of interest in the on-premises network and VCNs must not overlap.
- The VCN that acts as the *hub* uses a [dynamic routing gateway](#) (DRG) to communicate with the on-premises network. This *hub* VCN peers with each VCN that is acting as a spoke (referred to as *spoke* VCNs in this topic). The hub and spoke VCNs use [local peering gateways](#) (LPGs) to communicate.
- To enable the desired traffic from the on-premises network through the hub VCN to a peered spoke VCN, you implement route rules for the hub VCN's DRG attachment and LPG, and for the spoke VCN's subnets.
- If you like, you can set up transit routing *through a private IP in the hub VCN*. For example, you might want to filter or inspect the traffic between the on-premises network and a spoke VCN. In that case, you route the traffic to a private IP on an instance in the hub VCN for inspection, and the resulting traffic continues to its destination. This topic covers both situations: transit routing directly between gateways on the hub VCN, and transit routing through a private IP.
- By configuring route tables that reside in the hub VCN, you can control whether a particular subnet in a peered spoke VCN is advertised to the on-premises network, and whether a particular subnet in the on-premises network is advertised to a peered spoke VCN.



### Tip

There's another scenario that lets you connect your on-premises network to multiple VCNs. Instead of using a single DRG and hub-and-spoke layout, you set up a separate DRG for each VCN and a separate private virtual circuit over a single FastConnect. However, the scenario can be used only with FastConnect through a [third-party provider](#) or through [colocation with Oracle](#). The VCNs must be in the same region and same tenancy. For more information, see [FastConnect with Multiple DRGs and VCNs](#).

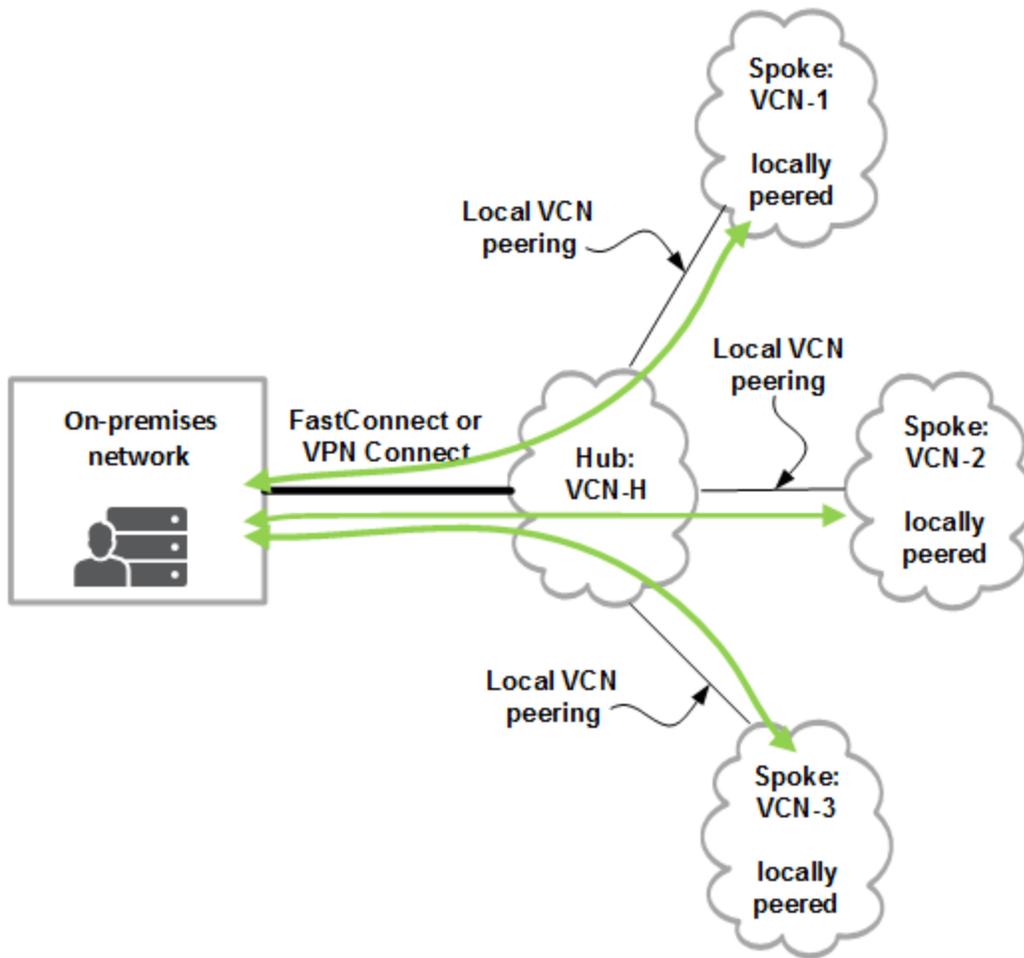
## Overview of Transit Routing

A basic networking scenario involves connecting your on-premises network to a VCN with either Oracle Cloud Infrastructure [FastConnect](#) or an [IPSec VPN](#). These two basic scenarios illustrate that layout: [Scenario B: Private Subnet with a VPN](#) and [Scenario C: Public and Private Subnets with a VPN](#).

There's an advanced networking scenario that lets you use your single FastConnect or IPSec VPN to communication with *multiple VCNs* from your on-premises network. The VCNs must be in the same region but can be in different tenancies.

Here's a basic example of why you might use transit routing: you have a large organization with different departments, each with their own VCN. Your on-premises network needs access to the different VCNs, but you don't want the administration overhead of maintaining a secure connection from each VCN to the on-premises network. Instead you want to use a single FastConnect or IPSec VPN.

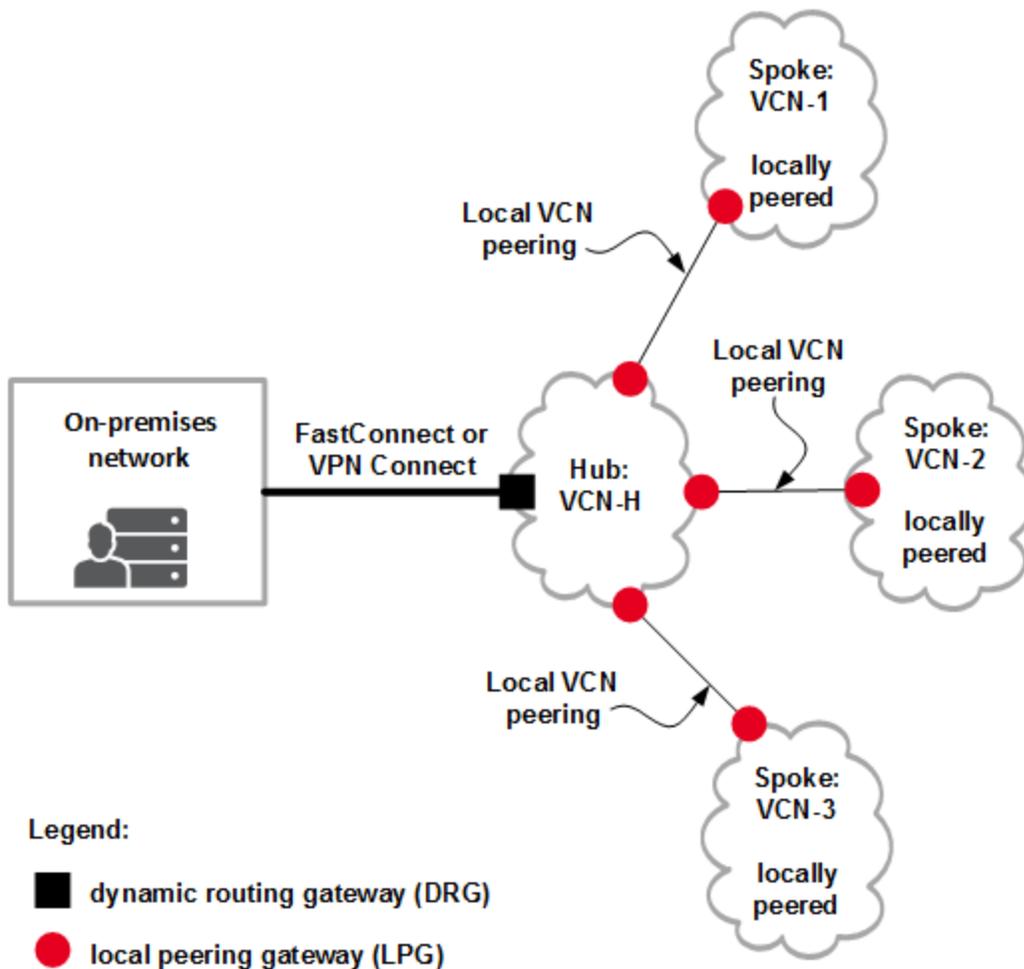
The scenario uses a *hub-and-spoke* layout, as illustrated in the following diagram. The term *hub VCN* here means only that a VCN is acting as the hub in this hub-and-spoke design.



One of the VCNs acts as the hub (VCN-H) and connects to your on-premises network by way of [FastConnect](#) or an [IPSec VPN](#). The other VCNs are [locally peered with the hub VCN](#). The traffic between the on-premises network and the peered VCNs transits through the hub VCN. The VCNs must be in the same region but can be in different tenancies.

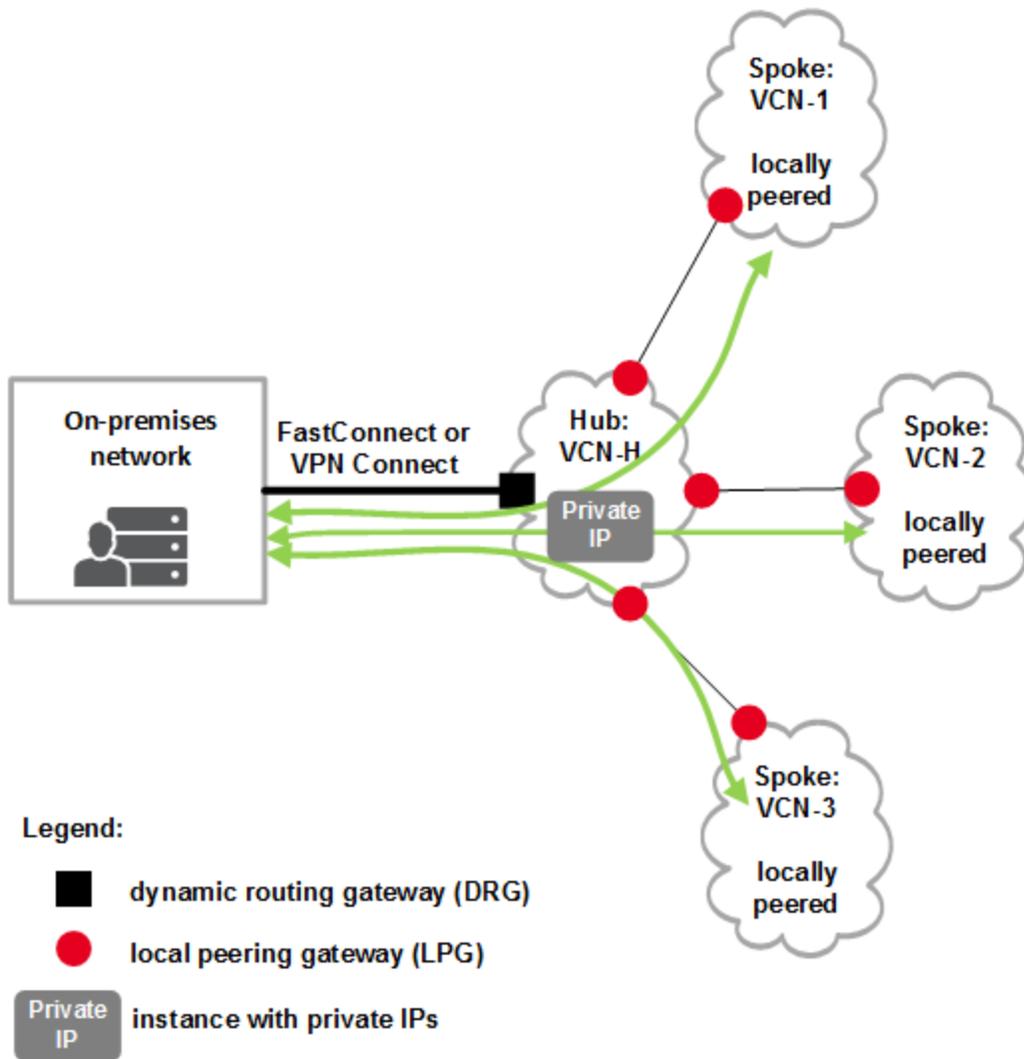
### Gateways Involved in Transit Routing

The next diagram shows the gateways on the VCNs. The hub VCN has a [dynamic routing gateway \(DRG\)](#), which is the communication path with the on-premises network. For each locally peered spoke VCN, there's a pair of [local peering gateways \(LPGs\)](#) that anchor the peering connection. One LPG is on the hub VCN, and the other is on the spoke VCN.



### **Transit Routing Through a Private IP in the Hub VCN**

If you want, you can route all the traffic through a private IP on an instance in the hub VCN. For example, you could set up a firewall or intrusion detection system on the instance in the hub VCN to filter or inspect the traffic between the on-premises network and spoke VCNs. Then you would configure the hub VCN's routing so that the transit traffic is routed to a private IP on that instance. You must [disable the source/destination check](#) for the private IP's VNIC. The resulting traffic would then travel on to its destination in either the spoke VCN or on-premises network. The following diagram illustrates the idea. Regardless of whether you choose to route through a private IP or not, the same gateways are required. However, you configure the hub VCN's routing a little differently.



### Summary of New Concepts for Experienced Networking Service Users

If you're already familiar with the Networking service and local VCN peering, these are the most important new concepts to understand:

- For each spoke VCN subnet that needs to communicate with the on-premises network, you must update the subnet's route table with a rule that sets the target (the next hop) as the spoke VCN's LPG for all traffic destined for the on-premises network.
- You must add a route table to the hub VCN, associate it with *the DRG attachment*, and add a route rule with a target that depends on your situation:
  - **Transit routing directly through gateways:** Set the target (the next hop) to the **hub VCN's LPG (for that spoke)** for all traffic destined for that spoke VCN (or a specific subnet in that VCN).
  - **Transit routing through a private IP:** Set the target (the next hop) to a [private IP on the instance](#), for all traffic destined for that spoke VCN (or a specific subnet in that VCN). Make sure to [disable the source/destination check](#) for the private IP's VNIC.
- You must add another route table to the hub VCN, associate it with *the hub VCN's LPG* (for that spoke), and add a route rule with a target that depends on your situation:
  - **Transit routing directly through gateways:** Set the target (the next hop) as the **DRG** for all traffic destined for the on-premises network (or a specific subnet in that network).
  - **Transit routing through a private IP:** Set the target (the next hop) to *another private IP* on that instance, for all traffic destined for the on-premises network (or a specific subnet in that network). Again, make sure to [disable the source/destination check](#) for the private IP's VNIC. In the example presented here, the private IP is on a *secondary VNIC* on the instance, in a different subnet. In the subsequent diagrams, the subnets are referred to as the *frontend* subnet and *backend* subnet.

For transit routing directly through gateways, see these specific tasks for more information:

- [Task 5: Add a route rule to the spoke VCN's subnet](#)
- [Task 6: Set up ingress routing for the DRG and LPG on the hub VCN](#)

For transit routing through a private IP: see these specific tasks for more information:

- [Task 5: Add a route rule to the spoke VCN's subnet](#)
- [Task 6: Set up the private IPs on an instance in the hub VCN](#)
- [Task 7: Set up ingress routing for the DRG and LPG on the hub VCN](#)

### Example: Components and Routing for a Hub and Single Spoke

The examples in this section show a VCN acting as a hub and only a single spoke VCN for simplicity.

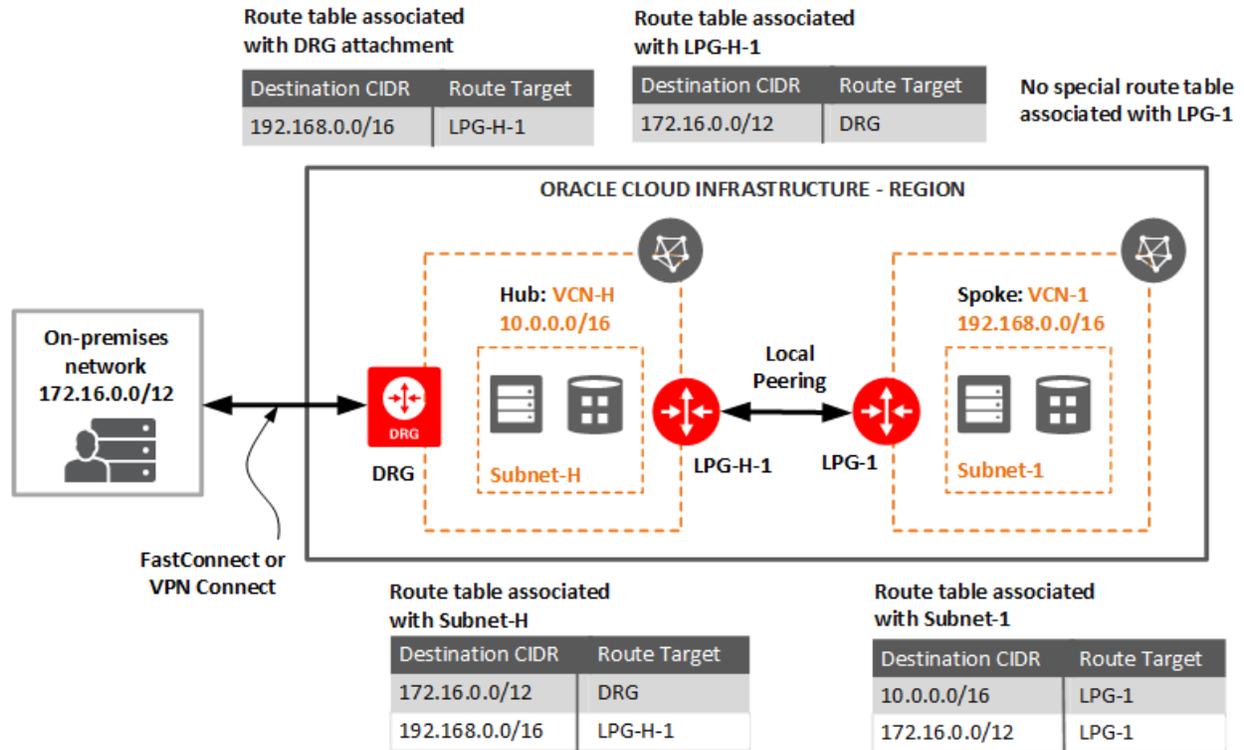


#### Note

In a hub-and-spoke model, the hub VCN can have multiple spokes and therefore multiple LPGs (one per spoke). This topic uses the phrase *the hub VCN's LPG*, which could therefore be ambiguous. When the phrase is used here, it means the hub LPG for the *particular spoke of interest*. In the following diagrams, it's LPG-H-1. Additional spokes would involve creation of an LPG-H-2, LPG-H-3, and so on.

### For transit routing directly through gateways

The following diagram shows the required Networking service route tables and route rules for transit routing directly through gateways. Although the hub VCN does not require a subnet to make transit routing work, the example presented here includes a subnet called Subnet-H.



The diagram shows four route tables, each associated with a different resource:

• **DRG attachment:**

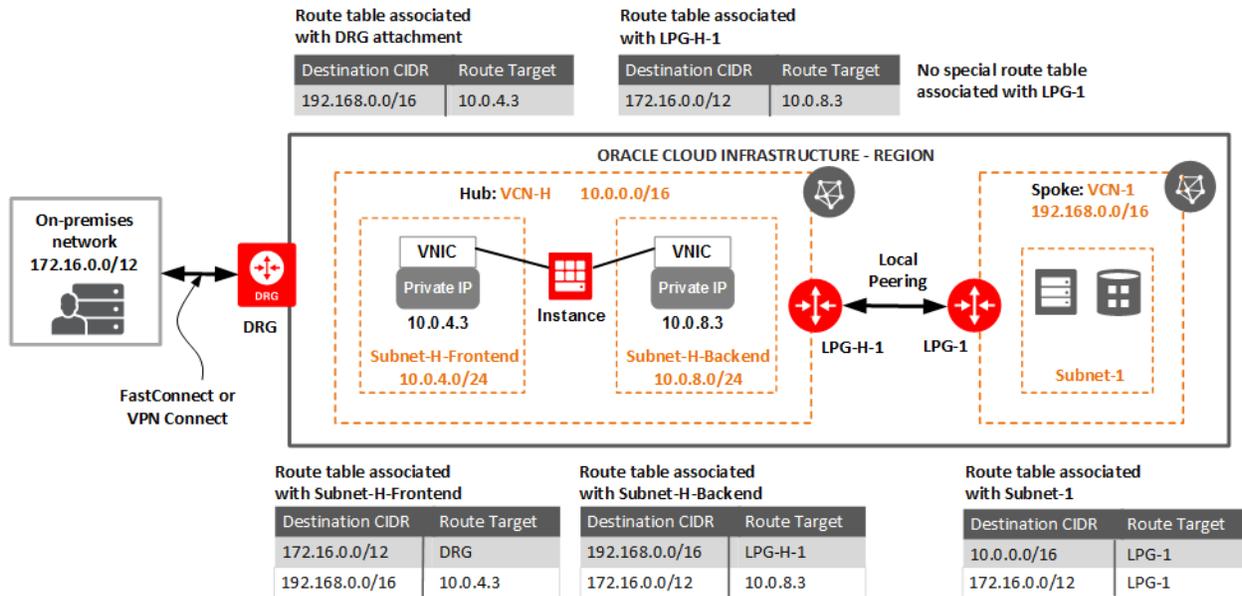
- The route table belongs to the hub VCN and is associated with the DRG *attachment*. Why the attachment and not the DRG itself? Because the DRG is a standalone resource that you can attach to any VCN in the same region and tenancy as the DRG. The attachment itself identifies which VCN.
- The route table routes the inbound traffic that is from the on-premises network and destined for the spoke VCN (VCN-1). You configure the rule to send that traffic to LPG-H-1.

• **LPG-H-1:**

- This route table belongs to the hub VCN and is associated with LPG-H-1.
- The route table routes inbound traffic that is from VCN-1 and destined for the on-premises network. You configure the rule to send that traffic to the DRG.
- **Subnet-H:**
  - This route table belongs to the hub VCN and is associated with Subnet-H.
  - This route table has a rule to route traffic that is destined for the on-premises network to the DRG. It has another rule to route traffic that is destined for the spoke VCN to LPG-H-1.
- **Subnet-1:**
  - This route table belongs to the spoke VCN and is associated with Subnet-1.
  - This route table has rules to route traffic that is destined for the hub VCN or the on-premises network to LPG-1.

### For transit routing through a private IP

The following diagram shows the required Networking service route tables and route rules for transit routing through a private IP on an instance in the hub VCN. You can choose to implement this scenario with either a single VNIC or multiple VNICs. The diagram and example here shows two VNICs: one in a subnet called Subnet-H-Frontend, and another in a subnet called Subnet-H-Backend. The frontend VNIC has private IP 10.0.4.3, and the backend VNIC has private IP 10.0.8.3.



The diagram shows five route tables, each associated with a different resource:

• **DRG attachment:**

- The route table belongs to the hub VCN and is associated with the DRG *attachment*. Why the attachment and not the DRG itself? Because the DRG is a standalone resource that you can attach to any VCN in the same region and tenancy as the DRG. The attachment itself identifies which VCN.
- The route table routes the inbound traffic that is from the on-premises network and destined for the spoke VCN (VCN-1). You configure the rule to send the traffic to the private IP in the frontend subnet.

• **LPG-H-1:**

- This route table belongs to the hub VCN and is associated with LPG-H-1.
- The route table routes inbound traffic that is from VCN-1 and destined for the on-premises network. You configure the rule to send that traffic to the private IP in the backend subnet.

- **Subnet-H-Frontend:**

- This route table belongs to the hub VCN and is associated with Subnet-H-Frontend.
- This route table has a rule to route traffic that is destined for the on-premises network to the DRG.
- Although Oracle does not recommend putting workloads in the hub VCN's subnets, the diagram also shows a route rule to route traffic that is destined for the spoke VCN to the private IP in the frontend subnet (10.0.4.3) for filtering by the instance. The second rule is shown here to give a more complete picture of routing for this example.

- **Subnet-H-Backend:**

- This route table belongs to the hub VCN and is associated with Subnet-H-Backend.
- This route table has a rule to route traffic that is destined for the spoke VCN (VCN-1) to LPG-H-1.
- Although Oracle does not recommend putting workloads in the hub VCN's subnets, the diagram also shows a route rule to route traffic that is destined for the on-premises network to the private IP in the backend subnet (10.0.8.3) for filtering by the instance. The second rule is shown here to give a more complete picture of routing for this example.

- **Subnet-1:**

- This route table belongs to the spoke VCN and is associated with Subnet-1.
- This route table has rules to route traffic that is destined for the hub VCN or the on-premises network to LPG-1.

### **Important Transit Routing Restrictions to Understand**

This section includes some additional important details about routing:

- **Route table for the DRG attachment:**

- A route table that is associated with a DRG attachment can have only rules that target an LPG or a private IP.
- A DRG attachment can exist without a route table associated with it. However, after you associate a route table with a DRG attachment, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

- **Route table for an LPG:**

- A route table that is associated with an LPG can have only rules that target a DRG or a private IP.
- An LPG can exist without a route table associated with it. However, after you associate a route table with an LPG, there must always be a route table associated with it. But, you can associate a different route table. You can also edit the table's rules, or delete some or all of the rules.

- **Traffic *through* the hub VCN:** The route tables discussed here are intended only for moving traffic *through* the hub VCN between locations in the on-premises network and locations in the spoke VCN. If you're using a private IP in the hub, you configure those route tables so that the private IP is placed in that traffic path going *through* the hub.
- **Inbound traffic to the hub VCN:** Even though the preceding statement is true (about traffic *through* the hub), inbound traffic to subnets *within the hub VCN* is always allowed. You do not need to set up explicit rules for this inbound traffic in the DRG attachment's route table or hub LPG's route table. When this kind of inbound traffic reaches the DRG or the hub LPG, the traffic is automatically routed to its destination in the hub VCN by the *VCN local routing*. Because of VCN local routing, for any route table belonging to a given VCN, you can't create a rule that lists that VCN's CIDR (or a subsection) as the rule's destination.
- **Hub VCN traffic when transit routing through a private IP:** The immediately preceding statement about VCN local routing means that you should use the hub VCN only for *transit* between the on-premises network and spoke VCNs. **Do not set up workloads in the hub VCN itself.** More explicitly, if you set up transit routing through a private IP in the hub VCN, you can't also route the *hub VCN's* traffic through

that private IP. For example, in the preceding diagram, if you were to change the route rule in the LPG-H-1 route table so that the destination CIDR is 0.0.0.0/0 instead of 172.16.0.0/12, only traffic coming from VCN-1 and destined for addresses *outside* the hub VCN's CIDR block would be routed through the private IP. Because of VCN local routing, any traffic destined for addresses within the VCN is automatically routed directly to the destination IP address. The VCN local routing takes precedence over the LPG-H-1 route table (in general, over *any* of the VCN's route tables).

### About CIDR Overlap

In this example, the various networks do not have overlapping CIDR blocks (172.16.0.0/12 versus 10.0.0.0/16 versus 192.168.0.0/16). The Networking service does not allow local VCN peering between two VCNs with overlapping CIDRs. That means each spoke must not overlap with the hub.

However, the Networking service does not validate whether the spoke VCNs themselves overlap with each other, or if any of the VCNs overlap with the on-premises network. You must ensure that CIDRs for all the subnets that need to communicate with each other don't overlap. Otherwise, traffic may be dropped.

A Networking service route table cannot contain two rules with the exact same destination CIDR. However, if two rules in the same route table have overlapping destination CIDRs, the most specific rule in the table is used to route the traffic (that is, the rule with the [longest prefix match](#)).

### Route Advertisement to the On-Premises Network and Spoke VCNs

From a security standpoint, you can control route advertisement so that only specific subnets in the on-premises network are advertised to the spoke VCNs. Similarly, you can control which subnets in the spoke VCNs are advertised to the on-premises network.

The routes advertised to the on-premises network consist of:

- The rules listed in the route table associated with the DRG attachment (192.168.0.0/16 in the preceding diagram)
- The individual subnets in the hub VCN

The routes advertised to the spoke VCN consist of:

- The individual subnets in the hub VCN
- The rules listed in the route table associated with the hub VCN's LPG for the spoke (172.16.0.0/12 in the preceding diagram)

Therefore, the administrator *of the hub VCN* alone can control which routes are advertised to the on-premises network and spoke VCNs.

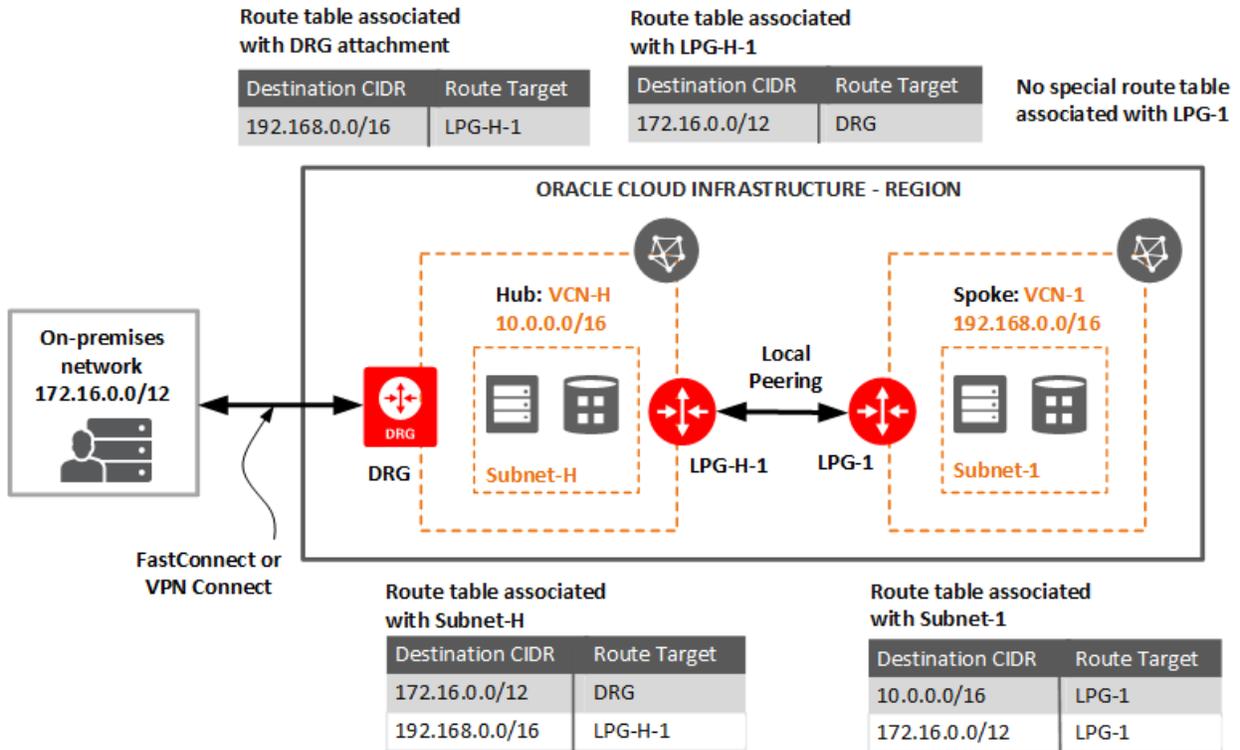
In the preceding example, the relevant routes use the full CIDR block of the on-premises network and spoke VCN as the destination (172.16.0.0/12 and 192.168.0.0/16, respectively), but they could instead use a subnet of those networks to restrict routing to specific subnets.

### **Details About Routing for Different Traffic Paths**

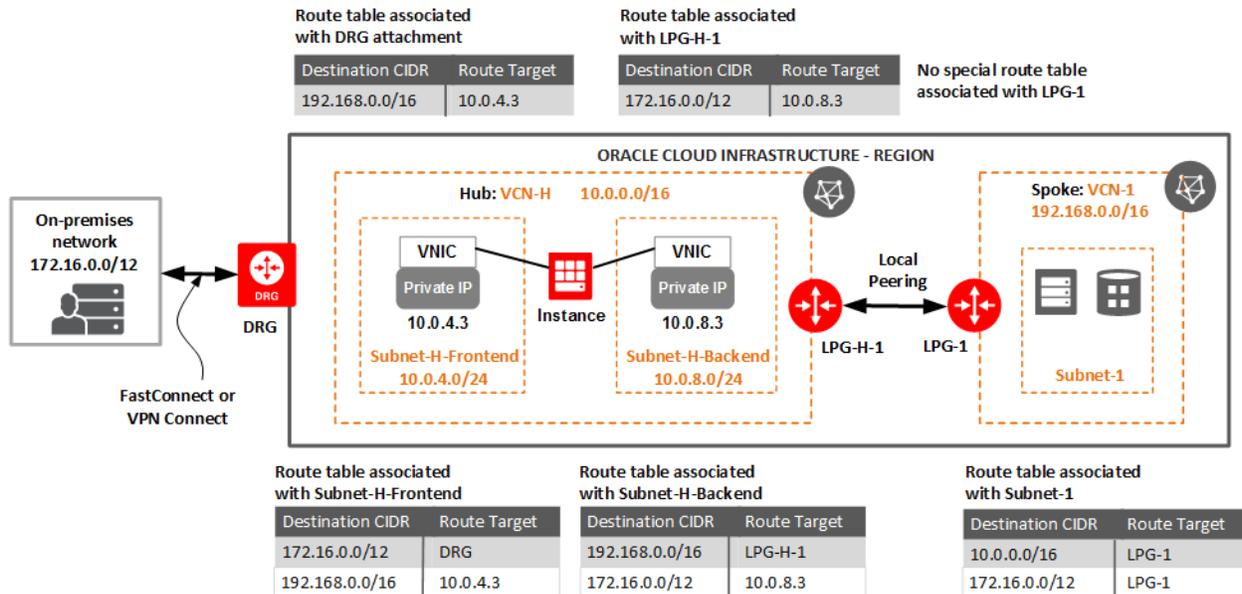
To further illustrate how routing takes place in the preceding example, let's look more closely at different paths of traffic. Here are the same diagrams again.

First, if you are transit routing directly through gateways on the hub VCN:

## CHAPTER 23 Networking



Second, if you are transit routing through a private IP in the hub VCN:



### Traffic from the on-premises network to the spoke VCN

1. Traffic leaves the on-premises network and reaches the DRG. The traffic's destination is in Subnet-1 (for example, 192.168.0.5).
2. The DRG attachment's associated route table has a rule for 192.168.0.0/16. It matches the destination and sends the traffic to the route target:
  - **Transit routing directly through gateways:** The rule's target is LPG-H-1.
  - **Transit routing through a private IP:** The rule's target is the private IP 10.0.4.3. The instance receives and processes the traffic and sends any resulting traffic out of the backend subnet's VNIC. The backend subnet's route table sends that traffic to LPG-H-1.

Remember that you can use the rules in the DRG attachment's route table to control which subnets in the spoke VCN are advertised to the on-premises network. You could instead set up the rule to list only a subnet of the spoke VCN.

3. LPG-H-1 receives the traffic.
4. Egress traffic leaving a VCN through an LPG is automatically routed to the LPG's peered LPG, which is LPG-1 in this situation. That routing occurs automatically because of the peering connection between the two LPGs.
5. LPG-1 receives the traffic.
6. Traffic coming in to a VCN through the LPG is automatically routed to the destination within the VCN because of VCN local routing. No explicit route rules are required.

### Traffic from the spoke VCN to the on-premises network

1. Traffic comes from an instance in Subnet-1 in the spoke VCN. The traffic's destination is in the on-premises network (for example, 172.16.0.3).
2. Subnet-1's associated route table has a rule for 172.16.0.0/12. It matches the destination and sends the traffic to the route target, LPG-1.
3. LPG-1 receives the traffic.
4. Egress traffic leaving a VCN through an LPG is automatically routed to the LPG's peered LPG, which is LPG-H-1 in this situation. That routing occurs automatically because of the peering connection between the two LPGs.
5. LPG-H-1 receives the traffic.
6. LPG-H-1's associated route table has a rule for 172.16.0.0/12. It matches the destination and sends the traffic to the route target:
  - **Transit routing directly through gateways:** The rule's target is the DRG.
  - **Transit routing through a private IP:** The rule's target is the private IP 10.0.8.3. The instance receives and processes the traffic and sends any resulting traffic out of the frontend subnet's VNIC. The frontend subnet's route table sends that to the DRG.

Remember that you can use the rules in the LPG's route table to control which subnets in the on-premises network are advertised to the spoke VCN. You could instead set up the rule to list only a subnet of the on-premises network.

7. The DRG receives the traffic.
8. Egress traffic leaving the VCN through the DRG is routed based on the IPsec VPN and FastConnect configuration. No explicit rules in the DRG attachment's route table are required.

Notice that Subnet-1 in the spoke VCN and LPG-H-1 both have route rules with 172.16.0.0/12 as the destination CIDR. Those rules don't have to use the exact same CIDR block. However, make sure both rules cover the traffic you want to route from the spoke to the on-premises network. The rule in Subnet-1's route table controls which traffic is routed from Subnet-1 to LPG-H-1. The rule in LPG-H-1's route table controls which traffic is routed from the spoke VCN to the on-premises network. If LPG-H-1's route rule is more restrictive than Subnet-1's route rule, some traffic leaving the subnet could ultimately be dropped and not reach the DRG.

### Traffic from the spoke VCN to a subnet in the hub VCN (routing directly between gateways only)

Depending on your situation, you might want to enable traffic between instances in the hub VCN and a spoke VCN, and not just traffic between the on-premises network and a spoke VCN. You can do this if you're routing directly between gateways. **You can't route the traffic from a spoke VCN through the private IP and on to other instances in the hub VCN.** The note at the end of this section explains why.

Here's how traffic would flow from the spoke VCN to a destination with an address in the hub VCN:

1. Traffic comes from an instance in Subnet-1 in the spoke VCN. The traffic's destination is in a subnet in the hub VCN (for example, 10.0.0.3).
2. Subnet-1's associated route table has a rule for 10.0.0.0/16. It matches the destination and sends the traffic to the route target, LPG-1.
3. LPG-1 receives the traffic.
4. Egress traffic leaving a VCN through an LPG is automatically routed to the LPG's peered LPG, which is LPG-H-1 in this situation. That routing occurs automatically because of the peering connection between the two LPGs.

5. LPG-H-1 receives the traffic.
6. Traffic coming in to a VCN through an LPG and destined for an address in the VCN is automatically routed to the destination by VCN local routing. No explicit route rules are required.

A similar series of routing steps occurs for traffic going from Subnet-H to Subnet-1, but in the reverse direction. Subnet-H's route table has a rule that matches the spoke VCN's CIDR (192.168.0.0/16) and sends the traffic to LPG-H-1, which forwards it on to LPG-1.



### Note

If you set up transit routing through a private IP in the hub VCN, remember that the LPG-H-1 route table only controls routing of traffic that is destined for addresses *outside the hub VCN*. Traffic destined for addresses within the VCN is instead handled by the hub VCN local routing, which takes precedence and always routes the traffic directly to the packet's destination address. This means that you cannot route traffic that is destined for addresses inside the hub VCN *through the private IP* that is being used for the transit traffic through the hub. Even if the LPG-H-1 route rule uses a destination = 0.0.0.0/0 and target = 10.0.8.3, the hub VCN local routing takes precedence and routes the traffic directly to the destination in the hub VCN instead of the private IP.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have

permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're a member of the Administrators group, you already have the required access to set up transit routing. Otherwise, you need access to the Networking service, and you need the ability to launch instances. See [IAM Policies for Networking](#).

### Setting Up VCN Transit Routing in the Console

This section shows how to use the Console to set up transit routing with a VCN to give your on-premises network access to multiple VCNs in the same region.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### For routing directly between gateways



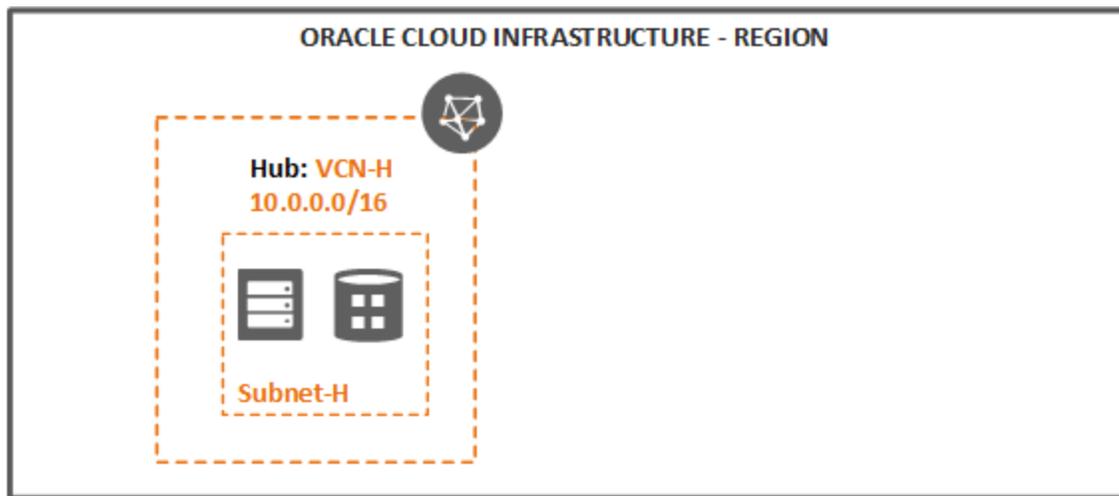
#### Tip

You might already have many of the necessary Networking components and connections in this advanced scenario already set up. So you might be able to skip some of the following tasks. **If you already have a network layout with a hub VCN connected to your on-premises network, and spoke VCNs locally peered with the hub VCN, then Task 5 and Task 6 are the most important.** They enable traffic



to be routed between your on-premises network and the spoke VCN.

### Task 1: Set up the hub VCN

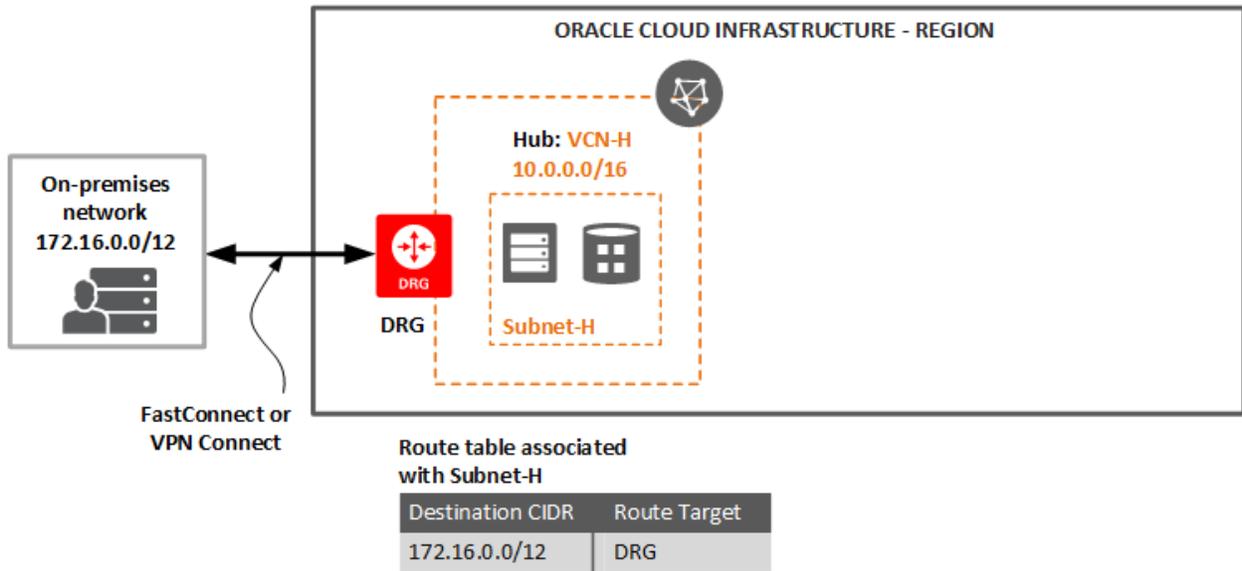


In this task, you set up the hub VCN. A subnet in the hub VCN is optional. However, this example includes one. The subnet can contain cloud resources that your on-premises network or the spoke VCN need to use.

For more information and instructions:

- [VCNs and Subnets](#)

Task 2: Connect the hub VCN with your on-premises network



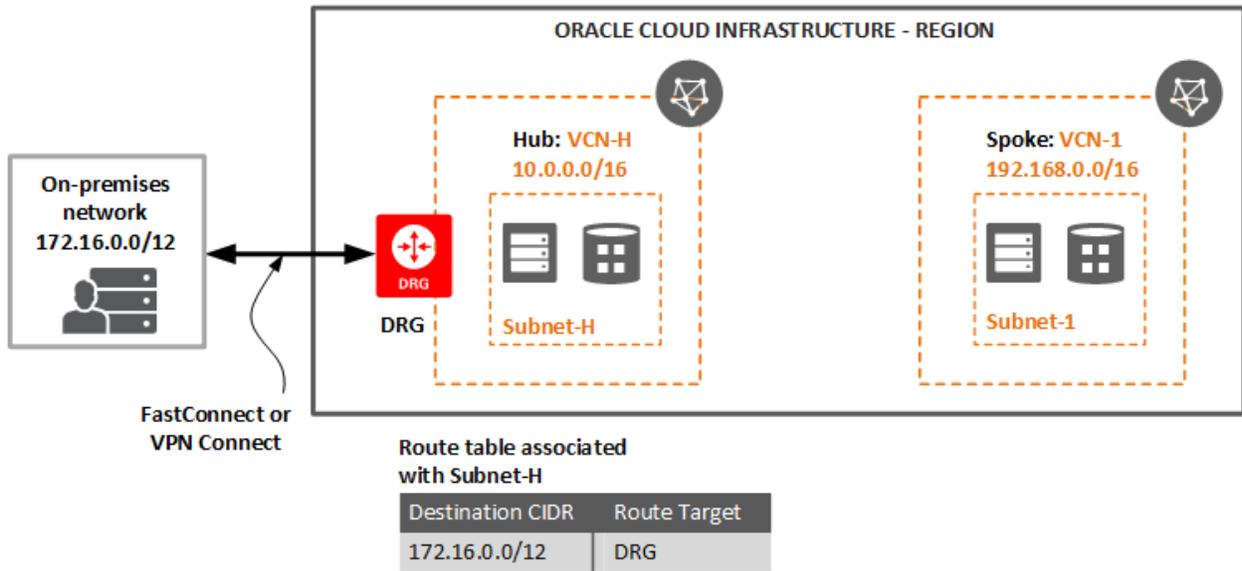
In this task, you set up either FastConnect or VPN Connect between your hub VCN and your on-premises network. As part of this process, you attach a DRG to the hub VCN and set up routing between the hub VCN and your on-premises network.

Notice that you do not yet create the route table that will be associated with the DRG attachment. That comes in a later step.

For more information and instructions:

- [FastConnect](#)
- [VPN Connect](#)
- [Dynamic Routing Gateways \(DRGs\)](#)

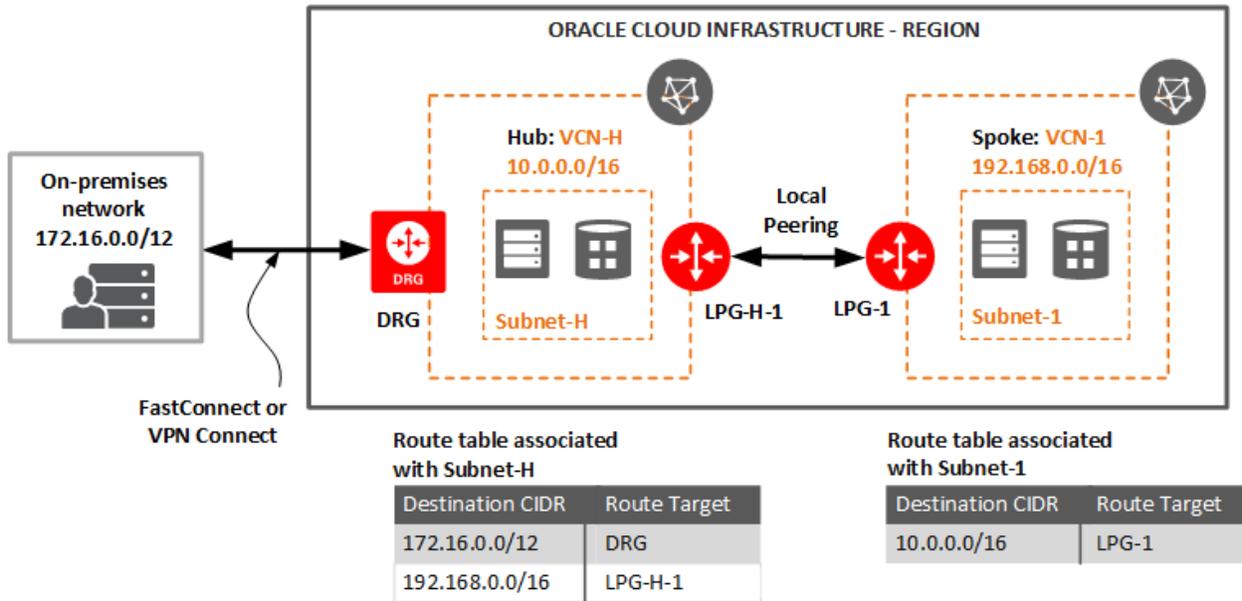
Task 3: Set up a spoke VCN with at least one subnet



In this task, you set up the spoke VCN with at least one subnet. For more information and instructions:

- [VCNs and Subnets](#)

Task 4: Set up a local peering between the hub VCN and the spoke VCN



In this task, you add an LPG to each VCN, establish a connection between the LPGs, and set up routing that enables resources in one VCN to communicate with resources in the other.

**✓ Important**

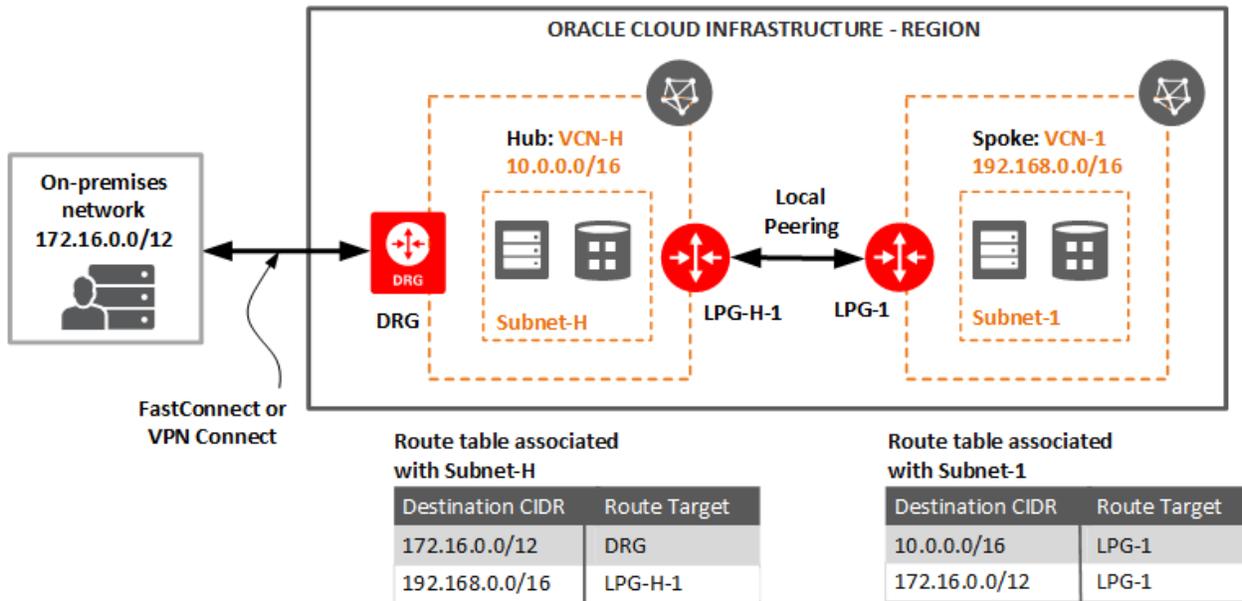
When setting up local peering between two VCNs, make sure to [establish the connection between the LPGs](#). It can be easy to overlook that part of the process.

Notice that you do not yet create the route table that will be associated with the LPG on the hub VCN (LPG-H-1). That comes in a later step.

For more information and instructions:

- [Setting Up a Local Peering](#)

Task 5: Add a route rule to the spoke VCN's subnet

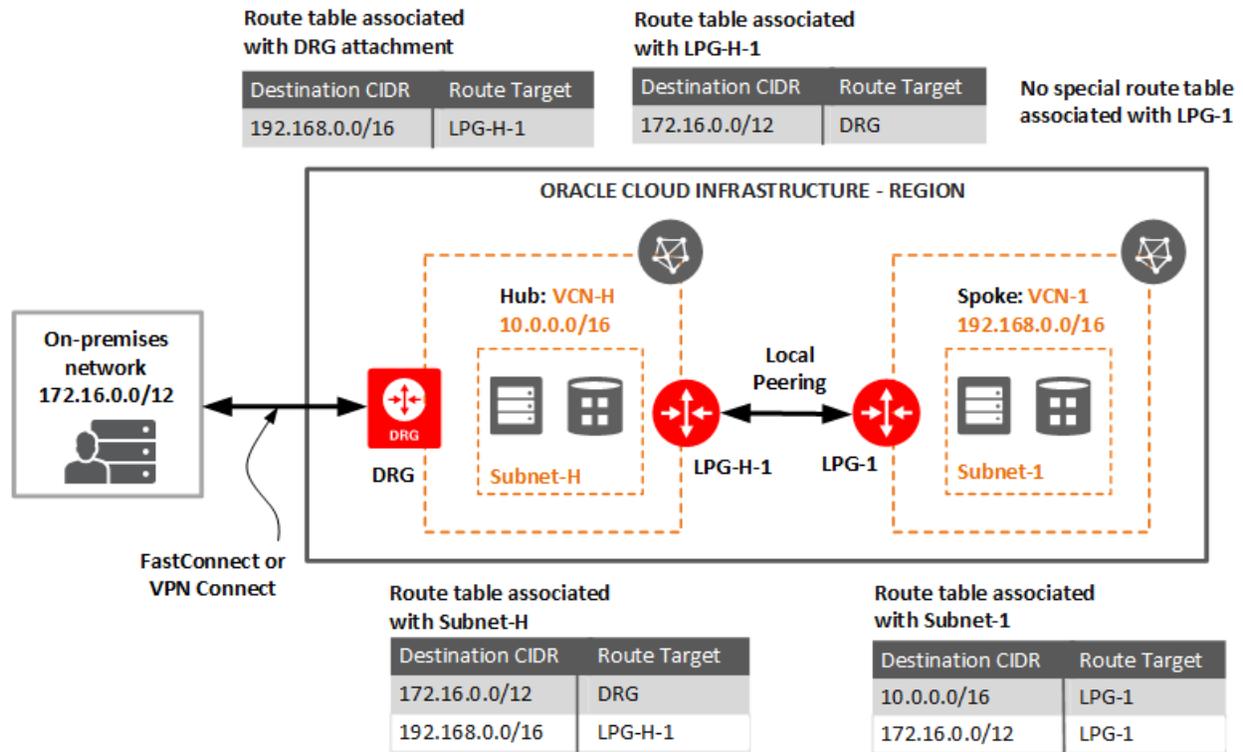


In this task, you add a rule to the route table associated with the spoke VCN's subnet. This rule routes traffic that is destined for the on-premises network to the spoke VCN's LPG (LPG-1 in the diagram).

Prerequisites: You already have an LPG for the spoke VCN, and a route table associated with the subnet (on the spoke VCN) that needs to communicate with the on-premises network.

1. For the spoke VCN, view the list of subnets.
2. For the subnet of interest, look at its details and click the link for its associated route table.
3. Edit the route table to include a rule that sends traffic to the on-premises network:
  - a. Click **Add Route Rules**.
  - b. Enter this information for the route rule:
    - **Target Type:** Local Peering Gateway.
    - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example).
    - **Compartment:** The compartment where the spoke VCN's LPG is located.
    - **Target Local Peering Gateway:** The spoke VCN's LPG.
  - c. Click **Add Route Rules**.

Task 6: Set up ingress routing for the DRG and LPG on the hub VCN



In this task, you set up the route tables for the DRG attachment and hub VCN's LPG for the spoke of interest (LPG-H-1).

Prerequisites:

- You already have a DRG attached to the hub VCN.
- You already have a hub VCN LPG for the spoke of interest.

1. Create a route table for the DRG attachment:
  - a. In the Console, view the hub VCN's details.
  - b. Under **Resources**, click **Route Tables** to view the VCN's route tables.
  - c. Click **Create Route Table**.
  - d. Enter the following:
    - **Name:** A descriptive name for the route table. Example: `DRG Route Table`. Avoid entering confidential information.
    - **Create in Compartment:** Leave as is.
  - e. Click **+ Additional Route Rule**, and enter this information for the route rule:
    - **Target Type:** Local Peering Gateway.
    - **Destination CIDR Block:** This spoke VCN's CIDR (192.168.0.0/16 in the earlier example). Remember that you can use the routes in this table to control which subnets in the spoke VCN are advertised to the on-premises network. You could instead set up the rule to list only a particular subnet of the spoke VCN that the on-premises network.
    - **Compartment:** The compartment where the hub VCN's LPG is located.
    - **Target:** The hub VCN's LPG.
  - f. Click **Create Route Table**.

The route table is created and displayed in the list.
2. Associate the route table (called *DRG Route Table* in this example) with the hub VCN's DRG attachment:
  - a. While still viewing the hub VCN's details, click **Dynamic Routing Gateways** to view the attached DRG.
  - b. Click the Actions icon (three dots), and then click **Associate Route Table**.
  - c. Select the route table.

- d. Click **Associate Route Table**.

The route table is associated with the DRG attachment.

3. Create a route table for the hub VCN's LPG for this spoke:

- a. While still viewing the hub VCN's details, click **Route Tables**.

- b. Click **Create Route Table**.

- c. Enter the following:

- **Create in Compartment:** Leave as is.
- **Name:** A descriptive name for the route table. Example: `Hub LPG-# Route Table` (where # indicates which spoke). Avoid entering confidential information.

- d. Click **+ Additional Route Rule**, and enter this information for the route rule:

- **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
- **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example). Remember that you can use the routes in this table to control which subnets in the on-premises network are advertised to this spoke VCN. You could instead set up the rule to list only a subnet of the on-premises network that needs to communicate with this spoke.

- e. Click **Create Route Table**.

The route table is created and displayed in the list.

4. Associate the route table (called *Hub LPG-# Route Table* in this example) with the hub VCN's LPG for the spoke of interest:

- a. While still viewing the hub VCN's details, click **Local Peering Gateways** to view the hub VCN's LPG for this spoke.

- b. For the LPG you're interested in, click the Actions icon (three dots), and then click **Associate With Route Table**.
- c. Enter the following:
  - **Route Table Compartment:** Select the compartment of the route table for the LPG.
  - **Route Table:** Select the route table for the LPG.
- d. Click **Associate**.

The route table is associated with the LPG.

### Later if you need more spoke VCNs

1. Repeat Tasks 3-5 for the new spoke VCN.
2. Repeat Task 6 with these changes:
  - For Step 1: Instead of creating a new route table for the DRG attachment, update the existing route table to include a new rule for the new spoke VCN. The destination CIDR is the spoke VCN's CIDR (or a subnet within). The target is the hub VCN's LPG for the new spoke.
  - For Step 2: Skip this step entirely because the DRG attachment is already associated with its route table.
  - For Step 3: Repeat as is. Name the new route table according to which spoke the route table is for (for example, **Hub LPG-2 Route Table** for the second spoke).
  - For Step 4: Repeat as is. Associate the new route table you created in Step 3 with the hub VCN's LPG for the new spoke.

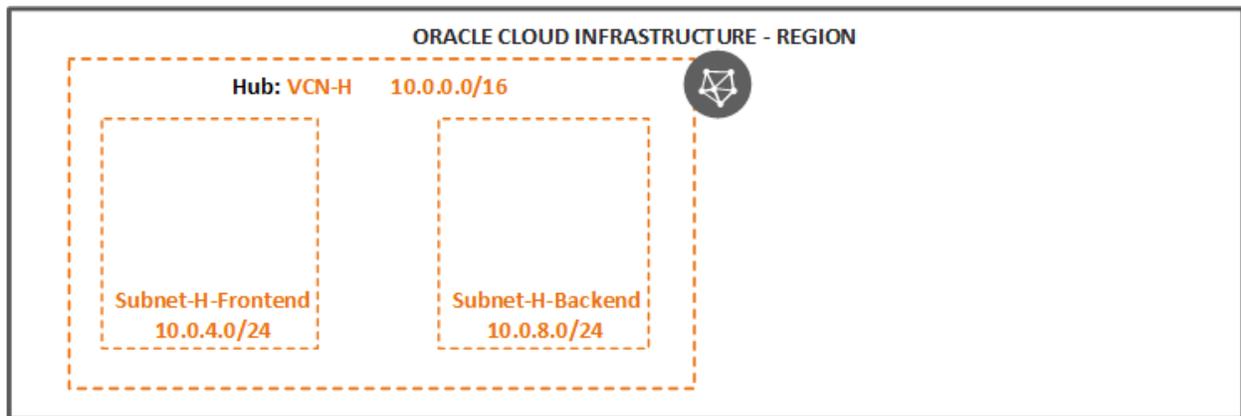
## For routing through a private IP



### Tip

You might already have many of the necessary Networking components and connections in this advanced scenario already set up. So you might be able to skip some of the following tasks. **If you already have a network layout with a hub VCN connected to your on-premises network, and spoke VCNs locally peered with the hub VCN, then Tasks 5 through 7 are the most important.** They enable traffic to be routed between your on-premises network and the spoke VCN.

## Task 1: Set up the hub VCN

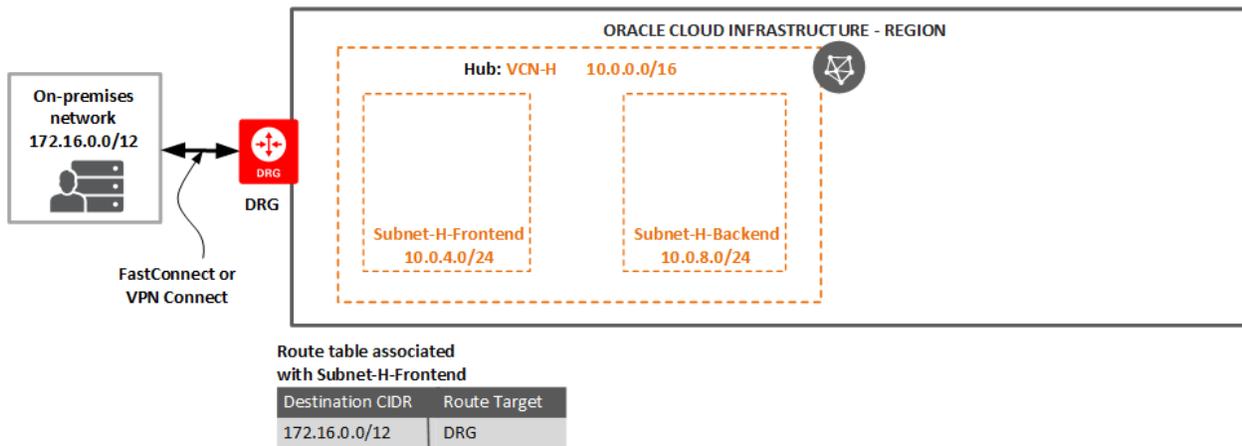


In this task, you set up the hub VCN. The hub VCN must have two subnets: one for the frontend VNIC on the instance, and one for the backend VNIC on the instance. Oracle recommends using regional *private* subnets, unless there will be resources in the frontend subnet that need internet access.

For more information and instructions:

- [VCNs and Subnets](#)

### Task 2: Connect the hub VCN with your on-premises network



In this task, you set up either FastConnect or VPN Connect between your hub VCN and your on-premises network. As part of this process, you attach a DRG to the hub VCN and set up routing between the hub VCN and your on-premises network.

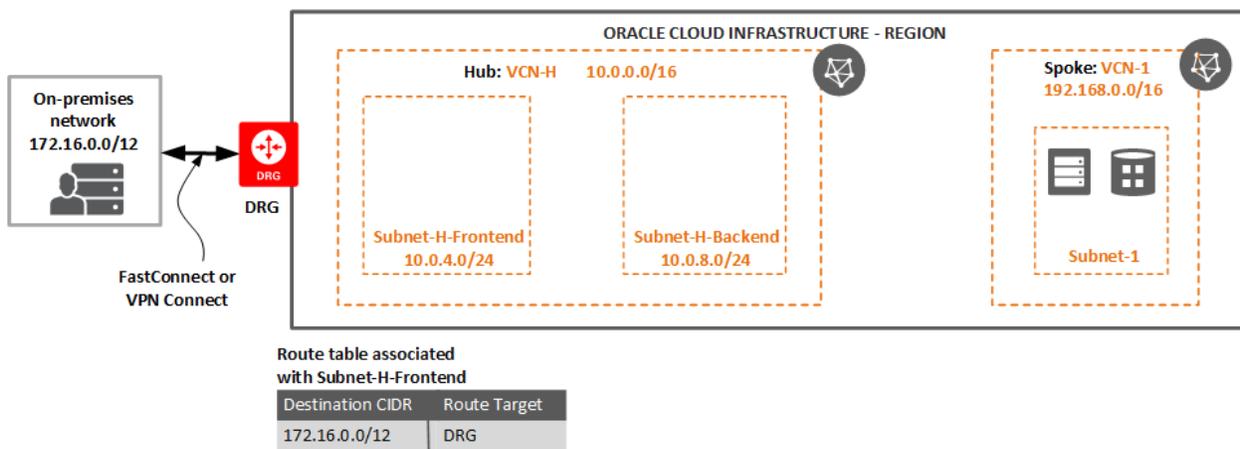
Notice that you do not yet create the route table that will be associated with the DRG attachment. That comes in a later step.

For more information and instructions:

## CHAPTER 23 Networking

- [FastConnect](#)
- [VPN Connect](#)
- [Dynamic Routing Gateways \(DRGs\)](#)

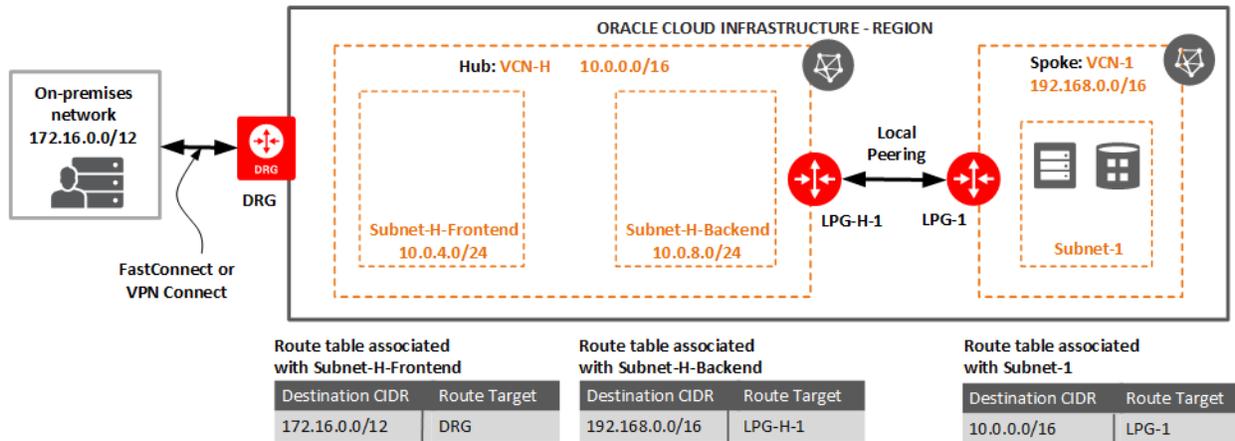
### Task 3: Set up a spoke VCN with at least one subnet



In this task, you set up the spoke VCN with at least one subnet. For more information and instructions:

- [VCNs and Subnets](#)

Task 4: Set up a local peering between the hub VCN and the spoke VCN



In this task, you add an LPG to each VCN, establish a connection between the LPGs, and set up routing that enables resources in one VCN to communicate with resources in the other.

**Important**

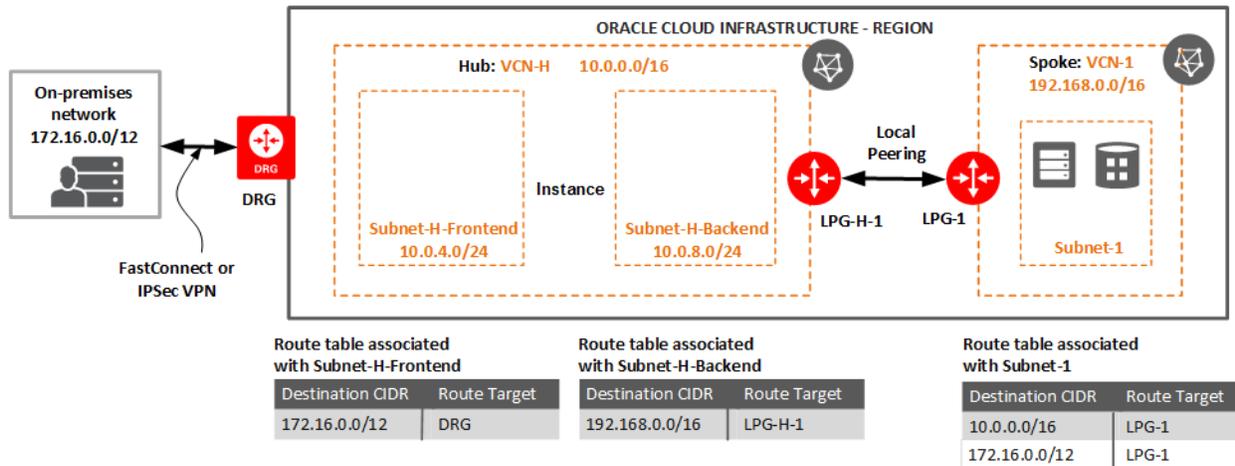
When setting up local peering between two VCNs, make sure to [establish the connection between the LPGs](#). It can be easy to overlook that part of the process.

Notice that you do not yet create the route table that will be associated with the LPG on the hub VCN (LPG-H-1). That comes in a later step.

For more information and instructions:

- [Setting Up a Local Peering](#)

Task 5: Add a route rule to the spoke VCN's subnet



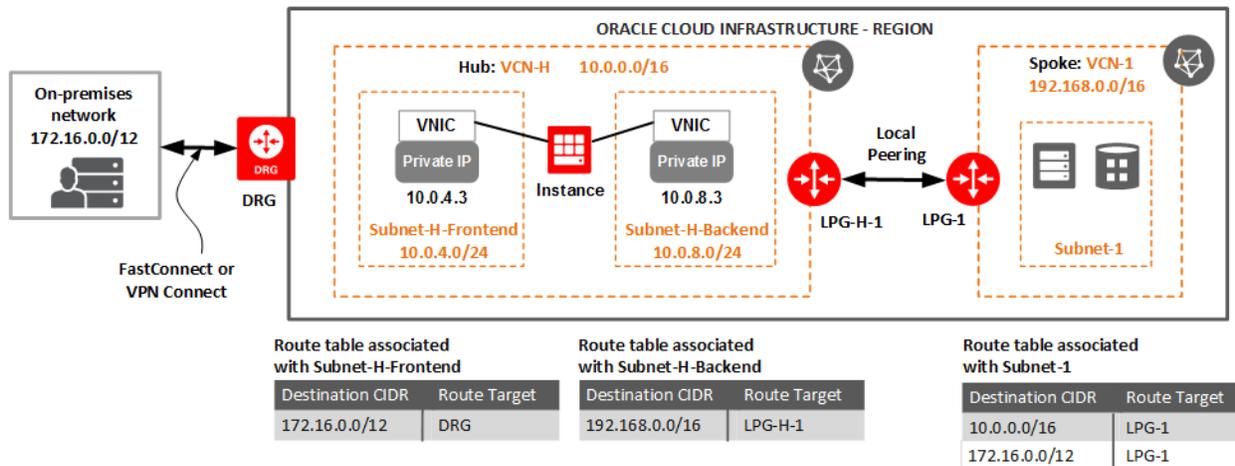
In this task, you add a rule to the route table associated with the spoke VCN's subnet. This rule routes traffic that is destined for the on-premises network to the spoke VCN's LPG (LPG-1 in the diagram).

Prerequisites: You already have an LPG for the spoke VCN, and a route table associated with the subnet (on the spoke VCN) that needs to communicate with the on-premises network.

1. For the spoke VCN, view the list of subnets.
2. For the subnet of interest, look at its details and click the link for its associated route table.
3. Edit the route table to include a rule that sends traffic to the on-premises network:
  - a. Click **Add Route Rules**.
  - b. Enter this information for the route rule:

- **Target Type:** Local Peering Gateway.
  - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example).
  - **Compartment:** The compartment where the spoke VCN's LPG is located.
  - **Target Local Peering Gateway:** The spoke VCN's LPG.
- c. Click **Add Route Rules**.

Task 6: Set up the private IPs on an instance in the hub VCN



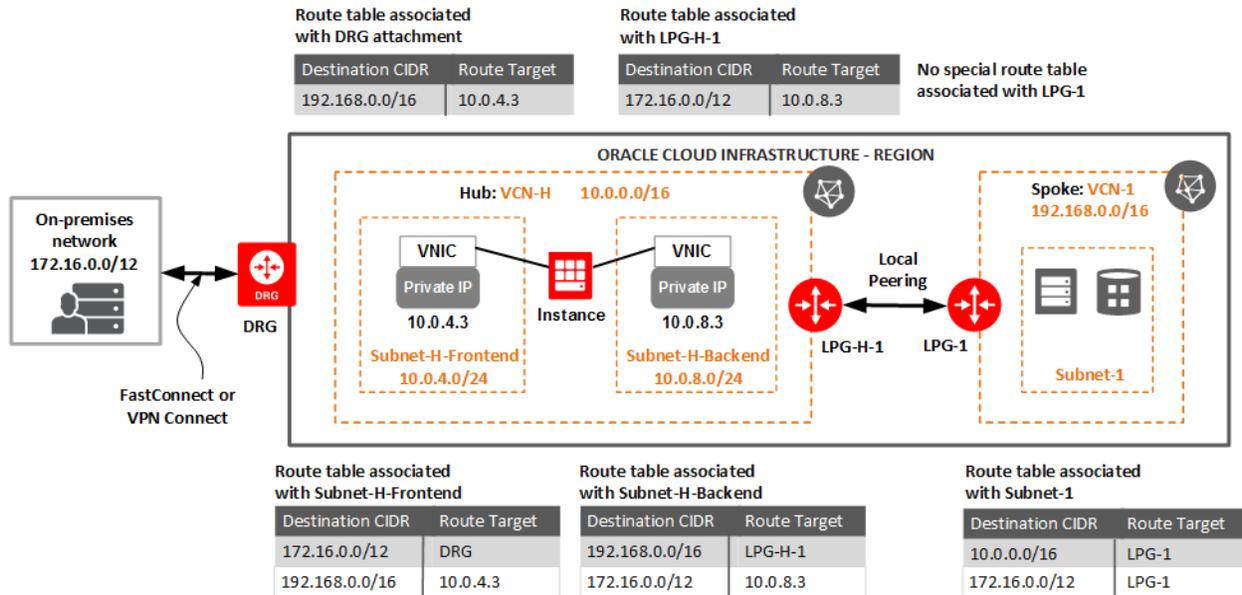
In this task, you set up the instance to have 2 private IPs.

Prerequisites:

- You already have a hub VCN with a subnet.
- Review this information: [Using a Private IP as a Route Target](#).

1. If you haven't already, create the instance in the hub VCN. See [Creating an Instance](#). The primary VNIC is created in the subnet you specify.
2. Create a secondary VNIC for the other subnet and configure the OS to use it. See [To create and attach a secondary VNIC](#).
3. Disable the source/destination check on each of the VNICs. See [Source/Destination Check](#).
4. For each VNIC, determine which private IP you want to use as the routing target. If you want to use a secondary private IP instead of the VNIC's primary private IP, assign that secondary private IP and configure the OS to use it. See [To assign a new secondary private IP to a VNIC](#).
5. For each of the private IPs you created, record the private IP address (for example: 10.0.4.3).
6. Configure the instance as necessary for the job it will perform (for example, configure the firewall or intrusion detection system on the instance).

Task 7: Set up ingress routing for the DRG and LPG on the hub VCN



In this task, you set up the route tables for the DRG attachment and hub VCN's LPG for the spoke of interest (LPG-H-1).

Prerequisites:

- You already have a DRG attached to the hub VCN.
- You already have a hub VCN LPG for the spoke of interest.
- You already have the two private IPs to use as the routing targets (see the preceding task).

1. Create a route table for the DRG attachment:
  - a. In the Console, view the hub VCN's details.
  - b. Under **Resources**, click **Route Tables** to view the VCN's route tables.
  - c. Click **Create Route Table**.
  - d. Enter the following:
    - **Name:** A descriptive name for the route table. Example: `DRG Route Table`. Avoid entering confidential information.
    - **Create in Compartment:** Leave as is.
  - e. Click **+ Additional Route Rule**, and enter this information for the route rule:
    - **Target Type:** Private IP.
    - **Destination CIDR Block:** This spoke VCN's CIDR (192.168.0.0/16 in the earlier example). Remember that you can use the routes in this table to control which subnets in the spoke VCN are advertised to the on-premises network. You could instead set up the rule to list only a particular subnet of the spoke VCN that the on-premises network.
    - **Compartment:** The compartment where the frontend subnet's private IP is located.
    - **Target:** The frontend subnet's private IP, which you recorded in the previous task (10.0.4.3 in the example).
  - f. Click **Create Route Table**.

The route table is created and displayed in the list.
2. Associate the route table (called *DRG Route Table* in this example) with the hub VCN's DRG attachment:
  - a. While still viewing the hub VCN's details, click **Dynamic Routing Gateways** to view the attached DRG.
  - b. Click the Actions icon (three dots), and then click **Associate With Route Table**.
  - c. Enter the following:

- **Route Table Compartment:** Select the compartment of the route table for the DRG attachment.
  - **Route Table:** Select the route table for the DRG attachment.
- d. Click **Associate**.
- The route table is associated with the DRG attachment.
3. Create a route table for the hub VCN's LPG for this spoke:
- a. While still viewing the hub VCN's details, click **Route Tables**.
  - b. Click **Create Route Table**.
  - c. Enter the following:
    - **Create in Compartment:** Leave as is.
    - **Name:** A descriptive name for the route table. Example: Hub LPG-# Route Table (where # indicates which spoke). Avoid entering confidential information.
  - d. Click **+ Additional Route Rule**, and enter this information for the route rule:
    - **Target Type:** Private IP.
    - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example). Remember that you can use the routes in this table to control which subnets in the on-premises network are advertised to this spoke VCN. You could instead set up the rule to list only a subnet of the on-premises network that needs to communicate with this spoke.
    - **Compartment:** The compartment where the private IP is located.
    - **Target:** The backend subnet's private IP, which you recorded in the previous task (10.0.8.3 in the example).
  - e. Click **Create Route Table**.
- The route table is created and displayed in the list.
4. Associate the route table (called *Hub LPG-# Route Table* in this example) with the hub

VCN's LPG for the spoke of interest:

- a. While still viewing the hub VCN's details, click **Local Peering Gateways** to view the hub VCN's LPG for this spoke.
- b. For the LPG you're interested in, click the Actions icon (three dots), and then click **Associate With Route Table**.
- c. Enter the following:
  - **Route Table Compartment:** Select the compartment of the route table for the LPG.
  - **Route Table:** Select the route table for the LPG.
- d. Click **Associate**.

The route table is associated with the LPG.

Although Oracle does not recommend putting workloads in the hub VCN's subnets, to give you a more complete picture of routing in the example, the diagram shows two additional route rules in the hub VCN's subnet route tables. For the frontend subnet, there's a route rule to route traffic that is destined for the spoke VCN to the private IP in the frontend subnet (10.0.4.3) for filtering by the instance. For the backend subnet, there's a route rule to route traffic that is destined for the on-premises network to the private IP in the backend subnet (10.0.8.3) for filtering by the instance. The following procedure adds those two route rules.

1. For the spoke VCN, view the list of subnets.
2. For the frontend subnet, look at its details and click the link for its associated route table.
3. Edit the frontend subnet's route table to include a rule that sends traffic destined for the spoke VCN to the private IP in the frontend subnet:
  - a. Click **Add Route Rules**.
  - b. Enter this information for the route rule:

- **Target Type:** Private IP.
  - **Destination CIDR Block:** This spoke VCN's CIDR (192.168.0.0/16 in the earlier example).
  - **Compartment:** The compartment where the frontend subnet's private IP is located.
  - **Target:** The frontend subnet's private IP, which you recorded in the previous task (10.0.4.3 in the example).
- c. Click **Add Route Rules**.
4. For the backend subnet, look at its details and click the link for its associated route table.
  5. Edit the backend subnet's route table to include a rule that sends traffic destined for the on-premises network to the private IP in the backend subnet:
    - a. Click **Add Route Rules**.
    - b. Enter this information for the route rule:
      - **Target Type:** Private IP.
      - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example).
      - **Compartment:** The compartment where the backend subnet's private IP is located.
      - **Target:** The backend subnet's private IP, which you recorded in the previous task (10.0.8.3 in the example).
    - c. Click **Add Route Rules**.

### Later if you need more spoke VCNs

1. Repeat Tasks 3-5 for the new spoke VCN.
2. Repeat task 7 with these changes:

- For Step 1: Instead of creating a new route table for the DRG attachment, update the existing route table to include a new rule for the new spoke VCN. The destination CIDR is the spoke VCN's CIDR (or a subnet within). The target is the frontend subnet private IP 10.0.4.3.
- For Step 2: Skip this step entirely because the DRG attachment is already associated with its route table.
- For Step 3: Repeat as is. Name the new route table according to which spoke the route table is for (for example, `Hub LPG-2 Route Table` for the second spoke).
- For Step 4: Repeat as is. Associate the new route table you created in Step 3 with the hub VCN's LPG for the new spoke.

### Turning Off Transit Routing

To turn off transit routing, remove the rules from:

- The route table associated with the DRG attachment.
- The route table associated with each LPG on the hub VCN.

A route table can be associated with a resource but have no rules. Without at least one rule, a route table does nothing.

A DRG attachment or LPG can exist without a route table associated with it. However, after you associate a route table with a DRG attachment or LPG, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

### Changes to the API

For information about changes to the Networking service API to support transit routing, see the [transit routing release notes](#).

## Transit Routing: Private Access to Oracle Services

*Transit routing* refers to a network setup in which your on-premises network uses a connected virtual cloud network (VCN) to reach Oracle resources or services beyond that VCN. You connect the on-premises network to the VCN with [FastConnect](#) or [VPN Connect](#), and then configure the VCN routing so that traffic *transits through the VCN* to its destination beyond the VCN.

There are two primary transit routing scenarios:

- **Private access to Oracle services:** The scenario covered in this topic. This scenario gives your on-premises network *private access* to Oracle services, so that your on-premises hosts can use their private IP addresses and the traffic does not go over the public internet. Instead, the traffic travels over a FastConnect private virtual circuit or VPN Connect, transits through a virtual cloud network (VCN), and then through a service gateway to the Oracle service of interest.
- **Access to multiple VCNs in the same region:** This scenario enables communication between an on-premises network and multiple VCNs in the same region over a single FastConnect private virtual circuit or VPN Connect. See [Transit Routing: Access to Multiple VCNs in the Same Region](#).

### Highlights

- You can set up a VCN so that your on-premises network has *private access* to Oracle services in the Oracle Services Network by way of the VCN. The hosts in your on-premises network communicate with their private IP addresses.
- The VCN uses a [dynamic routing gateway](#) (DRG) to communicate with the on-premises network. Access to Oracle services is through a [service gateway](#) on the VCN. The traffic from the VCN to the Oracle service travels over the Oracle network fabric and never traverses the public internet.
- The service gateway is regional and enables access only to supported Oracle services *in the same region* as the VCN.

- The supported Oracle services are Oracle Cloud Infrastructure Object Storage and others in the Oracle Services Network. For a list, see [Service Gateway: Supported Cloud Services in Oracle Services Network](#).
- The service gateway uses the concept of a *service CIDR label*, which is a string that represents all the regional public IP address ranges for the service or group of services of interest (for example, *OCI PHX Object Storage* is the string for Object Storage in US West (Phoenix)). You use that service CIDR label when you configure the service gateway and related route rules to control traffic to the service. You can optionally use it when configuring security rules. If the service's public IP addresses change in the future, you don't have to adjust those rules.
- To enable the desired traffic from the on-premises network through the VCN to Oracle services, you implement route rules for the VCN's DRG attachment and service gateway.
- If you want, you can set up transit routing *through a private IP in the VCN*. For example, you might want to filter or inspect the traffic between the on-premises network and the Oracle service. In that case, you route the traffic to a private IP on an instance in the VCN for inspection, and the resulting traffic continues to its destination. This topic covers both situations: transit routing directly between gateways on the VCN, and transit routing through a private IP.

### Overview of the Oracle Services Network

The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. These services have [public IP addresses](#) that you typically reach over the public internet. However, you can access the Oracle Services Network *without the traffic going over the public internet*. There are different ways, depending on which of your hosts need the access:

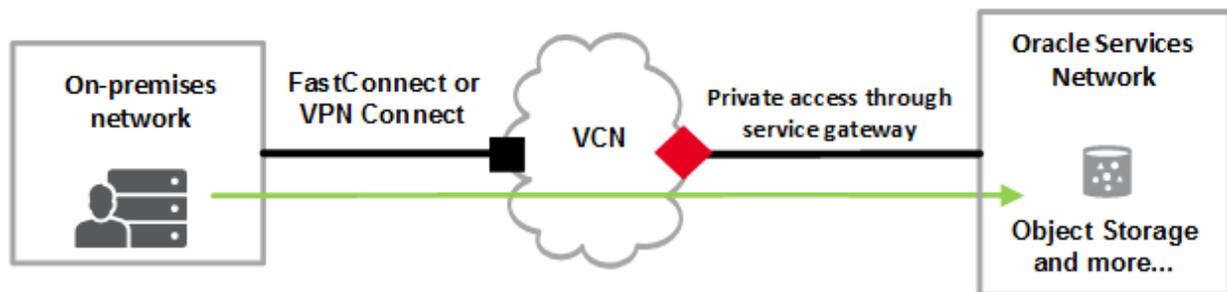
- **Hosts in your on-premises network:**
  - Private access through a VCN with FastConnect private peering or VPN Connect: This is the scenario covered in this topic. The on-premises hosts use private IP addresses and reach the Oracle Services Network by way of the VCN and the

VCN's service gateway.

- [Public access with FastConnect public peering](#): The on-premises hosts use public IP addresses.
- **Hosts in your VCN:**
  - [Private access through a service gateway](#): The VCN's hosts use private IP addresses.

### Overview of On-Premises Network Private Access to Oracle Services

The following diagram illustrates the basic layout for giving your on-premises network private access to Oracle services.



#### Legend:

- dynamic routing gateway (DRG)
- ◆ service gateway

Your on-premises network connects to the VCN by way of a [FastConnect](#) private virtual circuit or [VPN Connect](#). Each of these types of connections terminates on a dynamic routing gateway (DRG) that is attached to the VCN. The VCN also has a service gateway, which gives the VCN access to the Oracle Services Network. The traffic from your on-premises network transits through the VCN, through the service gateway, and to the Oracle service of interest. The responses return through the service gateway and VCN to your on-premises network.

When you set up a service gateway, you enable a *service CIDR label*, which is a string that represents all the regional public IP address ranges for the service or group of services that you want to access through the service gateway. For example, *All PHX Services in Oracle Services Network* is the service CIDR label for the Oracle services available in US West (Phoenix) through a service gateway. Oracle uses Border Gateway Protocol (BGP) on the DRG to advertise those regional public IP address ranges to the edge device (also called the customer-premises equipment or CPE) in your on-premises network. For a list of those ranges available through the service gateway, see [Public IP Addresses for VCNs and the Oracle Services Network](#).

### Multiple Connection Paths to Oracle Services

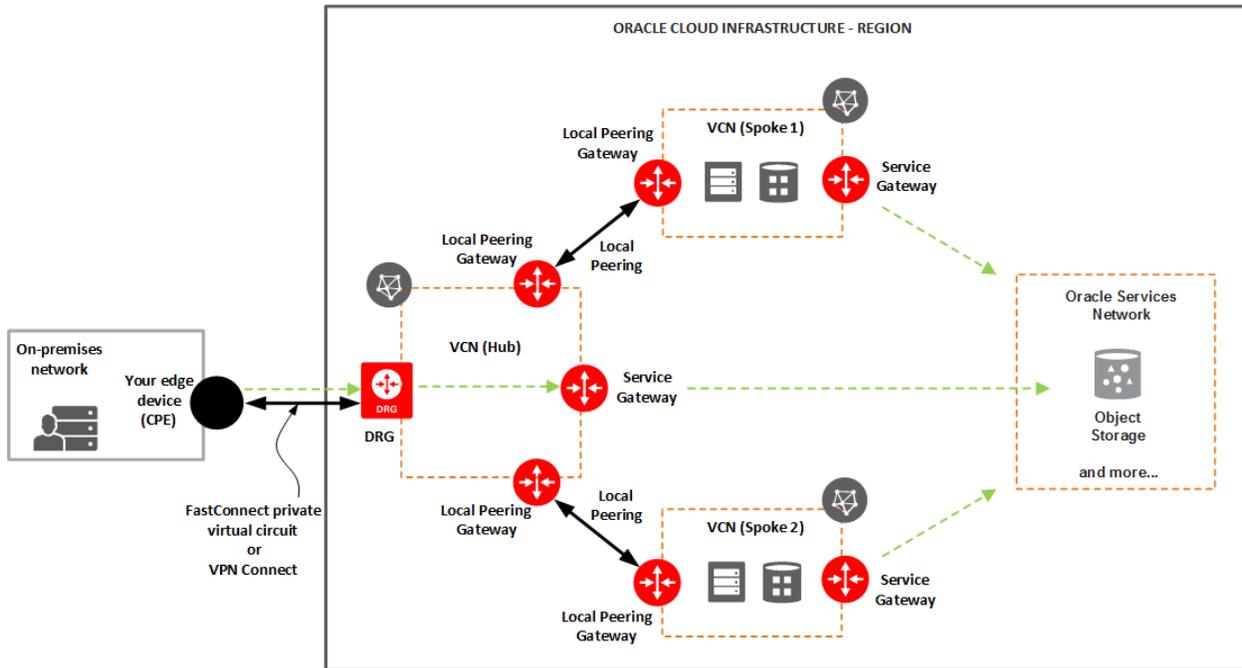
You can configure your on-premises network with multiple connection paths to Oracle Cloud Infrastructure and Oracle services for redundancy or other reasons. For example, you could use both FastConnect public peering and FastConnect private peering. If you have multiple paths, your edge device receives route advertisement of the Oracle services public IP address ranges over multiple paths. For important information about configuring your edge device correctly, see [Routing Details for Connections to Your On-Premises Network](#).

### Multiple VCNs with Private Access to Oracle Services

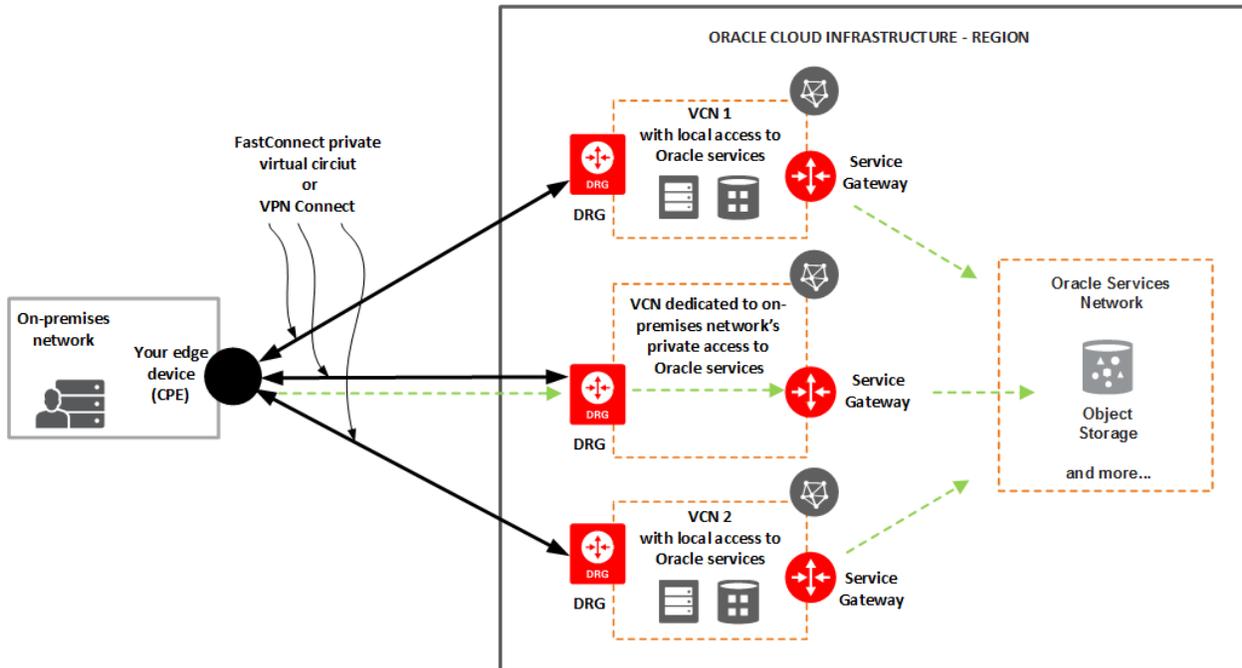
Your organization might choose to use multiple VCNs, each with a service gateway to give the VCN's resources access to Oracle services. For example, you might have a different VCN for each department in your organization.

If you *also* want to set up your on-premises network with private access to Oracle services through a VCN with a service gateway, this section describes two different network layouts you could use.

In the first layout, you set up a *single* DRG, with the VCNs in a hub-and-spoke layout as shown in the next diagram. The VCN that acts as the hub is dedicated to providing the on-premises network with private access to Oracle services. The other VCNs are [locally peered](#) with the hub VCN. You configure only the hub VCN according to instructions in [Setting Up Private Access to Oracle Services](#). This hub-and-spoke layout is recommended and described further in [Transit Routing: Access to Multiple VCNs in the Same Region](#).



In the second layout, there's a separate DRG for each VCN, with a separate FastConnect private virtual circuit or VPN Connect from your on-premises network to each DRG. You dedicate one DRG and VCN to providing your on-premises network with private access to Oracle services. In the next diagram, it's the VCN in the center. To configure that VCN, follow the instructions in [Setting Up Private Access to Oracle Services](#).



Notice that in both of these layouts, the on-premises network can reach the Oracle services only through a single VCN's service gateway (the one dedicated for this purpose) and not through the service gateways of the other VCNs. For those other VCNs, only the resources *inside* those VCNs can reach Oracle services through their VCN's service gateway.

Regardless of which layout you choose, you can write an IAM policy to restrict access to an Object Storage bucket so that *only requests that come through a specific VCN's service gateway* are allowed for that bucket. With either of these layouts, you might want to write the policy to allow requests from *multiple* VCNs. The following example policy lets resources in the example ObjectBackup group write objects to an existing bucket called db-backup that resides in a compartment called ABC. When writing a policy like this one, you can specify one or more VCN OCIDs. This example shows three.

```
Allow group ObjectBackup to read buckets in compartment ABC

Allow group ObjectBackup to manage objects in compartment ABC where
 all {target.bucket.name='db-backup',
```

```
any {request.vcn.id='<hub_VCN_OCID>', request.vcn.id='<spoke_1_VCN_OCID>',
request.vcn.id='<spoke_2_VCN_OCID>'},
any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}}
```

For more information, see [Task 4: \(Optional\) Update IAM Policies to Restrict Object Storage Bucket Access](#) in the procedure for setting up a service gateway.

### Requests from Oracle Services to Your Clients

The service gateway does not allow incoming connection requests to the VCN or your on-premises network. Any connection requests coming from an Oracle service to your on-premises network must come over a public path such as the internet or FastConnect public peering.

If you use Oracle Analytics Cloud so that it initiates connection requests to clients, and you *also* want to set up private access to Oracle services for your on-premises network, see this [known issue](#).

## Transit Routing Options for Private Access to Oracle Services

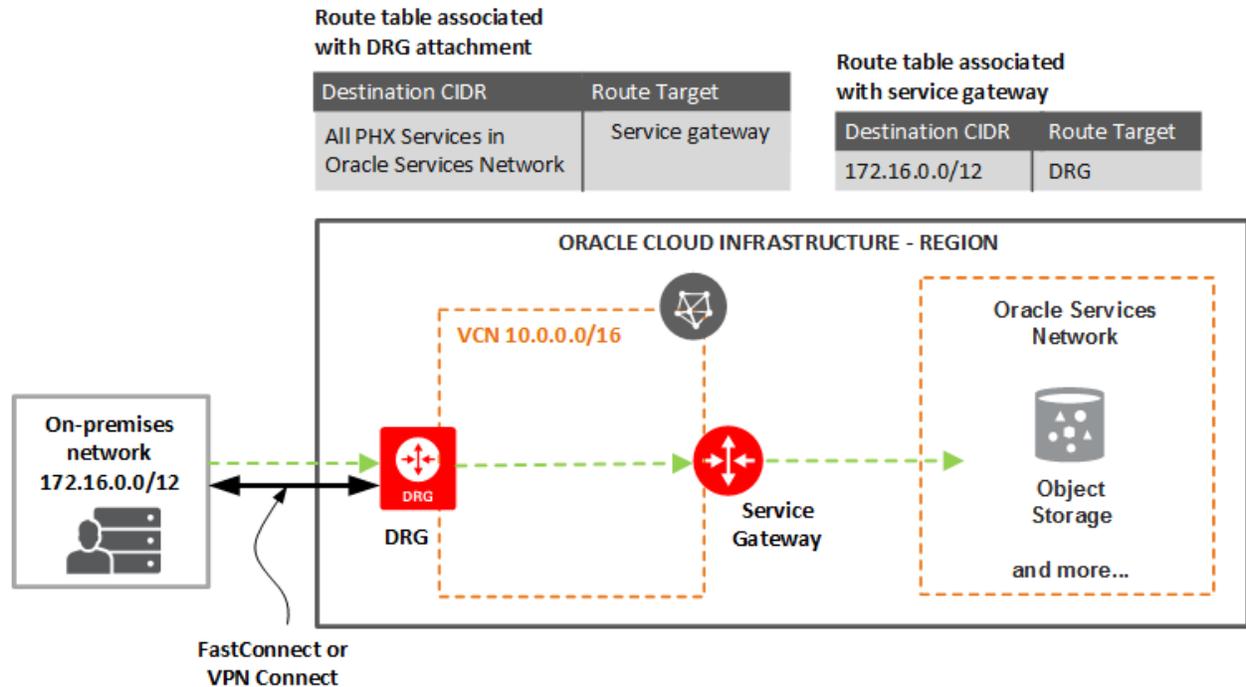
There are two options for routing through the VCN for private access to Oracle services:

- **Transit routing directly through gateways:** You route the traffic directly through the VCN, from one gateway to the other.
- **Transit routing through a private IP:** You set up an instance in the VCN to filter or inspect the traffic between the on-premises network and Oracle Services Network, and route traffic through a private IP on the instance.

The examples shown in the following sections assume that the VCN contains no workloads that need to access the on-premises network or Oracle Services Network. The VCN is being used only for transit routing of traffic *through the VCN*.

### Transit routing directly through gateways

In this example, you route directly through the two gateways on the VCN: the dynamic routing gateway (DRG) and the service gateway. See the following diagram.



The diagram shows two route tables, each associated with a different resource:

- **DRG attachment:**

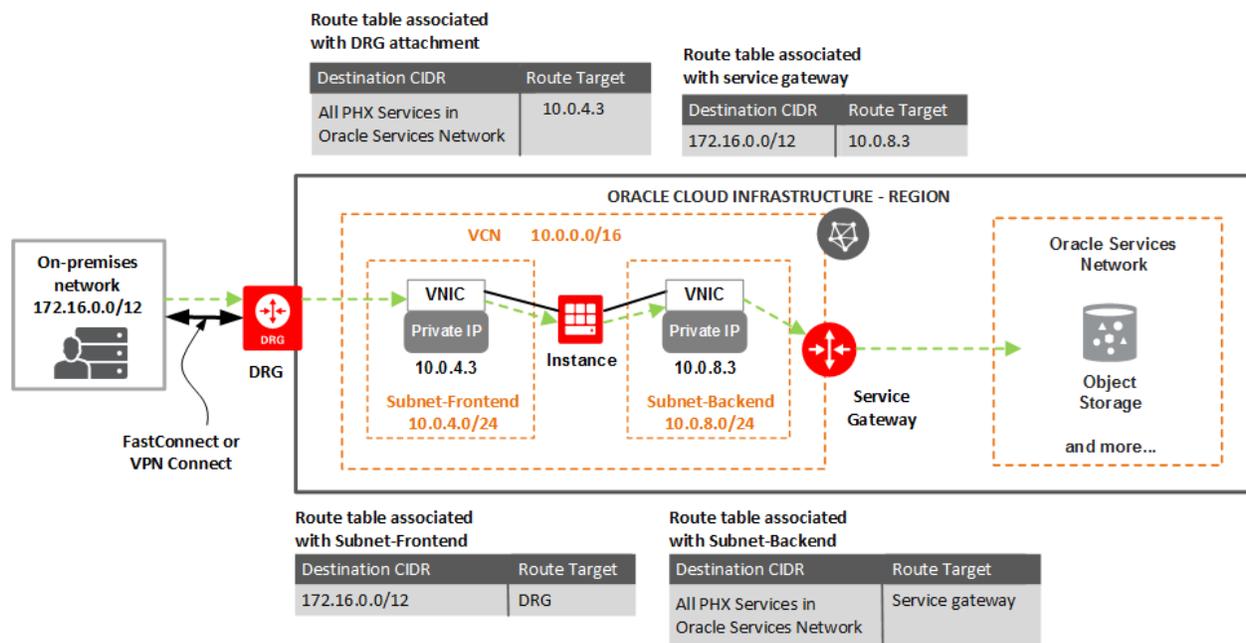
- The route table belongs to the VCN and is associated with the DRG *attachment*. Why the attachment and not the DRG itself? Because the DRG is a standalone resource that you can attach to any VCN in the same region and tenancy as the DRG. The attachment itself identifies which VCN.
- The route table routes the inbound traffic that is from the on-premises network and destined for a supported Oracle service. You configure the rule to send that traffic to the service gateway.

- **Service gateway:**

- This route table belongs to the VCN and is associated with the service gateway.
- The route table routes response traffic that is from a supported Oracle service and destined for the on-premises network. You configure the rule to send that traffic to the DRG.

### Transit routing through a private IP in the VCN

In this example, you set up an instance in the VCN to act as a firewall or intrusion detection system to filter or inspect the traffic between the on-premises network and Oracle Services Network. See the following diagram.



The instance has two VNICs, each with a private IP. One of the VNICs is in a subnet that faces the on-premises network (referred to here as the *frontend subnet*). The other VNIC is in a

subnet that faces the Oracle Services Network (referred to here as the *backend subnet*). The frontend VNIC has private IP 10.0.4.3, and the backend VNIC has private IP 10.0.8.3.

The diagram shows four route tables, each associated with a different resource:

- **DRG attachment:**

- The route table belongs to the VCN and is associated with the DRG *attachment*. Why the attachment and not the DRG itself? Because the DRG is a standalone resource that you can attach to any VCN in the same region and tenancy as the DRG. The attachment itself identifies which VCN.
- The route table routes the inbound traffic that is from the on-premises network and destined for a supported Oracle service. You configure the rule to send the traffic to the private IP in the frontend subnet.

- **Service gateway:**

- This route table belongs to the VCN and is associated with the service gateway.
- The route table routes response traffic that is from a supported Oracle service and destined for the on-premises network. You configure the rule to send that traffic to the private IP in the backend subnet.

- **Subnet-frontend:**

- This route table belongs to the VCN and is associated with Subnet-frontend.
- It includes a rule to enable traffic with the on-premises network.

- **Subnet-backend:**

- This route table belongs to the VCN and is associated with Subnet-backend.
- It includes a rule to enable traffic with the regional Oracle Services Network.

### **Important Transit Routing Restrictions to Understand**

This section includes some additional important details about routing:

- **Route table for the DRG attachment:**

- A route table that is associated with a DRG attachment can only have rules that target a service gateway, a private IP, or a local peering gateway.
- A DRG attachment can exist without a route table associated with it. However, after you associate a route table with a DRG attachment, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

- **Route table for a service gateway:**

- A route table that is associated with a service gateway can only have rules that target a DRG or a private IP.
- A service gateway can exist without a route table associated with it. However, after you associate a route table with a service gateway, there must always be a route table associated with it. But, you can associate a different route table. You can also edit the table's rules, or delete some or all of the rules.

- **Traffic *transiting through* the VCN:** The route tables discussed here are intended only for moving traffic *through* the VCN between locations in the on-premises network and the Oracle Services Network. If you're using a private IP in the VCN, you configure the route tables so that the private IP is placed in that traffic path going *through* the VCN.

- **Inbound traffic to the VCN:** Even though the preceding statement is true (about traffic *through* the VCN), inbound traffic to subnets *within the VCN* is always allowed. You do not need to set up explicit rules for this inbound traffic in the DRG attachment's route table or service gateway's route table. When this kind of inbound traffic reaches the DRG or the service gateway, the traffic is automatically routed to its destination in the VCN by the *VCN local routing*. Because of VCN local routing, for any route table belonging to a given VCN, you can't create a rule that lists that VCN's CIDR (or a subsection) as the rule's destination.

- **VCN traffic when transit routing through a private IP:** The immediately preceding statement about VCN local routing means that you should use the VCN only for *transit* between the on-premises network and spoke VCNs. **You should not set up workloads in the VCN itself.** More explicitly, if you set up transit routing through a

private IP in the VCN, you can't also route the *VCN's* traffic through that private IP. For example, in the preceding diagram, if you were to change the route rule in the service gateway's route table so that the destination CIDR is 0.0.0.0/0 instead of 172.16.0.0/12, only traffic coming from the Oracle Services Network and destined for addresses *outside* the VCN's CIDR block would be routed through the private IP. Because of VCN local routing, any traffic destined for addresses within the VCN is automatically routed directly to the destination IP address. The VCN local routing takes precedence over the service gateway's route table (in general, over *any* of the VCN's route tables).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're a member of the Administrators group, you already have the required access to set up transit routing. Otherwise, you need access to the Networking service, and you need the ability to launch instances. See [IAM Policies for Networking](#).

### Setting Up Private Access to Oracle Services

This section shows how to use the Console to set up transit routing with a VCN to give your on-premises network private access to Oracle services.



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

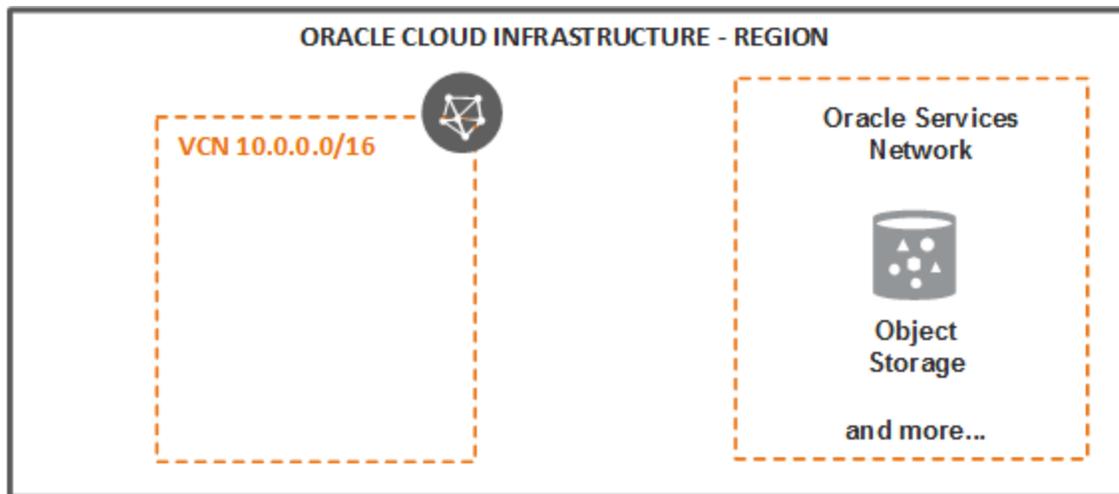
## For routing directly between gateways



### Tip

You might already have many of the necessary Networking components and connections in this advanced scenario already set up. So you might be able to skip some of the following tasks. **If you already have a network layout with a VCN connected to your on-premises network, and a service gateway for that VCN, then Task 4 is the most important.** It enables traffic to be routed between your on-premises network and the Oracle Services Network.

## Task 1: Set up the VCN

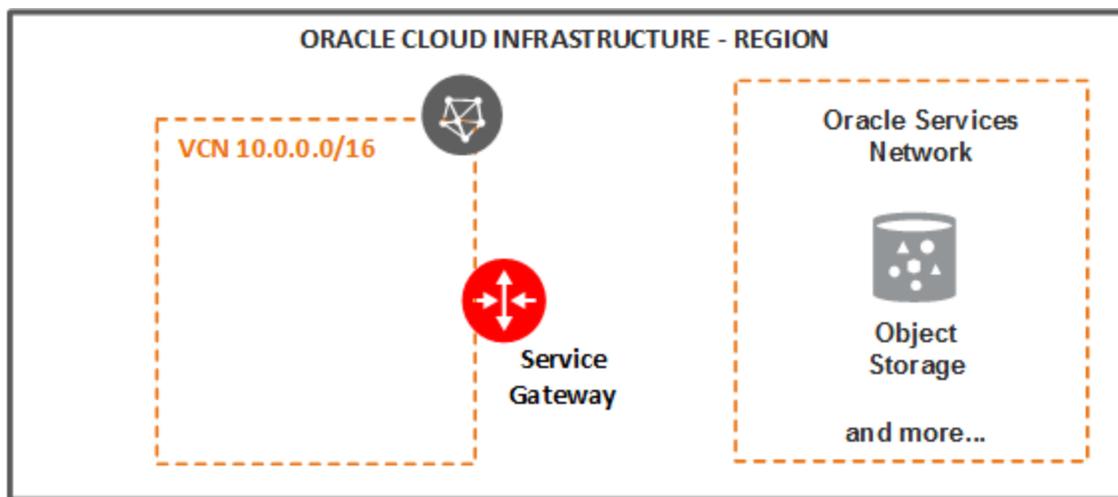


In this task, you set up the VCN. For this example, no subnet is required.

For more information and instructions:

- [VCNs and Subnets](#)

### Task 2: Add a service gateway to the VCN



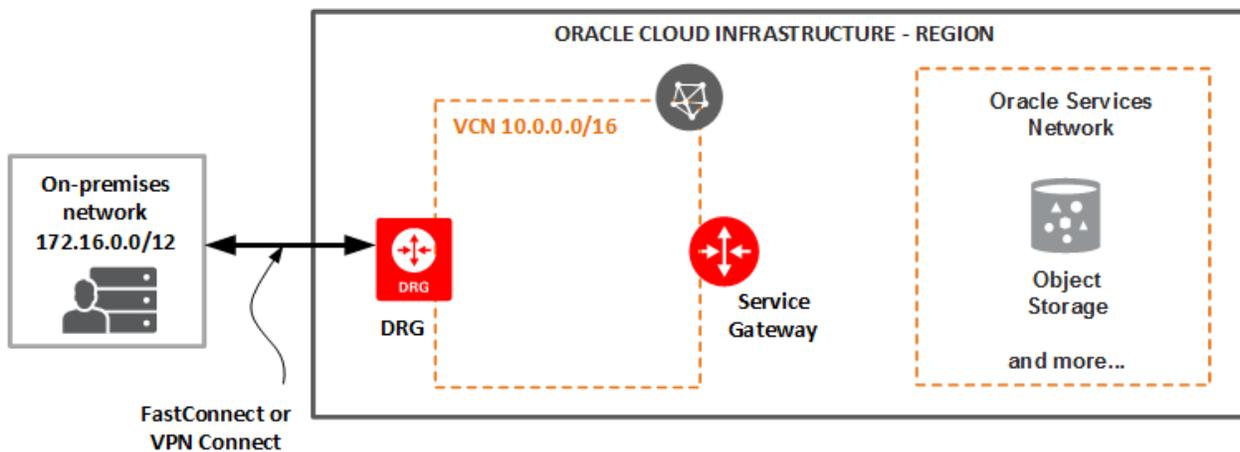
In this task, you add a service gateway to the VCN and enable the gateway for the regional Oracle Services Network.

Notice that you do not yet create the route table that will be associated with the service gateway. That comes in a later task.

1. In the Console, view the VCN's details.
2. Under **Resources**, click **Service Gateways**.
3. Click **Create Service Gateway**.
4. Enter the following values:

- **Name:** A descriptive name for the service gateway. It doesn't have to be unique. Avoid entering confidential information.
  - **Create in compartment:** The compartment where you want to create the service gateway, if different from the compartment you're currently working in.
  - **Services:** All *<region>* Services in Oracle Services Network.
5. Click **Create Service Gateway**.
- The service gateway is then created and displayed on the **Service Gateways** page in the compartment you chose.

### Task 3: Connect the VCN to your on-premises network



In this task, you set up either FastConnect or VPN Connect between your VCN and your on-premises network. As part of this process, you attach a DRG to the VCN and set up routing between the VCN and your on-premises network.

Notice that you do not yet create the route table that will be associated with the DRG attachment. That comes in a later task.

For more information and instructions:

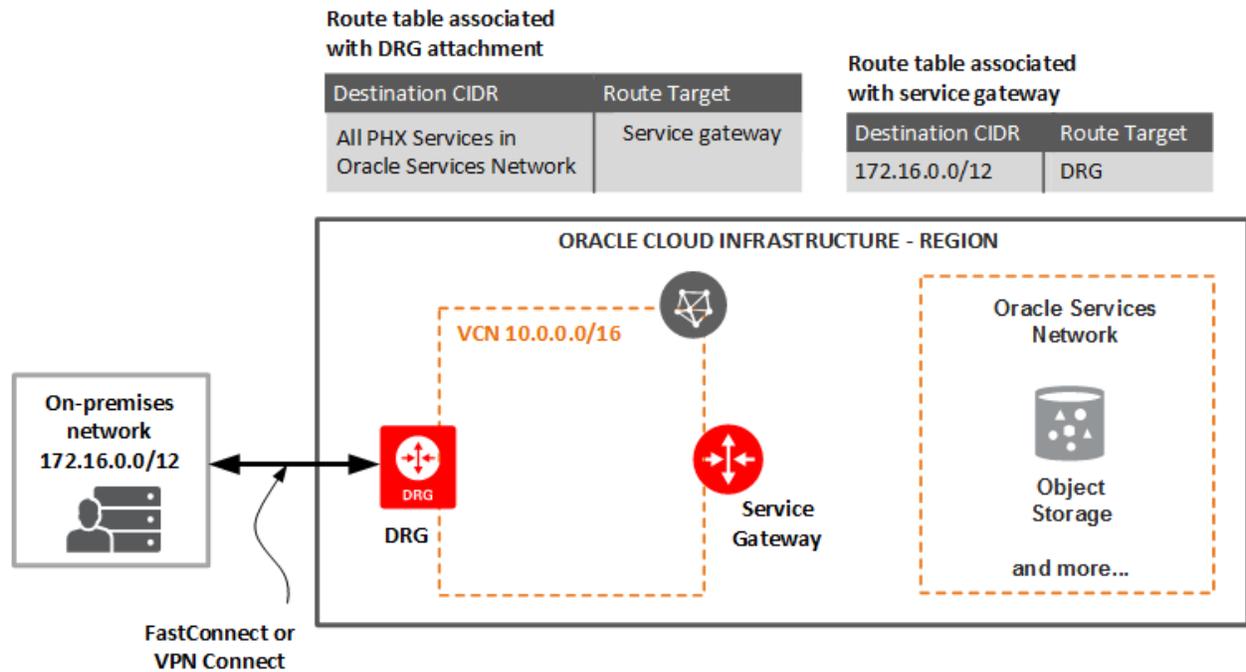
- [FastConnect](#)
- [VPN Connect](#)
- [Dynamic Routing Gateways \(DRGs\)](#)



### Important

If you're using VPN Connect with static routing, and you've configured the VCN to give your on-premises network private access to Oracle services, you must configure your edge device with the routes for the Oracle Services Network public IP ranges that are advertised by the DRG over the private path (through the service gateway). For a list of those ranges, see [Public IP Addresses for VCNs and the Oracle Services Network](#)

Task 4: Set up ingress routing for the DRG and service gateway



In this task, you set up the route tables for the DRG attachment and the service gateway.

Prerequisites:

- You already have a DRG attached to the VCN.
- You already have a service gateway.

1. Create a route table for the DRG attachment:
  - a. In the Console, view the VCN's details.
  - b. Under **Resources**, click **Route Tables** to view the VCN's route tables.

- c. Click **Create Route Table**.
  - d. Enter the following:
    - **Name:** A descriptive name for the route table. Example: `DRG Route Table`. Avoid entering confidential information.
    - **Create in Compartment:** Leave as is.
  - e. Click **+ Additional Route Rule**, and enter this information for the route rule:
    - **Target Type:** Service gateway.
    - **Destination Service:** All *<region>* Services in Oracle Services Network.
    - **Compartment:** The compartment where the service gateway is located.
    - **Target:** The service gateway.
  - f. Click **Create Route Table**.

The route table is created and displayed in the list.
2. Associate the route table (called *DRG Route Table* in this example) with the VCN's DRG attachment:
    - a. While still viewing the VCN's details, click **Dynamic Routing Gateways** to view the attached DRG.
    - b. Click the Actions icon (three dots), and then click **Associate With Route Table**.
    - c. Enter the following:
      - **Route Table Compartment:** Select the compartment of the route table for the DRG attachment.
      - **Route Table:** Select the route table for the DRG attachment.
    - d. Click **Associate**.

The route table is associated with the DRG attachment.
  3. Create a route table for the service gateway:
    - a. While still viewing the VCN's details, click **Route Tables**.
    - b. Click **Create Route Table**.

- c. Enter the following:
    - **Create in Compartment:** Leave as is.
    - **Name:** A descriptive name for the route table. Example: `Service Gateway Route Table`. Avoid entering confidential information.
  - d. Click **+ Additional Route Rule**, and enter this information for the route rule:
    - **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
    - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example).
  - e. Click **Create Route Table**.

The route table is created and displayed in the list.
4. Associate the route table (called *Service Gateway Route Table* in this example) with the service gateway:
    - a. While still viewing the VCN's details, click **Service Gateways**.
    - b. For the service gateway, click the Actions icon (three dots), and then click **Associate With Route Table**.
    - c. Enter the following:
      - **Route Table Compartment:** Select the compartment of the route table for the service gateway.
      - **Route Table:** Select the route table for the service gateway.
    - d. Click **Associate**.

The route table is associated with the service gateway.

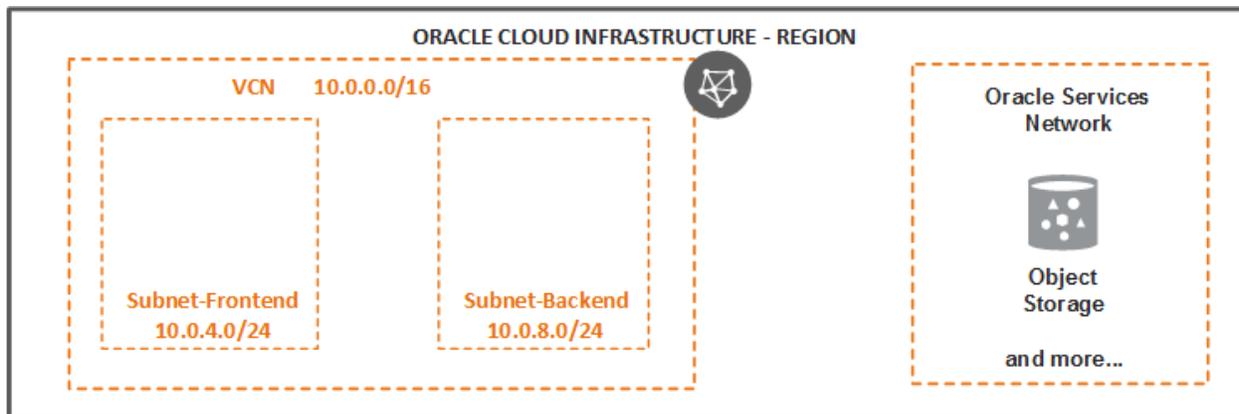
## For routing through a private IP



### Tip

You might already have many of the necessary Networking components and connections in this advanced scenario already set up. So you might be able to skip some of the following tasks. **If you already have a network layout with a VCN connected to your on-premises network, and a service gateway for that VCN, then Tasks 4 and 5 are the most important.** They enable traffic to be routed between your on-premises network and the spoke VCN.

## Task 1: Set up the VCN

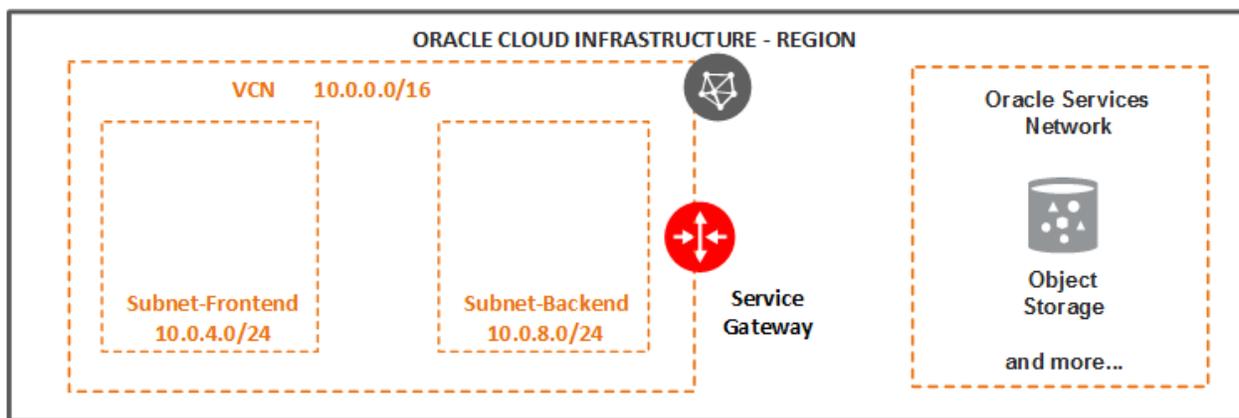


In this task, you set up the VCN. This example also has two subnets: one for the frontend VNIC on the instance, and one for the backend VNIC on the instance. Oracle recommends using regional *private* subnets.

For more information and instructions:

- [VCNs and Subnets](#)

### Task 2: Add a service gateway to the VCN



Route table associated with Subnet-Backend

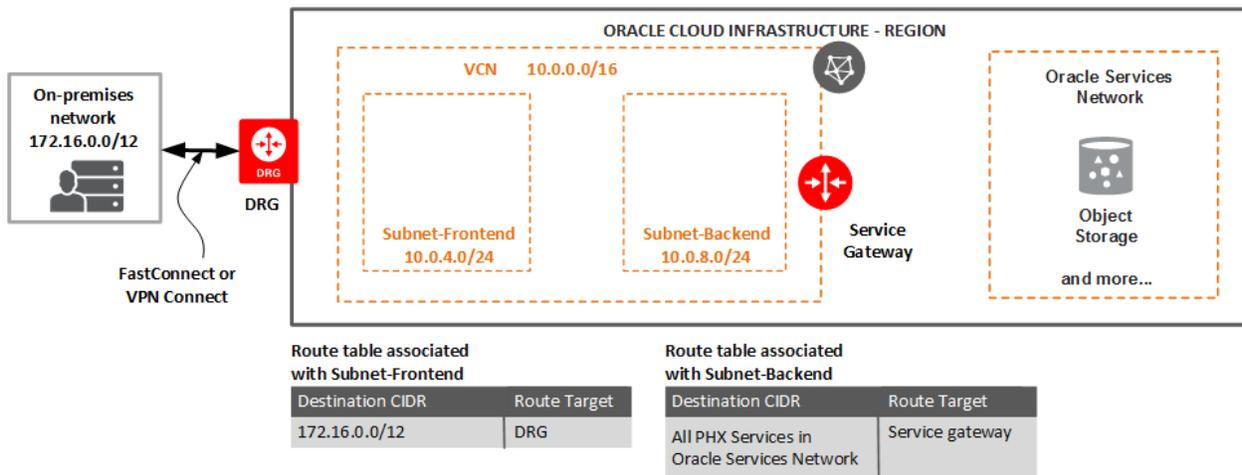
Destination CIDR	Route Target
All PHX Services in Oracle Services Network	Service gateway

In this task, you add a service gateway to the VCN and enable the gateway for the regional Oracle Services Network.

Notice that you do not yet create the route table that will be associated with the service gateway. That comes in a later task.

1. In the Console, view the VCN's details.
2. Under **Resources**, click **Service Gateways**.
3. Click **Create Service Gateway**.
4. Enter the following values:
  - **Name:** A descriptive name for the service gateway. It doesn't have to be unique. Avoid entering confidential information.
  - **Create in compartment:** The compartment where you want to create the service gateway, if different from the compartment you're currently working in.
  - **Services:** All *<region>* Services in Oracle Services Network.
5. Click **Create Service Gateway**.  
 The service gateway is then created and displayed on the **Service Gateways** page in the compartment you chose.

Task 3: Connect the VCN to your on-premises network



In this task, you set up either FastConnect or VPN Connect between your hub VCN and your on-premises network. As part of this process, you attach a DRG to the hub VCN and set up routing between the hub VCN and your on-premises network.

Notice that you do not yet create the route table that will be associated with the DRG attachment. That comes in a later step.

For more information and instructions:

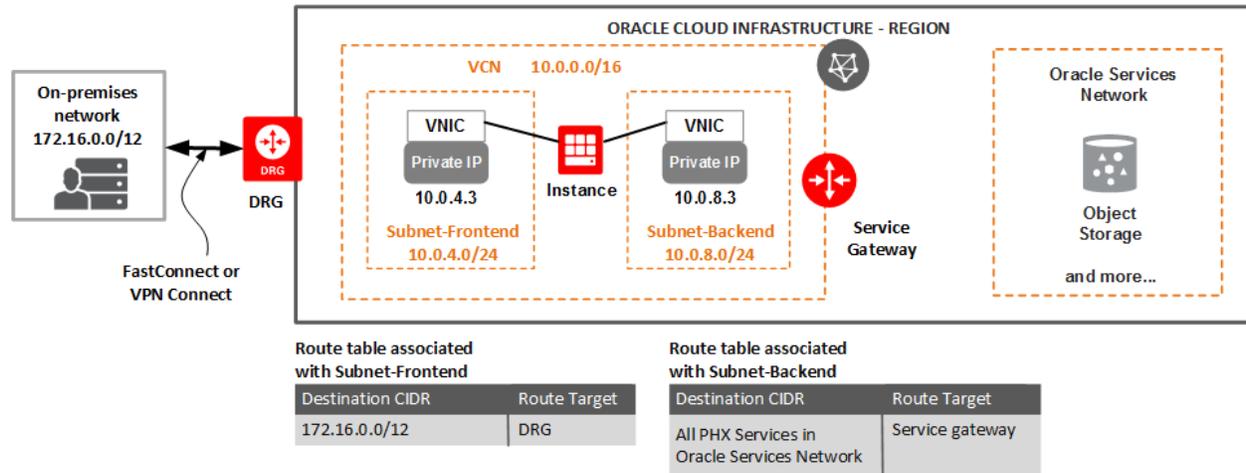
- [FastConnect](#)
- [VPN Connect](#)
- [Dynamic Routing Gateways \(DRGs\)](#)



### Important

If you're using VPN Connect with static routing, and you've configured the VCN to give your on-premises network private access to Oracle services, you must configure your edge device with the routes for the Oracle Services Network public IP ranges that are advertised by the DRG over the private path (through the service gateway). For a list of those ranges, see [Public IP Addresses for VCNs and the Oracle Services Network](#)

Task 4: Set up the private IPs on an instance in the VCN



In this task, you set up the instance to have 2 private IPs.

Prerequisites:

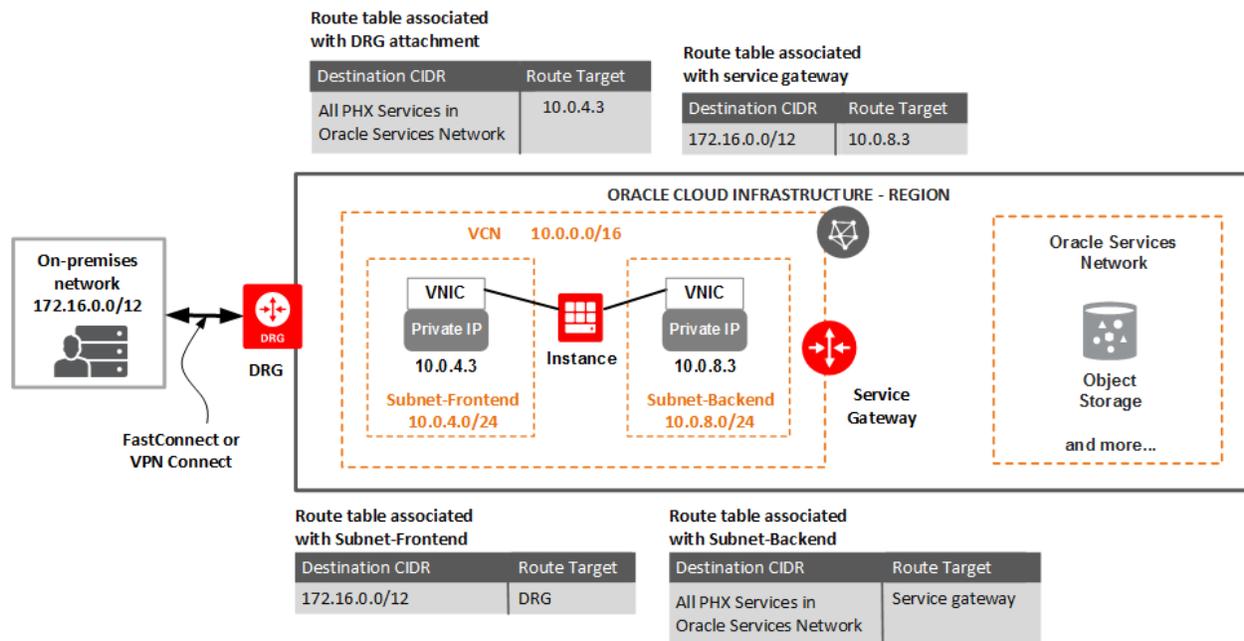
- You already have a VCN with two subnets.
  - Review this information: [Using a Private IP as a Route Target](#).
1. If you haven't already, create the instance in the VCN. See [Creating an Instance](#). The primary VNIC is created in the subnet you specify.
  2. Create a secondary VNIC for the other subnet and configure the OS to use it. See [To create and attach a secondary VNIC](#).
  3. Disable the source/destination check on each of the VNICs. See [Source/Destination Check](#).
  4. For each VNIC, determine which private IP you want to use as the routing target. If you want to use a secondary private IP instead of the VNIC's primary private IP, assign that

## CHAPTER 23 Networking

secondary private IP and configure the OS to use it. See [To assign a new secondary private IP to a VNIC](#).

5. For each of the private IPs you created, record the private IP address (for example: 10.0.4.3).
6. Configure the instance as necessary for the job it will perform (for example, configure the firewall or intrusion detection system on the instance).

### Task 5: Set up ingress routing for the DRG and service gateway



In this task, you set up the route tables for the DRG attachment and service gateway.

Prerequisites:

- You already have a DRG attached to the VCN.
  - You already have a service gateway.
  - You already have the 2 private IPs to use as the routing targets (see the preceding task).
1. Create a route table for the DRG attachment:
    - a. In the Console, view the VCN's details.
    - b. Under **Resources**, click **Route Tables** to view the VCN's route tables.
    - c. Click **Create Route Table**.
    - d. Enter the following:
      - **Name:** A descriptive name for the route table. Example: `DRG Route Table`. Avoid entering confidential information.
      - **Create in Compartment:** Leave as is.
    - e. Click **+ Additional Route Rule**, and enter this information for the route rule:
      - **Target Type:** Private IP.
      - **Destination:** Service.
      - **Destination Service:** All *<region>* Services in Oracle Services Network
      - **Compartment:** The compartment where the frontend subnet's private IP is located.
      - **Target:** The frontend subnet's private IP, which you recorded in the previous task (10.0.4.3 in the example).
    - f. Click **Create Route Table**.

The route table is created and displayed in the list.
  2. Associate the route table (called *DRG Route Table* in this example) with the VCN's DRG attachment:
    - a. While still viewing the VCN's details, click **Dynamic Routing Gateways** to view the attached DRG.

- b. Click the Actions icon (three dots), and then click **Associate Route Table**.
- c. Select the route table.
- d. Click **Associate Route Table**.

The route table is associated with the DRG attachment.

3. Create a route table for the service gateway:

- a. While still viewing the VCN's details, click **Route Tables**.
- b. Click **Create Route Table**.
- c. Enter the following:
  - **Create in Compartment:** Leave as is.
  - **Name:** A descriptive name for the route table. Example: `Service Gateway Route Table`. Avoid entering confidential information.
- d. Click **+ Additional Route Rule**, and enter this information for the route rule:
  - **Target Type:** Private IP.
  - **Destination:** CIDR Block.
  - **Destination CIDR Block:** The on-premises network's CIDR (172.16.0.0/12 in the earlier example).
  - **Compartment:** The compartment where the private IP is located.
  - **Target:** The backend subnet's private IP, which you recorded in the previous task (10.0.8.3 in the example).
- e. Click **Create Route Table**.

The route table is created and displayed in the list.

4. Associate the route table (called *Service Gateway Route Table* in this example) with the service gateway:
- a. While still viewing the VCN's details, click **Service Gateways**.
  - b. For the service gateway, click the Actions icon (three dots), and then click **Associate With Route Table**.

- c. Enter the following:
  - **Route Table Compartment:** Select the compartment of the route table for the service gateway.
  - **Route Table:** Select the route table for the service gateway.
- d. Click **Associate**.

The route table is associated with the service gateway.

### Turning Off Transit Routing

To turn off transit routing, remove the rules from:

- The route table associated with the DRG attachment.
- The route table associated with service gateway.

A route table can be associated with a resource but have no rules. Without at least one rule, a route table does nothing.

A DRG attachment or service gateway can exist without a route table associated with it. However, after you associate a route table with a DRG attachment or service gateway, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

### FastConnect with Multiple DRGs and VCNs

This topic summarizes an advanced networking scenario that enables communication between an on-premises network and multiple virtual cloud networks (VCNs) over a single [Oracle Cloud Infrastructure FastConnect](#). **Each VCN has its own dynamic routing gateway (DRG), and you set up a separate FastConnect private virtual circuit to each DRG.**

This scenario is supported for only certain FastConnect setups:

- Supported:
  - Using a [third-party provider](#)
  - [Colocation with Oracle in a FastConnect location](#)
- Not supported:
  - Using an [Oracle provider](#)

There's a scenario called *transit routing* that also involves using multiple VCNs, but only a single DRG. It can be used with VPN Connect or FastConnect. It involves setting up the VCNs in a hub-and-spoke layout and using the hub VCN for transit routing of traffic to the other VCNs. You might use this scenario if you need multiple VCNs for different parts of your organization, but you want to use one VCN for centralized services that all parts of the organization need. For more information, see [Transit Routing: Access to Multiple VCNs in the Same Region](#).

### Highlights

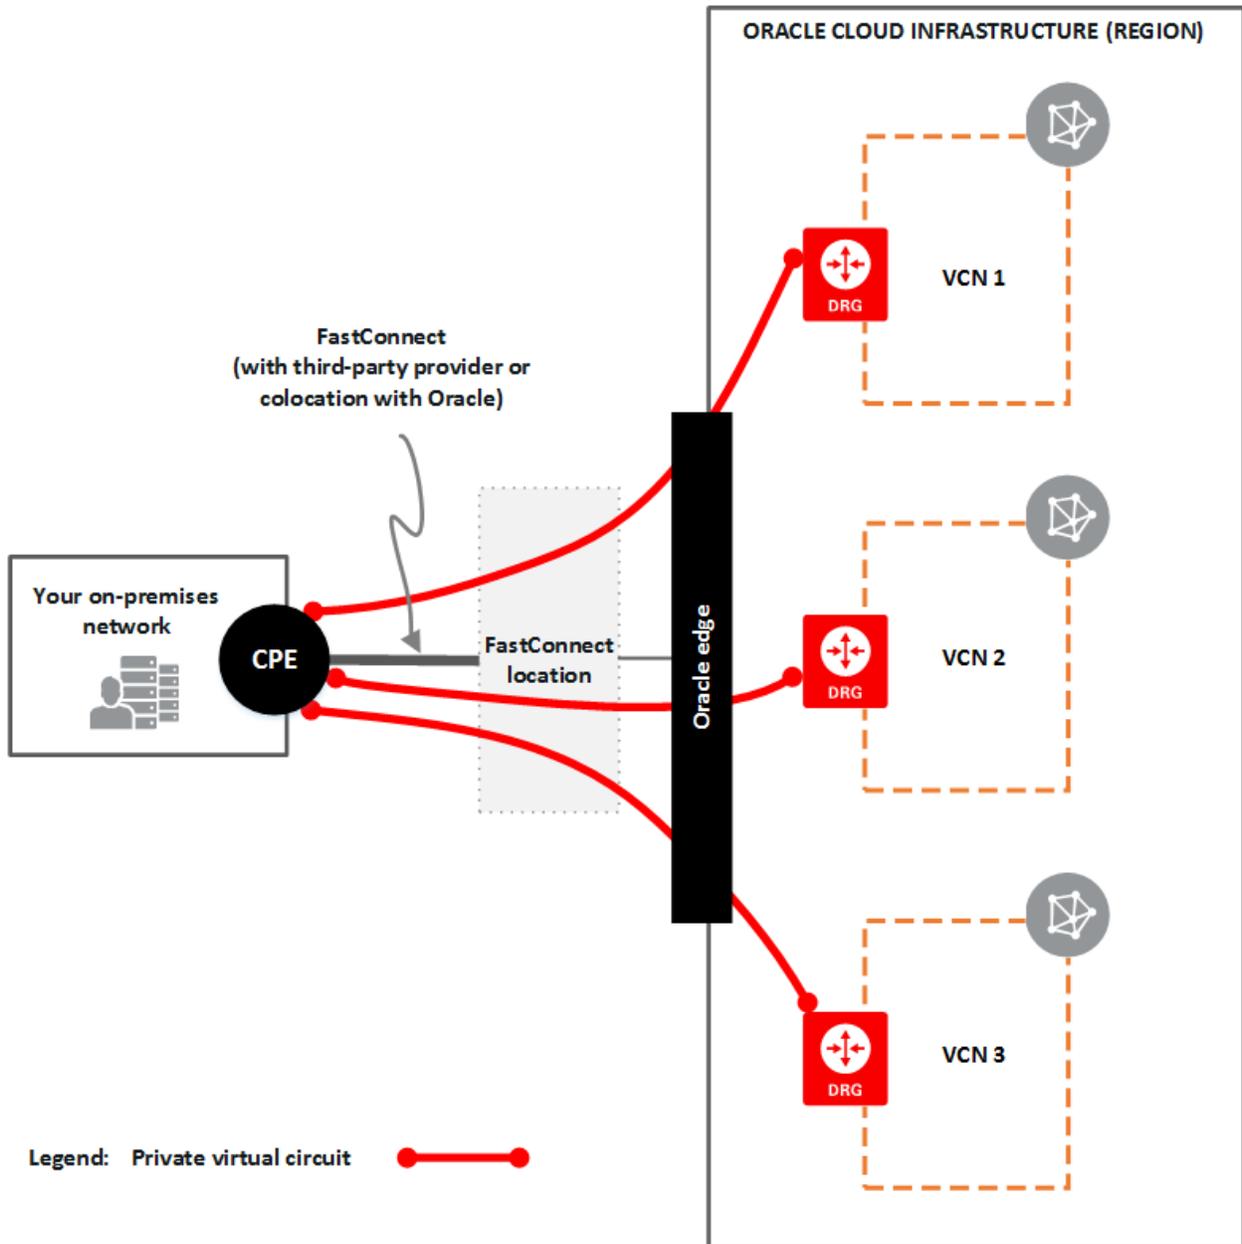
- You can use a single FastConnect to connect your on-premises network with *multiple* VCNs in the same region. The scenario is supported only for FastConnect through a [third-party provider](#) or through [colocation with Oracle](#). You need at least one physical connection (cross-connect) in your connection.
- The VCNs must be in the same region and same tenancy. The VCNs can be in the same compartment or different ones in the tenancy. For accurate routing, the CIDR blocks of the various subnets of interest in the on-premises network and VCNs must not overlap.
- Each VCN has its own dynamic routing gateway (DRG) and private virtual circuit. You must use a different VLAN and different set of BGP IP addresses for each private virtual circuit.
- You can also use FastConnect public peering to give your on-premises network access to public endpoints of Oracle services. In this case, you set up a single public virtual circuit, which covers all the VCNs. With public peering, make sure to configure your edge device (also known as your *customer-premises equipment* or *CPE*) to prefer FastConnect over your ISP for the Oracle Cloud Infrastructure public IP prefixes. Or, if you plan to also set up [private access to Oracle services through one of the VCNs](#), see

the important routing details in [Routing Details for Connections to Your On-Premises Network](#).

### Overview of the Scenario

In this scenario, you have a single FastConnect that connects your existing on-premises network to Oracle Cloud Infrastructure. That FastConnect has at least one physical connection, or cross-connect.

In Oracle Cloud Infrastructure, you have multiple VCNs, all in the same region. Each VCN has its own DRG. For each VCN, there's a private virtual circuit that runs on the FastConnect and terminates at your CPE on one end, and on the VCN's DRG on the other end. The private virtual circuit enables communication that uses private IP addresses between the VCN and the on-premises network. See the following diagram.



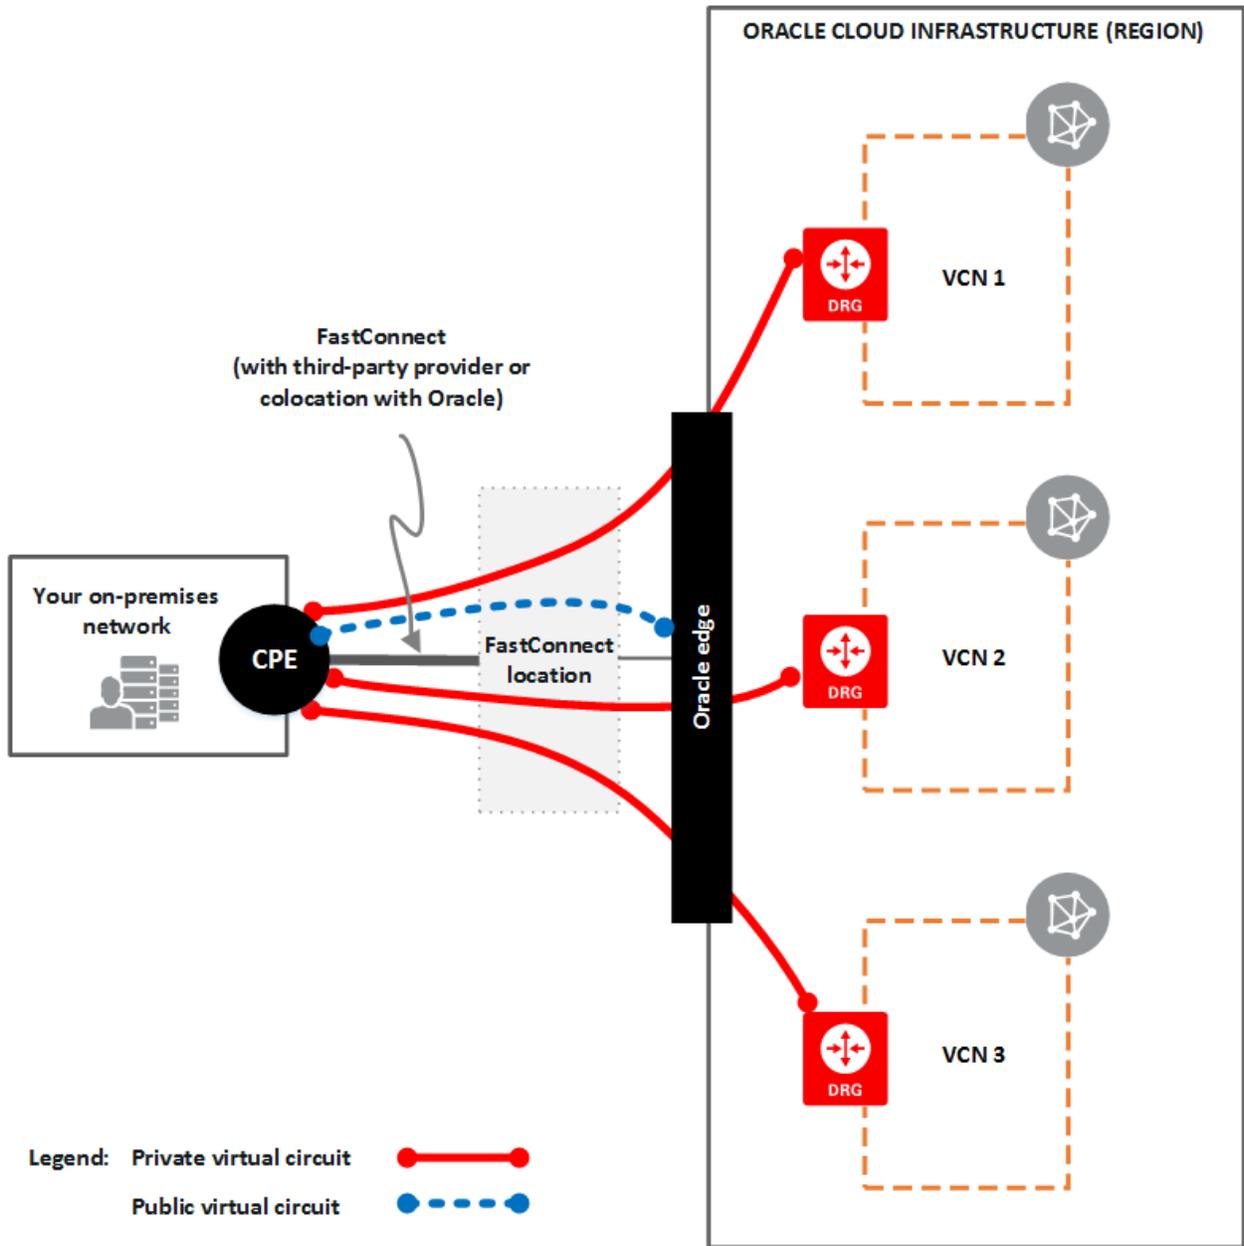
For example, imagine that each department in your organization has its own subnet in your on-premises network and a corresponding departmental VCN in Oracle Cloud Infrastructure. You want to enable private communication between each department's subnet and VCN over the FastConnect.

Or, perhaps all the departments need to communicate with all the VCNs. For example, instead perhaps the VCNs are for separate development, test, and production environments, and each department needs access to all three VCNs.

The FastConnect and virtual circuits give you the general private connection where none of the traffic traverses the internet. You can separately control which on-premises subnets and VCNs can communicate by configuring route rules in your on-premises network and VCN [route tables](#). You can optionally configure VCN [security rules](#) and other firewalls that you maintain to allow only certain types of traffic (such as SSH) between your on-premises network and VCN.

### Public Peering

You can also set up public peering on that same FastConnect by creating a public virtual circuit. In the following diagram, the public virtual circuit is shown separate from the private virtual circuits. It terminates at Oracle's edge. The public virtual circuit enables communication *that uses public IP addresses* but does not traverse the internet. If a given VCN happens to also have an internet gateway, Oracle's edge prefers the FastConnect route over the VCN's internet gateway. For other important details about how you can control route preferences when you have multiple connections between your on-premises network and Oracle, see [Routing Details for Connections to Your On-Premises Network](#).



When you set up public peering for your FastConnect, the public IP prefixes that you designate for the public virtual circuit are advertised to *all* the VCNs in your tenancy. The routes advertised to your on-premises network are all the Oracle Cloud Infrastructure public IP addresses (including the CIDRs for each of the VCNs in the tenancy).



### Important

Your network will receive Oracle's public IP addresses through both FastConnect and your Internet Service Provider (ISP). When configuring your edge, make sure to give higher preference to FastConnect over your ISP, or you will not receive the benefits of FastConnect. If you plan to also set up [private access to Oracle services through one of the VCNs](#), see the important routing details in [Routing Details for Connections to Your On-Premises Network](#).

For more information, see [Logical Connection: Public Virtual Circuit](#).

## General Setup Process

The setup process and instructions are in these topics, based on your particular FastConnect setup:

- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)

However, remember that:

- You set up a separate DRG for each VCN. A DRG can be attached to only a single VCN, and each VCN can be attached to only a single DRG.
- You set up a separate private virtual circuit for each DRG.

- For each private virtual circuit, **you must specify a different VLAN and a different set of BGP IP addresses.**
- When you configure your CPE, you can advertise the same on-premises routes to each VCN, or different ones, based on your own requirements.

## VCNs and Subnets

This topic describes how to manage virtual cloud networks (VCNs) and the subnets in them. This topic uses the terms *virtual cloud network*, *VCN*, and *cloud network* interchangeably. The Console uses the term *Virtual Cloud Network*, whereas for brevity the API uses *VCN*.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of VCNs and Subnets

A VCN is a software-defined network that you set up in the Oracle Cloud Infrastructure data centers in a particular region. A subnet is a subdivision of a VCN. For an overview of VCNs, allowed size, default VCN components, and scenarios for using a VCN, see [Overview of Networking](#).

You can privately connect a VCN to another VCN so that the traffic does not traverse the internet. The CIDRs for the two VCNs must not overlap. For more information, see [Access to Other VCNs: Peering](#). For an example of an advanced routing scenario that involves the peering of multiple VCNs, see [Transit Routing: Access to Multiple VCNs in the Same Region](#).

Each subnet in a VCN consists of a contiguous range of IPv4 addresses that do not overlap with other subnets in the VCN. Example: 172.16.1.0/24. The first two IPv4 addresses and the last in the subnet's CIDR are reserved by the Networking service. You can't change the size of

the subnet after creation, so it's important to think about the size of subnets you need before creating them.

Subnets act as a unit of configuration: all instances in a given subnet use the same route table, security lists, and DHCP options. For more information, see [Default Components that Come With Your VCN](#).

Subnets can be either public or private (see [Public vs. Private Subnets](#)). You choose this during subnet creation, and you can't change it later.

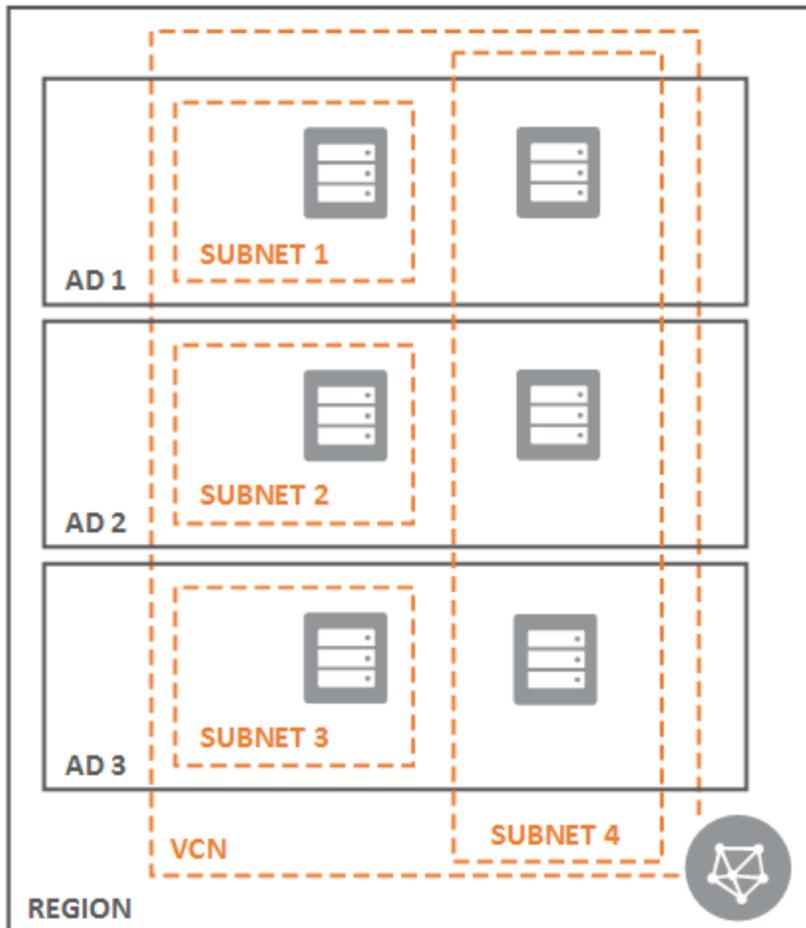
You can think of each Compute instance as residing in a subnet. But to be precise, each instance is actually attached to a [virtual network interface card \(VNIC\)](#), which in turn resides in the subnet and enables a network connection for that instance.

IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

### **About Regional Subnets**

Originally subnets were designed to cover only one availability domain (AD) in a region. They were all *AD-specific*, which means the subnet's resources were required to reside in a particular availability domain. Now subnets can be either *AD-specific* or *regional*. You choose the type when you create the subnet. Both types of subnets can co-exist in the same VCN. In the following diagram, subnets 1-3 are AD-specific, and subnet 4 is regional.

## Subnets can be regional or specific to an AD



Aside from the removal of the AD constraint, regional subnets behave the same as AD-specific subnets. **Oracle recommends using regional subnets** because they're more flexible. They make it easier to efficiently divide your VCN into subnets while also designing for availability domain failure.

When you create a resource such as a Compute instance, you choose which availability domain the resource will be in. From a virtual networking standpoint, you must also choose

which VCN and subnet the instance will be in. You can either choose a regional subnet, or choose an AD-specific subnet that matches the AD you chose for the instance.



### Warning

If anyone in your organization implements a regional subnet, be aware that you **may need to update any client code that works with Networking service subnets and private IPs**. There are possible breaking API changes. For more information, see the [regional subnet release note](#).

## Working with VCNs and Subnets

One of the first things you do when working with Oracle Cloud Infrastructure resources is create a VCN with one or more subnets. You can easily get started in the Console with a simple VCN and some related resources that enable you to launch and connect to an instance. See [Tutorial - Launching Your First Linux Instance](#) or [Tutorial - Launching Your First Windows Instance](#).

For the purposes of access control, when you create a VCN or subnet, you must specify the compartment where you want the resource to reside. Consult an administrator in your organization if you're not sure which compartment to use.

You may optionally assign friendly names to the VCN and its subnets. The names don't have to be unique, and you can change them later. Oracle automatically assigns each resource a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

You can also add a DNS label for the VCN and each subnet, which are required if you want the instances to use the *Internet and VCN Resolver* feature for DNS in the VCN. For more information, see [DNS in Your Virtual Cloud Network](#).

When you create a subnet, you may optionally specify a [route table](#) for the subnet to use. If you don't, the subnet uses the cloud network's default route table. You can [change which route table the subnet uses](#) at any time.

Also, you may optionally specify one or more [security lists](#) for the subnet to use (up to five). If you don't specify any, the subnet uses the cloud network's [default security list](#). You can [change which security list the subnet uses](#) at any time. Remember that the [security rules](#) are enforced at the instance level, even though the list is associated at the subnet level. [Network security groups](#) are an alternative to security lists and let you apply a set of security rules to a set of resources that all *have the same security posture*, instead of all the resources in a particular subnet.

You may optionally specify a [set of DHCP options](#) for the subnet to use. All instances in the subnet will receive the configuration specified in that set of DHCP options. If you don't specify a set, the subnet uses the cloud network's default set of DHCP options. You can [change which set of DHCP options the subnet uses](#) at any time.

To delete a subnet, it must contain no resources (no instances, [load balancers](#), [DB systems](#), and [orphaned mount targets](#)). For more details, see [Subnet or VCN Deletion](#).

To delete a VCN, its subnets must contain no resources. Also, the VCN must have no attached gateways. If you're using the Console, there's a "Delete All" process you can use after first ensuring the subnets are empty. See [To delete a VCN](#).

For information about the number of VCNs and subnets you can have, see [Service Limits](#).

### **Required IAM Policy**

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

### Using the Console

#### To create a VCN



#### Note

The following procedure creates a VCN without any subnets or gateways for access. You must manually create the subnets and other resources before you can use the VCN. For a quick procedure that creates a VCN that you can try out immediately (that is, with subnets and an internet gateway), see [Scenario A: Public Subnet](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Virtual Cloud Network**.
4. Enter the following:
  - **Name:** A friendly name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** Leave as is.
  - **Create Virtual Cloud Network Only:** Make sure this radio button is selected (the default).

- **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the Government Cloud. For more information, see [IPv6 Addresses](#).
- **CIDR Block:** A single, contiguous CIDR block for the VCN. For example: 172.16.0.0/16. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
- **Use DNS Hostnames in this VCN:** Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (`<VCN DNS label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create Virtual Cloud Network**.

The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose.

Next you'll typically want to create one or more subnets in the cloud network.

### To create a subnet

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Create Subnet**.
4. In the **Create Subnet** dialog box, you specify the resources to associate with the

subnet (for example, a route table, and so on). By default, the subnet will be created in the current compartment, and you'll choose the resources from the same compartment. Click the **click here** link in the dialog box if you want to enable compartment selection for the subnet and each of those resources.

Enter the following:

- **Create in Compartment:** If you've enabled compartment selection, specify the compartment where you want to put the subnet.
- **Name:** A friendly name for the subnet. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
- **Regional or AD-specific subnet:** Oracle recommends creating only [regional subnets](#), which means that the subnet can contain resources in any of the region's availability domains. If you instead choose **Availability Domain-Specific** (the only kind of subnet that Oracle originally offered), you must also specify an availability domain. This choice means that any instances or other resources later created in this subnet must also be in that availability domain.
- **CIDR Block:** A single, contiguous CIDR block for the subnet (for example, 172.16.0.0/24). Make sure it's within the cloud network's CIDR block and doesn't overlap with any other subnets. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
- **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
- **Route Table:** The route table to associate with the subnet. If you've enabled compartment selection, under **Route Table Compartment**, you must specify the compartment that contains the route table.
- **Private or public subnet:** This controls whether VNICs in the subnet can have public IP addresses. For more information, see [Access to the Internet](#).
- **Use DNS Hostnames in this Subnet:** This option is available only if you provided a DNS label for the VCN during creation. The option is required for

assignment of DNS hostnames to hosts in the subnet, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the subnet, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the subnet (`<subnet_DNS_label>.<VCN_DNS_label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).

- **DHCP Options:** The set of DHCP options to associate with the subnet. If you've enabled compartment selection, under **DHCP Options Compartment**, you must specify the compartment that contains the set of DHCP options.
  - **Security Lists:** One or more security lists to associate with the subnet. If you've enabled compartment selection, you must specify the compartment that contains the security list.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create**.
- The subnet is then created and displayed on the **Subnets** page in the compartment you chose.

### To edit a subnet

You can change these characteristics of a subnet:

- Name
- Which set of DHCP options the subnet uses
- Which route table the subnet uses
- Which security lists the subnet uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Subnets**.
4. Click the subnet you're interested in.
5. Click **Edit**.
6. Make your changes.
7. Click **Save Changes**.  
The changes take effect within a few seconds.

### To delete a subnet

Prerequisite: The subnet must have no instances, [load balancers](#), [DB systems](#), and [orphaned mount targets](#) in it. For more information, see [Subnet or VCN Deletion](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Subnets**.
4. Click the subnet you're interested in.
5. Click **Terminate**.
6. Confirm when prompted.

If the subnet is empty, its state changes to TERMINATING briefly and then TERMINATED. If the subnet is not empty, you get an error indicating that there are still instances or other resources in it that you must delete first.

### To delete a VCN



#### Important

The Console has an easy "Delete all" process that deletes a VCN and its related Networking resources (subnets, route tables, security lists, sets of DHCP options, internet gateway, and so on). If the VCN is attached to a dynamic routing gateway (DRG), the attachment is deleted, but the DRG remains.

The "Delete All" process deletes one resource at a time and takes a minute or two. A progress report is displayed to show you what's been deleted so far.

Before using the "Delete All" process, make sure there are no instances, [load balancers](#), [DB systems](#), or [orphaned mount targets](#) in any of the subnets. For more information, see [Subnet or VCN Deletion](#).

If there are still resources in any subnet, or if you don't have permission to delete a particular Networking resource, the "Delete All" process stops and an error message is displayed. Any resources deleted up to that point cannot be restored. You might need to contact your tenancy administrator to help you delete any remaining resources.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Terminate**.

The resulting dialog box displays a list of the resources to be deleted. The list doesn't include the [default components that come with the VCN](#), but they will be deleted along with the VCN.

4. Click **Delete All**.

The process begins. The progress is displayed as each resource is deleted.

5. When the process is complete, click **Close**.

### To manage tags for a VCN

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To manage tags for a subnet

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click the subnet you're interested in.
4. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To move a VCN to a different compartment

You can move a VCN from one compartment to another. When you move a VCN, its associated

VNICs, private IPs, and ephemeral IPs move with it to the new compartment. For more information, see [Moving Resources to a Different Compartment](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Find the VCN in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.
5. If there are alarms monitoring the VCN, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.

### To move a subnet to a different compartment

You can move a subnet from one compartment to another. For more information, see [Moving Resources to a Different Compartment](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Find the subnet in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your VCNs, use these operations:

- [ListVcns](#)
- [CreateVcn](#)
- [GetVcn](#)
- [UpdateVcn](#)
- [DeleteVcn](#): Deletes only the VCN and not its related resources. For more information, see [Subnet or VCN Deletion](#). Note that the Console offers a "Delete All" process that makes it easy to delete the VCN and its related resources. See [To delete a VCN](#).
- [ChangeVcnCompartment](#)

To manage a VCN's subnets, use these operations:

- [ListSubnets](#)
- [CreateSubnet](#)
- [GetSubnet](#)
- [UpdateSubnet](#)
- [DeleteSubnet](#)
- [ChangeSubnetCompartment](#)

## Ways to Secure Your Network

There are several ways you can control security for your cloud network and compute instances:

- **Public versus private subnets:** You can designate a subnet to be private, which means instances in the subnet cannot have public IP addresses. For more information, see [Public vs. Private Subnets](#).
- **Security rules:** To control packet-level traffic in and out of an instance. You configure security rules in the Oracle Cloud Infrastructure API or Console. To implement security rules, you can use network security groups or security lists. For more information, see [Security Rules](#).

- **Firewall rules:** To control packet-level traffic in/out of an instance. You configure firewall rules directly on the instance itself. Notice that Oracle Cloud Infrastructure images that run Oracle Linux automatically include default rules that allow ingress on TCP port 22 for SSH traffic. Also, the Windows images include default rules that allow ingress on TCP port 3389 for Remote Desktop access. For more information, see [Oracle-Provided Images](#).

**Important**

Firewall rules and security rules both operate at the instance level. However, you configure security lists at the subnet level, which means all resources in a given subnet have the same set of security list rules. Also, the security rules in a network security group apply only to the resources in the group.

When troubleshooting access to an instance, make sure all of the following items are set correctly: the network security groups that the instance is in, the security lists associated with the instance's subnet, and the instance's firewall rules.

If your instance is running Oracle Linux 7, you need to use [firewalld](#) to interact with the iptables rules.

For your reference, here are commands for opening a port (1521 in this example):

```
sudo firewall-cmd --zone=public --permanent --add-
port=1521/tcp
```

```
sudo firewall-cmd --reload
```

For instances with an iSCSI boot volume, the preceding `--reload` command can cause problems. For details and a workaround, see [Instances experience system hang after running firewall-cmd --reload](#).

- **Gateways and route tables:** To control general traffic flow from your cloud network to outside destinations (the internet, your on-premises network, or another VCN). You configure your cloud network's gateways and route tables in the Oracle Cloud

Infrastructure API or Console. For more information about the gateways, see [Networking Components](#). For more information about route tables, see [Route Tables](#).

- **IAM policies:** To control who has access to the Oracle Cloud Infrastructure API or Console itself. You can control the type of access, and which cloud resources can be accessed. For example, you can control who can set up your network and subnets, or who can update route tables, network security groups, or security lists. You configure policies in the Oracle Cloud Infrastructure API or Console. For more information, see [Access Control](#).

## Access Control

This topic gives basic information about using compartments and IAM policies to control access to your cloud network.

## Compartments and Your Cloud Network

Anytime you create a cloud resource such as a virtual cloud network (VCN) or compute instance, you must specify which IAM *compartment* you want the resource in. A compartment is a collection of related resources that can be accessed only by certain groups that have been given permission by an administrator in your organization. The administrator will create compartments and corresponding IAM policies to control which users in your organization have access to which compartments. Ultimately, the goal is to ensure that each person has access to only the resources they need.

If your company is starting to try out Oracle Cloud Infrastructure, only a few people need to create and manage the VCN and its components, launch instances into the VCN, and attach block storage volumes to those instances. Those few people need access to *all* these resources, so all those resources could be in the same compartment.

In an enterprise production environment with a VCN, your company will want to use multiple compartments to make it easier to control access to certain types of resources. For example, your administrator could create `Compartment_A` for your VCN and other networking components. The administrator could then create `Compartment_B` for all the compute instances and block storage volumes that the HR organization uses, and `Compartment_C` for

all the instances and block storage volumes that the Marketing organization uses. The administrator would then create IAM policies that give users only the level of access they need in each compartment. For example, the HR instance administrator is not entitled to modify the existing cloud network. So they would have full permissions for `Compartment_B`, but limited access to `Compartment_A` (just what's required to launch instances into the network). If they tried to modify other resources in `Compartment_A`, the request would be denied.

Network resources such as VCNs, subnets, route tables, security lists, service gateways, NAT gateways, VPNs, and FastConnect connections can be [moved from one compartment to another](#). When you move a resource to a new compartment, inherent policies apply immediately.

For more information about compartments and how to control access to your cloud resources, see "Setting Up Your Tenancy" in the *Oracle Cloud Infrastructure Getting Started Guide* and [Overview of Oracle Cloud Infrastructure Identity and Access Management](#).

### IAM Policies for Networking

The simplest approach to granting access to Networking is the policy listed in [Let network admins manage a cloud network](#). It covers the cloud network and all the other Networking components (subnets, security lists, route tables, gateways, and so on). To also give network admins the ability to launch instances (to test network connectivity), see [Let users launch Compute instances](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For reference material for writing more detailed policies for Networking, see [Details for the Core Services](#).

### Individual Resource-Types

If you'd like, you can write policies that focus on individual resource-types (for example, just security lists) instead of the broader `virtual-network-family`. Be aware that the `instance-family` resource-type also includes several permissions for VNICs, which reside in a subnet

but attach to an instance. For more information, see [For instance-family Resource Types and Virtual Network Interface Cards \(VNICs\)](#).

There's a resource-type called `local-peering-gateways` that is included within `virtual-network-family` and includes two other resource-types related to local VCN peering (within region):

- `local-peering-from`
- `local-peering-to`

The `local-peering-gateways` resource-type covers all permissions related to local peering gateways (LPGs). The `local-peering-from` and `local-peering-to` resource-types are for granting permission to *connect* two LPGs and form a peering relationship within a single region. For more information, see [Local VCN Peering \(Within Region\)](#).

Similarly, there's a resource-type called `remote-peering-connections` that is included within `virtual-network-family` and includes two other resource-types related to remote VCN peering (across regions):

- `remote-peering-from`
- `remote-peering-to`

The `remote-peering-connections` resource-type covers all permissions related to remote peering connections (RPCs). The `remote-peering-from` and `remote-peering-to` resource-types are for granting permission to *connect* two RPCs and form a peering relationship across regions. For more information, see [Remote VCN Peering \(Across Regions\)](#).

### Nuances of Different Verbs

If you'd like, you can write policies that limit the level of access by using a different policy verb ( `manage` versus `use`, and so on). If you do, there are some nuances to understand about the policy verbs for Networking.

Be aware that the `inspect` verb not only returns general information about the cloud network's components (for example, the name and OCID of a security list, or of a route table). It also includes the contents of the component (for example, the actual rules in the security list, the routes in the route table, and so on).

Also, the following types of abilities are available only with the `manage` verb, not the `use` verb:

- Update (enable/disable) `internet-gateways`
- Update `security-lists`
- Update `route-tables`
- Update `dhcp-options`
- Attach a dynamic routing gateway (DRG) to a virtual cloud network (VCN)
- Create an IPSec connection between a DRG and customer-premises equipment (CPE)
- Peer VCNs



### Important

Each VCN has various components that directly affect the behavior of the network (route tables, security lists, DHCP options, Internet Gateway, and so on). When you create one of these components, you establish a relationship between that component and the VCN, which means you must be allowed in a policy to both create the component and manage the VCN itself. However, the ability to *update* that component (to change the route rules, security list rules, and so on) does NOT require permission to manage the VCN itself, even though changing that component can directly affect the behavior of the network. This discrepancy is designed to give you flexibility in granting least privilege to users, and not require you to grant excessive access to the VCN just so the user can manage other components of the network. Be aware that by giving someone the ability to update a particular type of component, you're implicitly trusting them with controlling the network's behavior.

For more information about policy verbs, see [Verbs](#).

## Security Rules

The Networking service offers two virtual firewall features that both use *security rules* to control traffic at the packet level. The two features are:

- **Security lists:** The original virtual firewall feature from the Networking service.
- **Network security groups (NSGs):** A subsequent feature designed for application components that have different security postures. NSGs are supported only for [specific services](#).

These two features offer different ways to apply security rules to a set of virtual network interface cards (VNICs) in the virtual cloud network (VCN).

This topic summarizes basic differences between the two features. It also explains important security rule concepts that you need to understand. How you create, manage, and apply security rules varies between security lists and network security groups. For implementation details, see these related topics:

- [Security Lists](#)
- [Network Security Groups](#)

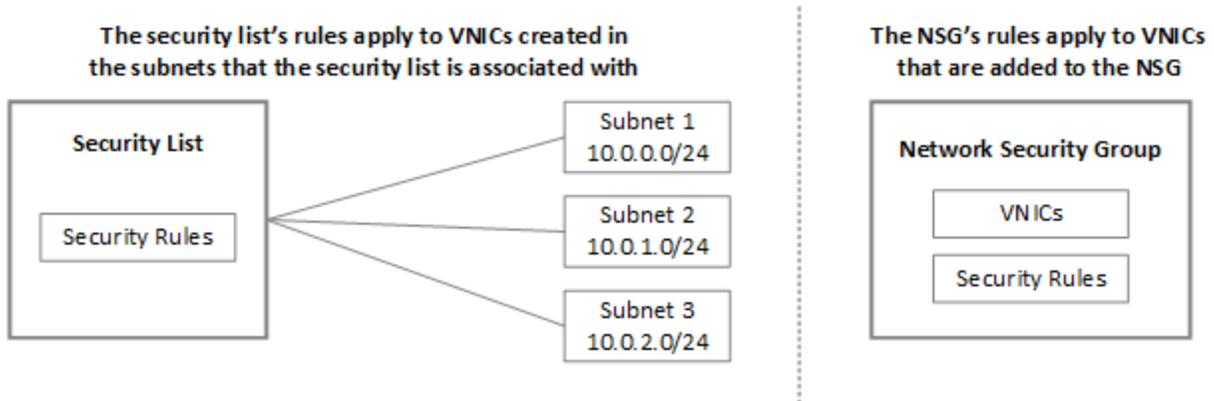
## Comparison of Security Lists and Network Security Groups

*Security lists* let you define a set of security rules that applies to all the VNICs in an entire *subnet*. To use a given security list with a particular subnet, you *associate* the security list with the subnet either during subnet creation or later. A subnet can be associated with a maximum of five security lists. Any VNICs that are created in that subnet are subject to the security lists associated with the subnet.

*Network security groups (NSGs)* let you define a set of security rules that applies to a *group of VNICs of your choice* (or the VNICs' [parent resources such as load balancers or DB systems](#)). For example: the VNICs that belong to a set of Compute instances that all have the same security posture. To use a given NSG, you add the VNICs of interest to the group. Any VNICs

added to that group are subject to that group's security rules. A VNIC can be added to a maximum of five NSGs.

The following diagram illustrates the concept.



**Oracle recommends using NSGs instead of security lists because NSGs let you separate the VCN's subnet architecture from your application security requirements.**

However, you can use both security lists and NSGs together if you want. For more information, see [If You Use Both Security Lists and Network Security Groups](#).

### About VNICs and Parent Resources

A [VNIC](#) is a Networking service component that enables a networked resource such as a Compute instance to connect to a virtual cloud network (VCN). The VNIC determines how the instance connects with endpoints inside and outside the VCN. Each VNIC resides in a subnet in a VCN.

When you create a Compute instance, a VNIC is automatically created for the instance in the instance's subnet. The instance is considered to be the *parent resource* for the VNIC. You can add additional (secondary) VNICs to a Compute instance. For this reason, an instance's VNICs are displayed prominently as part of a Compute instance's related resources in the Console.

There are other types of parent resources that you can create that also result in a VNIC automatically being created. For example: when you create a load balancer, the Load Balancing service automatically creates VNICs for balancing traffic across the backend set. Also, when you create a DB system, the Database service automatically creates VNICs as DB system nodes. Those services create and manage those VNICs for you. For this reason, those VNICs are not readily apparent in the Console the same way VNICs are for Compute instances.

To use an NSG, you put VNICs of your choice into the group. However, you typically work *with the parent resource* when you add a VNIC to the group, not the VNIC itself. For example, when you create a Compute instance, you can optionally specify an NSG for the instance. Although you conceptually put the instance in the group, you're actually putting the instance's *primary VNIC* in the network security group. The group's security rules apply to that VNIC, not the instance. Also, if you add a secondary VNIC to the instance, you can optionally specify an NSG for that VNIC, and the rules apply to that VNIC, not the instance. Note that all the VNICs in a given NSG must be in the VCN that the NSG belongs to.

Likewise, when you put a load balancer in a network security group, you conceptually put the load balancer in the group. But you're actually putting VNICs managed by the Load Balancing service into the network security group.

You manage the VNIC membership of an NSG *at the parent resource*, and not at the NSG itself. In other words, to add a parent resource to an NSG, you execute the action on the *parent resource* (by specifying which NSGs the parent resource should be added to). You do not execute the action on the NSG (by specifying which VNICs or parent resources should be added to the NSG). Similarly, to remove a VNIC from an NSG, you execute that action by updating the parent resource, not the NSG. For example, to add an existing Compute instance's VNIC to an NSG, you update that VNIC's properties and specify the NSG. For example, with the REST API, you call `UpdateVnic`. In the Console, you view the instance and then the VNIC of interest, and then edit the VNIC's properties there.

For a list of parent resources that support the use of NSGs, see [Support for Network Security Groups](#).

### Network Security Group as Source or Destination of a Rule

There's an important difference in how you can write security rules for NSGs compared to security lists.

When writing rules for an NSG, you can specify an *NSG* as the source of traffic (for ingress rules) or the traffic's destination (for egress rules). Contrast this with security list rules, where you specify a *CIDR* as the source or destination.

The ability to specify an NSG means that you can easily write rules to control traffic between two different NSGs. The NSGs must be in the same VCN.

Also, if you want to control traffic between *VNICs in a specific NSG*, you can write rules that specify the *rule's own NSG* as the source (for ingress rules) or destination (for egress rules).

For more information, see [Overview of Network Security Groups](#).

### REST API Differences

There are a few basic differences in the REST API model for NSGs compared to security lists. For more information, see [Using the API](#).

### Default Rules

Your VCN automatically comes with a [default security list](#) that contains several default security rules to help you get started using the Networking service. When you create a subnet, the default security list is associated with the subnet unless you specify a custom security list that you've already created in the VCN. For comparison, the VCN does NOT have a default network security group.

### Limits

The two features have different limits.

## Security List Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Security lists	VCN	300	300
Security lists	Subnet	5*	5*
Security rules	Security list	200 ingress rules* and 200 egress rules*	200 ingress rules* and 200 egress rules*
* Limit for this resource cannot be increased			

## Network Security Group Limits

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Network security groups	VCN	1000	1000
VNICs	Network security group	A given network security group can have as many VNICs as are in the VCN.  A given VNIC can belong to a maximum of 5 network security groups.*	A given network security group can have as many VNICs as are in the VCN.  A given VNIC can belong to a maximum of 5 network security groups.*

Resource	Scope	Monthly Universal Credits	Pay-as-You-Go or Promo
Security rules	Network security group	120 (total ingress plus egress)	120 (total ingress plus egress)
* Limit for this resource cannot be increased			

## Best Practices for Security Rules

### Use Network Security Groups

Oracle recommends using NSGs for components that all have the same security posture. For example, in a multi-tier architecture, you would have a separate NSG for each tier. A given tier's VNICs would all belong to that tier's NSG. Within a given tier, you might have a particular subset of the tier's VNICs that have additional, special security requirements. Therefore you would create another NSG for those additional rules, and place that subset of VNICs into both the tier's NSG and the additional NSG.

Oracle also recommends using NSGs because Oracle will prioritize NSGs over security lists when implementing future enhancements.

### Get Familiar with the Default Security List Rules

Your VCN automatically comes with a [default security list](#) that contains several default security rules to help you get started using the Networking service. Those rules exist because they enable basic connectivity. Even if you choose not to use security lists or the default security list, get familiar with the rules so you better understand the types of traffic that your networked cloud resources require. You might want to use those rules in your NSGs or any custom security lists that you set up.

The default security list does not include rules to enable ping. If you need to use ping, see [Rules to Enable Ping](#).

### Don't Delete Default Security Rules Indiscriminately

Your VCN might have subnets that use the default security list by default. Do not delete any of the list's default security rules unless you've first confirmed that resources in your VCN do not require them. Otherwise, you might disrupt your VCN's connectivity.

### Confirm That Your OS Firewall Rules Align with Your Security Rules

Your instances running Oracle-provided Linux images or Windows images also have OS firewall rules that control access to the instance. When troubleshooting access to an instance, make sure that all of the following items are set correctly:

- The rules in the network security groups that the instance is in
- The rules in the security lists associated with the instance's subnet
- The instance's OS firewall rules

For more information, see [Oracle-Provided Images](#).

If your instance is running Oracle Linux 7, you need to use [firewalld](#) to interact with the iptables rules. For your reference, here are commands for opening a port (1521 in this example):

```
sudo firewall-cmd --zone=public --permanent --add-port=1521/tcp

sudo firewall-cmd --reload
```

For instances with an iSCSI boot volume, the preceding `--reload` command can cause problems. For details and a workaround, see [Instances experience system hang after running firewall-cmd --reload](#).

### Use VNIC Metrics to Troubleshoot Packets Dropped Because of Security Rules

The Networking service offers [metrics for VNICs](#), which show the levels of VNIC traffic (packets and bytes). Two of the metrics are for ingress and egress packets that violate security rules and are therefore dropped. You can use these metrics to help you troubleshoot issues related to security rules and whether your VNICs are receiving the desired traffic.

## If You Use Both Security Lists and Network Security Groups

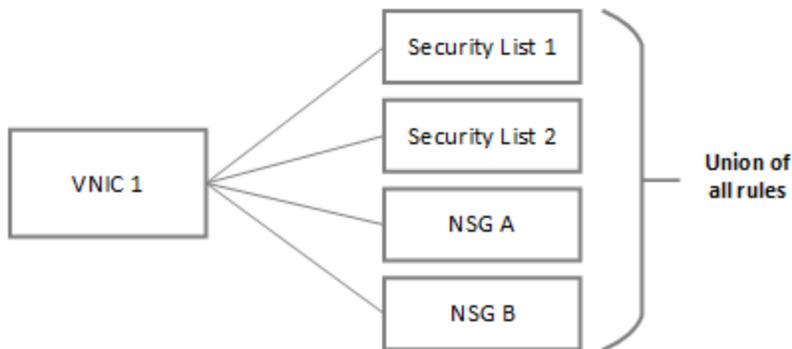
You can use security lists alone, network security groups alone, or both together. It depends on your particular security needs.

If you have security rules that you want to enforce for *all VNICs in a VCN*: the easiest solution is to put the rules in one security list, and then associate that security list with all subnets in the VCN. This way you can ensure that the rules are applied, regardless of who in your organization creates a VNIC in the VCN. If you like, you can use the VCN's [default security list](#), which automatically comes with the VCN and contains a set of essential rules by default.

If you choose to use *both* security lists and network security groups, the set of rules that applies to a given VNIC is the union of these items:

- The security rules in the security lists associated with the VNIC's subnet
- The security rules in all NSGs that the VNIC is in

The following diagram is a simple illustration of the idea.



A packet in question is allowed if *any rule in any of the relevant lists and groups* allows the traffic (or if the traffic is part of an existing [connection being tracked](#)). There's a caveat if the lists happen to contain both stateful and stateless rules that cover the same traffic. For more information, see [Stateful Versus Stateless Rules](#).

### Parts of a Security Rule

A security rule **allows** a particular type of traffic in or out of a VNIC. For example, a commonly used security rule allows ingress TCP port 22 traffic for establishing SSH connections to the instance (more specifically to the instance's VNICs). Without security rules, no traffic is allowed in and out of VNICs in the VCN.



#### Note

Security rules are not enforced for traffic involving the 169.254.0.0/16 CIDR block, which includes services such as iSCSI and instance metadata.

Each security rule specifies the following items:

- **Direction (ingress or egress):** Ingress is inbound traffic to the VNIC, and egress is outbound traffic from the VNIC. The REST API model for security lists is different from network security groups. With security lists, there is an `IngressSecurityRule` object and a separate `EgressSecurityRule` object. With network security groups, there is only a `SecurityRule` object, and the object's `direction` attribute determines whether the rule is for ingress or egress traffic.
- **Stateful or stateless:** If stateful, connection tracking is used for traffic matching the rule. If stateless, no connection tracking is used. See [Stateful Versus Stateless Rules](#).
- **Source type and source (ingress rules only):** The source you provide for an ingress rule depends on the source type you choose.

## Allowed source types

Source Type	Allowed Source
CIDR	The CIDR block where the traffic originates from. Use 0.0.0.0/0 to indicate all IP addresses. The prefix is required (for example, include the /32 if specifying an individual IP address).
Network Security Group	An NSG that is in the same VCN as this rule's NSG. This source type is available only if the rule belongs to an NSG and not a security list.
Service	Only for packets coming from an Oracle service through a <a href="#">service gateway</a> . The source service is the <a href="#">service CIDR label</a> that you're interested in.

- **Destination type and destination (egress rules only):** The destination you provide for an egress rule depends on the destination type you choose.

## Allowed destination types

Destination Type	Allowed Destination
CIDR	The CIDR block where the traffic is destined to. Use 0.0.0.0/0 to indicate all IP addresses. The prefix is required (for example, include the /32 if specifying an individual IP address).
Network Security Group	An NSG that is in the same VCN as this rule's NSG. This destination type is available only if the rule belongs to an NSG and not a security list.
Service	Only for packets going to an Oracle service through a <a href="#">service gateway</a> . The destination service is the <a href="#">service CIDR label</a> that you're interested in.

- **IP Protocol:** Either a single [IPv4 protocol](#) or "all" to cover all protocols.
- **Source port:** The port where the traffic originates from. For TCP or UDP, you can specify all source ports, or optionally specify a single source [port number](#), or a range.
- **Destination port:** The port where the traffic is destined to. For TCP or UDP, you can specify all destination ports, or optionally specify a single destination [port number](#), or a range.
- **ICMP type and code:** For ICMP, you can specify all types and codes, or optionally specify a single [type](#) with an optional code. If the type has multiple codes, create a separate rule for each code you want to allow.
- **Description** (NSG rules only): NSG security rules include an optional attribute where you can provide a friendly description of the rule. This is currently not supported for security list rules.

For examples of security rules, see [Networking Scenarios](#).

For the limit on the number of rules you can have, see [Limits](#).



### Note

If you're using NSGs, and two VNICs that are in the same VCN need to communicate *using their public IP addresses*, you must use the VNIC's public IP address and not the VNIC's NSG as the source (for ingress) or destination (for egress) in the relevant security rules. The packet is routed to the VCN's internet gateway, and at that point, the VNIC's NSG information is not available. Therefore a security rule that specifies the NSG as the source or destination will be ineffective in allowing that specific type of traffic.

## Stateful Versus Stateless Rules

When you create a security rule, you choose whether it's *stateful* or *stateless*. The difference is described in the next sections. The default is stateful. Stateless rules are recommended if you have a high-volume internet-facing website (for the HTTP/HTTPS traffic).

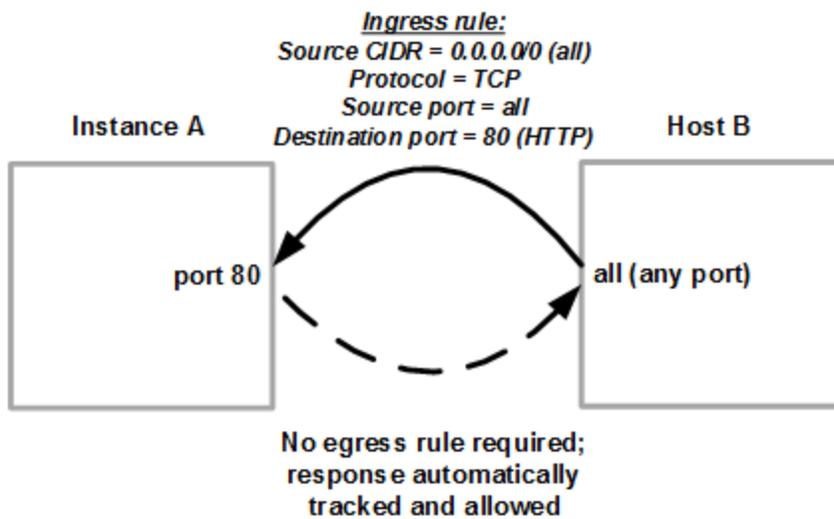
This section refers specifically to Compute instances and their traffic. However, the discussion is applicable to all types of resources with VNICs. See [About VNICs and Parent Resources](#).

### Stateful Rules

Marking a security rule as stateful indicates that you want to use connection tracking for any traffic that matches that rule. This means that when an instance receives traffic matching the stateful ingress rule, the response is tracked and automatically allowed back to the originating host, regardless of any egress rules applicable to the instance. And when an instance sends traffic that matches a stateful egress rule, the incoming response is automatically allowed, regardless of any ingress rules. For more details, see [Connection Tracking Details for Stateful Rules](#).

The figure below illustrates a stateful ingress rule for an instance that needs to receive and respond to HTTP traffic. Instance A and Host B are communicating (Host B could be any host, whether an instance or not). The stateful ingress rule allows traffic from any source IP address (0.0.0.0/0) to destination port 80 only (TCP protocol). No egress rule is required to allow the response traffic.

**Stateful Security Rule: Receive HTTP Traffic**



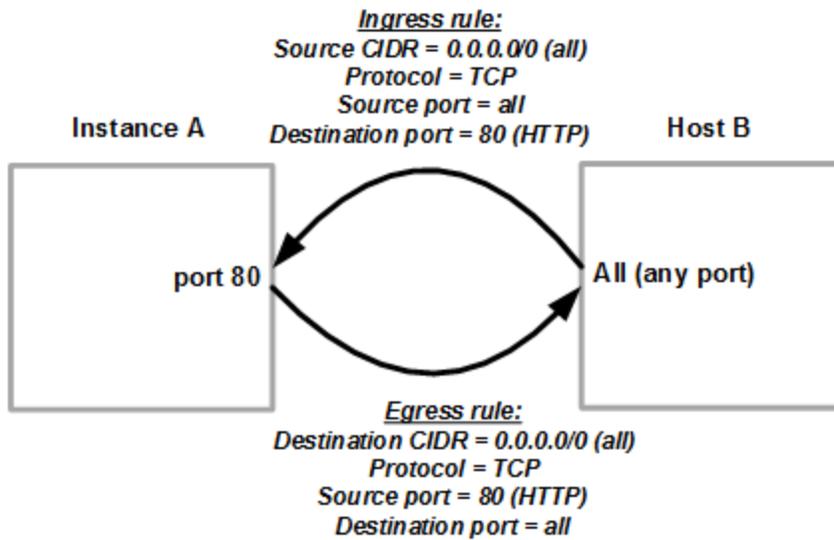
**Stateless Rules**

Marking a security rule as stateless indicates that you do NOT want to use connection tracking for any traffic that matches that rule. This means that response traffic is not automatically allowed. To allow the response traffic for a stateless ingress rule, you must create a corresponding stateless egress rule.

The next figure shows Instance A and Host B as before, but now with stateless security rules. As with the stateful rule in the preceding section, the stateless ingress rule allows traffic from all IP addresses and any ports, on destination port 80 only (using the TCP protocol). To allow the response traffic, there needs to be a corresponding stateless egress rule that allows traffic

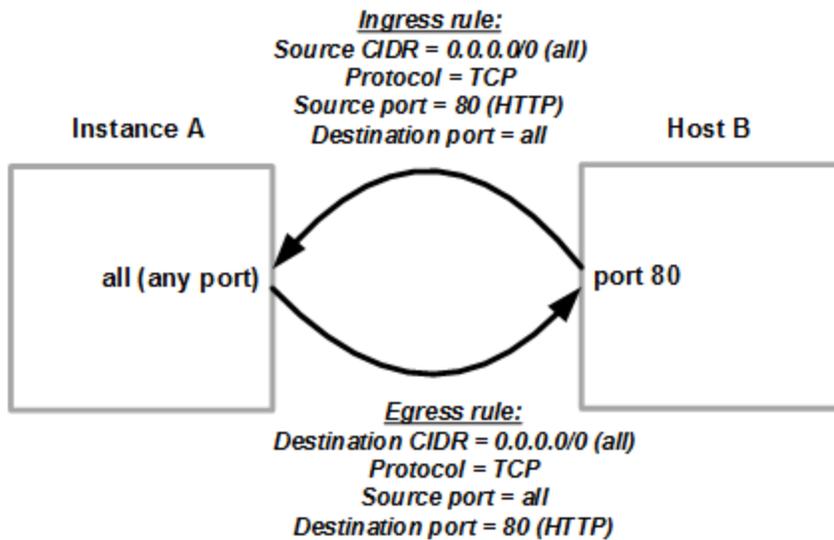
to any destination IP address (0.0.0.0/0) and any ports, from source port 80 only (using the TCP protocol).

**Stateless Security Rule: Receive HTTP Traffic**



If Instance A needs instead to *initiate* HTTP traffic and get the response, then a different set of stateless rules are required. As the next figure shows, the egress rule would have source port = all and destination port = 80 (HTTP). The ingress rule would then have source port 80 and destination port = all.

**Stateless Security Rule: Initiate HTTP Traffic**



If you were to use port binding on Instance A to specify exactly which port the HTTP traffic would come from, you could specify that as the source port in the egress rule and the destination port in the ingress rule.



### Note

If for some reason you use both stateful and stateless rules, and there's traffic that matches both a stateful and stateless rule in a particular direction (for example, ingress), the stateless rule takes precedence and the connection is not tracked. You would need a corresponding rule in the other direction (for example, egress, either stateless or stateful) for the response traffic to be allowed.

### Connection Tracking Details for Stateful Rules

Oracle uses connection tracking to allow responses for traffic that matches stateful rules (see [Stateful Versus Stateless Rules](#)). Each instance has a maximum number of concurrent connections that can be tracked, based on the instance's shape.

To determine response traffic for TCP, UDP, and ICMP, Oracle performs connection tracking on these items for the packet:

- Protocol
- Source and destination IP addresses
- Source and destination ports (for TCP and UDP only)



### Note

For other protocols, Oracle tracks only the protocol and IP addresses, and not the ports. This means that when an instance initiates traffic to another host and that traffic is allowed by egress security rules, any traffic that the instance subsequently receives from that host for a period is considered response traffic and is allowed.

### Enabling Path MTU Discovery Messages for Stateless Rules

If you decide to use stateless security rules to allow traffic to/from endpoints outside the VCN, it's important to add a security rule that allows ingress ICMP traffic type 3 code 4 from source 0.0.0.0/0 and any source port. This rule enables your instances to receive Path MTU Discovery fragmentation messages. This rule is critical for establishing a connection to your instances. Without it, you can experience connectivity issues. For more information, see [Hanging Connection](#).

### Rules to Enable Ping

The VCN's [default security list](#) contains several default rules, but not one to allow ping requests. If you want to ping an instance, make sure the instance's applicable security lists or NSGs include an *additional* stateful ingress rule to specifically allow ICMP traffic type 8 from the source network you plan to ping from. To allow ping access from the internet, use 0.0.0.0/0 for the source. Note that this rule for pinging is separate from the default ICMP-related rules in the default security list. Do not remove those rules.

### Rules to Handle Fragmented UDP Packets

Instances can send or receive UDP traffic. If a UDP packet is too large for the connection, it is fragmented. However, only the first fragment from the packet contains the protocol and port

information. If the security rules that allow this ingress or egress traffic specify a particular port number (source or destination), the fragments after the first one are dropped. If you expect your instances to send or receive large UDP packets, set both the source and destination ports for the applicable security rules to ALL (instead of a particular port number).

## Network Security Groups

The Networking service offers two virtual firewall features to control traffic at the packet level:

- **Network security groups:** Covered in this topic. Network security groups are supported only for [specific services](#).
- **Security lists:** The original type of virtual firewall offered by the Networking service. See [Security Lists](#).

Both of these features use *security rules*. For important information about how security rules work, and a general comparison of security lists and network security groups, see [Security Rules](#).

## Highlights

- Network security groups (NSGs) act as a virtual firewall for your Compute instances and [other kinds of resources](#). An NSG consists of a set of ingress and egress [security rules](#) that apply only to *a set of VNICs of your choice in a single VCN* (for example: all the Compute instances that act as web servers in the web tier of a multi-tier application in your VCN).
- Compared to security lists, NSGs let you separate your VCN's subnet architecture from your application security requirements. See [Comparison of Security Lists and Network Security Groups](#).
- At this time, you can use NSGs with Compute instances, load balancers, and DB systems. For more information, see [Support for Network Security Groups](#).
- NSG security rules function the same as security list rules. However, for an NSG security rule's source (for ingress rules) or destination (for egress rules), you can

specify an NSG instead of a CIDR. This means you can easily write security rules to control traffic between two NSGs *in the same VCN*, or traffic within a single NSG. See [Parts of a Security Rule](#).

- Unlike with security lists, the VCN does not have a default NSG. Also, each NSG you create is initially empty. It has no default security rules.
- Network security groups have separate and different limits compared to security lists. See [Limits](#).

### Support for Network Security Groups

NSGs are available for you to create and use. However, they are not yet supported by all the relevant Oracle Cloud Infrastructure services.

Currently, the following types of [parent resources](#) support the use of NSGs:

- **Compute instances:** When you create an instance, you can specify one or more NSGs for the instance's primary VNIC. If you add a secondary VNIC to an instance, you can specify one or more NSGs for that VNIC. You can also update existing VNICs on an instance so that they are in one or more NSGs.
- **Load balancers:** When you create a load balancer, you can specify one or more NSGs for the load balancer (not the backend set). You can also update an existing load balancer to use one or more NSGs.
- **DB systems:** When you create a DB system, you can specify one or more NSGs. You can also update an existing DB system to use one or more NSGs.
- **Autonomous Databases:** When you create an Autonomous Database using the [dedicated deployment](#) option, you can specify one or more NSGs. You can also update an existing dedicated deployment to use NSGs.

For resource types that do not yet support NSGs, continue to use security lists to control traffic in and out of those parent resources.

## Overview of Network Security Groups

A network security group (NSG) provides a virtual firewall for a set of cloud resources that all have the same security posture. For example: a group of Compute instances that all perform the same tasks and thus all need to use the same set of ports.

An NSG consists of two types of items (as illustrated in the following diagram):

- **VNICs:** One or more [VNICs](#) (for example, the VNICs attached to the set of Compute instances that all have the same security posture). All the VNICs must be in the VCN that the NSG belongs to. Also see [About VNICs and Parent Resources](#).
- **Security rules:** [Security rules](#) that define the types of traffic allowed in and out of the VNICs in the group. For example: ingress TCP port 22 SSH traffic from a particular source.

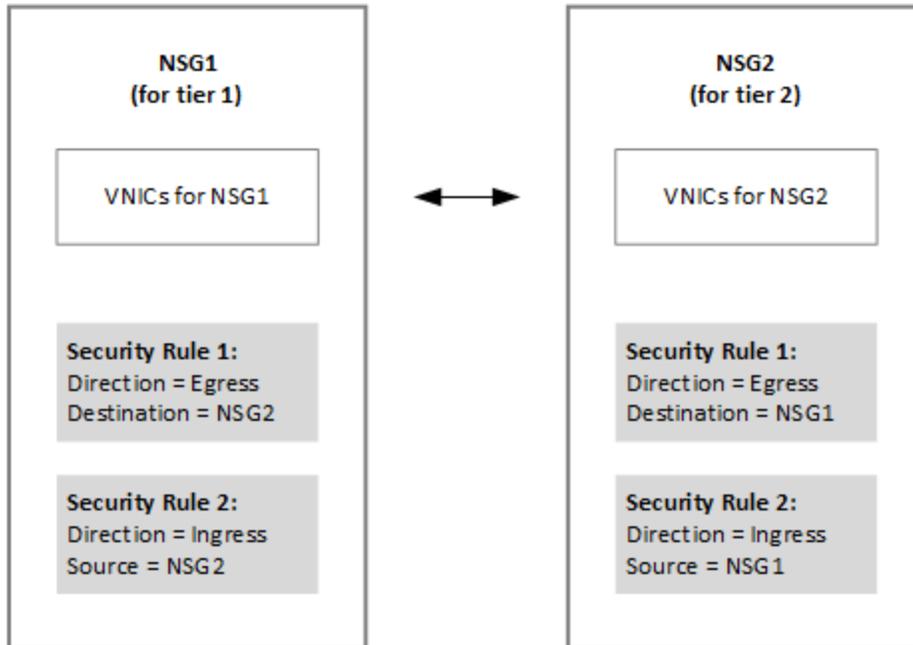
**The NSG's rules apply to VNICs  
that are added to the NSG**



If you have resources with different security postures *in the same VCN*, you can write NSG security rules to control traffic between the resources with one posture versus another. For example, in the following diagram, NSG1 has VNICs running in one tier of a multi-tier architecture application. NSG2 has VNICs running in a second tier. Both NSGs must belong to the same VCN. The assumption is that both NSGs need to initiate connections to the other NSG.

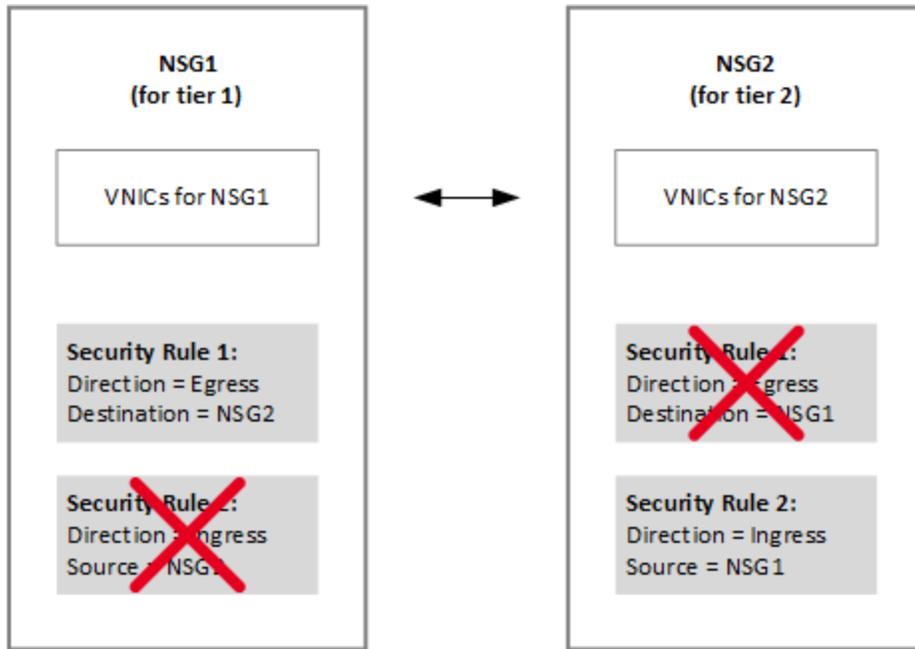
For NSG1, you set up egress security rules that specify NSG2 as the destination, and ingress security rules that specify NSG2 as the source. Likewise for NSG2, you set up egress security

rules that specify NSG1 as the destination, and ingress security rules that specify NSG1 as the source. The rules are assumed to be [stateful](#) in this example.



The preceding diagram assumes that each NSG needs to initiate connections to the other NSG.

The next diagram assumes that you instead want to only allow connections initiated from NSG1 to NSG2. To do that, remove the ingress rule from NSG1 and the egress rule from NSG2. The remaining rules do not allow connections initiated from NSG2 to NSG1.



The next diagram assumes that you want to control traffic between VNICs *in the same NSG*. To do that, set the rule's source (for ingress) or destination (for egress) as the rule's own NSG.



A VNIC can be in a maximum of five NSGs. A packet in question is allowed if *any rule in any of the VNIC's NSGs* allows the traffic (or if the traffic is part of an existing connection being tracked). There's a caveat if the lists happen to contain both stateful and stateless security rules that cover the same traffic. For more information, see [Stateful Versus Stateless Rules](#).

Network security groups are regional entities. For limits related to network security groups, see [Limits](#).

### Security Rules

If you're not yet familiar with the basics of NSG security rules, see these sections in the security rules topic:

- [Parts of a Security Rule](#)
- [Stateful Versus Stateless Rules](#)

### Working with Network Security Groups



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

#### General Process for Working with NSGs

1. Create an NSG.
2. Add security rules to the NSG.
3. Add parent resources (or more specifically, VNICs) to the NSG. You can do this when you create the parent resource, or you can update the parent resource and add it to one or more NSGs. When you create a Compute instance and add it to an NSG, the instance's primary VNIC is added to the NSG. Separately, you can create secondary VNICs and optionally add them to NSGs.

Before deleting an NSG, you must remove all VNICs from it.

See the next sections for more details.

#### Creating NSGs

Each VCN comes with a [default security list](#) that has default security rules in it to enable basic connectivity. However, there is no default NSG in the VCN.

When you create an NSG, it is initially empty, without any security rules or VNICs. If you're using the Console, you can add security rules to the NSG during creation.

You may optionally assign a friendly name to the NSG during creation. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the NSG a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

For the purposes of access control, you must specify the compartment where you want the NSG to reside. Consult an administrator in your organization if you're not sure which compartment to use. For more information, see [Access Control](#).

### Updating Security Rules and Group Membership

After the NSG is created, you can add or remove security rules to allow the types of ingress and egress traffic that the VNICs in the group require.

If you're familiar with security lists and use the REST API, note that the model for *updating* existing security rules is different between security lists and NSGs. With NSGs, each rule in a given group has a unique Oracle-assigned identifier (example: 04ABEC). When you call `UpdateNetworkSecurityGroupSecurityRules`, you provide the IDs of the specific rules that you want to update. For comparison, with security lists, the rules have no unique identifier. When you call `UpdateSecurityList`, you must pass in the *entire* list of rules, including rules that are not being updated in the call.

When you manage an NSG's VNIC membership, you do it as part of working with the parent resource, not the NSG itself. For more information, see [About VNICs and Parent Resources](#).

### Specifying an NSG in a Security Rule

As mentioned earlier in [Overview of Network Security Groups](#), you can specify an NSG as the source (for ingress rules) or destination (for egress rules) in a given NSG's security rule. The two NSGs must be in the same VCN. For example, if both NSG1 and NSG2 belong to the same VCN, you could add an ingress rule to NSG1 that lists NSG2 as the source. If someone deletes NSG2, the rule becomes invalid. The REST API uses an `isValid` Boolean in the `SecurityRule` object to convey that status.

### Deleting NSGs

To delete an NSG, it must not contain any VNICs or parent resources. When a parent resource (or a Compute instance VNIC) is deleted, it is automatically removed from the NSGs it was in. You might not have permission to delete a particular parent resource. Contact your administrator to determine who owns a given resource.

The Console displays a list of parent resources that are in an NSG, with a link to each parent resource. If the parent resource is a Compute instance, the Console also displays the instance's VNIC or VNICs that are in the NSG.

To remove a parent resource from its NSGs without deleting the resource, first view the parent resource's details in the Console. There you can see a list of the NSGs that the resource belongs to. From there, you can click **Edit** and remove the resource from all NSGs. If you're instead working with a Compute instance, view the details of the specific VNIC that you want to remove from the NSGs.

If you're using the REST API: the `ListNetworkSecurityGroupVnics` lists the parent resources and VNICs in an NSG. Use the resource's `Update` operation to remove the resource from NSGs. For example, for a Compute instance, use the `UpdateVnic` operation. For a load balancer, use the `UpdateNetworkSecurityGroups` operation, and so on.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let network admins manage a cloud network](#) covers management of all Networking components, including NSGs.

If you have security admins who need to manage NSGs but not other components in the Networking service, you could write a more restrictive policy:

```
Allow group NSGAdmins to manage network-security-groups in tenancy
```

```
Allow group NSGAdmins to manage vnics in tenancy
```

```
Allow group NSGAdmins to use VNICs in tenancy
```

The first statement lets the NSGAdmins group create and manage NSGs and their security rules.

The second statement is required because the creation of an NSG affects the VCN that the NSG is in. The scope of the second statement *also* allows the NSGAdmins group to create VCNs. However, the group can't create subnets or manage any other components related to any of those VCNs (except for the NSGs), because additional permissions would be required for those resources. The NSGAdmins group also can't delete any existing VCNs that already have subnets in them, because that action would require permissions related to subnets.

The third statement is required if the NSGAdmins need to put VNICs into an NSG. Whoever creates the parent resource of the VNIC (for example, a Compute instance) must already have this level of permission to create the parent resource.

For more information, see [IAM Policies for Networking](#).

### Using the Console

#### To view the security rules and resources in an NSG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Network Security Groups**.
4. Click the NSG you're interested in to view its details.

The NSG's security rules are displayed on the page. From there you can add, edit, or remove rules.

5. Under **Resources**, click **VNICs** to see the parent resources that belong to the NSG.

If the [parent resource](#) is a Compute instance, the corresponding VNICs from that instance are also listed on the page.

For other types of parent resources, the relevant service manages the VNICs on your behalf. Therefore only the parent resource (and not its corresponding VNICs) is listed on the page.

### To create an NSG

Prerequisite: Become familiar with the [parts of security rules](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Network Security Groups**.
4. Click **Create Network Security Group**.
5. Enter the following:
  - a. **Name:** A descriptive name for the network security group. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
  - b. **Create in Compartment:** The compartment where you want to create the network security group, if different from the compartment you're currently working in.
  - c. **Show Tagging Options:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Next**.

If you want to create the NSG without any rules yet, click **Create** and you're done. Otherwise proceed to the next step.
7. For the first security rule, enter the following items (for examples of rules, see [Networking Scenarios](#)):
  - **Stateful or stateless:** If stateful, connection tracking is used for traffic matching the rule. If stateless, no connection tracking is used. By default, rules

are stateful unless you specify otherwise. See [Stateful Versus Stateless Rules](#).

- **Direction (ingress or egress):** Ingress is inbound traffic to the VNIC, and egress is outbound traffic from the VNIC.
- **Source Type** and **Source** (for ingress rules only):

#### Allowed source types

Source Type	Allowed Source
CIDR	The CIDR block where the traffic originates from. Use 0.0.0.0/0 to indicate all IP addresses. The prefix is required (for example, include the /32 if specifying an individual IP address).
Service	Only for packets coming from an Oracle service through a <a href="#">service gateway</a> . The source service is the <a href="#">service CIDR label</a> that you're interested in.
Network Security Group	An NSG that is in the same VCN as this rule's NSG.

- **Destination Type** and **Destination** (for egress rules only):

## Allowed destination types

Destination Type	Allowed Destination
CIDR	The CIDR block where the traffic is destined to. Use 0.0.0.0/0 to indicate all IP addresses. The prefix is required (for example, include the /32 if specifying an individual IP address).
Service	Only for packets going to an Oracle service through a <a href="#">service gateway</a> . The destination service is the <a href="#">service CIDR label</a> that you're interested in.
Network Security Group	An NSG that is in the same VCN as this rule's NSG.

- **IP Protocol:** Either a single [IPv4 protocol](#) (for example, TCP or ICMP) or "all" to cover all protocols.
  - **Source port range:** The port where the traffic originates from. For TCP or UDP, you can specify all source ports, or optionally specify a single source [port number](#), or a range.
  - **Destination port range:** The port where the traffic is destined to. For TCP or UDP, you can specify all destination ports, or optionally specify a single destination [port number](#), or a range.
  - **ICMP type and code:** For ICMP, you can specify all types and codes, or optionally specify a single [type](#) with an optional code. If the type has multiple codes, create a separate rule for each code you want to allow.
8. To add another security rule, click **+ Another Rule** and enter the rule's information. Repeat for each rule you want to add.

9. When you're done, click **Create**.

The NSG is created and then displayed on the **Network Security Group** page in the compartment you chose. You can now specify this NSG when creating or managing instances or other types of [parent resources](#).

When you view all the security rules in an NSG, you can filter the list by ingress or egress.

### To add or remove a resource from an NSG

In general, you manage the resource membership of an NSG *at the [parent resource](#)*, and not at the NSG itself. In other words, to add a parent resource to an NSG, you execute the action on the *parent resource* (by specifying which NSGs the parent resource should be added to). You do not execute the action on the NSG (by specifying which VNICs or parent resources should be added to the NSG). Similarly, to remove a VNIC from an NSG, you execute that action by updating the parent resource, not the NSG. For a list of the parent resources that support the use of NSG, see [Support for Network Security Groups](#).

#### Example: Compute instances

- **When creating an instance:** In the **Configure networking** section, select the check box for **Use network security groups to control traffic**, and then specify one or more NSGs. The instance's primary VNIC is added to the NSGs. See the procedure in [Creating an Instance](#).
- **For an existing instance:** Adding an existing instance to an NSG means adding its *primary VNIC* to the NSG. You can also add a secondary VNIC to an NSG. See [To add or remove a VNIC from a network security group](#).

#### Example: Exadata DB systems

- **When creating an Exadata DB system:** In the **Network Information** section, you set up the client network and backup network. For each network, select the check box

for **Use network security groups to control traffic**, and then specify one or more NSGs for the specific network. See [To create an Exadata DB system](#). Also see [Network Setup for Exadata DB Systems](#).

- **For an existing Exadata DB system:** An Exadata's details include a list of the NSGs that the client network belongs to (if any), and a list of the NSGs that the backup network belongs to (if any). Next to the relevant **Network Security Groups**, click **Edit** to change that list. See [To edit the network security groups \(NSGs\) for your client or backup network](#).

### To delete an NSG

Prerequisite: To delete a security list, it must not be associated with a subnet . You can't delete the default security list in a VCN.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.
4. For the security list you want to delete, click the Actions icon (three dots), and then click **Terminate**.
5. Confirm when prompted.

### To manage security rules for an NSG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Network Security Groups**.
4. Click the NSG you're interested in to view its details.

The NSG's security rules are displayed on the page. From there you can add, edit, or remove rules.

### To manage tags for an NSG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Network Security Groups**.
4. Click the NSG you're interested in.
5. Click the **Tags** tab to view or edit the existing tags. Or click **Add tags** to add new ones.

For more information, see [Resource Tags](#).

### To move an NSG to a different compartment

You can move an NSG from one compartment to another. When you move an NSG to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Network Security Groups**.
4. Click the the Actions icon (three dots) for the NSG, and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage a VCN's network security groups, use these operations:

- [ListNetworkSecurityGroups](#)
- [GetNetworkSecurityGroup](#)
- [UpdateNetworkSecurityGroup](#)
- [CreateNetworkSecurityGroup](#)
- [DeleteNetworkSecurityGroup](#)
- [ChangeNetworkSecurityGroupCompartment](#)
- [ListNetworkSecurityGroupVnics](#)
- [ListNetworkSecurityGroupSecurityRules](#)
- [AddNetworkSecurityGroupSecurityRules](#)
- [RemoveNetworkSecurityGroupSecurityRules](#)
- [UpdateNetworkSecurityGroupSecurityRules](#)

There are some differences in the REST API model for NSGs compared to security lists:

- With security lists, there is an `IngressSecurityRule` object and a separate `EgressSecurityRule` object. With network security groups, there is only a `SecurityRule` object, and the object's `direction` attribute determines whether the rule is for ingress or egress traffic.
- With security lists, the rules are part of the `SecurityList` object, and you work with the rules by calling the security list operations (such as `UpdateSecurityList`). With NSGs, the rules are not part of the `NetworkSecurityGroup` object. Instead you use separate operations to work with the rules for a given NSG (example: `UpdateNetworkSecurityGroupSecurityRules`).

- The model for *updating* existing security rules is different between security lists and NSGs. With NSGs, each rule in a given group has a unique Oracle-assigned identifier (example: 04ABEC). When you call `UpdateNetworkSecurityGroupSecurityRules`, you provide the IDs of the specific rules that you want to update. For comparison, with security lists, the rules have no unique identifier. When you call `UpdateSecurityList`, you must pass in the *entire* list of rules, including rules that are not being updated in the call.
- There is a limit of 25 rules when calling the operations to add, remove, or update security rules.

## Security Lists

The Networking service offers two virtual firewall features to control traffic at the packet level:

- **Security lists:** Covered in this topic. This is the original type of virtual firewall offered by the Networking service.
- **Network security groups:** Another type of virtual firewall that Oracle recommends over security lists. See [Network Security Groups](#).

Both of these features use *security rules*. For important information about how security rules work, and a general comparison of security lists and network security groups, see [Security Rules](#).

## Highlights

- Security lists act as virtual firewalls for your Compute instances and [other kinds of resources](#). A security list consists of a set of ingress and egress [security rules](#) that apply to all the VNICs *in any subnet that the security list is associated with*. This means that all the VNICs in a given subnet are subject to the same set of security lists. See [Comparison of Security Lists and Network Security Groups](#).
- Security list rules function the same as network security group rules. For a discussion of rule parameters, see [Parts of a Security Rule](#).

- Each VCN comes with a [default security list](#) that has several default rules for essential traffic. If you don't specify a custom security list for a subnet, the default security list is automatically used with that subnet. You can add and remove rules from the default security list.
- Security lists have separate and different limits compared to network security groups. See [Limits](#).

### Overview of Security Lists

A security list acts as a virtual firewall for an instance, with ingress and egress rules that specify the types of traffic allowed in and out. Each security list is enforced at the VNIC level. However, you configure your security lists *at the subnet level*, which means that all VNICs in a given subnet are subject to the same set of security lists. The security lists apply to a given VNIC whether it's communicating with another instance in the VCN or a host outside the VCN.

Each subnet can have multiple security lists associated with it, and each list can have multiple rules (for the maximum number, see [Limits](#)). A packet in question is allowed if *any rule in any of the lists* allows the traffic (or if the traffic is part of an existing connection being tracked). There's a caveat if the lists happen to contain both stateful and stateless rules that cover the same traffic. For more information, see [Stateful Versus Stateless Rules](#).

Security lists are regional entities. For limits related to security lists, see [Limits](#).

Security lists can control both IPv4 and IPv6 traffic. However, IPv6 addressing and related security list rules are currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

### Default Security List

Each cloud network has a *default security list*. You can also create other security lists for the VCN. A given subnet automatically has the default security list associated with it if you don't specify one or more other security lists during subnet creation. At any time after you create a subnet, you can change which security lists are associated with it. And you can change the rules in the lists.

Unlike other security lists, the default security list comes with an initial set of stateful rules, which you can change:

- **Stateful ingress:** Allow TCP traffic on destination port 22 (SSH) from source 0.0.0.0/0 and any source port. This rule makes it easy for you to create a new cloud network and public subnet, launch a Linux instance, and then immediately use SSH to connect to that instance without needing to write any security list rules yourself.



### Important

The default security list does not include a rule to allow Remote Desktop Protocol (RDP) access. If you're using [Windows images](#), make sure to add a stateful ingress rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port.

See [To enable RDP access](#) for more information.

- **Stateful ingress:** Allow ICMP traffic type 3 code 4 from source 0.0.0.0/0. This rule enables your instances to receive Path MTU Discovery fragmentation messages.
- **Stateful ingress:** Allow ICMP traffic type 3 (all codes) from source = your VCN's CIDR. This rule makes it easy for your instances to receive connectivity error messages from other instances within the VCN.
- **Stateful egress:** Allow all traffic. This allows instances to initiate traffic of any kind to any destination. Notice that this means the instances with public IP addresses can talk to any internet IP address if the VCN has a configured internet gateway. And because stateful security rules use connection tracking, the response traffic is automatically allowed regardless of any ingress rules. For more information, see [Connection Tracking Details for Stateful Rules](#).

The default security list comes with no stateless rules. However, you can add or remove rules from the default security list as you like.

If your VCN is enabled for IPv6 addressing (which is currently supported in only the Government Cloud), the default security list contains some default rules for IPv6 traffic. For more information, see [IPv6 Addresses](#).

### Enabling Ping

The default security list does not include a rule to allow ping requests. If you plan to ping an instance, see [Rules to Enable Ping](#).

## Security Rules

If you're not yet familiar with the basics of security rules, see these sections in the security rules topic:

- [Parts of a Security Rule](#)
- [Stateful Versus Stateless Rules](#)

## Working with Security Lists



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### General Process for Working with Security Lists

1. Create a security list.
2. Add security rules to the security list.
3. Associate the security list with one or more subnets.

4. Create resources in the subnet (for example, create Compute instances in the subnet). The security rules apply to all the VNICs in that subnet. See [About VNICs and Parent Resources](#).

### Additional Details

When you create a subnet, you must associate at least one security list with it. It can be either the VCN's default security list or one or more other security lists that you've already created (for the maximum number, see [Service Limits](#)). You can [change which security lists the subnet uses](#) at any time.

You may optionally assign a friendly name to the security list during creation. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the security list a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

For the purposes of access control, you must specify the compartment where you want the security list to reside. Consult an administrator in your organization if you're not sure which compartment to use. For more information, see [Access Control](#).

You can move security lists from one compartment to another. Moving a security list doesn't affect its attachment to a subnet. When you move a security list to a new compartment, inherent policies apply immediately and affect access to the security list. For more information, see [Managing Compartments](#).

You can add and remove rules from the security list. A security list can have no rules. Notice that when you update a security list in the API, the new set of rules replaces the entire existing set of rules.

To delete a security list, it must not be associated with a subnet. You can't delete a VCN's default security list.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have

permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let network admins manage a cloud network](#) covers management of all Networking components, including security lists.

If you have security admins who need to manage security lists but not other components in Networking, you could write a more restrictive policy:

```
Allow group SecListAdmins to manage security-lists in tenancy
```

```
Allow group SecListAdmins to manage vcn in tenancy
```

Both statements are needed because the creation of a security list affects the VCN the security list is in. The scope of the second statement *also* allows the SecListAdmins group to create VCNs. However, the group can't create subnets or manage any other components related to any of those VCNs (except for the security lists), because additional permissions would be required for those resources. The group also can't delete any existing VCNs that already have subnets in them, because that action would require permissions related to subnets.

For more information, see [IAM Policies for Networking](#).

## Using the Console

### To view a VCN's default security list

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.
4. Click the default security list to view its details.  
Under **Resources**, you can click **Ingress Rules** or **Egress Rules** to switch between the different types of rules.

### To update rules in an existing security list



#### Important

When updating the default security list, be aware of the [default rules](#), the purpose each serves, and the consequences of deleting them. For example, deleting the default stateful egress rule can cause problems if the instances need to initiate connections.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.
4. Click the security list you're interested in.
5. Under **Resources**, click either **Ingress Rules** or **Egress Rules** depending on the type of rule you want to work with.
6. If you want to add a rule, click **Add Ingress Rule** (or **Add Egress Rule**). See details of adding a rule in [To create a security list](#).
7. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Remove**.
8. If you wanted to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.

### To create a security list

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.

3. Under **Resources**, click **Security Lists**.
4. Click **Create Security List**.
5. Enter the following:
  - a. **Name:** A descriptive name for the security list. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - b. **Create in Compartment:** The compartment where you want to create the security list, if different from the compartment you're currently working in.
6. Add either an ingress rule or egress rule (for examples of rules, see [Networking Scenarios](#)):
  - a. Click either **Add Ingress Rule** or **Add Egress Rule**.
  - b. Choose whether it's a stateful or stateless rule (see [Stateful Versus Stateless Rules](#)). By default, rules are stateful unless you specify otherwise.
  - c. Enter either the source CIDR (for ingress) or destination CIDR (for egress). For example, use 0.0.0.0/0 to indicate all IP addresses. Other typical CIDRs you might specify in a rule are the CIDR block for your on-premises network, or for a particular subnet. If you're setting up a security list rule to allow traffic with a service gateway, instead see [Task 3: \(Optional\) Update security rules](#).
  - d. Select the IP protocol (for example, TCP, UDP, ICMP, "All protocols", and so on).
  - e. Enter further details depending on the protocol:
    - If you chose TCP or UDP, enter a source port range and destination port range. You can enter "All" to cover all ports. If you want to allow a specific [port](#), enter the port number (for example, 22 for SSH or 3389 for RDP) or a port range (for example, 20–22).
    - If you chose ICMP, you can enter "All" to cover all types and codes. If you want to allow a specific [ICMP type](#), enter the type and an optional code separated by a comma (for example, 3,4). If the type has multiple codes you want to allow, create a separate rule for each code.

7. Repeat the preceding step for each rule you want to add to the list.
8. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
9. When you're done, click **Create Security List**.

The security list is created and then displayed on the **Security Lists** page in the compartment you chose. You can now specify this security list when creating or updating a subnet.

When you view all the rules in a security list, notice that any stateless rules in the list are shown above any stateful rules. Stateless rules in the list take precedence over stateful rules. In other words: If there's traffic that matches both a stateless rule and a stateful rule across all the security lists associated with the subnet, the stateless rule takes precedence and the connection is not tracked.

### To change which security lists a subnet uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Subnets**.
4. Click the subnet you're interested in.
5. Under **Resources**, click **Security Lists**.
6. If you want to add a security list, click **Add Security List**, and select the new security list you want the subnet to use.
7. If you want to remove a security list, click the Actions icon (three dots), and then click **Remove**. Remember that a subnet must always have at least one security list associated with it.

The changes take effect within a few seconds.

### To move a security list to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.
4. Find the security list, click the Actions icon (three dots), and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

### To delete a security list

Prerequisite: To delete a security list, it must not be associated with a subnet. You can't delete the default security list in a VCN.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.
4. For the security list you want to delete, click the Actions icon (three dots), and then click **Terminate**.
5. Confirm when prompted.

### To manage tags for a security list

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Security Lists**.

4. Click the security list you're interested in.
5. Click the **Tags** tab to view or edit the existing tags. Or click **Add tags** to add new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage a VCN's security lists, use these operations:

- [ListSecurityLists](#)
- [GetSecurityList](#)
- [UpdateSecurityList](#)
- [CreateSecurityList](#)
- [DeleteSecurityList](#)
- [ChangeSecurityListCompartment](#)

### Virtual Network Interface Cards (VNICs)

This topic describes how to manage the virtual network interface cards (VNICs) in a virtual cloud network (VCN).



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Overview of VNICs and Physical NICs

The servers in Oracle Cloud Infrastructure data centers have physical network interface cards (NICs). When you launch an instance on one of these servers, the instance communicates using Networking service *virtual* NICs (VNICs) associated with the physical NICs. The next sections talk about VNICs and NICs, and how they're related.

#### About VNICs

A VNIC enables an instance to connect to a VCN and determines how the instance connects with endpoints inside and outside the VCN. Each VNIC resides in a subnet in a VCN and includes these items:

- One *primary* private IPv4 address from the subnet the VNIC is in, chosen by either you or Oracle.
- Up to 31 optional [secondary private IPv4 addresses](#) from the same subnet the VNIC is in, chosen by either you or Oracle.
- An optional [public IPv4 address](#) for each private IP, chosen by Oracle but assigned by you at your discretion.
- An optional hostname for DNS for each private IP address (see [DNS in Your Virtual Cloud Network](#)).
- A MAC address.
- A VLAN tag assigned by Oracle and available when attachment of the VNIC to the instance is complete (relevant only for bare metal instances).
- A flag to enable or disable the source/destination check on the VNIC's network traffic (see [Source/Destination Check](#)).
- Optional membership in one or more [network security groups](#) (NSGs) of your choice. NSGs have [security rules](#) that apply only to the VNICs in that NSG.
- Up to 32 optional IPv6 addresses. IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

Each VNIC also has a friendly name you can assign, and an Oracle-assigned OCID (see [Resource Identifiers](#)).

Each instance has a *primary VNIC* that is automatically created and attached during launch. That VNIC resides in the subnet you specify during launch. The primary VNIC cannot be removed from the instance.

### How VNICs and Physical NICs Are Related

This section is relevant to bare metal instances.

The OS on a bare metal instance recognizes two physical network devices and configures them as two physical NICs, 0 and 1. Whether they're both active depends on the underlying hardware:

- **Oracle X5 servers (also called *first-generation*):** Only NIC 0 is active.
- **Oracle X6 servers:** Only NIC 0 is active.
- **Oracle X7 servers (also called *second-generation*):** Both NIC 0 and NIC 1 are active. Each physical NIC has 25 Gbps bandwidth.

NIC 0 is automatically configured with the primary VNIC's IP configuration (the IP addresses, DNS hostname, and so on).

If you add a *secondary VNIC* to a second-generation instance, you must specify which physical NIC the secondary VNIC should use. You must also configure the OS so that the physical NIC has the secondary VNIC's IP configuration. For Linux instances, see [Linux: Configuring the OS for Secondary VNICs](#). For Windows instances, see [Windows: Configuring the OS for Secondary VNICs](#).

### About Secondary VNICs

You can add secondary VNICs to an instance after it's launched. Each secondary VNIC can be in a subnet in the same VCN as the primary VNIC, or in a different subnet that is either in the same VCN or a different one. However, all the VNICs must be in the same availability domain as the instance.

Here are some reasons why you might use secondary VNICs:

- **Use your own hypervisor on a bare metal instance:** The virtual machines on the bare metal instance each have their own secondary VNIC, giving direct connectivity to other instances and services in the VNIC's VCN. For more information, see [Installing and Configuring KVM on Bare Metal Instances with Multi-VNIC](#).
- **Connect an instance to subnets in multiple VCNs:** For example, you might set up your own firewall to protect traffic between VCNs, so the instance needs to connect to subnets in different VCNs.

Here are more details about secondary VNICs:

- They're supported for these types of instances:
  - **Linux:** Both VM and bare metal instances. Also see [Linux: Configuring the OS for Secondary VNICs](#).
  - **Windows:** Both VM and bare metal instances, but only on X7/second-generation shapes ([shapes with "2" in the name](#), such as VM.Standard 2.16 and BM.Standard2.52). For bare metal, secondary VNICs are supported only on the second physical NIC. Remember that the first physical NIC is NIC 0, and the second is NIC 1. Also see [Windows: Configuring the OS for Secondary VNICs](#).
- There's a limit to how many VNICs can be attached to an instance, and it varies by shape. For those limits, see [Compute Shapes](#).
- They can be added only after the instance is launched.
- They must always be attached to an instance and cannot be moved. The process of creating a secondary VNIC automatically attaches it to the instance. The process of detaching a secondary VNIC automatically deletes it.
- They are automatically detached and deleted when you terminate the instance.
- The instance's bandwidth is fixed regardless of the number of VNICs attached. You can't specify a bandwidth limit for a particular VNIC on an instance.
- Attaching multiple VNICs from the same subnet CIDR block to an instance can introduce asymmetric routing, especially on instances using a variant of Linux. If you need this

type of configuration, Oracle recommends assigning multiple [private IP addresses](#) to one VNIC, or using policy-based routing as shown in the script later in this topic.

### Source/Destination Check

By default, every VNIC performs the source/destination check on its network traffic. The VNIC looks at the source and destination listed in the header of each network packet. If the VNIC is not the source or destination, then the packet is dropped.

If the VNIC needs to forward traffic (for example, if it needs to perform Network Address Translation (NAT)), you must disable the source/destination check on the VNIC. For instructions, see [To update an existing VNIC](#). For information about the general scenario, see [Using a Private IP as a Route Target](#).

### VNIC Information in the Instance Metadata

The [instance metadata](#) includes information about the VNICs at this URL:

```
http://169.254.169.254/opc/v1/vnics/
```

Here's an example response that shows the VNICs that are attached to an instance:

```
[{
 "vnicId" : "ocidl.vnic.oc1.phx.exampleuniqueID",
 "privateIp" : "10.0.3.6",
 "vlanTag" : 11,
 "macAddr" : "02:00:17:00:12:D3",
 "virtualRouterIp" : "10.0.3.1",
 "subnetCidrBlock" : "10.0.3.0/24",
 "nicIndex" : 0
}, {
 "vnicId" : "ocidl.vnic.oc1.phx.exampleuniqueID",
 "privateIp" : "10.0.4.3",
 "vlanTag" : 12,
 "macAddr" : "02:00:17:00:13:13",
 "virtualRouterIp" : "10.0.4.1",
 "subnetCidrBlock" : "10.0.4.0/24",
 "nicIndex" : 0
}]
```

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

VNICs reside in a subnet but attach to an instance. The VNIC's *attachment to the instance* is a separate object from the VNIC or the instance itself. Be aware that the VNIC and subnet always exist together in the same compartment, but the VNIC's *attachment to the instance* always exists in the instance's compartment. This distinction isn't important if you have a simple access control scenario where all the cloud resources are in the same compartment (for example, for a proof-of-concept). When you move to a production implementation, you might decide to have network administrators manage the network, and other personnel administer instances. That means you might put instances in a different compartment than the subnet.

For administrators: see [IAM Policies for Networking](#).

### Monitoring VNICs

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring the traffic coming in and out of VNICs, see [VNIC Metrics](#).

## Using the Console

### To view an instance's VNICs

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.

3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed. If the instance has two [active physical NICs](#), the VNICs are grouped by NIC 0 and NIC 1.

### To create and attach a secondary VNIC

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click **Create VNIC**.
6. In the **Create VNIC** dialog box, you specify which VCN and subnet to put the VNIC in. By default, the VNIC will be created in the current compartment, and you'll choose the VCN and subnet from the same compartment. Click the **click here** link in the dialog box if you want to enable compartment selection and choose a VCN or subnet in a different compartment.

Enter the following:

- **Name:** A friendly name for the secondary VNIC. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
- **Virtual Cloud Network Compartment:** The compartment that contains the VCN that in turn contains the subnet of interest.
- **Virtual Cloud Network:** The VCN that contains the subnet of interest.
- **Subnet Compartment:** The compartment that contains the subnet of interest.
- **Subnet:** The subnet of interest. The secondary VNIC must be in the same availability domain as the instance's primary VNIC, so the subnet list includes any

[regional subnets or AD-specific subnets](#) in the primary VNIC's availability domain.

- **Physical NIC:** Only relevant if this is a bare metal instance with two [active physical NICs](#). Select which one you want the secondary VNIC to use. When you later view the instance's details and the list of VNICs attached to the instance, they'll be grouped by NIC 0 and NIC 1.
- **Use network security groups to control traffic:** Select this check box to add the secondary VNIC to at least one [network security group](#) (NSG) of your choice. NSGs have security rules that apply only to the VNICs in that NSG.
- **Skip Source/Destination Check:** By default, this check box is NOT selected, which means the VNIC performs the source/destination check. Only select this check box if you want the VNIC to be able to forward traffic. See [Source/Destination Check](#).
- **Private IP Address:** Optional. An available private IP address of your choice from the subnet's CIDR (otherwise the private IP address is automatically assigned).
- **Assign public IP address:** Whether to assign an ephemeral public IP address to the VNIC's primary private IP. Available only if the subnet is public. For more information, see [Public IP Addresses](#).
- **Hostname:** Optional. A hostname to be used for DNS within the cloud network. Available only if the VCN and subnet both have DNS labels. For more information, see [DNS in Your Virtual Cloud Network](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Create VNIC**.

The secondary VNIC is created and then displayed on the **Attached VNICs** page for the instance. It can take several seconds before the secondary VNIC appears on the page.

8. Configure the OS to use the VNIC. See [Linux: Configuring the OS for Secondary VNICs](#) or [Windows: Configuring the OS for Secondary VNICs](#).

### To update an existing VNIC

You can update the VNIC's friendly name or hostname, or whether the VNIC performs the source/destination check.

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. For the VNIC you want to edit, click the Actions icon (three dots), and then click **Edit VNIC**.
6. Make your changes and click **Update VNIC**.

### To add or remove a VNIC from a network security group

You can change which [network security groups](#) (NSGs) a VNIC belongs to, or remove a VNIC from all NSGs.

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.

The primary VNIC and any secondary VNICs attached to the instance are displayed. Each VNIC's details include a list of the NSGs that the VNIC belongs to (if any).

5. For the VNIC you want to edit, next to **Network Security Groups** in the VNIC's details, click **Edit**.
6. Make your changes and click **Save**.

### To delete a secondary VNIC



#### Warning

If the VNIC has a private IP that is the [target of a route rule](#), deleting the VNIC causes the route rule to blackhole and traffic will be dropped.

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. For the VNIC you want to delete, click the Actions icon (three dots), and then click **Delete**.
6. Confirm when prompted.

It takes typically a few seconds before the VNIC is deleted.

If the secondary VNIC is on a Linux instance: If you then run the provided script in [Linux: Configuring the OS for Secondary VNICs](#), it removes the secondary VNIC from the OS configuration.

### To manage tags for a VNIC

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage VNICs on an instance, use these operations:

- [ListVnicAttachments](#): Use this to list the VNICs attached to an instance.
- [GetVnicAttachment](#): Use this to get the VNIC's VLAN tag and other properties.
- [GetVnic](#): Use this to get the VNIC's private IP address, MAC address, optional public IP address, optional DNS hostname, and other properties.
- [AttachVnic](#)
- [DetachVnic](#)
- [UpdateVnic](#)

### Linux: Configuring the OS for Secondary VNICs

This section gives details about OS configuration that is required for secondary VNICs on instances that run a variant of Linux.

At the end of the section is a script that you can use to configure secondary VNICs on either VM instances or bare metal instances.

#### Linux VM Instances

When you add a secondary VNIC to a Linux VM instance, a new interface (that is, an Ethernet device) is added to the instance and automatically recognized by the OS. However, DHCP is not active for the secondary VNIC, and you must configure the interface with the static IP address and default route. The script provided here handles that configuration for you.

#### Linux Bare Metal Instances

When you add a secondary VNIC to a Linux bare metal instance, the OS does not automatically recognize the secondary VNIC, so you must configure it in the OS. Depending on your requirements, you can configure it as either:

- An SR-IOV virtual function (see [Installing and Configuring KVM on Bare Metal Instances with Multi-VNIC](#)).
- A VLAN-tagged interface (see the script in the following section).

#### Configuration Script for Either Linux VM Instances or Linux Bare Metal Instances

The following script works for both VM instances and bare metal instances. It looks at the secondary VNIC information in the [instance metadata](#) and configures the OS accordingly. You could run the script periodically to bring the OS configuration up to date with the instance metadata.

For VM instances in particular, the OS automatically recognizes the secondary VNIC's interface, and the script just needs to configure the static IP address and default route.

For bare metal instances in particular, the script creates the interface for the secondary VNIC and configures it with the relevant information. If the instance has [two active physical NICs](#)

(an X7/second-generation shape with NIC 0 and NIC 1), the script configures the secondary VNIC to use whichever physical NIC you chose when you [added the VNIC to the instance](#). Note that for NIC 1, if a secondary VNIC has VLAN tag 0, it uses the NIC's interface. The script doesn't create an interface for that secondary VNIC.

Here are some additional notes about how the script works for both VM instances and bare metal instances:

- **Default namespace and policy-based routing:** By default, this script configures the secondary VNIC in the default namespace and with policy-based routing so applications can communicate through the VNIC with hosts outside the VNIC's subnet. **This policy-based routing has effect only if the packets are sourced from the IP address of the secondary VNIC.** The ability to bind to a specific source IP address or source interface exists in most tools (such as `ssh`, `ping`, and `wget`) and libraries that initiate connections. For example, the `ssh -b` option lets you bind to the private IP address of the secondary VNIC:

```
ssh -b <secondary_VNIC_IP_address> <destination_IP_address>
```

Be aware that if traffic comes in to a service on the instance through a secondary VNIC's interface and the service replies, the reply packets automatically have the VNIC's interface IP address as the source IP address. Policy-based routing is required for that reply to go back out on the same interface and find the correct default gateway.

- **A separate namespace:** If you're familiar with namespaces, you can instead configure the secondary VNIC in another namespace of your choice by running the script with the `-n` option. A separate namespace is required when an instance has secondary VNICs that are attached to subnets in different VCNs, and those subnets have overlapping CIDR blocks.
- **Secondary private IPs:** The instance metadata does not include information about any [secondary private IPs](#) assigned to the instance. To include that as part of the script's OS configuration, you must provide the secondary private IP address and OCID at the command line when you run the script.
- **Removal of a secondary VNIC:** After [deleting a secondary VNIC from an instance](#), running the script removes the VNIC's information from the OS configuration.



### Important

The script uses a simple configuration process that does not persist if you reboot the instance. If you use the script, make sure to rerun it after each reboot.

Here are basic examples of how to run the script:

- `<script_name> -c`: Configure (adds or deletes) secondary VNIC host IP configuration
- `<script_name> -c -n`: Same but uses separate namespaces
- `<script_name> -d`: Force removes all secondary VNIC host IP configuration

See the script's help for more information.



### Tip

Download the script from the online version of this user guide at <https://docs.cloud.oracle.com/iaas/Content/Network/Tasks/managingVNICs.htm#linux>.

## Windows: Configuring the OS for Secondary VNICs

Secondary VNICs are supported on VM and bare metal instances, **but only on X7/second-generation shapes** ([shapes with "2" in the name](#), such as VM.Standard2.16 and BM.Standard2.52). For bare metal, secondary VNICs are supported **only on the second physical NIC**.



### Tip

The first physical NIC is NIC 0, and the second is NIC 1.

You must configure the secondary VNIC within the OS. There's an Oracle-provided PowerShell script that performs configuration. When running the script, you can optionally provide the secondary VNIC's OCID (which you can get from the [instance's VNIC metadata](#)):

```
.\secondary_vnic_windows_configure.ps1 "<secondary_VNIC_OCID>"
```

Otherwise, the script shows a list of the secondary VNICs on the instance and asks you to choose the one you want to configure. Here's generally what the script does:

1. The script checks if the network interface has an IP address and a default route.
2. To enable the OS to recognize the secondary VNIC, the **script must overwrite the IP address and default route with static settings (which effectively disables DHCP)**. The script prompts you with a choice: to overwrite with the static settings, or exit.



### Tip

Download the script from the online version of this user guide at <https://docs.cloud.oracle.com/iaas/Content/Network/Tasks/managingVNICs.htm#windows>.

The overall process for configuration varies slightly depending on the type of instance (VM or bare metal) and how many secondary VNICs you add to the instance.

## Windows VM instances

Here's the overall process:

1. [Add one or more secondary VNICs](#) to the instance. Keep each VNIC's OCID handy so you can provide it later when you run the configuration script. You can also get the OCID from the instance's [VNIC metadata](#).
2. [Connect to the instance](#) with Remote Desktop.

3. Run the script:
  - a. Open PowerShell as an administrator.
  - b. Run the script with the secondary VNIC's OCID:

```
.\secondary_vnic_windows_configure.ps1 "<secondary_VNIC_OCID>"
```

4. Repeat the preceding step for each additional secondary VNIC.

### Windows bare metal instances: adding the first secondary VNIC

If you're adding only a single secondary VNIC to the bare metal instance, here's the overall process:

1. [Add the secondary VNIC](#) to your instance. Keep the VNIC's OCID handy so you can provide it when later running the configuration script. You can also get the OCID from the instance's [VNIC metadata](#).
2. [Connect to the instance](#) with Remote Desktop.
3. Enable the second physical NIC on the instance:
  - a. Open the Device Manager, and then click **Network adapters**.
  - b. Right-click the adapter that corresponds to the instance's second physical NIC, and click **Enable**.
4. Run the script:
  - a. Open PowerShell as an administrator.
  - b. Run the script with the secondary VNIC's OCID:

```
.\secondary_vnic_windows_configure.ps1 "<secondary_VNIC_OCID>"
```

- c. When the script prompts you to overwrite the network interface's settings, say yes.

### Windows bare metal instances: adding additional secondary VNICs

If you have one secondary VNIC on the second physical NIC of a bare metal instance, and you

want to one or more additional VNICs, here's the overall process. It includes a task for setting up NIC teaming, which is required if the instance has more than one VNIC on the second physical NIC.



### Note

If you increase the number of secondary VNICs on the second physical NIC from one to two or more, you must enable *NIC teaming* for the second physical NIC (see instructions that follow). In your NIC "team," you create a separate interface for *each* secondary VNIC on that physical NIC, including the initial one. This means that the original interface for that first secondary VNIC will no longer work, and any subsequent configuration you want to do for that VNIC must be done instead on the VNIC's new interface that's part of the "team".

1. [Add one or more additional secondary VNICs](#) to your instance. Keep each VNIC's OCID and VLAN tag handy so you can provide them when later running the configuration script. You can also get the values from the instance's [VNIC metadata](#).
2. [Connect to the instance](#) with Remote Desktop.
3. Set up NIC teaming on the instance:
  - a. Open the Server manager, and then click **Local Server**.
  - b. In the list of properties, locate **NIC Teaming**, and then click **Disabled** to enable and set up NIC teaming.
  - c. In the **Teams** section on the lower-left side of the screen, click **Tasks**, and then click **New Team**.
  - d. Enter a name for the team, select the check box for the second physical NIC on the instance, and click **OK**.

The new team is created and appears in the list of teams in the **Teams** section.

- e. Click the new team so it's selected, and then in the **Adapters and Interfaces** section on the right side of the screen, click the **Team Interfaces** tab.
  - f. Click **Tasks**, and then click **Add Interface** (you will create a separate interface for each secondary VNIC on the second physical NIC).
  - g. Click the radio button for **Specific VLAN**, and then enter the Oracle-assigned VLAN tag number for the VLAN (for example, 1). You can get the VLAN tag from the Console or the instance's [VNIC metadata](#).
  - h. Click **OK**.
  - i. Repeat the four preceding steps (e-h) for each of the other secondary VNICs. You create a separate interface for each secondary VNIC.
4. Run the script:
- a. Open PowerShell as an administrator.
  - b. For the first secondary VNIC, run the script with the secondary VNIC's OCID:

```
.\secondary_vnic_windows_configure.ps1 "<secondary_VNIC_OCID>"
```
  - c. When the script prompts you to overwrite the network interface's settings, say yes.
  - d. Repeat the preceding two steps (b and c) for each of the additional secondary VNICs.

## VNIC Metrics

You can monitor the health, capacity, and performance of your Networking service VNICs by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_vcn` (the Networking service).

Resources: virtual network interface cards (VNICs).

### Overview of Metrics for an Instance and Its Network Devices

If you're not already familiar with the different types of metrics available for an instance and its storage and network devices, see [Compute Instance Metrics](#).

#### Overview of Metrics: `oci_vcn`

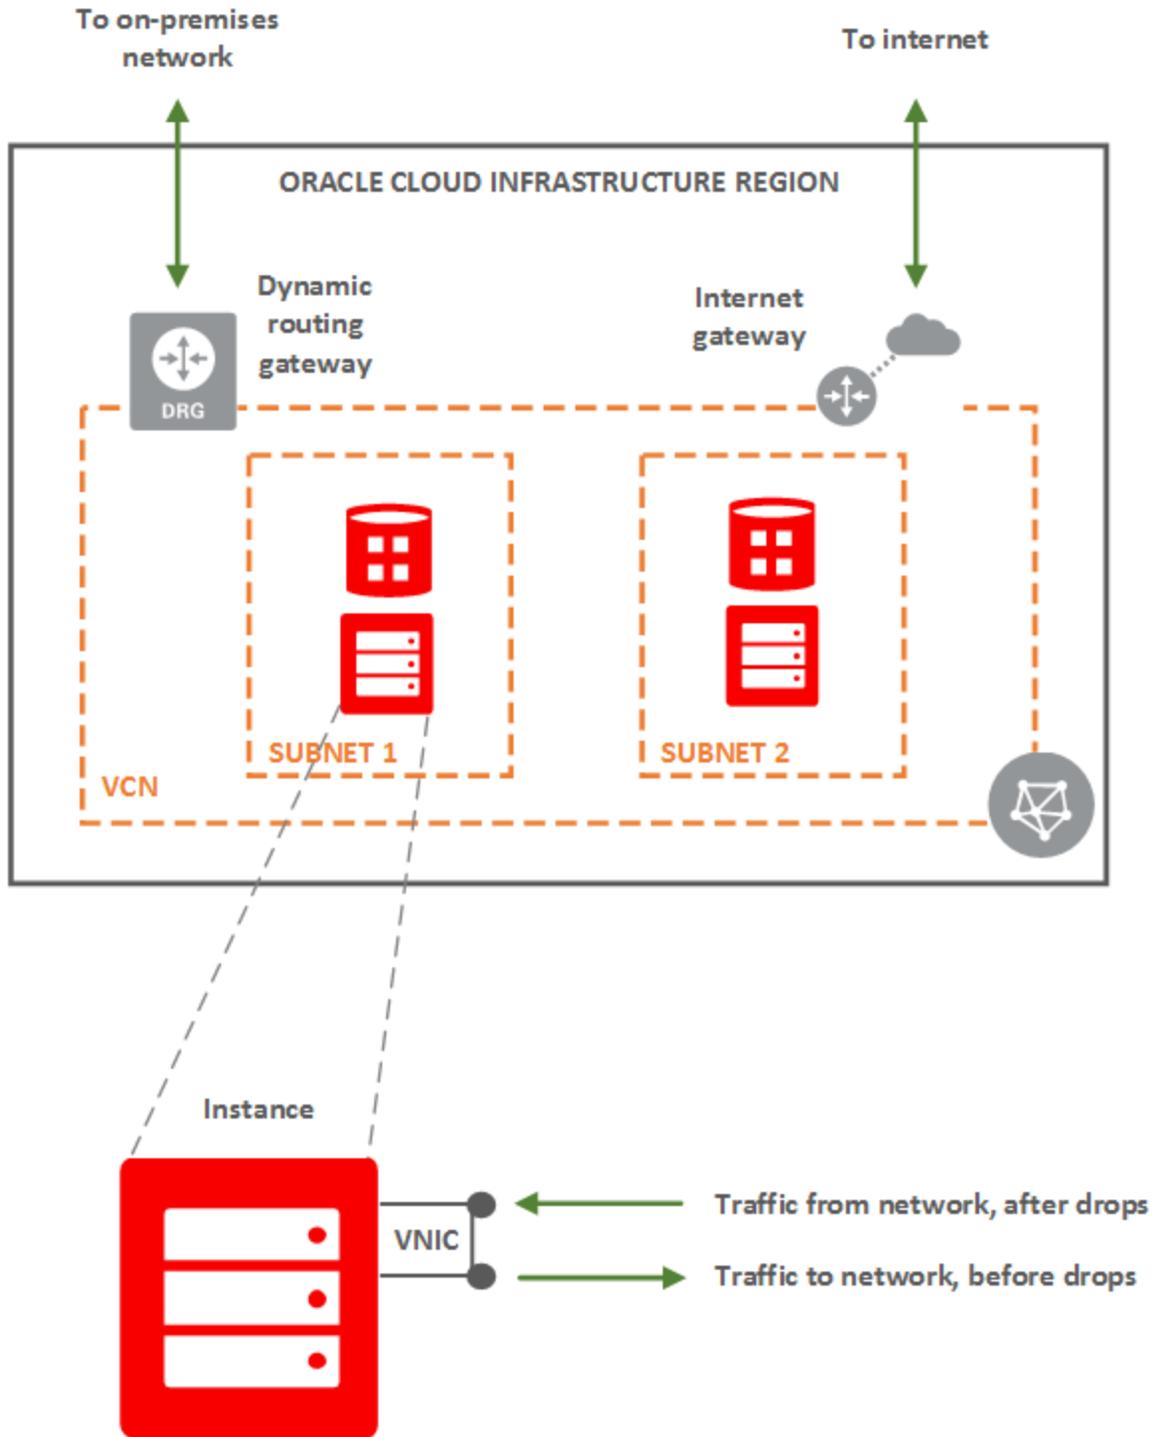
Each Compute instance has one or more Networking service [VNICs](#). A VNIC connects the instance to a subnet in a virtual cloud network (VCN). A given VNIC controls how the instance communicates with endpoints inside the VCN (other instances) and endpoints outside the VCN (hosts on the internet, in your on-premises network, in another VCN, and so on).

With the Networking service metrics (in metric namespace `oci_vcn`), you can get this information for a VNIC:

- **Traffic to and from the network:** Per-VNIC traffic levels (packets and bytes), which can help you identify meaningful increases or decreases in traffic coming in and out of your instances
- **Packets dropped due to security list violations:** Per-VNIC *drops* (dropped packets), which can help you identify changes in traffic caused by security list changes

The following diagram illustrates the general concept. A given instance resides in a subnet within a VCN that has one or more gateways to communicate with other networks. The instance is enlarged to show its VNIC, which the instance uses to communicate with the network. In this context, the term *network* means both the other instances in the VCN and hosts outside the VCN available through the gateways.

The VNIC receives traffic from the network and sends traffic to the network. The Networking service drops packets according to security list rules you set up for the instance's subnet. Traffic coming to the VNIC from the network is measured *after* the Networking service drops the packets that violate the subnet's security list rules. Traffic leaving the VNIC is measured *before* the Networking service drops the packets that violate the subnet's security list rules.



The Compute service separately reports network-related metrics *as measured on the instance itself and aggregated across all the attached VNICs*. Those metrics are available in the `oci_computeagent` metric namespace. For more information, see [Compute Instance Metrics](#).

### Raw Data Point Frequency

For every 1-minute interval, the Networking service posts one raw data point to the Monitoring service. The Monitoring service charts show data points at 1-minute, 5-minute, and 60-minute intervals. The available statistics are calculated by using the count of 1-minute data points in the select interval. For example, for a given metric:

- The mean for each 5-minute interval is calculated over 5 raw data points.
- The mean for each 60-minute interval is calculated over 60 raw data points.

### Required IAM Policy

When writing an IAM policy for viewing VNIC metrics, it's important to remember that:

- The VNIC and the VNIC's metrics (emitted by the `oci_vcn` metric namespace) reside in the *subnet's compartment*, and not the instance's compartment.
- The *VNIC attachment* (which is an object different from the VNIC itself) resides in the *instance's compartment*.

If the instance and subnet are in the same compartment, these details aren't so important when you write the IAM policy.

### Minimum required policy for getting VNIC metrics

The following policy contains the one statement required to get VNIC metrics, which are part of the `oci_vcn` metric namespace.

If you're using the Console, this policy lets you go to the **Monitoring** tab in the Console and view the metrics for one or more VNICs in the specified compartment. The policy uses an example group called `VnicMetricReaders`. The condition at the end (*where*

## CHAPTER 23 Networking

---

`target.metrics.namespace='oci_vcn')` allows the group to view only the metrics in the `oci_vcn` metric namespace.

```
Allow group VnicMetricReaders to read metrics in compartment <subnet_compartment> where
target.metrics.namespace='oci_vcn'
```

### Policy for viewing a VNIC's details and metrics in the Console

The following policy lets you view an instance in the Console, click through to a specific VNIC, and then view that VNIC's details and metrics.

```
Allow group VnicMetricReaders to read metrics in compartment <subnet_compartment> where
target.metrics.namespace='oci_vcn'
```

```
Allow group VnicMetricReaders to read instance-family in compartment <instance_compartment>
```

```
Allow group VnicMetricReaders to inspect virtual-network-family in compartment <subnet_compartment>
```

The second and third statements let you view the instance's details and the VNIC's details, respectively.

### Available Metrics: `oci_vcn`

The metrics listed in the following table are automatically available for any VNIC on any instance you create. You do not need to enable monitoring on the instance to get these metrics for the VNIC or VNICs on the instance.

You also can use the Monitoring service to create [custom queries](#).

Each metric includes the following dimension:

#### **RESOURCEID**

The OCID of the VNIC.

Metric	Metric Display Name	Unit	Description	Dimensions
VnicEgressDropsSecurityList	<b>Egress Packets Dropped by Security List</b>	packets	Packets sent by the VNIC, destined for the network, dropped due to security rule violations.	resourceId
VnicIngressDropsSecurityList	<b>Ingress Packets Dropped by Security List</b>	packets	Packets received from the network, destined for the VNIC, dropped due to security rule violations.	
VnicFromNetworkBytes*	<b>Bytes from Network</b>	bytes	Bytes received at the VNIC from the network, after drops.	
VnicFromNetworkPackets*	<b>Packets from Network</b>	packets	Packets received at the VNIC from the network, after drops.	

Metric	Metric Display Name	Unit	Description	Dimensions
VnicToNetworkBytes*	<b>Bytes to Network</b>	bytes	Bytes sent from the VNIC to the network, before drops.	
VnicToNetworkPackets*	<b>Packets to Network</b>	packets	Packets sent from the VNIC to the network, before drops.	

\* The Compute service separately reports network-related metrics *as measured on the instance itself and aggregated across all the attached VNICs*. Those metrics are available in the `oci_computeagent` metric namespace. For more information, see [Compute Instance Metrics](#).

## Tips for Working with VNIC Metrics

Here are some tips to help you use VNIC metrics.

### Default Metric Charts for One VNIC Versus Multiple VNICs

The default charts for VNIC metrics use these default settings:

- Time range = the last hour
- Interval = 1 minute
- Statistic displayed: Sum
- Aggregation of metric streams = not selected (which means each VNIC is displayed as a separate line on the chart)

You can [view the default charts with data for only a single VNIC](#) by viewing the VNIC's details in the Console. When looking at a single VNIC, these statistics are the most useful: sum, mean, max, and min.

You can [view the default charts with data for multiple VNICs](#) by going to the **Service Metrics** page in the Console. Make sure to select the desired compartment and metric namespace (`oci_vcn`) at the top of the page. For all the charts, you can either show each VNIC as a separate line, or show a single line that aggregates the data for all the VNICs in your selected compartment. To aggregate the data, select the check box for **Aggregate Metric Streams**.

When viewing aggregated data, you can use the P90 - P99.9 statistics to help identify typical behavior of your instance fleet and outliers. To view these statistics over an even larger number of data points, expand the chart's start and end time (for example, view the last 7 days instead of the last hour), and set the interval to 1 hour.

For general information about how to work with and modify the default metric charts, see [Using the Console](#) in the Monitoring documentation.

### Alarms for VNIC Metrics

You can set up [alarms](#) for a given metric. For VNICs, an alarm makes the most sense for the egress security list drops metric (`VnicEgressDropsSecurityList`). In a normal situation, you shouldn't have egress security list drops. If you do, it most likely means that one or both of these is true:

- An application is behaving in an unexpected manner
- Your security list is configured incorrectly

In either case, an alarm is warranted.

### Using the Console

#### To view default metric charts for a single VNIC

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance to view its details.
3. Click **Attached VNICs**.
4. Click the VNIC to view its details.
5. Under **Resources**, click **Metrics**.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

#### To view default metric charts for multiple VNICs

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that contains the VNICs you're interested in. Remember that a given VNIC resides in its *subnet's* compartment.
3. For **Metric Namespace**, select **oci\_vcn**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.



### Tip

If there are multiple VNICs in the compartment, the charts default to show a separate line for each VNIC. You can instead show a single line aggregated across all the VNICs by selecting the check box for **Aggregate Metric Streams** on the right side of the page.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

## Private IP Addresses

This topic describes how to manage the IPv4 addresses assigned to an instance in a virtual cloud network (VCN).

IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of IP Addresses

Instances use IP addresses for communication. Each instance has at least one private IP address and optionally one or more public IP addresses. A private IP address enables the instance to communicate with other instances inside the VCN, or with hosts in your on-premises network (via an IPSec VPN or Oracle Cloud Infrastructure FastConnect). A public IP address enables the instance to communicate with hosts on the internet. For more information, see these related topics:

- [Public vs. Private Subnets](#)
- [How IP Addresses Are Assigned](#)
- [Public IP Addresses](#)

### About the Private IP Object

The Networking service defines an object called a *private IP*, which consists of:

- Private IPv4 address, assigned by either you or Oracle.
- Optional hostname for DNS (see [DNS in Your Virtual Cloud Network](#)).

Each private IP object has an Oracle-assigned OCID (see [Resource Identifiers](#)). If you're using the API, you can also assign each private IP object a friendly name.

Each instance receives a *primary private IP* object during launch. The Networking service uses the Dynamic Host Configuration Protocol (DHCP) to pass the object's private IP address to the instance. This address does not change during the instance's lifetime and cannot be

removed from the instance. The private IP object is terminated when the instance is terminated.

If an instance has any [secondary VNICs](#) attached, each of those VNICs also has a primary private IP.

A private IP can have a [public IP](#) assigned to it at your discretion.

A private IP can be the target of a route rule in your VCN. For more information, see [Using a Private IP as a Route Target](#).

### About Secondary Private IP Addresses

You can add a *secondary private IP* to an instance after it's launched. You can add it to either the primary VNIC or a secondary VNIC on the instance. The secondary private IP address must come from the CIDR of the VNIC's subnet. You can move a secondary private IP from a VNIC on one instance to a VNIC on another instance if both VNICs belong to the same subnet.

Here are a few reasons why you might use secondary private IPs:

- **Instance failover:** You assign a secondary private IP to an instance. Then if the instance has problems, you can easily reassign that secondary private IP to a standby instance in the same subnet. If the secondary private IP has a public IP assigned to it, that public IP moves along with the private IP.
- **Running multiple services or endpoints on a single instance:** For example, you could have multiple container pods running on a single instance, and each uses an IP address from the VCN's CIDR. The containers have direct connectivity to other instances and services in the VCN. Another example: you could run multiple SSL websites with each one using its own IP address.

Here are more details about secondary private IP addresses:

- They're supported for all shapes and OS types, for both bare metal and VM instances.
- A VNIC can have a maximum of 31 secondary private IPs.
- They can be assigned only after the instance is launched (or the secondary VNIC is created/attached).

- A secondary private IP that is assigned to a VNIC in a [regional subnet](#) has a null availability domain attribute. Compare this with the VNIC's *primary* private IP, which always has its availability domain attribute set to the instance's availability domain, regardless of whether the instance's subnet is regional or AD-specific.
- Deleting a secondary private IP from a VNIC returns the address to the pool of available addresses in the subnet.
- They are automatically deleted when you terminate the instance (or detach/delete the secondary VNIC).
- The instance's bandwidth is fixed regardless of the number of private IP addresses attached. You can't specify a bandwidth limit for a particular IP address on an instance.
- A secondary private IP can have a [reserved public IP](#) assigned to it at your discretion.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

## Using the Console

### To view an instance's private IPs

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.

The [primary VNIC and any secondary VNICs](#) assigned to the instance are displayed.

5. Click the VNIC you're interested in.

6. Under **Resources**, click **IP Addresses**.

The VNIC's primary private IP and any secondary private IPs are displayed.

### To assign a new secondary private IP to a VNIC

1. Confirm you're viewing the compartment that contains the instance you're interested in.

2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.

3. Click the instance to view its details.

4. Under **Resources**, click **Attached VNICs**.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. Click the VNIC you're interested in.

6. Under **Resources**, click **IP Addresses**.

The VNIC's primary private IP and any secondary private IPs are displayed.

7. Click **Assign Private IP Address**.

8. Enter the following:

- **Private IP Address:** Optional. An available private IP address of your choice from the subnet's CIDR (otherwise the private IP address is automatically assigned).
- **Unassign if already assigned to another VNIC:** Select this check box to force reassignment of the IP address if it's already assigned to another VNIC in the subnet. Relevant only if you specify a private IP address in the preceding field.
- **Hostname:** Optional. A hostname to be used for DNS within the cloud network. Available only if the VCN and subnet both have DNS labels. See [DNS in Your Virtual Cloud Network](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
  - **Public IP address:** Whether to assign a public IP address. Available only if the VNIC is in a public subnet. See [Public IP Addresses](#).
9. Click **Assign**.  
The secondary private IP is created and then displayed on the **IP Addresses** page for the VNIC.
  10. Configure the IP address:
    - For instances running a variant of Linux, see [Linux: Details about Secondary IP Addresses](#).
    - For Windows instances, see [Windows: Details about Secondary IP Addresses](#).

### To move a secondary private IP to another VNIC in the same subnet

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. Click **Assign Private IP Address**.

8. Enter the following:
  - **Private IP Address:** The secondary private IP address you want to move.
  - **Unassign if already assigned to another VNIC:** Select this check box to move the secondary IP address from the VNIC it's currently assigned to.
  - **Hostname:** Optional. The hostname to be used for DNS within the cloud network. Available only if the VCN and subnet both have DNS labels. See [DNS in Your Virtual Cloud Network](#).
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
  - **Public IP address:** Whether to assign a public IP address. Available only if the VNIC is in a public subnet. See [Public IP Addresses](#).
9. Click **Assign**.

The private IP address is moved from the original VNIC to the new VNIC.

### To update the hostname for an existing private IP

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.

The VNIC's primary private IP and any secondary private IPs are displayed.

7. For the IP address you're interested in, click the Actions icon (three dots), and then click **Edit**.
8. Make your changes and click **Update**.

### To delete a secondary private IP from a VNIC



#### Warning

If the private IP is the [target of a route rule](#), deleting it from the VNIC causes the route rule to blackhole and the traffic will be dropped.

Prerequisite: Oracle recommends removing the IP address from the OS configuration before deleting it from the VNIC. See [Linux: Details about Secondary IP Addresses](#) or [Windows: Details about Secondary IP Addresses](#).

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the private IP you want to delete, click the Actions icon (three dots), and then click **Delete Private IP**.
8. Confirm when prompted.

The private IP address is returned to the pool of available addresses in the subnet.

### To manage tags for a private IP

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the private IP you're interested in, click the Actions icon (three dots), and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage private IPs on a VNIC, use these operations:

- [GetPrivateIp](#): Use this to get a single `privateIp` object by specifying its OCID.
- [ListPrivateIps](#): Use this to get a single `privateIp` object by specifying the private IP address (for example, 10.0.3.3) and the subnet's OCID. Or you can list all the `privateIp` objects in a given subnet, or just the ones assigned to a given VNIC.
- [CreatePrivateIp](#): Use this to assign a new secondary private IP to a VNIC.

- [UpdatePrivateIp](#): Use this to reassign a secondary private IP to a different VNIC in the same subnet, or to update the hostname or display name of a secondary private IP.
- [DeletePrivateIp](#): Use this to delete a secondary private IP from a VNIC. The private IP address is returned to the subnet's pool of available addresses.

### Linux: Details about Secondary IP Addresses

After assigning a secondary private IP to a VNIC, you must configure the OS to use it.

#### Basic Commands (Not Persistent Through a Reboot)

On the instance, run the following command. It works on all variants of Linux, for both bare metal and VM instances:

```
ip addr add <address>/<subnet_prefix_len> dev <phys_dev> label <phys_dev>:<addr_seq_num>
```

where:

- `<address>`: The secondary private IP address.
- `<subnet_prefix_len>`: The subnet's prefix length. For example, if the subnet is 192.168.20.0/24, the subnet prefix length is 24.
- `<phys_dev>`: The interface to add the address to (for example, ens2f0).
- `<addr_seq_num>`: The sequential number in the stack of addresses on the device (for example, 0).

For example:

```
ip addr add 192.168.20.50/24 dev ens2f0 label ens2f0:0
```

Later if you want to delete the address, you can use:

```
ip addr del 192.168.20.50/24 dev ens2f0:0
```

Also make sure to [delete the secondary IP from the VNIC](#). You can do that before or after executing the above command to delete the address from the OS configuration.



### Note

If you've assigned a secondary IP to a [secondary VNIC](#), and you're using policy-based routing for the secondary VNIC, make sure to configure the route rules to look up the same route table for the secondary IP address.

### Configuration File (Persistent Through a Reboot)

You can make the configuration persistent through a reboot by adding the information to a configuration file.

### For Oracle Linux and CentOS

Create an `ifcfg` file named `/etc/sysconfig/network-scripts/ifcfg-<phys_dev>:<addr_seq_num>`. To continue with the preceding example, the file name would be `/etc/sysconfig/network-scripts/ifcfg-ens2f0:0`, and the contents would be:

```
DEVICE="ens2f0:0"
BOOTPROTO=static
IPADDR=192.168.20.50
NETMASK=255.255.255.0
ONBOOT=yes
```



### Note

If you've assigned a secondary IP to a [secondary VNIC](#), and you're using policy-based routing for the secondary VNIC, make sure to configure the route rules to look up the same route table for the secondary IP address.

### For Ubuntu

Add the following information to the end of the `/etc/network/interfaces` file:

```
auto <phys_dev>:<addr_seq_num>
iface <phys_dev>:<addr_seq_num> inet static
 address <address>
 netmask <address_netmask>
```

Where the netmask is not the prefix but the 255.255.x.x. address.

To continue with the earlier example:

```
auto ens2f0:0
iface ens2f0:0 inet static
 address 192.168.20.50
 netmask 255.255.255.0
```



#### Note

If you've assigned a secondary IP to a [secondary VNIC](#), and you're using policy-based routing for the secondary VNIC, make sure to configure the route rules to look up the same route table for the secondary IP address.

### Windows: Details about Secondary IP Addresses

After assigning a secondary private IP to a VNIC, you must configure the OS to use it. Here are instructions for using a PowerShell script or the Network and Sharing Center UI.

#### Using a PowerShell Script

You must run PowerShell as an administrator. The script configures two things: static IP addressing on the instance and the secondary private IP. The configuration persists through a

reboot of the instance.

1. In your browser, go to the Console, and note the secondary private IP address that you want to configure on the instance.
2. Connect to the instance, and run the following command at a command prompt:

```
ipconfig /all
```

3. Note the values for the following items so you can enter them into the script in the next step:
  - Default Gateway
  - DNS Servers
4. Replace the variables in the following PowerShell script with your own values:

```
netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-NetIPInterface -DHCP Disabled
$netadapter | New-NetIPAddress -AddressFamily IPv4 -IPAddress <secondary_IP_address> -
PrefixLength <subnet_prefix_length> -Type Unicast -DefaultGateway <default_gateway>
Set-DnsClientServerAddress -InterfaceAlias Ethernet -ServerAddresses <DNS_server>
```

For example:

```
netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-NetIPInterface -DHCP Disabled
$netadapter | New-NetIPAddress -AddressFamily IPv4 -IPAddress 192.168.11.14 -PrefixLength 24 -
Type Unicast -DefaultGateway 192.168.11.1
Set-DnsClientServerAddress -InterfaceAlias Ethernet -ServerAddresses 169.254.169.254
```

5. Save the script with the name of your choice and a `.ps1` extension, and run it on the instance.

```

Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2014 Microsoft Corporation. All rights reserved.

PS C:\Windows\system32> cd
PS C:\Windows\system32> cd C:\Users\opc\Documents
PS C:\Users\opc\Documents> .\set-secondary-ip.ps1

IPAddress : 192.168.11.14
InterfaceIndex : 12
InterfaceAlias : Ethernet
AddressFamily : IPv4
Type : Unicast
PrefixLength : 24
PrefixOrigin : Manual
SuffixOrigin : Manual
AddressState : Tentative
ValidLifetime : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)
SkipAsSource : False
PolicyStore : ActiveStore

IPAddress : 192.168.11.14
InterfaceIndex : 12
InterfaceAlias : Ethernet
AddressFamily : IPv4
Type : Unicast
PrefixLength : 24
PrefixOrigin : Manual
SuffixOrigin : Manual
AddressState : Invalid
ValidLifetime : Infinite ([TimeSpan]::MaxValue)
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)
SkipAsSource : False
PolicyStore : PersistentStore

PS C:\Users\opc\Documents> gc .\set-secondary-ip.ps1
$netadapter = Get-NetAdapter -Name Ethernet
$netadapter | Set-NetIPInterface -DHCP Disabled
$netadapter | New-NetIPAddress -AddressFamily IPv4 -IPAddress 192.168.11.14 -PrefixLength 24 -Type Unicast -DefaultGate
way 192.168.11.1
Set-NetClientServerAddress -InterfaceAlias Ethernet -ServerAddresses 169.254.169.254
PS C:\Users\opc\Documents>

```

If you run `ipconfig /all` again, you'll see that DHCP has been disabled and the secondary private IP address is included in the list of IP addresses.

Later if you want to delete the address, you can use this command:

```
Remove-NetIPAddress -IPAddress 192.168.11.14 -InterfaceAlias Ethernet
```

Also make sure to [delete the secondary IP from the VNIC](#). You can do that before or after executing the above command to delete the address from the OS configuration.

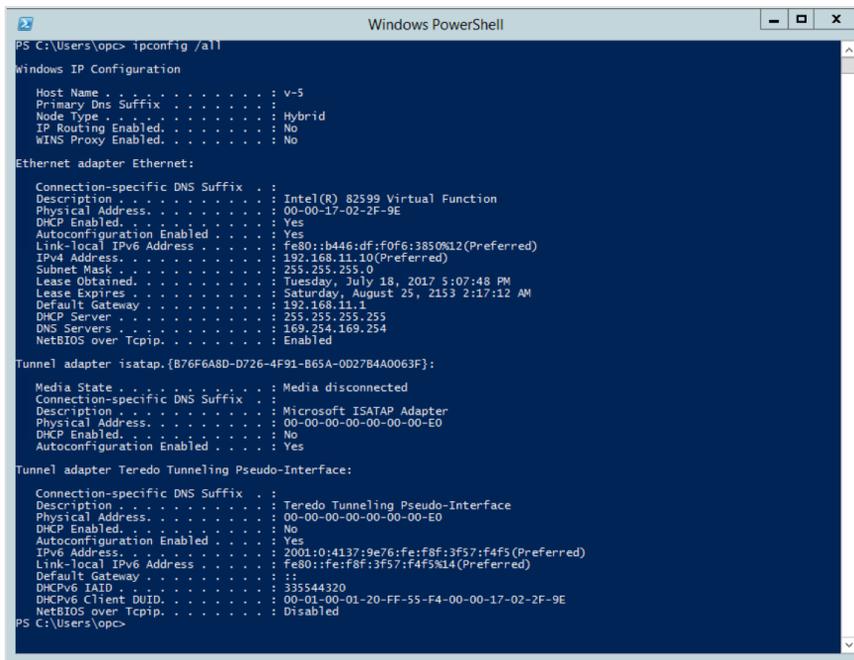
## Using the Network and Sharing Center UI

The following instructions configure two things: static IP addressing on the instance and the secondary private IP. The configuration persists through a reboot of the instance.

1. In your browser, go to the Console, and note the secondary private IP address that you want to configure on the instance.

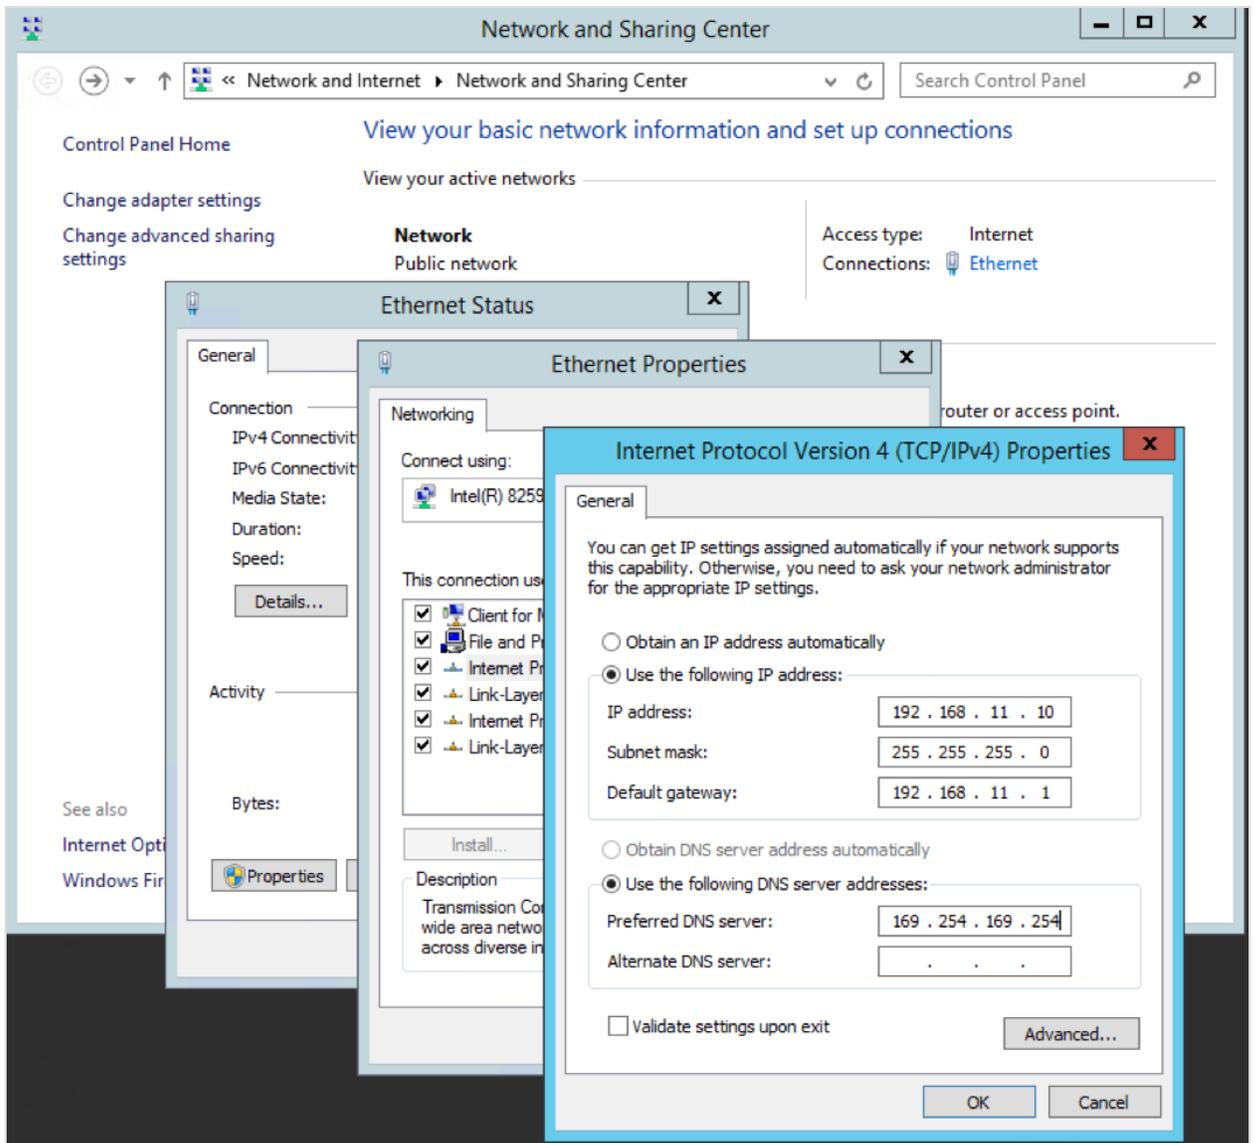
2. Connect to the instance, and run the following command at a command prompt:

```
ipconfig /all
```



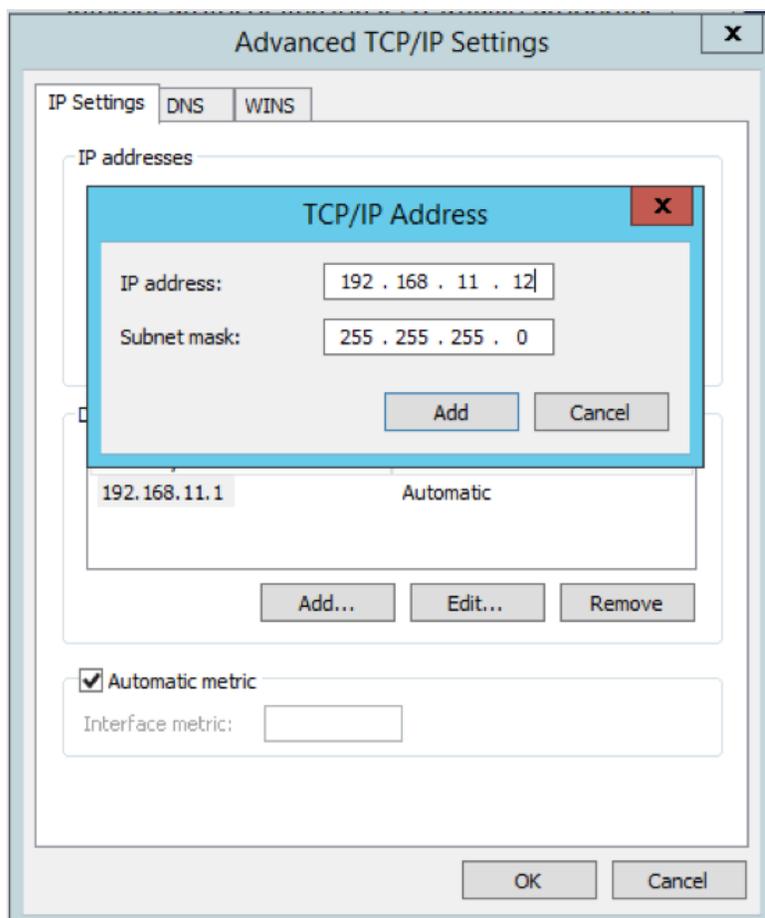
3. Note the values for the following items so you can enter them elsewhere in a later step:
  - IPv4 Address
  - Subnet Mask
  - Default Gateway
  - DNS Servers
4. In the instance's **Control Panel**, open the **Network and Sharing Center** (see the image below for the set of dialog boxes you'll see in these steps).
5. For the active networks, click the connection (**Ethernet**).

6. Click **Properties**.
7. Click **Internet Protocol Version 4 (TCP/IPv4)**, and then click **Properties**.
8. Select the radio button for **Use the following IP address**, and then enter the values you noted earlier for the IP address, subnet mask, default gateway, and DNS servers.



9. Click **Advanced...**
10. Under **IP addresses**, click **Add...**

11. Enter the secondary private IP address and the subnet mask you used earlier and click **Add**.



12. Click **OK** until the Network and Sharing Center is closed.
13. Verify the changes by returning to the command prompt and running `ipconfig /all`.

You should now see that DHCP is disabled (static IP addressing is enabled), and the secondary private IP address is in the list of addresses displayed. The address is now configured on the instance and available to use.

```
Windows PowerShell
PS C:\Users\opc> ipconfig /all

Windows IP Configuration

Host Name : v-5
Primary Dns Suffix :
Node Type : Hybrid
IP Routing Enabled. : No
WINS Proxy Enabled. : No

Ethernet adapter Ethernet:

Connection-specific DNS Suffix . . :
Description : Intel(R) 82599 Virtual Function
Physical Address. : 00-00-17-02-2F-9E
DHCP Enabled. : No
Autoconfiguration Enabled : Yes
Link-local IPv6 Address : fe80::b446:df:f0f6:3850%12(Preferred)
IPv4 Address. : 192.168.11.10(Preferred)
Subnet Mask : 255.255.255.0
Lease Obtained. : Tuesday, July 18, 2017 5:07:48 PM
Lease Expires : Saturday, August 25, 2153 2:30:54 AM
IPv4 Address. : 192.168.11.12(Preferred)
Subnet Mask : 255.255.255.0
Default Gateway : 192.168.11.1
DNS Servers : 169.254.169.254
NetBIOS over Tcpip. : Enabled

Tunnel adapter isatap.{B7F6A8D-D726-4F91-B65A-0D27B4A0063F}:

Media State : Media disconnected
Connection-specific DNS Suffix . . :
Description : Microsoft ISATAP Adapter
Physical Address. : 00-00-00-00-00-00-E0
DHCP Enabled. : No
Autoconfiguration Enabled : Yes

Tunnel adapter Teredo Tunneling Pseudo-Interface:

Connection-specific DNS Suffix . . :
Description : Teredo Tunneling Pseudo-Interface
Physical Address. : 00-00-00-00-00-00-E0
DHCP Enabled. : No
Autoconfiguration Enabled : Yes
IPv6 Address. : 2001:0:4137:9e76:2044:3b51:3f57:f4f5(Preferred)
Link-local IPv6 Address : fe80::2044:3b51:3f57:f4f5%14(Preferred)
Default Gateway : ::
DHCPv6 Iaid : 335544320
DHCPv6 Client DUID. : 00-01-00-01-20-ff-55-f4-00-00-17-02-2f-9e
NetBIOS over Tcpip. : Disabled

PS C:\Users\opc>
```



### Note

You might not see the primary private IP address when you again view the properties for Internet Protocol Version 4 (TCP/IPv4) in the Network and Sharing Center UI. The best way to confirm your changes is to use `ipconfig /all` at the command line.

## Public IP Addresses

This topic describes how to manage public IPv4 addresses on instances in a virtual cloud network (VCN).

IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of Public IP Addresses

A public IP address is an IPv4 address that is reachable from the internet. If a resource in your tenancy needs to be directly reachable from the internet, it must have a public IP address. Depending on the type of resource, there might be other requirements.

Certain types of resources in your tenancy are designed to be directly reachable from the internet and therefore automatically come with a public IP address. For example: a NAT gateway or a public load balancer. Other types of resources are directly reachable only if you configure them to be. For example: instances in your VCN.

This topic focuses on these subjects:

- The types of public IP addresses and their characteristics
- How to control whether an instance has a public IP address

For more information about resources that automatically get a public IP address, see [Resources That Always Get a Public IP](#).

### Instances and Public IP Addresses

You can assign a public IP address to an instance to enable communication with the internet. The instance is assigned a public IP address from the Oracle Cloud Infrastructure address pool.

The assignment is actually to a [private IP](#) object on the instance. The [VNIC](#) that the private IP is assigned to must be in a [public subnet](#). A given instance can have multiple secondary VNICs, and a given VNIC can have multiple secondary private IPs. So you can assign a given instance multiple public IPs across one or more VNICs if you like.

For an instance to communicate directly with the internet, all of the following are required:

- The instance must be in a [public subnet](#).
- The instance must have a public IP address.
- The instance's VCN must have an [internet gateway](#).
- The public subnet must have [route tables](#) and [security lists](#) configured accordingly.



#### Tip

Oracle Cloud Infrastructure FastConnect public peering lets your on-premises network access the public IP addresses of resources in Oracle Cloud Infrastructure *without the traffic traversing the internet*. For more information, see [FastConnect](#).

### The Public IP Object

The Networking service defines an object called a *public IP*, which consists of these items:

- Public IPv4 address (chosen by Oracle)
- Properties that further define the public IP's type and behavior

Each public IP object has an Oracle-assigned OCID (see [Resource Identifiers](#)). If you're using the API, you can also assign each public IP object a friendly name.

## Types of Public IPs

There are two types of public IPs:

- **Ephemeral:** Think of it as temporary and existing for the lifetime of the instance.
- **Reserved:** Think of it as persistent and existing beyond the lifetime of the instance it's assigned to. You can unassign it and then reassign it to another instance whenever you like. Exception: reserved public IPs on public load balancers. See [Resources That Always Get a Public IP](#).

The following table summarizes the differences between the two types.

Characteristic	Ephemeral Public IPs	Reserved Public IPs
<b>Allowed assignment</b>	To a VNIC's <a href="#">primary private IP</a> only Limits: <ul style="list-style-type: none"> <li>• One per <a href="#">VNIC</a></li> <li>• Two per VM instance, and 16 per bare metal instance</li> </ul>	To either a primary or <a href="#">secondary private IP</a> Limit: 32 per <a href="#">VNIC</a>
<b>Creation</b>	Optionally created and assigned during instance launch or secondary VNIC creation. You can create and assign one later if the VNIC doesn't already have one.	You create one at any time. You can then assign it when you like. Limit: You can create 50 per region

Characteristic	Ephemeral Public IPs	Reserved Public IPs
<b>Unassignment</b>	<p>You can unassign it at any time, which deletes it. You might do this if whoever launched the instance included a public IP, but you don't want the instance to have one.</p> <p>When you stop an instance, its ephemeral public IPs remain assigned to the instance.</p>	<p>You can unassign it at any time, which returns it to your tenancy's pool of reserved public IPs.</p>
<b>Moving to a different resource</b>	<p>You cannot move an ephemeral public IP to a different private IP.</p>	<p>If assigned to a secondary private IP: If you move the private IP to a different VNIC (must be in the same subnet), the reserved public IP goes with it.</p> <p>You can move it (unassign and then reassign it) at any time to another private IP in the same region. Can be in a different VCN or availability domain.</p>
<b>Automatic deletion</b>	<p>Its lifetime is tied to the private IP's lifetime. Automatically unassigned and deleted when:</p> <ul style="list-style-type: none"> <li>• Its private IP is deleted</li> <li>• Its VNIC is detached or terminated</li> <li>• Its instance is terminated</li> </ul>	<p>Never. Exists until you delete it.</p>

Characteristic	Ephemeral Public IPs	Reserved Public IPs
<b>Scope</b>	Availability domain	Regional (can be assigned to a private IP in any availability domain in the region)
<b>Compartment and availability domain</b>	Same as the private IP's	Can be different from the private IP's

When you launch an instance in a public subnet, by default, the instance gets a public IP unless you say otherwise. See [To choose whether an ephemeral public IP is assigned when launching an instance](#).

After you create a given public IP, you can't change which type it is. For example, if you launch an instance that is assigned an ephemeral public IP with address 129.146.1.9, you can't convert the ephemeral public IP to a reserved public IP with address 129.146.1.9.

The preceding table notes the public IP limits per VNIC and instance. If you try to perform any operation that assigns or moves a public IP to a VNIC or instance that has already reached its public IP limit, an error is returned. The operations include:

- Assigning a public IP
- Creating a new secondary VNIC with a public IP
- Moving a private IP with a public IP to another VNIC
- Moving a public IP to another private IP

### Resources That Always Get a Public IP

As mentioned earlier, certain types of resources are designed to be directly reachable from the internet. Examples: a NAT gateway or a public load balancer. These resources automatically get a public IP address upon creation. Oracle chooses the public IP address from the Oracle pool. You can't remove or change the address.

For public load balancers, the address is a regional reserved public IP that is assigned to a private IP on the load balancer. This public IP appears in the list of your tenancy's reserved public IPs, which [you can view in the Console](#). However, unlike other reserved public IPs that you create, you have no control over this public IP. You can't edit it or unassign it from the load balancer yourself. It's automatically unassigned and deleted from your tenancy when you terminate the load balancer.

For NAT gateways, the address is a regional ephemeral public IP that is assigned to the NAT gateway. Like other ephemeral public IPs, it's automatically unassigned and deleted when you terminate its assigned resource (the NAT gateway). However, unlike other ephemeral public IPs, you can't edit it or unassign it yourself.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

## Ephemeral Public IPs: Using the Console

To choose whether an ephemeral public IP is assigned when launching an instance

When you launch an instance into a [public subnet](#), there's an **Assign public IP address** check box in the **Launch Instance** dialog box (in the **Advanced Options**, on the **Networking** tab). By default, the check box is selected, which means the instance gets an ephemeral public IP.

If you do NOT want an ephemeral public IP assigned, you can either:

- Clear the **Assign public IP address** check box
- [Delete the ephemeral public IP after instance launch](#)

### To assign an ephemeral public IP when creating a secondary VNIC

When you add a secondary VNIC to an instance, you choose whether the primary private IP on the new VNIC gets an ephemeral public IP. This choice is available only if the secondary VNIC is in a [public subnet](#).

In the **Create VNIC** dialog box, there's an **Assign public IP address** check box. By default, the check box is NOT selected, which means the secondary VNIC does not get an ephemeral public IP. You must select the check box.

For instructions, see [To create and attach a secondary VNIC](#).

### To assign an ephemeral public IP to an existing primary private IP

Prerequisite: The primary private IP must not have a reserved or ephemeral public IP already assigned to it. If it does, first [delete the ephemeral public IP](#), or [unassign the reserved public IP](#).

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the VNIC's primary private IP, click the Actions icon (three dots), and then click **Edit**.

8. In the **Public IP Address** section, for **Public IP type**, select the radio button for **Ephemeral Public IP**.
9. In the **Ephemeral Public IP Name** field, enter an optional friendly name for the public IP. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
10. Click **Update**.

### To delete an ephemeral public IP from an instance

Deleting an ephemeral public IP automatically unassigns it from its private IP.

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the VNIC's primary private IP, click the Actions icon (three dots), and then click **Edit**.
8. In the **Public IP Address** section, for **Public IP Type**, select the radio button for **No Public IP**.
9. Click **Update**.

### To change the display name for an ephemeral public IP

1. Confirm you're viewing the compartment that contains the instance you're interested in.

2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the VNIC's primary private IP, click the Actions icon (three dots), and then click **Edit**.
8. In the **Public IP Address** section, edit the **Ephemeral Public IP Name**. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
9. Click **Update**.

### Reserved Public IPs: Using the Console

#### To view your reserved public IPs

1. Confirm you're viewing the region and compartment you're interested in.
2. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.

The details of the reserved public IPs in the selected region and compartment are displayed. If the reserved public IP is assigned, there's a link to the relevant VNIC.

#### To create a new reserved public IP in your pool

1. Confirm you're viewing the region and compartment where you want to create the

reserved public IP.

2. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.
3. Click **Create Reserved Public IP**.
4. Enter the following:
  - **Name:** An optional friendly name for the reserved public IP. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
  - **Compartment:** Leave as is.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create Reserved Public IP**.

The new reserved public IP is created and displayed on the page. You can now [assign it to an existing private IP](#) if you like.

### To delete a reserved public IP from your pool

The reserved public IP can be in the "Assigned" state. Deleting a reserved public IP automatically unassigns it from its private IP.

1. Confirm you're viewing the region and compartment that contains the reserved public IP.
2. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.

3. For the reserved public IP you want to delete, click the Actions icon (three dots), and then click **Terminate**.
4. Confirm when prompted.

After a few seconds, the reserved public IP is unassigned (if it was assigned) and deleted from your pool.

### To assign a reserved public IP to a private IP

Prerequisite: The private IP must not have an ephemeral or reserved public IP already assigned to it. If it does, first [delete the ephemeral public IP](#), or [unassign the reserved public IP](#).

1. Confirm you're viewing the compartment that contains the instance with the private IP you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the private IP you're interested in, click the Actions icon (three dots), and then click **Edit**.
8. In the **Public IP Address** section, for **Public IP Type**, select the radio button for **Reserved Public IP**.
9. Enter the following:
  - **Compartment:** The compartment that contains the reserved public IP you want to assign.

- **Reserved Public IP:** The reserved public IP you want to assign. You have three choices:
    - Create a new reserved public IP. You may optionally provide a friendly name for it. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
    - Assign a reserved public IP that is currently unassigned.
    - Move a reserved public IP from another private IP.
10. Click **Update**.

### To unassign a reserved public IP and return it to the pool

1. Confirm you're viewing the compartment that contains the instance with the reserved public IP you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For the private IP you're interested in, click the Actions icon (three dots), and then click **Edit**.
8. In the **Public IP Address** section, for **Public IP Type**, select the radio button for **No Public IP**.
9. Click **Update**.

The reserved public IP is unassigned and returned to your pool.

### To move a reserved public IP from one private IP to another

Let's say you want to move a reserved public IP from private IP 1 to private IP 2. In summary: Make sure private IP 2 doesn't have a public IP already assigned to it. Then assign the reserved public IP to private IP 2. It will be automatically unassigned from private IP 1 first. Detailed instructions:

1. Confirm you're viewing the compartment that contains the instance with private IP 2.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IP Addresses**.  
The VNIC's primary private IP and any secondary private IPs are displayed.
7. For private IP 2, click the Actions icon (three dots), and then click **Edit**.
8. If private IP 2 already has a public IP assigned to it:
  - a. In the **Public IP Address** section, select the radio button for **No Public IP**.
  - b. Click **Update**.
  - c. Again for private IP 2, click the Actions icon (three dots), and then click **Edit**.
9. In the **Public IP Address** section, select the radio button for **Reserved Public IP**.
10. Enter the following:
  - **Compartment:** The compartment that contains the reserved public IP you want to assign.
  - **Reserved Public IP:** The reserved public IP you want to assign. It will be moved from the public IP it's currently assigned to.
11. Click **Update**.

### To change the display name for a reserved public IP

1. Confirm you're viewing the region and compartment that contains the reserved public IP.
2. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.
3. For the reserved public IP you want to edit, click the Actions icon (three dots), and then click **Edit**.
4. Make your changes to the friendly name. The name doesn't have to be unique, and you can change it later. Avoid entering confidential information.
5. Click **Save**.

### To manage tags for a reserved public IP

1. Confirm you're viewing the region and compartment that contains the reserved public IP.
2. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.
3. For the reserved public IP you're interested in, click the Actions icon (three dots), and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move a reserved public IP to a different compartment

You can move a reserved public IP from one compartment to another. When you move a reserved public IP to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under Core Infrastructure, go to **Networking** and click **Public IPs**.

2. For the reserved public IP you're interested in, click the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage public IPs, use these operations:

- [GetPublicIp](#): Use this to get a `publicIp` object by specifying its OCID.
- [GetPublicIpByIpAddress](#): Use this to get a `publicIp` object by specifying its public IP address.
- [GetPublicIpByPrivateIpId](#): Use this to get a `publicIp` object by specifying the OCID of the private IP it's assigned to.
- [ListPublicIps](#): Use this to list either your ephemeral or reserved `publicIp` objects.
- [CreatePublicIp](#): Use this to create a new reserved public IP in your pool.
- [UpdatePublicIp](#): Use this to assign, reassign, or unassign a reserved public IP, or to update the display name of an ephemeral or reserved public IP. You can also update a reserved public IP's tags.
- [DeletePublicIp](#): Use this to delete an ephemeral public IP from its private IP, or delete a reserved public IP from your pool. The operation first unassigns the public IP.

- [ChangePublicIpCompartment](#): Use this to move a reserved public IP from one compartment to another. This operation applies only to reserved public IPs. Ephemeral public IPs always belong to the same compartment as their VNIC and move accordingly.

## IPv6 Addresses

This topic describes support for IPv6 addressing in your VCN.

### Highlights

- IPv6 addressing is currently supported only in the US Government Cloud. See [For All Government Cloud Customers](#).
- During VCN creation, you choose whether the VCN is enabled for IPv6. You also choose whether each subnet in an IPv6-enabled VCN is enabled for IPv6. **You cannot change whether a VCN or subnet is IPv6-enabled after creation.**
- IPv6-enabled VCNs use a /48 IPv6 CIDR block. Oracle assigns a /48 *public* IPv6 CIDR block to the VCN for internet communication. You can either let the *private* IPv6 CIDR block be the same as the public CIDR, or provide your own value (in which case it's referred to as a *custom* IPv6 CIDR). All subnets are /64.
- You also choose whether a given VNIC in an IPv6-enabled subnet has IPv6 addresses (up to 32 maximum per VNIC), and whether each address can be used for internet communication.
- You can choose which particular IPv6 address in the subnet is assigned to a VNIC. This means you can plan how the VCN's private and public address space is allocated within your organization.
- Only these Networking gateways support IPv6 traffic: dynamic routing gateway (DRG) and internet gateway.
- Both inbound- and outbound-initiated IPv6 connections are supported between your VCN and the internet, and between your VCN and your on-premises network.
- Private IPv6 traffic between resources within a region (intra- and inter-VCN) is not yet supported. See other important details in [Routing for IPv6 Traffic](#).

- Both FastConnect and IPsec VPN support IPv6 traffic between your VCN and on-premises network. You must configure the FastConnect or IPsec VPN for IPv6.

### Overview of IPv6 Addresses

Oracle supports dual-stack IPv4/IPv6 addressing for VCNs. Every VCN always supports IPv4, and you can optionally enable IPv6 during VCN creation. Enabling IPv6 for the VCN means that when you create a subnet, you can optionally enable it to also have IPv6 addresses. Therefore a VCN can have a mix of IPv4-only subnets and IPv6-enabled subnets.

After you create a Compute instance, you may optionally add an *IPv6* to the VNIC. You can add up to 32 IPv6s to a given VNIC. You can remove an IPv6 from a VNIC at any time.

### **CIDRs Assigned to an IPv6-Enabled VCN**

An IPv6-enabled VCN has 3 CIDR blocks assigned to it. The following table summarizes them.

IPv4 or IPv6	Use and Size	Who Assigns the CIDR Block	Allowed Values
Private IPv4 CIDR	Private communication /16 to /30	You	Typically RFC 1918 range
Private IPv6 CIDR *	On-premises communication Only /48	<p>Optionally, you can assign it. If you do, it's referred to in this documentation as a <i>custom IPv6 CIDR</i>.</p> <p>Or, you can let Oracle assign it.</p> <p><b>Important:</b> You must assign this value if you want instances in the same VCN to communicate with each other using public IPv6 addresses. For more information, see <a href="#">Routing for IPv6 Traffic</a>.</p>	If you assign it, see <a href="#">Allowed Custom IPv6 CIDR Ranges</a> .

IPv4 or IPv6	Use and Size	Who Assigns the CIDR Block	Allowed Values
Public IPv6 CIDR	Internet communication Only /48	Oracle	<p>If you assign the VCN's private IPv6 CIDR, it will be different from the <i>public</i> IPv6 CIDR that Oracle assigns.</p> <p>But if you let Oracle assign the VCN's private IPv6 CIDR, Oracle uses the <i>same</i> CIDR for both the private and public IPv6 CIDRs. That means the private address and public address for a given IPv6 <i>are the same</i>.</p>
<p>* Oracle assigns IPv6 CIDR blocks that are NOT in the IPv6 unique local address (ULA) range. This range is analogous to the IPv4 RFC 1918 private ranges. Therefore, all Oracle-assigned IPv6 CIDRs can be considered <i>public</i> ranges by this definition.</p>			

### Allowed Custom IPv6 CIDR Ranges

Your custom IPv6 CIDR block can be in these general ranges:

- Global unicast: 2000::/3
- ULA: fc00::/7

But it cannot be in these [IANA special registry](#) ranges:

- IETF protocol assignments: 2001::/23
- Documentation: 2001:db8::/32
- 6to4: 2002::/16
- Direct Delegation AS112 Service: 2620:4f:8000::/48

### Internet Communication

Regardless of whether you or Oracle assigns the VCN's private IPv6 CIDR, Oracle also assigns the VCN an IPv6 CIDR block for the *public* IP address space (the *public IPv6 CIDR*). These addresses are used for internet communication. If you do not assign a custom CIDR, Oracle uses the same Oracle-assigned public IPv6 CIDR for the private address space. **This means that a given VNIC might use the same IPv6 IP address for both private and internet communication.**

You control whether a given IPv6 address can be used for internet communication. If the IPv6 is in a private subnet, it can never be used for internet communication. If it's in a public subnet, you can [enable or disable internet access](#) for that IPv6 at any time. If internet access is enabled, the IPv6 uses its public IPv6 address for communication.

### Assignment of IPv6 Addresses to a VNIC

To enable IPv6 for a given VNIC, you [assign an IPv6 to the VNIC](#). You can assign up to 32 IPv6s to a VNIC.

As with IPv4, when assigning an IPv6, you can specify the particular address you want to use, or let Oracle choose one for you. By choosing the IPv6 addresses yourself, you can plan how the VCN's private and public address space is allocated within your organization.

You also choose whether the IPv6 has internet access enabled (it is enabled by default if the VNIC is in a public subnet). A VNIC with an internet-enabled IPv6 is not required to have a public IPv4 address.

You can [move an IPv6 address from one VNIC to another in the same subnet](#).

After adding an IPv6 to a VNIC, you must [configure the instance's OS to use the IPv6](#).

### Format of IPv6 Addresses

IPv6 addresses have 128 bits.

An IPv6 CIDR block for a VCN must be /48 in size. The left 48 bits identify the VCN portion of the address. For example:

*2001:0db8:0123::/48*

An IPv6 CIDR block for a subnet must be /64 in size. The right 16 bits in a subnet's CIDR identify the subnet portion of the address. In the following example, the **1111** is the unique portion for the subnet:

```
2001:0db8:0123:1111::/64
```

The right-most 64 bits of an IPv6 address identify the unique portion specific to the particular IPv6 address. For example:

```
2001:0db8:0123:1111:abcd:ef01:2345:6789
```

For a given IPv6, those right-most 64 bits **are identical for both the private and public address for an IPv6**. When you assign an IPv6 to a VNIC, you can specify which specific IPv6 address to use (those 64 bits). Therefore you can control how the private and public address space is allocated within your organization.

### Example 1: You assign a custom CIDR



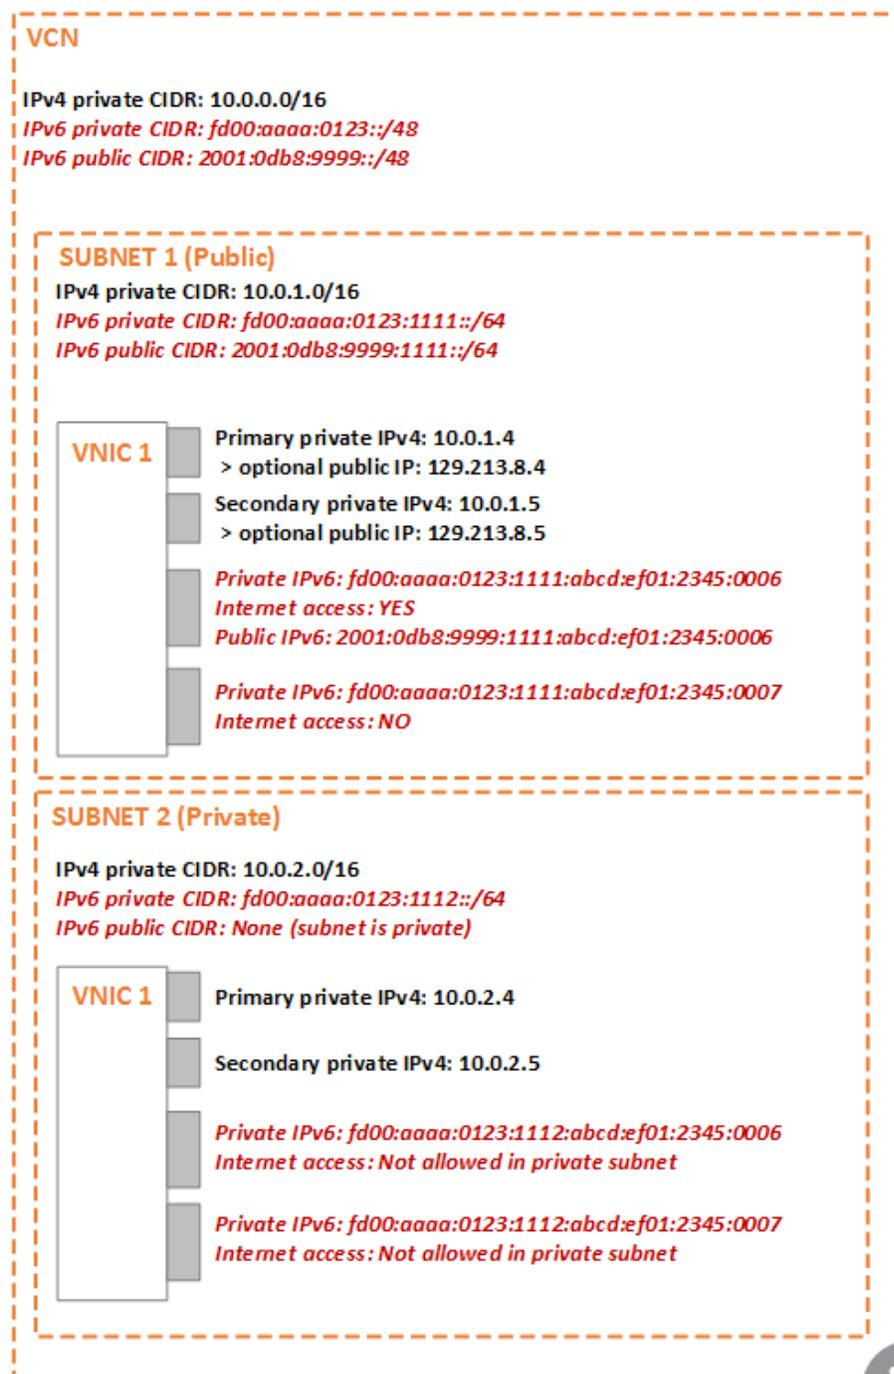
#### Important

Oracle recommends this option if you want instances within the same IPv6-enabled VCN to communicate with each other using public IPv6 addresses. For more information, see [Routing for IPv6 Traffic](#).

Let's say you provide this custom IPv6 CIDR: **fd00:aaaa:0123**::/48.

Oracle assigns a separate CIDR block for the VCN's public CIDR: **2001:0db8:9999**::/48.

The following diagram illustrates the VCN and includes two subnets: public subnet 1111 and private subnet 1112.



The VNIC in the public subnet has a primary private IPv4 (10.0.1.4) with an optional public IP address assigned. The VNIC has a secondary private IPv4 (10.0.1.5), also with an optional public IP address assigned.

The VNIC also has two IPv6s. The first one has internet access enabled and therefore has both private and public IPv6 addresses, which are the following:

- Private IPv6 address: fd00:aaaa:0123:1111:abcd:ef01:2345:0006
- Public IPv6 address: 2001:0db8:9999:1111:abcd:ef01:2345:0006

**Notice that the right-most 64 bits are the same for both the private and public IP address, as are the subnet's 16 bits. Only the left 48-most bits differ.**

The second IPv6 in the public subnet does not have internet access enabled and therefore has only a private IP address, which is fd00:aaaa:0123:1111:abcd:ef01:2345:0007.

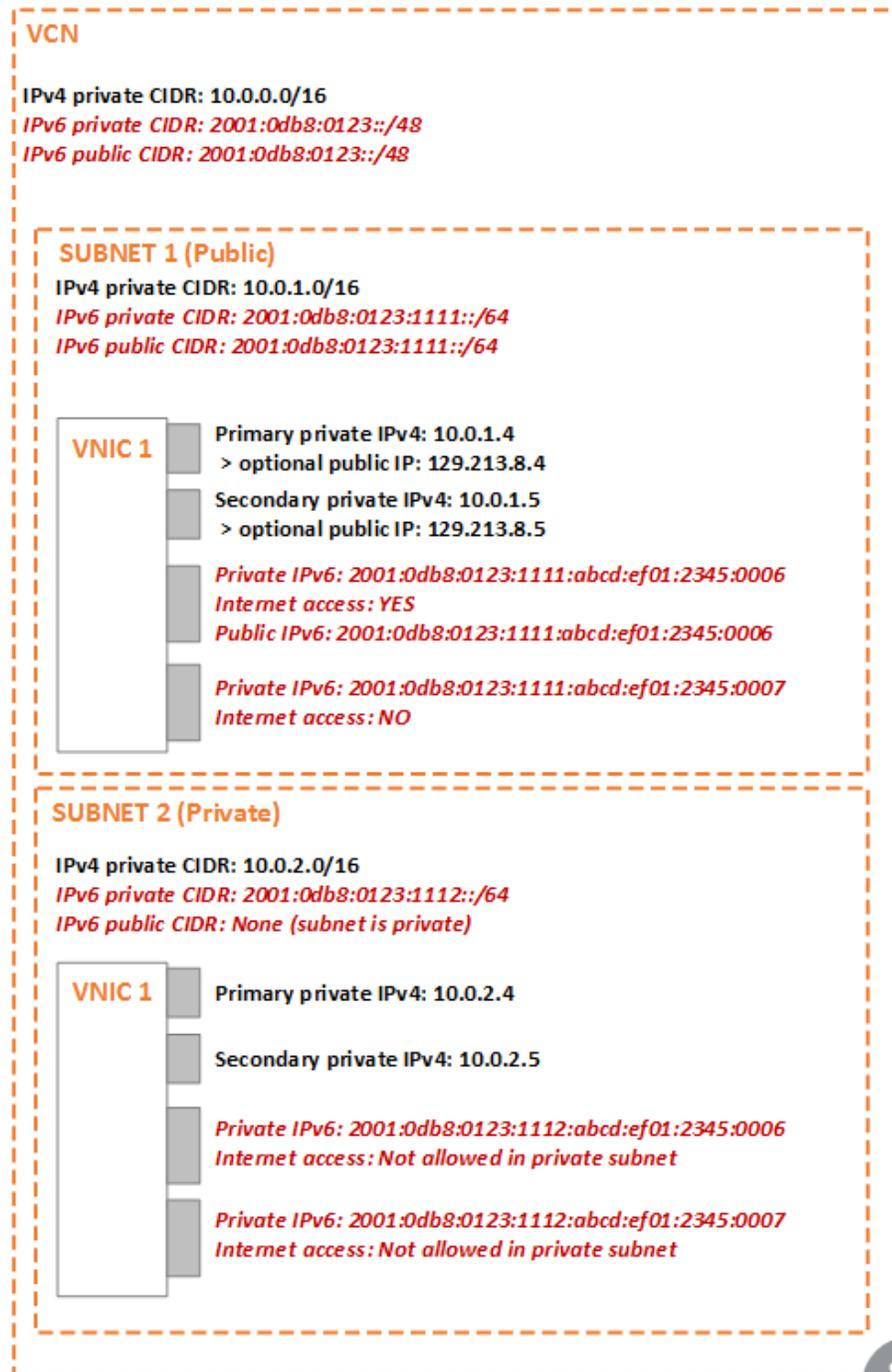
The second subnet is private, which means the VNICs can never have public IPv4 or IPv6 addresses. In this case, there's one VNIC that has a primary and secondary IPv4 with addresses 10.0.2.4 and 10.0.2.5, respectively.

The VNIC also has two IPv6s. The first has private address fd00:aaaa:0123:1112:abcd:ef01:2345:0006, and the second IPv6 has private address fd00:aaaa:0123:1112:abcd:ef01:2345:0007.

### Example 2: You let Oracle assign the VCN's CIDR

You do not assign a custom CIDR, and Oracle assigns this CIDR: *2001:0db8:0123::/48*. Oracle uses this same CIDR for both the private and public IP addresses.

The following diagram illustrates the VCN and includes two subnets: public subnet 1111 and private subnet 1112.



The VNIC in the public subnet has a primary private IPv4 (10.0.1.4) with an optional public IP address assigned. The VNIC has a secondary private IPv4 (10.0.1.5), also with an optional public IP address assigned.

The VNIC also has two IPv6s. The first one has internet access enabled and therefore has both private and public IPv6 addresses, which are the following:

- Private IPv6 address: 2001:0db8:0123:1111:abcd:ef01:2345:0006
- Public IPv6 address: 2001:0db8:0123:1111:abcd:ef01:2345:0006

**Notice that the two addresses are the same.**

The second IPv6 does not have internet access enabled and therefore has only a private IP address, which is 2001:0db8:0123:1111:abcd:ef01:2345:0007.

The second subnet is private, which means the VNICs can never have public addresses, IPv4 or IPv6. In this case, there's one VNIC that has a primary and secondary IPv4 with addresses 10.0.2.4 and 10.0.2.5, respectively.

The VNIC also has two IPv6s. The first has private address 2001:0db8:0123:1112:abcd:ef01:2345:0006, and the second IPv6 has private address 2001:0db8:0123:1112:abcd:ef01:2345:0007.

### Routing for IPv6 Traffic

Both inbound- and outbound-initiated IPv6 connections are supported between your VCN and the internet, and between your VCN and your on-premises network.

Here are other important details about routing of IPv6 traffic:

- Currently IPv6 traffic is supported only through these gateways:
  - [Dynamic routing gateway \(DRG\)](#): For access to your on-premises network or other networks outside the region. Both Oracle Cloud Infrastructure [FastConnect](#) and [IPSec VPN](#) support IPv6 traffic. For more details about IPv6 configuration,

see the upcoming sections.

- [Internet gateway](#): For access to the internet.
- Traffic between instances on their public IPv6 addresses is supported and must traverse the VCN's internet gateway. Exception: if the given IPv6 uses the [same address for both private and public communication](#), traffic between instances on their public IPv6 address is not supported. **Therefore, if you want instances in the same VCN to communicate with each other using public IPv6 addresses, specify your own private IPv6 CIDR when creating the VCN.** This means the private address for an IPv6 in the VCN will be different than its public address. For more information, see [CIDRs Assigned to an IPv6-Enabled VCN](#).
- Private IPv6 traffic between resources within a region (intra- and inter-VCN) is not yet supported.

### VCN Route Tables and IPv6

The VCN's [route tables](#) support both IPv4 rules and IPv6 rules that use a DRG or internet gateway as the target. For example, the route table for a given subnet could have these rules:

- Rule to route traffic that matches a certain IPv4 CIDR to the VCN's attached DRG
- Rule to route traffic that matches a certain IPv4 CIDR to the VCN's service gateway
- Rule to route traffic that matches a certain IPv4 CIDR to the VCN's NAT gateway
- Rule to route traffic that matches a certain **IPv6** CIDR to the VCN's attached DRG
- Rule to route traffic that matches a certain **IPv6** CIDR to the VCN's attached internet gateway

### Security Rules for IPv6 Traffic

Like route tables, the VCN's [network security groups](#) and [security lists](#) support both IPv4 and IPv6 [rules](#). For example, a network security group or security list could have these security rules:

- Rule to allow SSH traffic from the on-premises network's IPv4 CIDR
- Rule to allow ping traffic from the on-premises network's IPv4 CIDR
- Rule to allow SSH traffic from the on-premises network's **IPv6** CIDR
- Rule to allow ping traffic from the on-premises network's **IPv6** CIDR

The [default security list](#) in an IPv6-enabled VCN includes default IPv4 rules and the following default IPv6 rules:

- **Stateful ingress:** Allow IPv6 TCP traffic on destination port 22 (SSH) from source `::/0` and any source port. This rule makes it easy for you to create a new VCN with a public subnet and internet gateway, create a Linux instance, add an internet-access-enabled IPv6, and then immediately connect with SSH to that instance without needing to write any security rules yourself.



### Important

The default security list does not include a rule to allow Remote Desktop Protocol (RDP) access. If you're using [Windows images](#), make sure to add a stateful ingress rule for TCP traffic on destination port 3389 from source `::/0` and any source port. See [To enable RDP access](#) for more information.

- **Stateful ingress:** Allow ICMPv6 traffic type 2 code 0 (Packet Too Big) from source `::/0` and any source port. This rule enables your instances to receive Path MTU Discovery fragmentation messages.
- **Stateful egress:** Allow all IPv6 traffic. This allows instances to initiate IPv6 traffic of any kind to any destination. Notice that this means the instances with an internet-access-enabled IPv6 can talk to any internet IPv6 address if the VCN has a configured internet gateway. And because stateful security rules use connection tracking, the response traffic is automatically allowed regardless of any ingress rules. For more information, see [Connection Tracking Details for Stateful Rules](#).

### FastConnect and IPv6

If you use FastConnect, you can configure it so that on-premises hosts with IPv6 addresses can communicate with an IPv6-enabled VCN. In general, you must ensure that the FastConnect virtual circuit has IPv6 BGP addresses, and update the VCN's routing and security rules for IPv6 traffic.

#### About the IPv6 BGP Addresses

A FastConnect virtual circuit always requires IPv4 BGP addresses, but IPv6 BGP addresses are optional and only required for IPv6 traffic. Depending on how you're using FastConnect, you might be asked to provide all of the virtual circuit's BGP addresses yourself (both IPv4 and IPv6).

The addresses consist of a pair: one for your end of the BGP session, and another for the Oracle end of the BGP session.

When you specify a BGP address pair, you must include a subnet mask that contains both of the addresses. Specifically for IPv6, the allowed subnet masks are:

- /64
- /96
- /126
- /127

For example, you could specify 2001:db8::6/127 for the address at your end of the BGP session, and 2001:db8::7/127 for the Oracle end.

#### Process to Enable IPv6

In general, here's how to enable IPv6 for a FastConnect virtual circuit:

- **Virtual circuit BGP:** Ensure the FastConnect virtual circuit has IPv6 BGP addresses. If you're responsible for providing the BGP IP addresses, when you set up a new virtual circuit or edit an existing one, there's a place for the two IPv4 BGP addresses. There's a

separate check box for **Enable IPv6 Address Assignment** and a place to provide the two IPv6 addresses. Be aware that if you're editing an existing virtual circuit to add support for IPv6, it will go down while it's being reprovisioned to use the new BGP information.

- **VCN route tables:** For each IPv6-enabled subnet in the VCN, update its [route table](#) to include rules that route the IPv6 traffic from the VCN to the desired IPv6 subnets in your on-premises network. For example, the **Destination CIDR Block** for a route rule would be an IPv6 subnet in your on-premises network, and the **Target** would be the dynamic routing gateway (DRG) attached to the IPv6-enabled VCN.
- **VCN security rules:** For each IPv6-enabled subnet in the VCN, update its security lists or relevant network security groups to allow the desired IPv6 traffic between the VCN and your on-premises network. See [Security Rules for IPv6 Traffic](#).

If you do not yet have a FastConnect connection, see these topics to get started:

- [FastConnect Overview](#)
- [FastConnect Requirements](#)

### VPN Connect and IPv6

If you use [VPN Connect](#), you can configure it so that on-premises hosts with IPv6 addresses can communicate with an IPv6-enabled VCN. Here's how to enable IPv6 for the connection:

- **IPSec connection static routes:** Configure the IPSec connection with the IPv6 static routes of your on-premises network. Currently the Oracle IPSec VPN does not support BGP dynamic routing.
- **VCN route tables:** For each IPv6-enabled subnet in the VCN, update its [route table](#) to include rules that route the IPv6 traffic from the VCN to the desired IPv6 subnets in your on-premises network. For example, the **Destination CIDR Block** for a route rule would be an IPv6 static route for your on-premises network, and the **Target** would be the dynamic routing gateway (DRG) attached to the IPv6-enabled VCN.

- **VCN security rules:** For each IPv6-enabled subnet in the VCN, update its security lists or relevant network security groups to allow the desired IPv6 traffic between the VCN and your on-premises network. See [Security Rules for IPv6 Traffic](#).

If you have an existing IPsec VPN that uses static routing, you can update the list of static routes to include ones for IPv6. Be aware that changing the list of static routes causes the IPsec VPN to go down while it's being reprovisioned. See [Changing the Static Routes](#).

If you do not yet have an IPsec VPN, see these topics to get started:

- [VPN Connect Overview](#)
- [Setting Up VPN Connect](#)
- [Working with VPN Connect](#)

### DHCP

Currently DHCPv6 auto-configuration of IP addresses is not supported.

### DNS

The VCN's [Internet Resolver](#) supports IPv6, which means resources in your VCN can resolve IPv6 addresses of hosts outside the VCN. IPv6 traffic between resources within the VCN is not yet supported, and assignment of a hostname to an IPv6 address is not supported.

### Load Balancers

When you create a [load balancer](#), you can optionally choose to have an IPv4/IPv6 dual-stack configuration. When you choose the IPv6 option, the Load Balancing service assigns both an IPv4 and an IPv6 address to the load balancer. The load balancer receives client traffic sent to the assigned IPv6 address. The load balancer uses only IPv4 addresses to communicate with backend servers. There is no IPv6 communication between the load balancer and the backend servers.

IPv6 address assignment occurs only at load balancer creation. You cannot assign an IPv6 address to an existing load balancer.

## Comparison of IPv4 and IPv6 for Your VCN

The following table summarizes the differences between IPv4 and IPv6 addressing in a VCN.

Characteristic	IPv4	IPv6
Addressing type supported	IPv4 addressing is always required, regardless of whether IPv6 is enabled.	IPv6 addressing is optional per VCN, optional per subnet in an IPv6-enabled VCN, and optional per VNIC in an IPv6-enabled subnet.
Supported traffic types	IPv4 traffic is supported for all gateways. IPv4 traffic between instances within the VCN is supported (east/west traffic).	IPv6 traffic is supported only with these gateways: internet gateway and DRG. Both inbound- and outbound-initiated IPv6 connections are supported between your VCN and the internet, and between your VCN and your on-premises network. Private IPv6 traffic between resources within a region (intra- and inter-VCN) is not yet supported (east/west traffic). Also see the caveats in <a href="#">Routing for IPv6 Traffic</a> .
VCN size	/16 to /30	/48 only
Subnet size	/16 to /30, with 3 addresses reserved in each subnet by Oracle (first 2 and last 1).	/64 only, with 8 addresses in the subnet reserved by Oracle (first 4 and last 4).

Characteristic	IPv4	IPv6
<p>Private and public IP address space</p>	<p>Private: A VCN's private IPv4 CIDR can be from an RFC 1918 range or a publicly routable range (in which case, it's treated as private). You must specify the range, unless you use the Console's VCN creation wizard, which always uses 10.0.0.0/16.</p> <p>Public: The VCN does not have a dedicated public IPv4 address space. Any public addresses in your VCN are always chosen by Oracle.</p>	<p>You can specify a /48 from the list of supported ranges for the private IPv6 CIDR (see <a href="#">CIDRs Assigned to an IPv6-Enabled VCN</a>). If you don't specify a range, Oracle assigns a /48 CIDR that is used for both the private and public IP address space.</p> <p><b>Important:</b> You must assign this value if you want instances in the same VCN to communicate with each other using public IPv6 addresses. For more information, see <a href="#">Routing for IPv6 Traffic</a>.</p> <p>Unlike with IPv4, your VCN has a <i>dedicated</i> public IPv6 address space, which is always /48 in size. When you assign an IPv6 to a VNIC, you can choose the address, or you can let Oracle chose it.</p>

Characteristic	IPv4	IPv6
IP address assignment	<p>Private: Each VNIC gets a private IPv4 address. You can choose the address or let Oracle choose it.</p> <p>Public: You determine whether the private IPv4 address has a public IP address associated with it (assuming the VNIC is in a public subnet). Oracle chooses the public IP address.</p> <p>From an API standpoint: the <code>PrivateIp</code> object is separate from the <code>PublicIp</code> object. You can remove the public IP address from the private IPv4 address at any time.</p>	<p>You decide whether a VNIC in an IPv6-enabled subnet gets an IPv6. You can choose the private IPv6 address or let Oracle choose it.</p> <p>You also decide whether that IPv6 has internet access enabled (assuming the VNIC is in a public subnet). You can remove the internet access for that IPv6 at any time. When an IPv6 is internet enabled, it has a public IPv6 address. The public IPv6 address always has the same right-most 64 bits as the private IPv6 address.</p> <p>Recall that if Oracle assigns the VCN's private IPv6 CIDR, then the private and public CIDRs for the VCN are the same. In that case, each IPv6 uses the same address (all 128 bits) for both its private IP address and public IP address.</p> <p>From an API standpoint: both the private and public IP addresses are included in the <code>Ipv6</code> object and always exist together.</p>
Internet access	You control whether a subnet is public or private. You add or remove a public IP address from a private IPv4 address on a VNIC (assuming the VNIC is in a public subnet).	You control whether a subnet is public or private. You do not add or remove a public IP address to or from the VNIC as you do with IPv4. Instead you enable or disable the internet access for a given IPv6 that you've added to a VNIC (assuming the VNIC is in a public subnet).

Characteristic	IPv4	IPv6
Primary and secondary labels	Each VNIC automatically has a primary private IP address, and you can assign up to 31 secondary private IPs per VNIC.	You choose to add an IPv6 to a VNIC. There is no <i>primary</i> or <i>secondary</i> label for it. You can assign up to 32 IPv6s per VNIC.
Hostnames	You can assign hostnames to IPv4 addresses.	You cannot assign hostnames to IPv6 addresses.
Route rule limits	See <a href="#">Service Limits</a> .	IPv4 and IPv6 route rules can reside together in the same route table. IPv6 route rules can target only an internet gateway or DRG. Limit on number of IPv6 route rules in a route table: 8.
Security rule limits	See <a href="#">Service Limits</a> .	IPv4 and IPv6 security rules can reside together in same network security group or security list. IPv6 security rules can use only IPv6 CIDR ranges for source or destination, and not a service CIDR label used for a service gateway. Limit on number of IPv6 security rules in a security list: 8 ingress and 8 egress. Limit on number of IPv6 security rules in a network security group: 16 total.

Characteristic	IPv4	IPv6
Reserved public IP addresses	Supported.	Not supported.
Regional or AD-specific	Primary private IPv4 addresses are AD-specific. Secondary private IPv4 addresses are AD-specific unless assigned to a VNIC in a regional subnet. Public IP addresses can be AD-specific or regional depending on the type (ephemeral or reserved). See <a href="#">Public IP Addresses</a> .	IPv6 addresses are regional.

### Setting Up an IPv6-Enabled VCN with Internet Access

Use the following process if you want to set up an IPv6-enabled VCN with internet access so you can easily launch an instance and connect to it by using its public IPv6 address.



**Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Task 1: Create the IPv6-enabled VCN

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Virtual Cloud Network**.
4. Enter the following:
  - **Name:** A descriptive name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** Leave as is.
  - **Create Virtual Cloud Network Only:** Make sure this radio button is selected (the default).
  - **CIDR Block:** A single, contiguous IPv4 CIDR block for the VCN. For example: 172.16.0.0/16. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** Select the check box and optionally provide the private IPv6 CIDR in the field labeled **Private IPv6 CIDR Block**. You must provide the value **if you want the instances in this IPv6-enabled VCN to communicate with each other using their public IP addresses**. Leave the field blank if you want Oracle to assign the private IPv6 CIDR for you. You *cannot* later disable IPv6 for the VCN or change the CIDR. All IPv6-enabled VCNs are always /48 in size.
  - **Use DNS Hostnames in this VCN** (supported for IPv4 only): Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the

check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (<*VCN DNS label*>.oraclevcn.com). For more information, see [DNS in Your Virtual Cloud Network](#).

- **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
5. Click **Create Virtual Cloud Network**.

The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose.

### Task 2: Create a regional IPv6-enabled public subnet

1. While still viewing the VCN, click **Create Subnet**.
2. Enter the following:
  - **Name:** A descriptive name for the subnet (for example, Regional Public Subnet). It doesn't have to be unique, and you can change it later. Avoid entering confidential information.
  - **Regional or Availability Domain-specific subnet:** Oracle recommends creating only [regional subnets](#), which means that the subnet can contain resources in any of the region's availability domains. If you instead choose **Availability Domain-Specific** (the only kind of subnet that Oracle originally offered), you must also specify an availability domain. This choice means that any instances or other resources later created in this subnet must also be in that availability domain.
  - **CIDR Block:** A single, contiguous IPv4 CIDR block for the subnet (for example, 172.16.0.0/24). Make sure it's within the VCN's IPv4 CIDR block and doesn't overlap with any other subnets. You *cannot* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, here's a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** Select the check box and provide your choice of 16 bits for the subnet (example: 1111). You *cannot* later disable IPv6 for

the subnet or change the CIDR. All IPv6-enabled subnets are always /64 in size. For more information about IPv6 address format, see [Format of IPv6 Addresses](#).

- **Route Table:** Select the default route table.
  - **Private or public subnet:** Select **Public Subnet**, which means instances in the subnet can optionally have public IP addresses. For more information, see [Access to the Internet](#).
  - **Use DNS Hostnames in this Subnet** (supported for IPv4 only): This option is available only if you provided a DNS label for the VCN during creation. The option is required for assignment of DNS hostnames to hosts in the subnet, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the subnet, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the subnet (`<subnet_DNS_label>.<VCN_DNS_label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).
  - **DHCP Options:** Select the default set of DHCP options.
  - **Security Lists:** Make sure the default security list is selected (the default).
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
3. Click **Create Subnet**.  
The subnet is then created and displayed on the **Subnets** page.

### Task 3: Create the internet gateway

1. Under **Resources**, click **Internet Gateways**.
2. Click **Create Internet Gateway**.
3. Enter the following:
  - **Name:** A descriptive name for the internet gateway. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API).

Avoid entering confidential information.

- **Create in Compartment:** Leave as is.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
4. Click **Create Internet Gateway**.

Your internet gateway is created and displayed on the **Internet Gateways** page. It's already enabled, but you must add route rules that allow IPv4 and IPv6 traffic to flow to the gateway.

### Task 4: Update the default route table to use the internet gateway

The default route table starts out with no rules. Here you add rules that route all IPv4 and IPv6 traffic destined for addresses outside the VCN to the internet gateway. The existence of these rules also enables inbound connections to come from the internet to the subnet, through the internet gateway. You use security rules to control the *types of traffic* that are allowed in and out of the instances in the subnet (see the next task).

No route rule is required in order to route traffic within the VCN itself.

1. Under **Resources**, click **Route Tables**.
2. Click the default route table to view its details.
3. Click **Add Route Rules**.
4. Enter the following:
  - **Target Type:** Internet Gateway
  - **Destination CIDR block:** 0.0.0.0/0 (which means that all IPv4 non-intra-VCN traffic that is not already covered by other rules in the route table goes to the target specified in this rule).
  - **Compartment:** The compartment where the internet gateway is located.
  - **Target:** The internet gateway you created.
5. Click **+ Additional Route Rule**.

6. Enter the following:
  - **Target Type:** Internet Gateway
  - **Destination CIDR block:** `::/0` (for the IPv6 traffic).
  - **Compartment:** The compartment where the internet gateway is located.
  - **Target:** The internet gateway you created.
7. Click **Add Route Rules**.

The default route table now has two rules for the internet gateway, one for IPv4 traffic and one for IPv6 traffic. Because the subnet was set up to use the default route table, the resources in the subnet can now use the internet gateway. The next step is to specify the types of traffic you want to allow in and out of the instances you later create in the subnet.

### Task 5: Update the default security list (optional)



#### Note

This task is about configuring [security rules](#) to allow the desired traffic to and from your instances. Although this task uses a [security list](#) to implement those rules, you can also use [network security groups](#) to implement security rules.

Earlier you set up the subnet to use the VCN's [default security list](#). This list already includes basic rules that allow essential IPv4 and IPv6 traffic. In this task, you add any *additional* security rules that allow the types of connections that the instances in the VCN will need.

For example: This is a public subnet with an internet gateway, so the instances you create might need to receive inbound HTTPS connections from the internet (if they're web servers). Here's how to add another rule to the default security list to enable that traffic:

1. Under **Resources**, click **Security Lists**.
2. Click the default security list to view its details. By default, you land on the **Ingress Rules** page.
3. Click **Add Ingress Rule**.
4. To enable inbound connections for HTTPS (TCP port 443), enter the following:
  - **Stateless**: Unselected (this is a [stateful rule](#))
  - **Source Type**: CIDR
  - **Source CIDR**: 0.0.0.0/0 (or `::/0` if you want to enable IPv6 traffic with this rule)
  - **IP Protocol**: TCP
  - **Source Port Range**: All
  - **Destination Port Range**: 443
5. Click **Add Ingress Rule**.



### Important

#### *Security List Rule for Windows Instances*

If you're going to create Windows instances, you need to add a security rule to enable Remote Desktop Protocol (RDP) access. Specifically, you need a stateful ingress rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 (and a separate rule with `::/0` for IPv6 traffic) and any source port. For more information, see [Security Rules](#).

For a production VCN, you typically set up one or more *custom* security lists for each subnet. If you like, you can edit the subnet to [use different security lists](#). If you choose not to use the default security list, do so only after carefully assessing which of its default rules you want to duplicate in your custom security list. For example: the [default ICMP rules in the default security list](#) are important for receiving connectivity messages for IPv4.

### Task 6: Create an instance

Your next step is to create an instance in the subnet. When you create the instance, you choose the availability domain, which VCN and subnet to use, and several other characteristics.

Each instance automatically gets a private IPv4 address. When you create an instance in a *public subnet*, you choose whether the instance gets a public IP address. A public IPv4 address is NOT required for public IPv6 traffic. But if you want to connect to the instance from an IPv4 host, you *must* give the instance a public IP address, or else you can't access them through the internet gateway. The default (for a public subnet) is for the instance to get a public IP address.

For more information and instructions, see [Launching an Instance](#).

### Task 7: Add an internet-enabled IPv6 to the instance

1. While viewing the instance you just created, click **Attached VNICs**.
2. Click the VNIC.
3. Under **Resources**, click **IPv6 Addresses**.
4. Click **Assign Private IPv6 Address**.
5. Enter the following:
  - **Private IPv6 Address:** Optional. An available private IPv6 address of your choice from the subnet's private IPv6 CIDR (otherwise the private IP address is automatically assigned).
  - **Unassign if already assigned to another VNIC:** Leave this check box as is (cleared). Use this only to force reassignment of an IPv6 address if it's already assigned to another VNIC in the subnet. Relevant only if you specify a private IP address in the preceding field.
  - **Enable Internet Access:** Select this check box. This enables internet access and assigns the IPv6 a public address (if it's in a public subnet, which is the case here). See [Internet Communication](#).

- **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
6. Click **Assign**.
- The IPv6 is created and then displayed on the **IPv6 Addresses** page for the VNIC. Notice that it has both a private and public IPv6 address.

### Task 8: Configure the instance's OS to use the IPv6

You must configure the instance's OS to use the IPv6. For more information, see [Configuring the OS to Use an IPv6](#).

After performing this task, you can connect to the instance over the internet with SSH or RDP from your on-premises network or a location on the internet. The host connecting to the instance must be using a public IPv6 address.

## Managing IPv6s in the Console

This section includes basic tasks for working with IPv6-related resources.

### To create an IPv6-enabled VCN



#### Important

You can't enable IPv6 for an existing VCN. You can only enable IPv6 when creating a VCN. After enabling IPv6 for a VCN, you cannot disable it.

**Summary:** When creating a VCN, you choose between these two options:

- Creating *only* a VCN
- Creating a VCN plus related resources

The following table summarizes the difference between the two options. Regardless of which option you select, make sure to select the check box for **Enable IPv6 Address Assignment** when you create the VCN. Further instructions follow the table.

Choice for creating the VCN	What Is Created	Notes
VCN only	<p>VCN only</p> <p>You must manually create and configure other components required so you can create instances in the VCN and connect to them. For instructions, see <a href="#">Setting Up an IPv6-Enabled VCN with Internet Access</a>.</p>	<p><b>Recommended choice if you want the instances in the same IPv6-enabled VCN to communicate with each other using their public IP addresses.</b></p> <p>This type of communication is supported only if the private address for an IPv6 in the VCN is different than its public address. That means you must specify the value for the VCN's private IPv6 CIDR and NOT let Oracle choose it (Oracle uses the same IPV6 CIDR for both private and public address space).</p> <p>For more information, see <a href="#">CIDRs Assigned to an IPv6-Enabled VCN</a> and also <a href="#">Routing for IPv6 Traffic</a>.</p>
VCN plus related resources	<p>VCN</p> <p>One public subnet per availability domain</p> <p>Internet gateway</p> <p>Default route table with IPv4 rule and IPv6 rule for internet gateway traffic</p> <p>Default security list with rules for essential IPv4 traffic and IPv6 traffic</p>	<p>Oracle automatically chooses the IPv6 CIDR for the VCN and uses that same CIDR for the private and public IPv6 address space.</p> <p>For more information, see <a href="#">CIDRs Assigned to an IPv6-Enabled VCN</a>.</p>

### Instructions to create an IPv6-enabled VCN only

See the instructions in [Task 1: Create the IPv6-enabled VCN](#).

### Instructions to create an IPv6-enabled VCN plus related resources

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Virtual Cloud Network**.
4. Enter the following:
  - **Name:** A descriptive name for the VCN. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** Leave as is.
  - **Create Virtual Cloud Network Plus Related Resources:** Make sure this radio button is selected.
  - **Enable IPv6 Address Assignment:** Select the check box. Oracle chooses the private IPv6 CIDR and uses the same value for the public IPv6 CIDR. For more information, see [CIDRs Assigned to an IPv6-Enabled VCN](#).
  - **Use DNS Hostnames in this VCN** (supported for IPv4 only): Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (`<VCN DNS label>.oraclevcn.com`). For more information, see [DNS in Your Virtual Cloud Network](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create Virtual Cloud Network**.

The VCN is then created and displayed on the **Virtual Cloud Networks** page in the compartment you chose. The VCN has one public subnet per availability domain, an internet gateway, routing set up to use the internet gateway, and basic security rules in the default security list.

### To create an IPv6-enabled subnet



#### **Important**

After enabling IPv6 for a subnet, you cannot disable it.

**Summary:** Creating an IPv6-enabled subnet is similar to creating an IPv4 subnet. The only difference is that you must select the check box for **Enable IPv6 Address Assignment** and provide 16 bits for the subnet's portion of the IPv6 CIDR. See [Format of IPv6 Addresses](#).

For general instructions, see [Task 2: Create a regional IPv6-enabled public subnet](#). If you want a private subnet, make sure to select the radio button for **Private Subnet** when creating the subnet.

### To assign an IPv6 to a VNIC

**Summary:** The process for adding an IPv6 to a VNIC is similar to adding a [secondary private IPv4 address](#). You can specify the particular IPv6 address to use or let Oracle choose it from the subnet. You can enable internet access if you like. The resulting public IPv6 address uses

the same right-most 64 bits as the private IPv6 address. In certain situations, the entire IPv6 address is the same. For more information, see [Format of IPv6 Addresses](#). After assigning the IPv6 to the VNIC, you must [configure the OS to use the IPv6](#).

### Instructions

1. Assign the IPv6. For general instructions, see [Task 7: Add an internet-enabled IPv6 to the instance](#). If you want an IPv6 without internet access, do not select the check box for **Enable Internet Access** when assigning the IPv6.
2. Configure the OS to use the IPv6 address. For more information, see [Configuring the OS to Use an IPv6](#).

### To move an IPv6 to another VNIC in the subnet

**Summary:** The process is similar to [moving a secondary private IPv4 address](#) from one VNIC to another (let's call them the *original VNIC* and the *new VNIC*). You assign the IPv6 to the new VNIC, specify the private IPv6 address, and select the check box for **Unassign if already assigned to another VNIC**. Oracle automatically unassigns it from original VNIC and assigns it to the new VNIC. The public address for the IPv6 stays the same regardless of which VNIC the IPv6 is assigned to.

### Instructions

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.

6. Under **Resources**, click **IPv6 Addresses**.
7. Click **Assign Private IP Address**.
8. Enter the following:
  - **Private IPv6 Address:** The private IPv6 address that you want to move.
  - **Unassign if already assigned to another VNIC:** Select this check box to move the IPv6 address from the VNIC it's currently assigned to.
  - **Enable Internet Access:** Whether to assign a public IPv6 address. Available only if the VNIC is in a public subnet. See [Internet Communication](#).
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
9. Click **Assign**.

The private IP address is moved from the original VNIC to the new VNIC.

### To delete an IPv6 from a VNIC

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IPv6 Addresses**.

7. For the IPv6 you want to delete, click the Actions icon (three dots), and then click **Delete IPv6**.
8. Confirm when prompted.

The IPv6 address is returned to the pool of available addresses in the subnet.

### To enable or disable internet access for an IPv6

**Summary:** Internet access for an IPv6 is controlled by the IPv6's **Enable Internet Access** check box. When you enable internet access, the IPv6 is assigned a public IPv6 address. That address's right-most 64 bits are the same as the private IPv6 address. In certain situations, the entire IPv6 address is the same. For more information, see [Format of IPv6 Addresses](#).

### Instructions

1. Confirm you're viewing the compartment that contains the instance you're interested in.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
3. Click the instance to view its details.
4. Under **Resources**, click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.
6. Under **Resources**, click **IPv6 Addresses**.
7. For the IPv6 you're interested in, click the Actions icon (three dots), and then click **Edit**.
8. Either select or clear the check box for **Enable Internet Access**.
9. Click **Update**.

When you disable internet access, the public IPv6 address becomes null. If you re-enable internet access, the public IPv6 address is again assigned to the IPv6.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

For IPv6 addressing, there's an [Ipv6](#) object with the following operations:

- [ListIpv6s](#)
- [GetIpv6](#)
- [UpdateIpv6](#)
- [CreateIpv6](#)
- [DeleteIpv6](#)

### Configuring the OS to Use an IPv6

After assigning an IPv6 to a VNIC, you must configure the OS to use the IPv6.

#### Getting the IPv6 Virtual Router IP (Default Gateway)

You need the *IPv6 virtual router IP* (called the *default gateway* in Windows), which is included in the [instance metadata](#) available at the following URL:

```
http://169.254.169.254/opc/v1/vnics/
```

Here's an example response:

```
[{
 "vnicId" : "ocid1.vnic.oc1.phx.examplelvq7kncmdtfr23dznohdkd2cywjcem33eg3dxa",
 "privateIp" : "10.0.3.7",
 "vlanTag" : 3396,
 "macAddr" : "00:00:17:01:14:0C",
 "virtualRouterIp" : "10.0.3.1",
 "subnetCidrBlock" : "10.0.3.0/24",
 "ipv6SubnetCidrBlock" : "2001:0db8:95f4::/64",
 "ipv6VirtualRouterIp" : "2001:0db8::200:17ff:fee3:c491"
}]
```

### Oracle Linux 7 Configuration

The following commands are for Oracle Linux 7. They are NOT persistent through a reboot. You need the IPv6 virtual router IP from the instance metadata (see the previous section).

```
sysctl net.ipv6.conf.all.disable_ipv6=0
ip -6 addr add <private_IPv6_address>/64 dev <interface_name>
ip -6 route add default via <IPv6_virtual_router_IP> dev <interface_name>
```

For example:

```
sysctl net.ipv6.conf.all.disable_ipv6=0
ip -6 addr add 2001:0db8:95f4::abcd:1234/64 dev ens3
ip -6 route add default via 2001:0db8::200:17ff:fee3:c491 dev ens3
```

If you haven't yet, make sure the VCN's route table and security rules are configured for the desired IPv6 traffic. See [Routing for IPv6 Traffic](#) and [Security Rules for IPv6 Traffic](#).

### Windows Configuration

You can use a command line or the Network Connections UI.

### Command Line

If you use PowerShell, you must run it as an administrator. The following configuration persists through a reboot of the instance.

1. In your browser, go to the Console, and note the private IPv6 address that you want to configure on the instance.
2. Connect to the instance, and run the following command at a command prompt:

```
http://169.254.169.254/opc/v1/vnics/
```

3. Note the value for the `ipv6VirtualRouterIp`, which is the *<default\_gateway>* to use in the next step.
4. Run the following 2 commands:

## CHAPTER 23 Networking

---

```
netsh interface ipv6 add address interface="Ethernet" address=<private_IPv6_address>/64
netsh interface ipv6 add route prefix>::/0 interface="Ethernet" nexthop=<default_gateway>
publish=Yes
```

For example:

```
netsh interface ipv6 add address interface="Ethernet" address=2001:0db8:95f4::abcd:1234/64
netsh interface ipv6 add route prefix>::/0 interface="Ethernet"
nexthop=2001:0db8::200:17ff:fee3:c491 publish=Yes
```

If you haven't yet, make sure the VCN's route table and security rules are configured for the desired IPv6 traffic. See [Routing for IPv6 Traffic](#) and [Security Rules for IPv6 Traffic](#).

You can run the following command to see that the IPv6 address has been configured for the interface:

```
netsh interface ipv6 show addresses
```

Later if you want to delete the address, you can use this command:

```
netsh interface ipv6 delete address interface="Ethernet" address=<private_IPv6_address>
```

For example:

```
netsh interface ipv6 delete address interface="Ethernet" address=2001:0db8:95f4::abcd:1234
```

Also make sure to [delete the IPv6 from the VNIC](#). You can do that before or after executing the earlier command to delete the address from the OS configuration.

## Network Connections UI

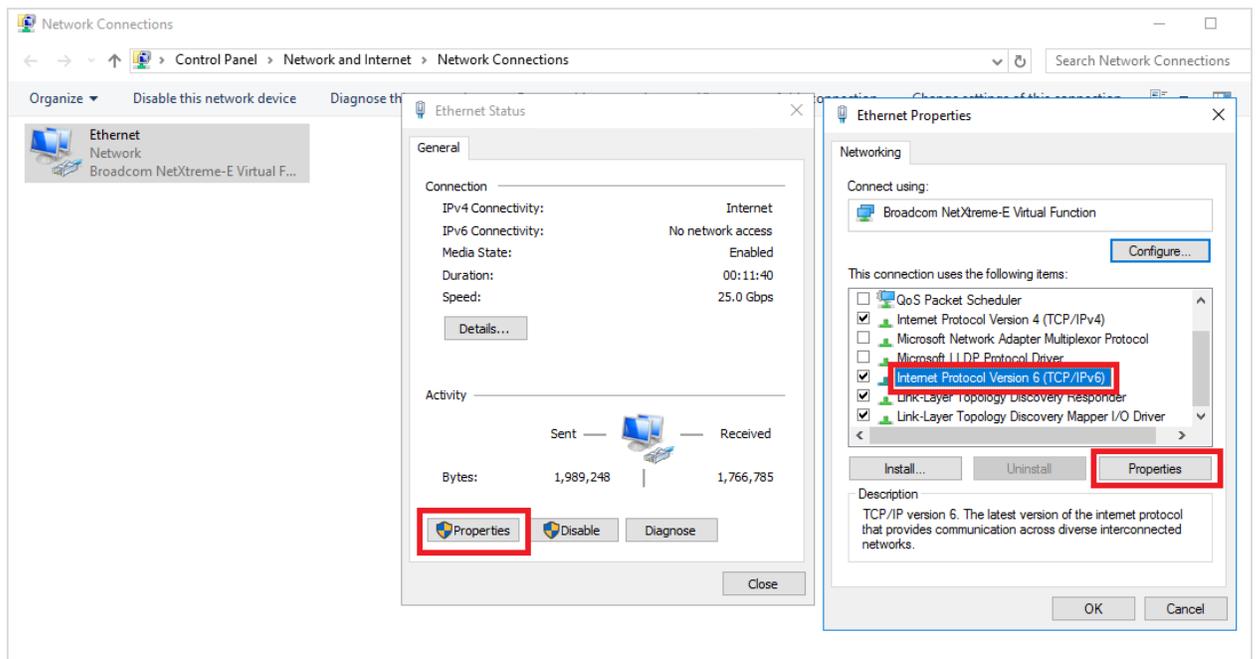
The following configuration persists through a reboot of the instance.

1. In your browser, go to the Console, and note the private IPv6 address that you want to configure on the instance.
2. Connect to the instance, and run the following command at a command prompt:

```
http://169.254.169.254/opc/v1/vnics/
```

## CHAPTER 23 Networking

3. Note the value for the `ipv6VirtualRouterIp`, which is the default gateway to use in a later step.
4. In the instance's **Control Panel**, go to **Network and Internet**, and view your network connections (see the image that follows for the set of dialog boxes you see in these steps).
5. For the active networks, click the connection (**Ethernet**).
6. Click **Properties**.
7. Click **Internet Protocol Version 6 (TCP/IPv6)**, and then click **Properties**.



8. Select the radio button for **Use the following IP address**, and then enter the values you noted earlier for the private IPv6 address and default gateway. Use 64 for the subnet prefix length.

The image shows a 'General' tab in a network configuration dialog box. At the top, there is a text box stating: 'You can get IPv6 settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IPv6 settings.' Below this, there are two radio button options: 'Obtain an IPv6 address automatically' (unselected) and 'Use the following IPv6 address:' (selected). The 'Use the following IPv6 address:' option is expanded to show three input fields: 'IPv6 address:' with the value '2001:0db8:95f4:1::abcd:1234', 'Subnet prefix length:' with the value '64', and 'Default gateway:' with the value '2001:0db8::200:17ff:fee3:c491'. Below these fields, there are two more radio button options: 'Obtain DNS server address automatically' (unselected) and 'Use the following DNS server addresses:' (selected). The 'Use the following DNS server addresses:' option is expanded to show two empty input fields: 'Preferred DNS server:' and 'Alternate DNS server:'. At the bottom left, there is a checkbox labeled 'Validate settings upon exit' which is unchecked. At the bottom right, there is an 'Advanced...' button. At the very bottom of the dialog box, there are 'OK' and 'Cancel' buttons.

9. Click **OK** until the dialog boxes are closed.

If you haven't yet, make sure the VCN's route table and security rules are configured for the desired IPv6 traffic. See [Routing for IPv6 Traffic](#) and [Security Rules for IPv6 Traffic](#).

## DNS in Your Virtual Cloud Network

The Domain Name System (DNS) lets computers use hostnames instead of IP addresses to communicate with each other.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Choices for DNS in Your VCN

Following are the choices for DNS name resolution for the instances in your VCN. You make this choice for *each subnet* in the VCN, using the subnet's set of DHCP options. This is similar to how you configure which route table and security lists are associated with each subnet. For more information, see [DHCP Options](#).



### Note

You use the Domain Name Server DHCP option to specify the DNS Type for the associated subnet. If you change the option's value, either restart the DHCP client on the instance or reboot the instance. Otherwise, the change does not get picked up until the DHCP client refreshes the lease (within 24 hours).

### DEFAULT CHOICE: INTERNET AND VCN RESOLVER

This is an Oracle-provided option that includes two parts:

- **Internet Resolver:** Lets instances resolve hostnames that are publicly published on the internet. The instances do not need to have internet access by way of either an internet gateway or a connection to your on-premises network (such as an IPsec VPN connection through a DRG).

- **VCN Resolver:** Lets instances resolve hostnames (which you can assign) of other instances in the same VCN. For more information, see [About the DNS Domains and Hostnames](#).

By default, new VCNs you create use the Internet and VCN Resolver. If you're using the Networking API, this choice refers to the `VcnLocalPlusInternet` enum in the [DhcpDnsOption object](#).



### Note

The Internet and VCN Resolver does not let instances resolve the hostnames of hosts in your on-premises network connected to your VCN by IPsec VPN connection or FastConnect. Use your own custom DNS resolver to enable that.

### CUSTOM RESOLVER

Use DNS servers of your choice for resolution (maximum three). They could be DNS servers that are:

- Available through the internet. For example, 216.146.35.35 for Dyn's Internet Guide.
- In your VCN.
- In your on-premises network, which is connected to your VCN by way of an IPsec VPN connection or FastConnect (through a DRG).

## About the DNS Domains and Hostnames

When you initially create a VCN and subnets, you may specify DNS labels for each. The labels, along with the parent domain of `oraclevcn.com` form the VCN domain name and subnet domain name:

- **VCN domain name:** `<VCN DNS label>.oraclevcn.com`
- **Subnet domain name:** `<subnet DNS label>.<VCN DNS label>.oraclevcn.com`

When you then launch an instance, you may assign a hostname. It's assigned to the VNIC that's automatically created during instance launch (that is, the *primary VNIC*). Along with the subnet domain name, the hostname forms the instance's fully qualified domain name (FQDN):

- **Instance FQDN:** `<hostname>.<subnet DNS label>.<VCN DNS label>.oraclevcn.com`

For example: `database1.privatesubnet1.abccorpvcn1.oraclevcn.com`.

The FQDN resolves to the instance's private IP address. The Internet and VCN Resolver also enables reverse DNS lookup, which lets you determine the hostname corresponding to the private IP address.

If you add a [secondary VNIC](#) to the instance, you can specify a hostname. The resulting FQDN resolves to the VNIC's private IP address (that is, the *primary private IP*).

If you add a [secondary private IP](#) to a VNIC, you can specify a hostname. The resulting FQDN resolves to that private IP address.

### Requirements for DNS Labels and Hostnames

- VCN and subnet labels: Max 15 alphanumeric characters and must start with a letter. **Notice that hyphens and underscores are NOT allowed.** The value cannot be changed later.
- Hostnames: Max 63 characters and must be compliant with RFCs [952](#) and [1123](#). The value can be changed later.



### Important

The Networking service allows hostnames up to 63 characters. However, some older operating systems enforce shorter hostnames. In Linux, here's how to determine the maximum allowed hostname length:

```
getconf HOST_NAME_MAX
```

If an instance has a hostname longer than the OS-specific maximum, the instance's FQDN is not resolvable within the VCN. You can use the Networking service to [update the VNIC](#) and change the hostname to a shorter value.

Uniqueness:

- VCN DNS label should be unique across your VCNs (not required, but a best practice)
- Subnet DNS labels must be unique within the VCN
- Hostnames must be unique within the subnet



### Tip

Don't confuse the DNS label or hostname with the friendly name you can assign to the object (that is, the *display name*), which doesn't have to be unique.

### Validation and Generation of the Hostname

If you've set DNS labels for the VCN and subnets, Oracle validates the hostname for DNS compliance and uniqueness during instance launch. If either of these requirements isn't met, the launch request fails.

If you don't specify a hostname during instance launch, Oracle tries to use the instance's display name as the hostname. If the display name does not pass the validation, Oracle automatically generates a DNS-compliant hostname that's unique across the subnet. You can see the generated hostname on the instance's page in the Console. In the API, the hostname is part of the [VNIC object](#).

If you don't provide a hostname or display name during instance launch, Oracle automatically generates a display name but NOT a hostname. This means the instance won't be resolvable using the Internet and VCN Resolver.



### Note

The Linux OS hostname on the instance is automatically set to the hostname you set during instance launch (or the one generated by Oracle). If you change the hostname directly on the instance, the FQDN of the instance does not get updated.

If you add a [secondary VNIC](#) to an instance, or add a [secondary private IP](#) to a VNIC, Oracle never tries to generate a hostname. Provide a valid hostname if you want the private IP address to be resolvable using the Internet and VCN Resolver.

### DHCP Options for DNS

There are two [DHCP options](#) related to DNS in your VCN:

- **Domain Name Server:** To specify your choice for DNS type (either Internet and VCN Resolver, or Custom Resolver).
  - **Default value in the default set of DHCP options:** Internet and VCN Resolver

- **Search Domain:** To specify a single search domain. When resolving a DNS query, the OS appends this search domain to the value being queried. You can specify only one search domain for the set of DHCP options.
  - **Default value in the default set of DHCP options:** The VCN domain name (`<VCN DNS label>.oraclevcn.com`), if you specified a DNS label for the VCN during creation but did not specify a search domain value. If you specified a search domain value, then that value is used for the Search Domain option. If you did NOT specify a DNS label, the default set of DHCP options does not include a Search Domain option.

The default set of DHCP options that you can associate with all the subnets in the VCN automatically uses the default values. This means you can simply use the `<hostname>.<subnet DNS label>` to communicate with any of the instances in the VCN. If the VCN uses a set of DHCP options that does not contain a Search Domain option, the instances must use the entire FQDN to communicate.



### Important

In general, when *any* set of DHCP options is initially created (the default set or a custom set you create), the Networking service automatically adds the Search Domain option and sets it to the VCN domain name (`<VCN DNS label>.oraclevcn.com`) *if all of these are true:*

- The VCN has a DNS label
- DNS Type = Internet and VCN Resolver
- You did NOT specify a search domain of your choice during creation of the set of DHCP options

After the set of DHCP options is created, you can always remove the Search Domain option or set it to a different value.

### How to Enable DNS Hostnames in Your VCN

Only new VCNs created after the release of the Internet and VCN Resolver feature have automatic access to it. How to enable DNS hostnames for a new VCN depends on which interface you're using.

#### If you create a new VCN and subnets with the Console

1. When creating the VCN:
  - Check the checkbox for **Use DNS Hostnames in this VCN**
  - Specify a DNS label of your choice for the VCN. If you check the checkbox but don't specify a DNS label, the Console assumes that you want to use the Internet and VCN Resolver in your VCN and automatically generates a DNS label for the VCN. The Console takes the VCN name you provided, removes non-alphanumeric characters, ensures that the first character is a letter, and truncates the label to 15 characters. The Console displays the result, and if you don't like it, you can instead enter your own value in the **DNS Label** field. See [Requirements for DNS Labels and Hostnames](#).
2. When creating the subnets:
  - Again, check the checkbox for **Use DNS Hostnames in this Subnet**
  - Specify a DNS label of your choice for each subnet. If you check the checkbox but don't specify the DNS label for a given subnet, the Console assumes you want to use the Internet and VCN Resolver for the subnet and automatically generates a DNS label for the subnet. The Console takes the subnet name you provided, removes non-alphanumeric characters, ensures that the first character is a letter, and truncates the label to 15 characters. The Console displays the result, and if you don't like it, you can instead enter your own value in the **DNS Label** field. See [Requirements for DNS Labels and Hostnames](#).
  - Associate any set of DHCP options that has DNS type = Internet and VCN Resolver. The default set of DHCP options in the VCN uses the Internet and VCN Resolver by default.

### 3. When launching instances:

- Specify a hostname (or at least a display name) for each instance. For more information, see [Validation and Generation of the Hostname](#).

If you don't check the checkbox for **Use DNS Hostnames in this VCN** when creating the VCN, you can't set the DNS label for the VCN or subnets, and you can't specify a hostname during instance launch.



#### Note

The above procedure assumes you create the VCN and subnets one at a time in the Console. The Console has a feature that automatically creates a VCN with subnets and an internet gateway all at the same time. If you use that feature to create the VCN and subnets, the Console automatically generates DNS labels for them.

### If you create a new VCN and subnets with the API

#### 1. When creating the VCN:

- Specify a DNS label for the VCN. See [Requirements for DNS Labels and Hostnames](#). If you don't set a value (if it's null), Oracle assumes you don't want to use the Internet and VCN Resolver, even if the DHCP options have [DhcpDnsOption](#) `serverType = VcnLocalPlusInternet`.

#### 2. When creating the subnets:

- Specify a DNS label for each subnet. See [Requirements for DNS Labels and Hostnames](#). If you specified a DNS label for the VCN, but you don't specify a DNS label for the subnet, Oracle assumes you don't want the instances in the subnet to use the Internet and VCN Resolver. They will not be able to use hostnames to communicate with instances in the VCN.

- Associate any set of DHCP options that has `DhcpDnsOption``serverType = VcnLocalPlusInternet`. The default set of DHCP options in the VCN uses this by default.
3. When launching instances:
- Specify a hostname (or at least a display name) for each instance. For more information, see [Validation and Generation of the Hostname](#).

If you don't specify a DNS label when creating the VCN, you can't set the DNS label for the subnets (the `CreateSubnet` call will fail), nor specify a hostname during instance launch (the `LaunchInstance` call will fail). You also cannot assign a hostname to a [secondary VNIC](#) or a [secondary private IP](#).

### Scenario 1: Use Internet and VCN Resolver with DNS Hostnames Across the VCN

The typical scenario is to enable the Internet and VCN Resolver *across your entire VCN*. This means all instances in the VCN can communicate with each other without knowing their IP addresses. To do that, follow the instructions for creating a new VCN in [How to Enable DNS Hostnames in Your VCN](#), and make sure to assign a DNS label to the VCN and every subnet. Then make sure to assign every instance a hostname (or at least a display name) at launch. If you add a [secondary VNIC](#) or [secondary private IP](#), also assign it a hostname. The instances can then communicate with each other using FQDNs instead of IP addresses. If you also set the Search Domain DHCP option to the VCN domain name (`<VCN DNS label>.oraclevcn.com`), the instances can then communicate with each other using just `<hostname>.<subnet DNS label>` instead of the FQDN.

### Scenario 2: Use Custom DNS Servers to Resolve DNS Hostnames

You can set up an instance to be a custom DNS server within your VCN and configure it to resolve the hostnames that you set when launching the instances. You must configure the servers to use 169.254.169.254 as the forwarder for the VCN domain (that is, `<VCN DNS label>.oraclevcn.com`).



### Note

The custom DNS servers must be located in a subnet that uses Internet and VCN Resolver for DNS.

For an example of an implementation of this scenario with the Oracle Terraform provider, see [Hybrid DNS Configuration](#).

### Scenario 3: Use Different DHCP Options Per Subnet

[Scenario 1](#) assumes you want to use the Internet and VCN Resolver the same way across all subnets, and thus all instances in the VCN. You could, however, configure different DNS settings *for each subnet*, because the DHCP options are configured *at the subnet level*. The important thing to understand is this: the subnet where you want to generate the DNS query is where you need to configure the corresponding Internet and VCN Resolver settings.

For example, if you want instance A in subnet A to resolve the hostname of instance B in subnet B, you must configure subnet A to use the Internet and VCN Resolver. Conversely, if you want instance B to resolve the hostname of instance A, you must configure subnet B to use the Internet and VCN Resolver.

You can configure a different set of DHCP options for each subnet. For example, you could set subnet A's Search Domain to `subneta.vcn1.oraclevcn.com`, which means all instances in subnet A could use just hostnames to communicate with each other. You could similarly set subnet B's Search domain to `subnetb.vcn1.oraclevcn.com` to enable Subnet B's instances to communicate with each other with just hostnames. But that means any instances in a given subnet would need to use FQDNs to communicate with instances in other subnets.

## DHCP Options

This topic describes how to manage the Dynamic Host Configuration Protocol (DHCP) options in a virtual cloud network (VCN).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of DHCP Options

The Networking service uses DHCP to automatically provide configuration information to instances when they boot up. Although DHCP lets you change some settings dynamically, others are static and never change. For example, when you launch an instance, either you or Oracle specifies the instance's private IP address. Each time the instance boots up or you restart the instance's DHCP client, DHCP passes that same private IP address to the instance. The address never changes during the instance's lifetime.

The Networking service provides *DHCP options* to let you control certain types of configuration on the instances in your VCN. You can change the values of these options at your discretion, unlike the static information that DHCP provides to the instance. The changes take effect the next time you restart a given instance's DHCP client or reboot the instance. For more details, see [Important Notes about Your Instances and DHCP Options](#).

Each subnet in a VCN can have a single set of DHCP options associated with it. That set of options applies to all instances in the subnet. Each VCN comes with a *default set of DHCP options* with initial values that you can change. If you don't specify otherwise, every subnet uses the VCN's default set of DHCP options.

The following table summarizes the available DHCP options you can configure.

DHCP Option	Possible Values	Initial Value in the Default DHCP Options	Notes
Domain Name Server	<p data-bbox="334 478 477 512"><u>DNS Type:</u></p> <ul data-bbox="383 541 532 737" style="list-style-type: none"><li data-bbox="383 541 532 646">• Internet and VCN Resolver</li><li data-bbox="383 674 532 737">• Custom Resolver</li></ul>	DNS Type = Internet and VCN resolver. For more information, see <a href="#">Choices for DNS in Your VCN</a> .	If you set DNS Type = Custom Resolver, you can specify up to three DNS servers of your choice. For more information, see <a href="#">Choices for DNS in Your VCN</a> .

DHCP Option	Possible Values	Initial Value in the Default DHCP Options	Notes
Search Domain	A single search domain	If you've set up your VCN with a DNS label, the default value for the Search Domain option is the <a href="#">VCN domain name</a> (<VCN DNS label>.oraclevcn.com). Otherwise, the Search Domain option is not present in the default set of DHCP options.	<p>In general, when <i>any</i> set of DHCP options is initially created (the default set or a custom set you create), the Networking service automatically adds the Search Domain option and sets it to the VCN domain name (&lt;VCN DNS label&gt;.oraclevcn.com) <i>if all of these are true</i>:</p> <ul style="list-style-type: none"> <li>• The VCN has a DNS label</li> <li>• DNS Type = Internet and VCN Resolver</li> <li>• You did NOT specify a search domain of your choice during creation of the set of DHCP options</li> </ul> <p>After the set of DHCP options is created, you can always remove the Search Domain option or set it to a different value.</p> <p>You can specify only a single search domain in a set of DHCP options.</p>

## Working with DHCP Options

When you create a subnet, you specify which set of DHCP options to associate with the subnet. If you don't, the default set of DHCP options for the VCN is used. You can [change which set of](#)

[DHCP options the subnet uses](#) at any time.

When creating a new set of DHCP options, you may optionally assign it a friendly name. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the set of options a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

You can change the values of an individual DHCP option in a set, but notice that when you use the REST API to update a single option in a set, the new set of options replaces the entire existing set.

To delete a set of DHCP options, it must not be associated with a subnet yet. You can't delete a VCN's default set of DHCP options.

For information about the maximum number of DHCP options allowed, see [Service Limits](#).

### **Required IAM Policy**

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

## Important Notes about Your Instances and DHCP Options

Whenever you change the value of one of the DHCP options, you must do one of the following for the change to take effect on existing instances in the subnets associated with that set of DHCP options: either restart the DHCP client on the instance, or reboot the instance.

Make sure to keep the DHCP client running so you can always access the instance. If you stop the DHCP client manually or disable NetworkManager (which stops the DHCP client on Linux instances), the instance can't renew its DHCP lease and will become inaccessible when the lease expires (typically within 24 hours). Do not disable NetworkManager unless you use another method to ensure renewal of the lease.

Stopping the DHCP client might remove the host route table when the lease expires. Also, loss of network connectivity to your iSCSI connections might result in loss of the boot drive.

Any changes you make to the `/etc/resolv.conf` file are overwritten whenever the DHCP lease is renewed or the instance is rebooted.

Changes you make to the `/etc/hosts` file are overwritten whenever the DHCP lease is renewed or the instance is rebooted. To persist your changes to the `/etc/hosts` file, add the following line to `/etc/oci-hostname.conf`:

```
PRESERVE_HOSTINFO=2
```

If the `/etc/oci-hostname.conf` file does not exist, create it.

### Using the Console

#### To view a VCN's set of default DHCP options

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **DHCP Options**.  
The default set and its details are displayed in the list.

#### To update options in an existing set of DHCP options

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **DHCP Options**.
4. For the set you're interested in, click the Actions icon (three dots), and then click **Edit**:

- For **DNS Type**: If you want instances in the subnet to resolve internet hostnames and hostnames of instances in the VCN, select **Internet and VCN Resolver**. Or to use a DNS server of your choice, select **Custom Resolver** and then enter the server's IP address (three servers maximum). For more information, see [DNS in Your Virtual Cloud Network](#).
  - For **Search Domain**: If you want instances in the subnet to append a particular search domain when resolving DNS queries, enter it here. If the Search Domain option is already set to the [VCN domain name](#) and you're not sure why, see the details in [Overview of DHCP Options](#).
5. When you're done, click **Save Changes**.
  6. If you have any existing instances in a subnet that uses this set of DHCP options, make sure to restart the DHCP client on each affected instance, or reboot the instance itself so that it picks up the new setting.

### To create a new set of DHCP options

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **DHCP Options**.
4. Click **Create DHCP Options**.
5. Enter the following:
  - **Name**: A friendly name for the set of options. It doesn't have to be unique, and you can change it later. Avoid entering confidential information.
  - **Create in Compartment**: The compartment where you want to create the set of DHCP options, if different from the compartment you're currently working in.
  - **DNS Type**: If you want instances in the subnet to resolve internet hostnames and hostnames of instances in the VCN, select **Internet and VCN Resolver**. Or to use a DNS server of your choice, select **Custom Resolver** and then enter the

server's IP address (three servers maximum). For more information, see [DNS in Your Virtual Cloud Network](#).

- **Search Domain:** If you want instances in the subnet to append a particular search domain when resolving DNS queries, enter it here. Be aware that the Networking service automatically sets the Search Domain option in certain situations. See the details in [Overview of DHCP Options](#).
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

6. When you're done, click **Create DHCP Options**.

The set of options is created and then displayed on the **DHCP Options** page of the compartment you chose. You can now specify this set of options when creating or updating a subnet.

### To change which set of DHCP options a subnet uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Subnets**.
4. Click the subnet you're interested in.
5. Click **Edit**.
6. In the **DHCP Options** section, select the new set of DHCP options you want the subnet to use.
7. Click **Save Changes**.

The changes take effect within a few seconds.

### To delete a set of DHCP options

Prerequisite: To delete a set of DHCP options, it must not be associated with a subnet yet. You can't delete the default set of DHCP options in a VCN.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **DHCP Options**.
4. For the set you want to delete, click the Actions icon (three dots), and then click **Terminate**.
5. Confirm when prompted.

### To manage tags for a set of DHCP options

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **DHCP Options**.
4. For the set you're interested in, click the Actions icon (three dots), and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move a set of DHCP options to a different compartment

You can move a set of DHCP options from one compartment to another. When you move a set of DHCP options to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.

3. Under **Resources**, click **DHCP Options**.
4. For the set you're interested in, click the Actions icon (three dots), and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage a VCN's DHCP options, use these operations:

- [ListDhcpOptions](#)
- [GetDhcpOptions](#)
- [UpdateDhcpOptions](#)
- [CreateDhcpOptions](#)
- [DeleteDhcpOptions](#)
- [ChangeDhcpOptionsCompartment](#)

### Route Tables

This topic describes how to manage the route tables in a virtual cloud network (VCN).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of Routing for Your VCN

Your VCN uses virtual route tables to send traffic out of the VCN (for example, to the internet, to your on-premises network, or to a peered VCN). These virtual route tables have rules that look and act like traditional network route rules you might already be familiar with. Each rule specifies a destination CIDR block and the target (the next hop) for any traffic that matches that CIDR.

Here are basics about routing in your VCN:

- The primary routing scenario is for sending a subnet's traffic to destinations outside the VCN. A subnet has a single route table of your choice associated with it. All VNICs in that subnet are subject to the rules in the route table. The rules govern how the traffic leaving the subnet is routed.
- Traffic within the VCN is automatically handled by the VCN *local routing*. No route rules are required to enable that traffic. And in general: for any route table that belongs to a given VCN, you can't create a rule that lists that VCN's CIDR (or a sub-section) as the rule's destination. Oracle uses a subnet's route table *only* if the destination IP address is not within the VCN's CIDR block.
- If a route table has overlapping rules, Oracle uses the most specific rule in the table to route the traffic (that is, the rule with the [longest prefix match](#)).
- If there is no route rule that matches the network traffic you intend to route outside the VCN, the traffic is dropped (blackholed).
- IPv6 addressing and routing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

For important details about routing between your VCN and on-premises network, see [Routing Details for Connections to Your On-Premises Network](#).

### Working with Route Tables and Route Rules

Each VCN automatically comes with a default route table that has no rules. If you don't specify otherwise, every subnet uses the VCN's default route table. When you add route rules to your VCN, you can simply add them to the default table if that suits your needs. However, if you need both a public subnet and a private subnet (for example, see [Scenario C: Public and Private Subnets with a VPN](#)), you instead create a separate (custom) route table for each subnet.

Each subnet in a VCN uses a single route table. When you create the subnet, you specify which one to use. You can [change which route table the subnet uses](#) at any time. You can also edit a route table's rules, or remove all the rules from the table.

You may optionally assign a descriptive name to a custom route table during creation. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the route table a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

A route rule specifies a destination CIDR block and the target (the next hop) for any traffic that matches that CIDR. Here are the allowed types of targets for a route rule:

- [dynamic routing gateway \(DRG\)](#): For subnets that need private access to networks connected to your VCN (for example, your on-premises network connected with an [IPSec VPN](#) or [FastConnect](#), or a [peered VCN in another region](#)).
- [internet gateway](#): For public subnets that need direct access to the internet.
- [NAT gateway](#): For subnets with instances that do not have public IP addresses but need outbound access to the internet.
- [service gateway](#): For subnets that need private access to Oracle services such as Object Storage.
- [local peering gateway \(LPG\)](#): For subnets that need private access to a peered VCN in the same region.

- [private IP](#): For subnets that need to route traffic to an instance in the VCN. For more information, see [Using a Private IP as a Route Target](#). Also see [Advanced Scenarios: Transit Routing](#).



### Note

You can't delete a particular resource if it's the target in a route rule. For example, you can't delete an internet gateway that has traffic routed to it. You must first delete all rules (in all route tables) with that internet gateway as the target.

When adding a route rule to a route table, you provide the destination CIDR block and target (plus the compartment where the target resides). Exception: if the target is a service gateway, instead of a destination CIDR block, you specify an Oracle-provided string that represents the public endpoints for the service of interest. That way you don't need to know all the service's CIDR blocks, which might change over time.

If you misconfigure a rule (for example, enter the wrong destination CIDR block), the network traffic you intended to route might be dropped (blackholed) or sent to an unintended target.

You can [move route tables from one compartment to another](#). Moving a route table doesn't affect its attachment to VCNs or subnets. When you move a route table to a new compartment, inherent policies apply immediately and affect access to the route table. For more information, see [Access Control](#).

You can't delete a VCN's default route table. To delete a custom route table, it must not be associated with a subnet yet. Or, in the advanced scenario of [transit routing](#), it must not be associated with a DRG attachment or LPG in the VCN.

For information about the maximum number of route tables and route rules, see [Service Limits](#).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

### Advanced Scenarios: Transit Routing

This documentation includes a few basic networking scenarios to help you understand the Networking service and generally how the components work together. See scenarios A, B, and C in [Networking Scenarios](#).

Scenarios A–C show your on-premises network connected to a VCN by way of [FastConnect](#) or [VPN Connect](#), and accessing only the resources in that VCN.

The following advanced routing scenarios give your on-premises network additional access beyond the resources in the connected VCN. Traffic travels from your on-premises network to the VCN, and then *transits through* the VCN to its destination. See these topics:

- [Transit Routing: Access to Multiple VCNs in the Same Region](#): Your on-premises network has access to *multiple* VCNs in the same region over a single FastConnect private virtual circuit or VPN Connect. The VCNs are in a hub-and-spoke layout, with the on-premises network connected to the VCN that acts as the hub. The spoke VCNs are peered with the hub VCN.
- [Transit Routing: Private Access to Oracle Services](#): Your on-premises network has *private access* to Oracle services in the [Oracle Services Network](#) by way of the connected VCN and the VCN's service gateway. The traffic does not go over the internet.

Both of the transit routing scenarios involve creating route tables that you associate *with a gateway instead of a subnet*.

For example:

- For [Transit Routing: Access to Multiple VCNs in the Same Region](#), you create a route table for a DRG attachment, and then multiple route tables, each for different LPGs.
- For [Transit Routing: Private Access to Oracle Services](#), you create a route table for a DRG attachment and a route table for a service gateway.

The Networking service imposes restrictions on how the route tables can be used:

- DRG attachment: If you associate a route table with the DRG attachment on a VCN, the route table can contain only rules that use an LPG on the VCN, a service gateway on the VCN, or a private IP in the VCN as the target.
- LPG or service gateway: If you associate a route table with an LPG or a service gateway, the route table can contain only rules that use the VCN's attached DRG or a private IP in the VCN as the target.

A DRG attachment or LPG can exist without a route table associated with it. However, after you associate a route table with a DRG attachment or LPG, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

### Using a Private IP as a Route Target

If you're not familiar with the definition of a private IP, see [Private IP Addresses](#). In short: a private IP is an object that contains a private IP address and related properties and has its own [OCID](#).

### General Use Cases

As mentioned earlier, Oracle uses a given subnet's route table *only* for traffic with a destination IP address outside the VCN. Typically you set up one or more rules to route that traffic to a gateway on the VCN (for example, a DRG connected to your on-premises network, or an internet gateway connected to the internet). However, you might want to route that traffic (going to destinations outside the VCN) through an instance in the VCN first. In that case, you can use a private IP in the VCN as the target instead of a gateway on the VCN. Here are a few reasons you might do this:

- To implement a virtual network appliance (NVA) such as a firewall or intrusion detection that filters outgoing traffic from instances.
- To manage an overlay network on the VCN, which lets you run container orchestration workloads.
- To implement Network Address Translation (NAT) in the VCN. Note that Oracle instead recommends using a [NAT gateway](#) with your VCN. In general, NAT enables outbound internet access for instances that don't have direct internet connectivity.

To implement these use cases, there's more to do than simply route traffic to the instance. There's also configuration required on the instance itself.



#### Tip

You can enable high availability of the private IP route target by using a [secondary private IP address](#). In the event of a failure, you can move the secondary private IP from an existing VNIC to another VNIC in the same subnet. See [To move a secondary private IP to another VNIC in the same subnet](#) (Console instructions) and [UpdatePrivateIp](#) (API instructions).

### Requirements for Using a Private IP as a Target

- The private IP must be in the same VCN as the route table.
- The private IP's VNIC must be configured to skip the source/destination check so that the VNIC can forward traffic. By default, VNICs are configured to perform the check. For more information, see [Source/Destination Check](#).
- You must configure the instance itself to forward packets. For more information, see [NAT Instance Configuration](#).
- The route rule must specify the OCID of the private IP as the target, and not the IP address itself. Exception: If you use the Console, you can instead specify the private IP address itself as the target, and the Console determines and uses the corresponding OCID in the rule.



### Important

A route rule with a private IP target can result in blackholing in these cases:

- The instance the private IP is assigned to is stopped or terminated
- The VNIC the private IP is assigned to is updated to enable the source/destination check or is deleted
- The private IP is unassigned from the VNIC

When a target private IP is terminated, in the Console, the route rule displays a note that the target OCID no longer exists.

For failover: If your target instance is terminated before you can move the secondary private IP to a standby, you must update the route rule to use the OCID of the new target private IP on the standby. The rule uses the target's OCID and not the private IP address itself.

### General Setup Process

1. Determine which instance will receive and forward the traffic (the NAT instance, for example).
2. Choose a private IP on the instance (can be on the instance's primary VNIC or a secondary VNIC). If you want to implement failover, set up a [secondary private IP](#) on one of the VNICs on the instance.
3. Disable the source/destination check on the private IP's VNIC. See [Source/Destination Check](#).

4. Get the OCID for the private IP. If you're using the Console, you can get either the OCID or the private IP address itself, along with the name of the private IP's compartment.
5. For the subnet that needs to route traffic to the private IP, view the subnet's route table. If the table already has a rule with the same destination CIDR but a different target, delete that rule.
6. [Add a route rule](#) with the following:
  - **Destination CIDR block:** If all traffic leaving the subnet needs to go to the private IP, use 0.0.0.0/0.
  - **Target type:** Private IP.
  - **Compartment:** The compartment of the private IP.
  - **Target:** The OCID of the private IP. If you're using the Console and instead enter the private IP address itself, the Console determines the corresponding OCID and uses it as the target for the route rule.

As mentioned earlier, you must configure the instance itself to forward packets. For more information, see [NAT Instance Configuration](#).

### Using the Console

#### To view a VCN's default route table

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Route Tables**.  
The default route table is displayed in the list of tables.
4. Click the default route table to view its details.

### To update rules in an existing route table

You can add, edit, or delete rules.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Route Tables**.
4. Click the route table you're interested in.
5. If you want to create a new route rule, click **Add Route Rule** and enter the following:
  - **Target Type:** See the list of target types in [Route Tables](#). If the target type is a DRG, the VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself. If the target is a private IP, make sure you've first disabled the source/destination check on the private IP's VNIC. For more information, see [Using a Private IP as a Route Target](#).
  - **Destination CIDR Block:** Only if the target is not a [service gateway](#). The value is the destination CIDR block for the traffic. A value of 0.0.0.0/0 means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule.
  - **Destination Service:** Only if the target is a [service gateway](#). The value is the [service CIDR label](#) that you're interested in.
  - **Compartment:** The compartment where the target is located.
  - **Target:** The target. If the target is a private IP, enter its OCID. Or you can enter the private IP address itself, in which case the Console determines the corresponding OCID and uses it as the target for the route rule.
6. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Remove**.
7. If you wanted to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.

### To create a route table

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Route Tables**.
4. Click **Create Route Table**.
5. Enter the following:
  - **Name:** A friendly name for the route table. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in Compartment:** The compartment where you want to create the route table, if different from the compartment you're currently working in.
6. Optionally, click **+Additional Route Rule** to add one or more route rules, each with the following information (remember, a route table can exist with no rules until you're ready to add them):
  - **Target Type:** See the list of target types in [Route Tables](#). If the target type is a DRG, the VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself. If the target is a private IP, make sure you've first disabled the source/destination check on the private IP's VNIC. For more information, see [Using a Private IP as a Route Target](#).
  - **Destination CIDR Block:** Only if the target is not a [service gateway](#). The value is the destination CIDR block for the traffic. A value of 0.0.0.0/0 means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule.
  - **Destination Service:** Only if the target is a [service gateway](#). The value is the [service CIDR label](#) that you're interested in.
  - **Compartment:** The compartment where the target is located.

- **Target:** The target. If the target is a private IP, enter its OCID. Or you can enter the private IP address itself, in which case the Console determines the corresponding OCID and uses it as the target for the route rule.
7. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
  8. Click **Create Route Table**.  
The route table is created and then displayed on the **Route Tables** page in the compartment you chose. You can now specify this route table when creating or updating a subnet.

### To change which route table a subnet uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Subnets**.
4. Click the subnet you're interested in
5. Click **Edit**.
6. In the **Route Table** section, select the new route table you want the subnet to use.
7. Click **Save Changes**.  
The changes take effect within a few seconds.

### To move a route table to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.

2. Click the VCN you're interested in.
3. Click **Route Tables**.
4. Find the route table in the list, click the the Actions icon (three dots), and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

### To delete a route table

Prerequisite: To delete a route table, it must not be associated with a subnet yet. You can't delete the default route table in a VCN.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Route Tables**.
4. Click the route table you're interested in.
5. Click **Terminate**.
6. Confirm when prompted.

### To manage tags for a route table

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Route Tables**.
4. Click the route table you're interested in.

5. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage a VCN's route tables, use these operations:

- [ListRouteTables](#)
- [GetRouteTable](#)
- [UpdateRouteTable](#)
- [CreateRouteTable](#)
- [DeleteRouteTable](#)
- [ChangeRouteTableCompartment](#)

## Dynamic Routing Gateways (DRGs)

This topic describes how to manage a dynamic routing gateway (DRG). This topic uses the terms *dynamic routing gateway* and *DRG* interchangeably. The Console uses the term *Dynamic Routing Gateway*, whereas for brevity the API uses *DRG*.

You use a DRG when connecting your existing on-premises network to your virtual cloud network (VCN) with one (or both) of these:

- [IPSec VPN](#)
- [Oracle Cloud Infrastructure FastConnect](#)

You also use a DRG when peering a VCN with a VCN in a different region:

- [Remote VCN Peering \(Across Regions\)](#)



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of Dynamic Routing Gateways

You can think of a DRG as a virtual router that provides a path for private traffic (that is, traffic that uses private IPv4 addresses) between your VCN and networks outside the VCN's region.

For example, if you use an [IPSec VPN](#) or [Oracle Cloud Infrastructure FastConnect](#) (or both) to connect your on-premises network to your VCN, that private IPv4 address traffic goes through a DRG that you create and attach to your VCN. For scenarios for using a DRG to connect a VCN to your on-premises network, see [Networking Scenarios](#). For important details about routing to your on-premises network, see [Routing Details for Connections to Your On-Premises Network](#).

Also, if you decide to peer your VCN with a VCN in another region, your VCN's DRG routes traffic to the other VCN over a private backbone that connects the regions (without traffic traversing the internet). For information about connecting VCNs in different regions, see [Remote VCN Peering \(Across Regions\)](#).

### Working with DRGs and DRG Attachments

For the purposes of access control, when creating a DRG, you must specify the compartment where you want the DRG to reside. If you're not sure which compartment to use, put the DRG in the same compartment as the VCN. For more information, see [Access Control](#).

You may optionally assign a friendly name to the DRG. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the DRG a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

A DRG is a standalone object. To use it, you must attach it to a VCN. A VCN can be attached to only one DRG at a time, and a DRG can be attached to only one VCN at a time. You can detach a DRG and reattach it at any time. In the API, the process of attaching creates a `DrgAttachment` object with its own OCID. To detach the DRG, you delete that attachment object.

After attaching a DRG, you must update the routing in the VCN to use the DRG. Otherwise, traffic from the VCN will not flow to the DRG. See [To route a subnet's traffic to a DRG](#).

To delete a DRG, it must not be attached to a VCN or connected to another network by way of IPSec VPN, Oracle Cloud Infrastructure FastConnect, or remote VCN peering. Also, there must not be a route rule that lists that DRG as a target.

For information about the number of DRGs you can have, see [Service Limits](#).

### Routing a Subnet's Traffic to a DRG

The basic routing scenario sends traffic from a subnet in the VCN to the DRG. For example, if you're sending traffic from the subnet to your on-premises network, you [set up a rule in the subnet's route table](#). The rule's destination CIDR is the CIDR for the on-premises network (or a subnet within), and the rule's target is the DRG. For more information, see [Route Tables](#).

### Advanced Scenarios: Transit Routing

This documentation includes a few basic networking scenarios to help you understand the Networking service and generally how the components work together. See scenarios A, B, and C in [Networking Scenarios](#).

Scenarios A–C show your on-premises network connected to a VCN by way of [FastConnect](#) or [VPN Connect](#), and accessing only the resources in that VCN.

The following advanced routing scenarios give your on-premises network additional access beyond the resources in the connected VCN. Traffic travels from your on-premises network to the VCN, and then *transits through* the VCN to its destination. See these topics:

- [Transit Routing: Access to Multiple VCNs in the Same Region](#): Your on-premises network has access to *multiple* VCNs in the same region over a single FastConnect private virtual circuit or VPN Connect. The VCNs are in a hub-and-spoke layout, with the on-premises network connected to the VCN that acts as the hub. The spoke VCNs are peered with the hub VCN.
- [Transit Routing: Private Access to Oracle Services](#): Your on-premises network has *private access* to Oracle services in the [Oracle Services Network](#) by way of the connected VCN and the VCN's service gateway. The traffic does not go over the internet.

In the transit routing scenarios, the VCN has a route table *associated with its DRG attachment* (typically route tables are associated with a VCN's subnets). That route table lets you manage routing of traffic *through the VCN* that is connected to the on-premises network.

When you attach a DRG to a VCN, you can optionally associate a route table with the attachment. Or if you already have a DRG attachment, you [can associate a route table with it](#). The route table must belong to the attached VCN. A route table associated with a DRG attachment can contain only rules that use one of the following as a target:

- A [local peering gateway \(LPG\)](#) on the attached VCN
- A [service gateway](#) on the attached VCN
- A [private IP](#) in the attached VCN

A DRG attachment can exist without a route table associated with it. However, after you associate a route table with a DRG attachment, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

### Using the Console

In general, to use a DRG, you must complete these steps:

1. Create the DRG.
2. Attach the DRG to your VCN.
3. Route subnet traffic to the DRG. This involves updating the route table associated with each subnet that must send traffic to the DRG. If all the subnets use the VCN's default route table, you must only update that one table.

### To create a DRG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Dynamic Routing Gateway**.
4. Enter the following items:
  - **Create in Compartment:** The compartment where you want to create the DRG, if different from the compartment you're currently working in.
  - **Name:** A descriptive name for the DRG. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create Dynamic Routing Gateway**.

The resource is created and then displayed on the **Dynamic Routing Gateways** page of the compartment you chose. It will be in the "Provisioning" state for a short period. You can connect it to other parts of your network only after provisioning is complete.

### To update a DRG's name

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you're interested in.
3. Click **Edit**.
4. Edit the name and click **Save Changes**.

### To attach a DRG to a VCN

**Note:** A VCN can be attached to only one DRG at a time, and a DRG can be attached to only one VCN at a time. The attachment is automatically created in the compartment that holds the VCN.

The following instructions have you navigate to the DRG and then choose which VCN to attach. You could instead navigate to the VCN and then choose which DRG to attach.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you want to attach.
3. Under **Resources**, click **Virtual Cloud Networks**. If you want to attach the DRG to a VCN in a different compartment than the one you're working in, choose that compartment from the list on the left side of the page.
4. Click **Attach to Virtual Cloud Network**.
5. Select the VCN.

6. (Optional) Only if you're setting up an [advanced scenario for transit routing](#), you can associate a route table with the DRG attachment (you can do this later if you want to):
  - a. Click **Show Advanced Options**.
  - b. Select the route table that you want to associate with the DRG attachment.
7. Click **Attach to Virtual Cloud Network**.

The attachment will be in the "Attaching" state for a short period before it's ready.

After it's ready, make sure to create a route rule that directs subnet traffic to this DRG. See [To route a subnet's traffic to a DRG](#).

### To route a subnet's traffic to a DRG

For each subnet that must send traffic to the DRG, you must add a route rule to the route table associated with that subnet. If all the subnets in the VCN use the default route table, you must add a rule to only that one table.

If all non-intra-VCN traffic that's not covered by another rule in the table must be routed to the DRG, then this is the new rule to add:

- **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
- **Destination CIDR Block** = 0.0.0.0/0. If you want to limit the rule to a specific network (for example, your on-premises network), then use that network's CIDR instead of 0.0.0.0/0.

For step-by-step instructions, see [To update rules in an existing route table](#).

To associate a route table with an existing DRG attachment



**Important**

Perform this task only if you're setting up an advanced scenario for transit routing. See [Transit Routing: Access to Multiple VCNs in the Same Region](#) and [Transit Routing: Private Access to Oracle Services](#).

A DRG attachment can exist without a route table associated with it. However, after you associate a route table with a DRG attachment, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

**Prerequisites:** The route table must exist and belong to the VCN that the DRG is already attached to.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG that is attached to the VCN that has the route table.
3. Click the Actions icon (three dots), and then click either:
  - **Associate Route Table:** If the DRG attachment has no route table associated with it yet.
  - **Associate Different Route Table:** If you're changing which route table is associated with the DRG attachment.
4. Select the route table.
5. Click **Associate Route Table**.

The route table is associated with the DRG attachment.

### To detach a DRG from a VCN

**Note:** You do not need to remove the route rule that routes traffic to the DRG before you detach the DRG from the VCN.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you want to detach.
3. Under **Resources**, click **Virtual Cloud Networks** to see the VCN the DRG is attached to. If the VCN is in a different compartment than the one you're working in, choose that compartment from the list on the left side of the page.
4. Click the Actions icon (three dots), and then click **Detach**.
5. Confirm when prompted.

The attachment will be in the "Detaching" state for a short period.

### To delete a DRG

#### Prerequisites:

- The DRG must not be attached to a VCN.
  - The DRG must not be connected to another network by way of an IPSec VPN, FastConnect, or remote VCN peering.
  - There must not be a [route rule](#) that lists the DRG as a target.
1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
  2. Click the DRG you're interested in.
  3. Click **Terminate**.
  4. Confirm when prompted.

The DRG will be in the "Terminating" state for a short period while it's being deleted.

### To manage tags for a DRG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Add Tags** to add new ones.

For more information, see [Resource Tags](#).

### To move a dynamic routing gateway to a different compartment

You can move a dynamic routing gateway from one compartment to another. When you move a dynamic routing gateway to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you're interested in.
3. Find the DRG in the list, click the the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your DRGs, use these operations:

- [ListDrgs](#)
- [CreateDrg](#)
- [GetDrg](#)
- [UpdateDrg](#)
- [DeleteDrg](#)
- [ChangeDrgCompartment](#)
- [ListDrgAttachments](#)
- [CreateDrgAttachment](#): This attaches a DRG to a VCN and results in a `DrgAttachment` object with its own OCID. You may optionally specify a route table if you're setting up the advanced routing scenario called [transit routing](#).
- [GetDrgAttachment](#)
- [UpdateDrgAttachment](#): Among other things, this associates a route table with an existing DRG attachment for [transit routing](#).
- [DeleteDrgAttachment](#): This detaches a DRG from a VCN by deleting the `DrgAttachment`.

For information about route table operations, see [Route Tables](#).

## Routing Details for Connections to Your On-Premises Network

You might use multiple connections between your on-premises network and virtual cloud network (VCN) for redundancy or other reasons.

For example, you might use both [FastConnect private peering](#) and [VPN Connect](#) to the dynamic routing gateway (DRG) attached to your VCN. Or perhaps you use redundant VPN Connect connections to the DRG (for an example scenario, see [Example Layout with Multiple Geographic Areas](#)). Or perhaps you use FastConnect public peering, FastConnect private peering, and VPN Connect.

This topic covers important details about route advertisement and path preferences when you have multiple connections.

### DRG Route Advertisements to Your On-Premises Network

FastConnect private peering and VPN Connect provide your on-premises network with private access to a VCN. Both types of connections terminate on a single DRG that is attached to the VCN. Remember that VPN Connect can use either Border Gateway Protocol (BGP) or static routing, or a combination. FastConnect always uses BGP for route advertisements.

The DRG advertises the routes for the individual subnets in the DRG's attached VCN. A DRG can be attached to only a single VCN, and a VCN can be attached to only a single DRG.

If you set up [transit routing to multiple VCNs](#) for your on-premises network, the DRG advertises additional routes. Transit routing is an advanced routing scenario that involves a single FastConnect or VPN Connect and multiple *peered* VCNs in a hub-and-spoke layout. With transit routing, the DRG also advertises routes for the VCNs that are peered with the DRG's attached VCN (the hub).

If you set up your on-premises network with [private access to Oracle services](#) through the VCN's service gateway, the DRG advertises more routes. They are routes for the Oracle Services Network, which is available with the service gateway. For a list of those ranges, see [Public IP Addresses for VCNs and the Oracle Services Network](#).



#### **Important**

If you're using VPN Connect with static routing, and you've configured the VCN to give your on-premises network private access to Oracle services, you must configure your edge device with the routes for the Oracle Services Network public IP ranges that are advertised by the DRG over the private path (through the service gateway). For a list of those ranges, see [Public IP Addresses for VCNs and the Oracle Services Network](#)

### Routing Preferences for Traffic from Oracle to Your On-Premises Network

This section describes how Oracle chooses which path to use when sending traffic to your on-premises network. The traffic can be for responding to a request or initiating new connections.

In general, routers use the most specific route (the one with the [longest prefix match](#)).

However, if the routes for the different paths are the same, **Oracle uses the shortest AS path** when sending traffic to your on-premises network, regardless of which path was used to initiate the connection to Oracle. **This means asymmetric routing is allowed.**

*Asymmetric routing* here means that Oracle's response to a request can follow a different path than the request. For example, depending on how your edge device (also called your *customer-premises equipment*, or CPE) is configured, you could send a request over VPN Connect, but the Oracle response could come back over FastConnect. If you want to force routing to be symmetric, Oracle recommends using BGP and AS path prepending with your routes to influence which path Oracle uses when responding to and initiating connections.

Oracle implements AS path prepending to determine which path to use if your edge device advertises the same route over multiple connections between your on-premises network and VCN. The details are summarized in the following table. Assuming that you're not influencing routing in some way, when the same route is advertised over multiple paths to the DRG at the Oracle end of the connections, Oracle prefers the paths in the following order:

Oracle preference	Path	Details of how Oracle prefers the path	Resulting AS path for the route
1	FastConnect	Oracle prepends no ASNs to the routes that your edge device advertises.	Your ASN
2	VPN Connect with BGP routing	Oracle prepends a single private ASN on all the routes that your edge device advertises over VPN Connect with BGP.	Private ASN, Your ASN
3	VPN Connect with static routing	Oracle prepends 3 private ASNs on the static routes that you've provided (Oracle advertises those routes to the dynamic routing gateway (DRG) at the Oracle end of the IPsec VPN) .	Private ASN, Private ASN, Private ASN

If you have two connections of the same type (for example, two IPsec VPNs that both use BGP), and you advertise the *same* routes across both connections, **Oracle prefers the oldest established route** when responding to requests or initiating connections.

## Routing Preferences for Traffic from Your On-Premises Network to Oracle

You can configure your edge device to prefer a specific path when sending traffic from your on-premises network to Oracle. This section describes a particular situation where you *must* do that to ensure a consistent traffic path if your on-premises hosts use Oracle services.

Your on-premises network can access Oracle services such as Object Storage over multiple paths. You can use public paths, such as the internet or [FastConnect public peering](#). With these public paths, the on-premises hosts communicate with Oracle services by using public IP addresses.

You can also set up your on-premises network with [private access to Oracle services](#) through the VCN's service gateway. You might do this if hosts in your on-premises network use any of the services listed in [Service Gateway: Supported Cloud Services in Oracle Services Network](#). This implementation lets your on-premises hosts communicate with those Oracle services from your private IP addresses.

If you've configured your on-premises network with *multiple* connection paths to Oracle services, your edge device may receive route advertisement of the Oracle services' public IP address routes over multiple paths. Here are the possible paths you can use with your on-premises network:

- Public access paths:
  - Internet service provider (ISP)
  - FastConnect public peering
- Private access paths by way of the VCN's DRG and service gateway:
  - FastConnect private peering
  - VPN Connect

Your edge device receives route advertisements from the DRG and possibly from routers over public paths. Most of the routes for Oracle services that the DRG advertises have a longer prefix (they are more specific) than the routes for Oracle services that are advertised over the public access paths. Therefore, if you set up your network with both public access and private access to Oracle services, **you must configure your edge device to prefer the private access path to the DRG** for traffic traveling from the on-premises network to Oracle services. This ensures a consistent path for all your access to Oracle services.

For a list of the public IP ranges advertised over FastConnect public peering, see [FastConnect Public Peering Advertised Routes](#).

For a list of the regional public IP ranges advertised over the private paths (for a VCN with a service gateway), see [Public IP Addresses for VCNs and the Oracle Services Network](#).

### Related Resources

For additional information, see these related resources:

- [Connectivity Redundancy Guide \(PDF\)](#)
- [VPN Connect Best Practices \(PDF\)](#)
- [FastConnect Redundancy Best Practices](#)

## VPN Connect

The following topics have information about setting up VPN Connect (also referred to as an IPSec VPN) between your on-premises network and virtual cloud network (VCN):

- [VPN Connect Overview](#)
- [VPN Connect Quickstart](#)
- [Routing Details for Connections to Your On-Premises Network](#)
- [Supported IPSec Parameters](#)
- [Supported Encryption Domain or Proxy ID](#)
- [Setting Up VPN Connect](#)
- [CPE Configuration](#)
  - [Verified CPE Devices](#)
  - Checkpoint:
    - [Check Point: Route-Based](#)
    - [Check Point: Policy-Based](#)
  - Cisco ASA:
    - [Cisco ASA: Route-Based](#)
    - [Cisco ASA: Policy-Based](#)
  - [Cisco IOS](#)
  - [FortiGate](#)
  - [Juniper MX](#)
  - [Juniper SRX](#)

- [Libreswan](#)
- [NEC IX Series](#)
- [Openswan](#)
- [Palo Alto](#)
- [WatchGuard](#)
- [Yamaha RTX Series](#)
- [Working with VPN Connect](#)
- [VPN Connect FAQ](#)
- [Using the API for VPN Connect](#)
- [VPN Connect Metrics](#)
- [VPN Connect Troubleshooting](#)

## VPN Connect Overview

One way to connect your on-premises network and your virtual cloud network (VCN) is to use VPN Connect, which is an IPsec VPN. IPsec stands for *Internet Protocol Security* or *IP Security*. IPsec is a protocol suite that encrypts the entire IP traffic before the packets are transferred from the source to the destination.

This topic gives an overview of an IPsec VPN for your VCN. For scenarios that include an IPsec VPN, see [Scenario B: Private Subnet with a VPN](#) and [Scenario C: Public and Private Subnets with a VPN](#).

## Required Personnel and Knowledge

Typically the following types of personnel are involved in setting up an IPsec VPN with Oracle Cloud Infrastructure:

- **Dev Ops team member** (or similar function) who uses the Oracle Cloud Infrastructure Console to set up the cloud components required for the virtual network and IPsec VPN.

- **Network engineer** (or similar function) who configures the customer-premises equipment (CPE) device with information provided by the Dev Ops team member.



### Tip

The Dev Ops team member must have the required permission to create and manage the cloud components. If the person is the default administrator for your Oracle Cloud Infrastructure tenancy or a member of the [Administrators group](#), then they have the required permission. For information about restricting access to your networking components, see [Access Control](#).

The personnel should be familiar with the following concepts and definitions:

- [The fundamentals of Oracle Cloud Infrastructure](#)
- [The basic Networking service components](#)
- [General IPSec VPN tunnel functionality](#)

### CLOUD RESOURCES

Anything you provision on a cloud platform. For example, with Oracle Cloud Infrastructure, a cloud resource can refer to a VCN, compute instance, user, compartment, load balancer, or any other service component on the platform.

### ON-PREMISES

A widely used term in cloud technologies that refers to your traditional data center environments. On-premises can refer to a colocation scenario, a dedicated floor space, a dedicated data center building, or a desktop running under your desk.

### ORACLE CLOUD IDENTIFIER (OCID)

A unique identifier assigned to each resource that you provision on Oracle Cloud Infrastructure. The OCID is a long string that Oracle automatically generates. You can't

choose the value for an OCID or change a resource's OCID. For more information, see [Resource Identifiers](#).

### About the Oracle IPSec VPN

In general, IPSec can be configured in the following modes:

- **Transport mode:** IPSec encrypts and authenticates only the actual payload of the packet, and the header information stays intact.
- **Tunnel mode (supported by Oracle):** IPSec encrypts and authenticates the entire packet. After encryption, the packet is then encapsulated to form a new IP packet that has different header information.

Oracle Cloud Infrastructure supports only the tunnel mode for IPSec VPNs.

Each Oracle IPSec VPN consists of multiple redundant IPSec tunnels. For a given tunnel, you can use either Border Gateway Protocol (BGP) dynamic routing or static routing to route that tunnel's traffic. More details about routing follow.

IPSec VPN site-to-site tunnels offer the following advantages:

- Public internet lines are used to transmit data, so dedicated, expensive lease lines from one site to another aren't necessary.
- The internal IP addresses of the participating networks and nodes are hidden from external users.
- The entire communication between the source and destination sites is encrypted, significantly lowering the chances of information theft.



#### Note

IPv6 addressing is currently supported only in the US Government Cloud. For more information, see [IPv6 Addresses](#).

### Routing for the Oracle IPSec VPN

When you create an IPSec VPN, it has two redundant IPSec tunnels. Oracle encourages you to configure your CPE device to use both tunnels (if your device supports it). Note that in the past, Oracle created IPSec VPNs that had up to four IPSec tunnels.

The following two routing types are available, and you choose the routing type separately *for each tunnel* in the IPSec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPSec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

### Important Routing Details for an Oracle IPSec VPN

Here are important details to understand about routing for your IPSec VPN:

- **Routing choices:**
  - Originally, the Oracle IPSec VPN supported only static routing, and you were required to provide at least one static route for the overall IPSec connection.
  - Now two different types of routing are available (BGP and static routing), and you configure the routing type *per tunnel*. *Only one type of routing at a time is supported for a given tunnel.*
  - In general, Oracle encourages you to use the same routing type for all tunnels in your IPSec connection. Exception: if you're in the process of transitioning between static routing and BGP, then one tunnel might temporarily still use static routing while the other has already been switched to BGP.
  - When you create an IPSec connection, static routing is the default type of routing for all tunnels *unless you explicitly configure each tunnel to use BGP.*
- **Routing information required:**

- If you choose BGP, for each tunnel you must provide two IP addresses (one for each of the two BGP speakers in the tunnel's BGP session). The addresses must be in the encryption domain for the IPsec connection. You must also provide the BGP autonomous system number (BGP ASN) for your network.
  - If you choose static routing, you must provide at least one static route (maximum 10). The static routes are configured with the *overall IPsec connection*, so the same set of static routes are used for *all* tunnels in the IPsec connection that are configured to use static routing. You can change the static routes at any time after creating the IPsec connection. If you're doing PAT between your CPE device and VCN, the static route for the IPsec connection is the PAT IP address. See [Example Layout with PAT](#).
  - If you choose static routing, you may optionally provide an IP address for each end of the tunnel for the purposes of tunnel troubleshooting or monitoring.
  - If the tunnel is configured to use BGP, the IPsec connection's static routes are ignored. Any static routes associated with the IPsec connection are used for routing a given tunnel's traffic *only if that tunnel is configured to use static routing*. This is especially relevant if you have an IPsec VPN that uses static routing, but want to switch to using BGP.
- **Changing the routing:**
- If you want to change a tunnel from BGP to static routing, you must first ensure that the IPsec connection itself has at least one static route associated with it.
  - You can change an existing tunnel's routing type at any time (unless the tunnel is currently being provisioned by Oracle). While you change the routing, the tunnel remains up (its IPsec status does not change). However, traffic flowing through the tunnel is disrupted temporarily during reprovisioning and while you reconfigure your CPE device. For information about making changes to an existing IPsec VPN, see [Working with VPN Connect](#).
  - Because you configure the routing type separately for each tunnel, if you want to switch your IPsec VPN from static routing to BGP, you can do it one tunnel at a time. This avoids the entire IPsec VPN being down. For instructions, see [Changing from Static Routing to BGP Dynamic Routing](#).

### Route Advertisements and Path Preferences When You Have Multiple Connections

When you use BGP, the DRG attached to your VCN advertises routes to your CPE.

If you set up multiple connections between your on-premises network and VCN, you must understand what routes the DRG advertises and how to set path preferences to use your desired connection.

For important information, see [Routing Details for Connections to Your On-Premises Network](#).

### Preferring a Specific Tunnel in the IPSec VPN

Within an IPSec VPN, you can influence *which tunnel* is preferred. Here are items you can configure:

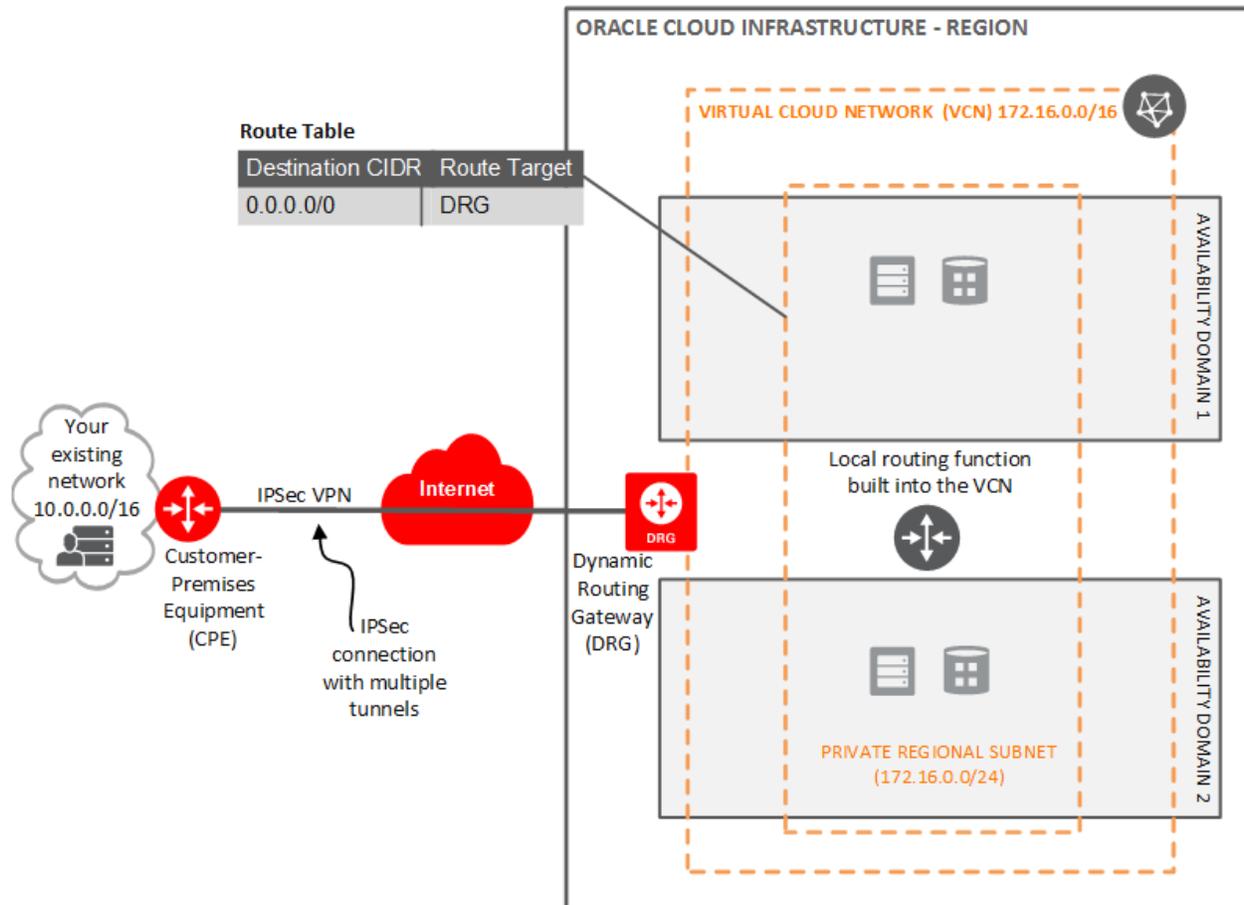
- **Your CPE's BGP local preference:** If you use BGP, you can configure the BGP local preference attribute on your CPE device to control which tunnel is preferred for connections initiated from your on-premises network to your VCN. Because Oracle generally uses asymmetric routing, you must configure other attributes if you want Oracle to respond on that same tunnel. See the next two items.
- **More specific routes on the preferred tunnel:** You can configure your CPE to advertise more specific routes for the tunnel that you want to prefer. Oracle uses the route with the [longest prefix match](#) when responding to or initiating connections.
- **AS path prepending:** BGP prefers the shortest AS path, so if you use BGP, you can use AS path prepending to control which tunnel has the shortest path for a given route. Oracle uses the shortest AS path when responding to or initiating connections.

### Overview of the IPSec VPN Components

If you're not already familiar with the basic Networking service components, see [Overview of Networking](#) before proceeding.

## CHAPTER 23 Networking

When you set up an IPsec VPN for your VCN, you must create several Networking components. You can create the components with either the Console or the API. See the following diagram and description of the components.



### CPE OBJECT

At your end of the IPsec VPN is the actual device in your on-premises network (whether hardware or software). The term *customer-premises equipment (CPE)* is commonly used in some industries to refer to this type of on-premises equipment. When setting up the

VPN, you must create a *virtual representation* of the device. Oracle calls the virtual representation a CPE, but this documentation typically uses the term *CPE object* to help distinguish the virtual representation from the actual CPE device. The CPE object contains basic information about your device that Oracle needs.

### **DYNAMIC ROUTING GATEWAY (DRG)**

At Oracle's end of the IPsec VPN is a virtual router called a dynamic routing gateway, which is the gateway into your VCN from your on-premises network. Whether you're using an IPsec VPN or [Oracle Cloud Infrastructure FastConnect private virtual circuits](#) to connect your on-premises network and VCN, the traffic goes through the DRG. For more information, see [Dynamic Routing Gateways \(DRGs\)](#).

A network engineer might think of the DRG as the *VPN headend*. After creating a DRG, you must *attach* it to your VCN, using either the Console or API. You must also add one or more route rules that route traffic from the VCN to the DRG. Without that DRG attachment and the route rules, traffic does not flow between your VCN and on-premises network. At any time, you can detach the DRG from your VCN but maintain all the remaining VPN components. You can then reattach the DRG, or attach it to another VCN.

### **IPSEC CONNECTION**

After creating the CPE object and DRG, you connect them by creating an IPsec connection, which you can think of as a parent object that represents the overall IPsec VPN. The IPsec connection has its own OCID. When you create this component, you configure the type of routing to use for each tunnel, and you provide any related routing information.

### **TUNNEL**

An IPsec tunnel is used to encrypt traffic between secure IPsec endpoints. Oracle creates two tunnels in each IPsec connection for redundancy. Each tunnel has its own OCID. Oracle recommends that you configure your CPE device to support both tunnels in case one fails or Oracle takes one offline for maintenance. Each tunnel has configuration information that your network engineer needs when [configuring your CPE device](#). This information includes an IP address and shared secret, as well as ISAKMP and IPsec parameters. For more information, see [Supported IPsec Parameters](#) and [Verified CPE Devices](#).

### **Access Control for the Components**

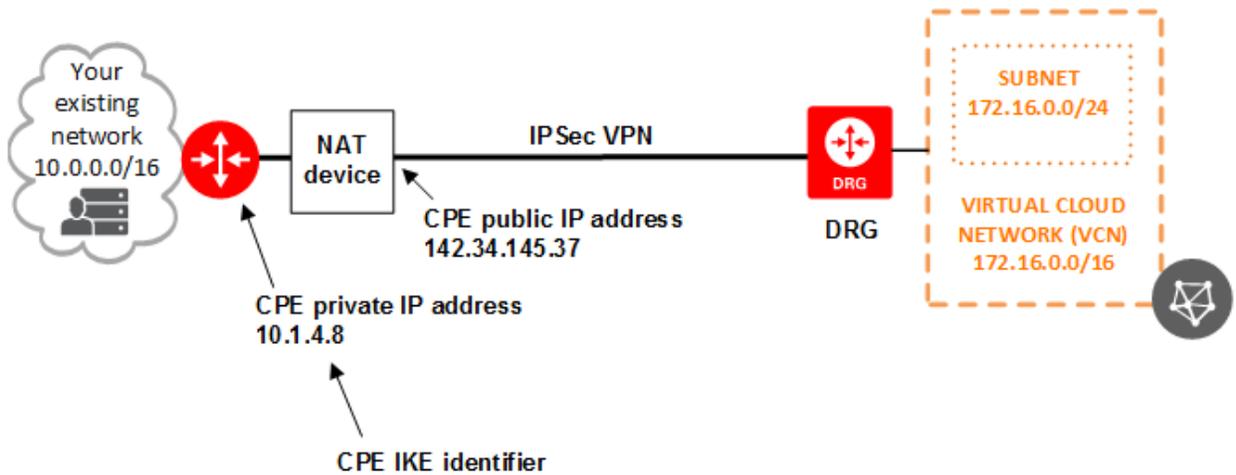
For the purposes of access control, when you set up the IPsec VPN, you must specify the compartment where you want each of the components to reside. If you're not sure which compartment to use, put all the components in the same compartment as the VCN. Note that the IPsec tunnels always reside in the same compartment as the parent IPsec connection. For information about compartments and restricting access to your networking components, see [Access Control](#).

### **Component Names and Identifiers**

You can optionally assign a descriptive name to each of the components when you create them. These names don't have to be unique, although it's a best practice to use unique names across your tenancy. Avoid including confidential information in the names. Oracle automatically assigns each component an OCID. For more information, see [Resource Identifiers](#).

### **If Your CPE Is Behind a NAT Device**

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



### Note

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

### About the Tunnel Shared Secret

Each tunnel has a shared secret. By default, Oracle assigns the shared secret to the tunnel unless you provide a shared secret yourself. You can provide a shared secret for each tunnel when you create the IPsec connection, or later after the tunnels are created. For the shared

secret, only letters, numbers, and spaces are allowed. If you change an existing tunnel's shared secret, the tunnel goes down while it is being reprovisioned.

For instructions, see [Changing the Shared Secret That an IPSec Tunnel Uses](#)

### Monitoring Your Connection

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring your connection, [VPN Connect Metrics](#).

### What's Next?

See these related topics:

- [VPN Connect Quickstart](#)
- [Setting Up VPN Connect](#)
- [Supported IPSec Parameters](#)
- [CPE Configuration](#)
- [Verified CPE Devices](#)
- [Routing Details for Connections to Your On-Premises Network](#)
- [Working with VPN Connect](#)
- [VPN Connect FAQ](#)
- [Using the API for VPN Connect](#)

### VPN Connect Quickstart

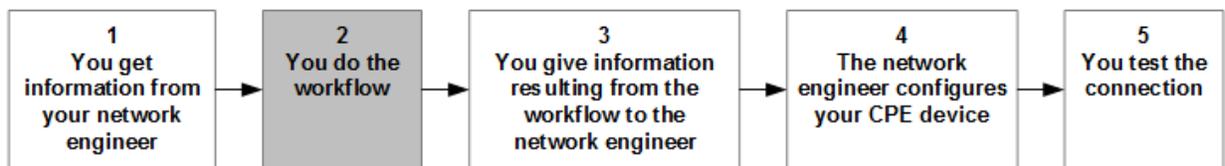
Using the VPN Connect workflow is the quickest way to set up an IPSec VPN between your on-premises network and your virtual cloud network (VCN). The workflow is a guided, step-by-step process in the Console that sets up the IPSec VPN plus related Networking service components.

### Purpose of the Workflow

VPN Connect involves setting up and configuring several Networking service components. The purpose of the workflow is to set up those components for you. In general, the workflow does the following:

- Uses a template with [assumptions](#) that will help you get started.
- Asks you for some basic network information.
- Sets up the Networking service components for you.

The workflow is a task within the overall process of setting up VPN Connect, which is illustrated in the following diagram. The workflow is the shaded box.



Notice that the overall process includes work by a network engineer in your organization. That engineer provides information that you, in turn, must supply during the workflow. The workflow returns information that the network engineer needs when configuring your CPE device.

The following short sections summarize each task.

#### Task 1: Information to get from your network engineer

- CPE device's public IP address.
- If the CPE is behind a NAT device, get the CPE IKE identifier. For more information, see [If Your CPE Is Behind a NAT Device](#).
- On-premises network routes.

- If you use BGP dynamic routing with the VPN:
  - Your network's BGP ASN
  - For each of the two IPsec tunnels that will be created, the pair of BGP IP addresses (with subnet mask) that you want to use for the inside tunnel interfaces at the ends of each tunnel. For example:
    - Tunnel 1: Inside tunnel interface - CPE: 10.0.0.8/31
    - Tunnel 1: Inside tunnel interface - Oracle: 10.0.0.9/31
    - Tunnel 2: Inside tunnel interface - CPE: 10.0.0.16/31
    - Tunnel 2: Inside tunnel interface - Oracle: 10.0.0.17/31
- The CIDR to use for the VCN. For the workflow, the allowed VCN size is /16 to /24. The CIDR must not overlap with your on-premises network.

### Task 2: Workflow

You walk through the workflow in the Console. For more information, see these sections:

- [Where to Access the Workflow](#)
- [What the Workflow Creates for You](#)

### Task 3: Information to give to your network engineer

- For each IPsec tunnel, the Oracle VPN IP address and shared secret.
- The supported IPsec parameter values.
- CPE-specific configuration information.

### Task 4: CPE configuration

Your network engineer takes the information you provide and configures your CPE device.

### Task 5: Testing

You and the network engineer test the connection and confirm that traffic is flowing.

### Workflow Assumptions

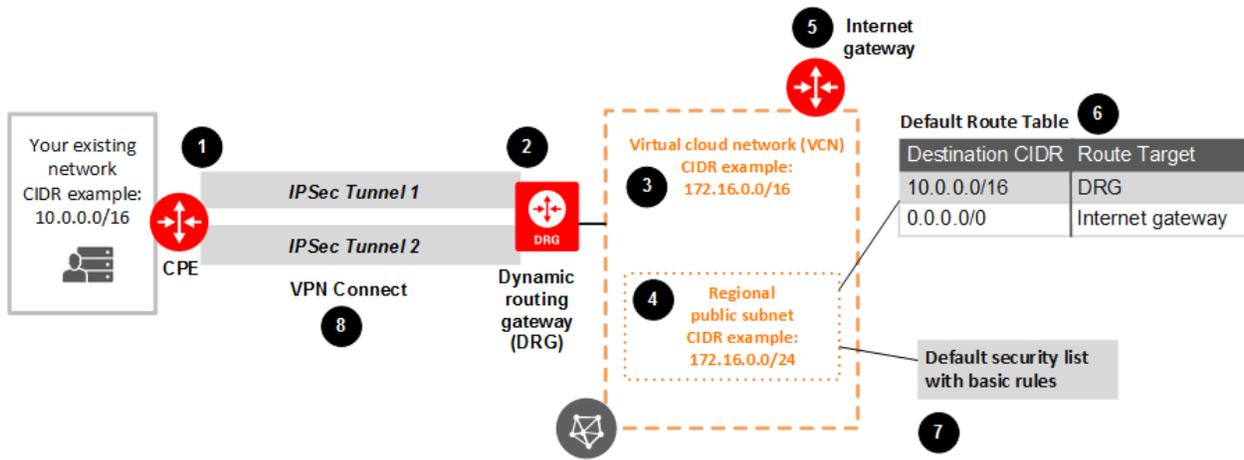
- **You do not already have a VCN:** The workflow automatically creates a [VCN](#) for you, along with related resources. If you instead have an existing VCN that you want to set up VPN Connect for, you can follow the step-by-step instructions in [Setting Up VPN Connect](#).
- **You want an internet gateway for easy initial access to the VCN:** The workflow automatically adds an [internet gateway](#) to make it easy for you to quickly create an instance in the VCN and connect to it over the internet. You can [delete this internet gateway](#) later if you don't want it.

### Alternative to the Workflow

If the workflow does not meet your specific needs (for example, if you already have a VCN), you can manually set up VPN Connect yourself. For step-by-step instructions, see [Setting Up VPN Connect](#).

### What the Workflow Creates for You

The workflow assumes that you start with only on-premises network and a CPE device. The workflow creates the numbered components in the diagram for you. The table describes each component.



Number	Component	Description	Can Use Existing One or Create New One?
1	CPE	A CPE is a <i>virtual representation</i> of your actual CPE device. This virtual representation contains basic information such as the CPE device's public IP address.	Yes, you can either use an existing CPE or the workflow creates a new one.
2	Dynamic routing gateway (DRG)	A <a href="#">DRG</a> is a <i>virtual representation</i> of the actual router at the Oracle end of your VPN Connect.	Yes. If you use an existing one, it must not already be attached to a VCN.
3	VCN	A <a href="#">VCN</a> is the extension of your on-premises network into the cloud. You can later <a href="#">add Compute instances</a> and other cloud resources to your VCN.	No. The workflow automatically creates a new VCN.
4	Subnet	A <a href="#">subnet</a> is a subdivision within the VCN. The workflow creates a regional public subnet. You can later add more subnets later if you like.	No. The workflow automatically creates the subnet in the new VCN.

Number	Component	Description	Can Use Existing One or Create New One?
5	Internet gateway	<p>An <a href="#">internet gateway</a> is a <i>virtual representation</i> of the actual router that gives your VCN access to the internet. Although this gateway is not necessary for VPN Connect, the workflow creates it to make it easy for you to quickly access any instances or other cloud resources you later create in the VCN. You can delete the internet gateway later if you like.</p>	No. The workflow automatically creates an internet gateway for the new VCN.
6	Default route table with rules	<p>The VCN automatically comes with a default <a href="#">route table</a>. The workflow configures the subnet to use this route table and adds two types of rules:</p> <ul style="list-style-type: none"><li>• One or more rules to route the desired VCN traffic to your on-premises network by way of the DRG. There's one rule per on-premises network route that you provide in the workflow.</li><li>• One rule to route the remaining VCN traffic to the internet by way of the internet gateway.</li></ul> <p>You can edit the rules or add more later if you want.</p>	No. The new VCN automatically comes with this component.

Number	Component	Description	Can Use Existing One or Create New One?
7	Default security list with rules	The VCN automatically comes with a <a href="#">default security list</a> . The workflow configures the subnet to use this security list, which automatically comes with default rules to enable basic traffic flow. The workflow also adds one or more rules to allow all types of traffic from your on-premises network. There's one rule per on-premises network route that you provide in the workflow. Notice that the security list does not include <a href="#">rules to allow ping</a> .	No. The new VCN automatically comes with this component.
8	VPN Connect IPSec tunnels	The workflow creates two <a href="#">IPSec tunnels</a> , each with specific configuration information that you must provide to your network engineer.  <b>Note:</b> The workflow uses IKEv1 for the tunnels. If you want to use IKEv2 instead, after creating the IPSec connection, edit each tunnel in the Oracle Console to use IKEv2. Then configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For more information, see <a href="#">Using IKEv2</a> .	No. The workflow automatically creates the tunnels.

## Where to Access the Workflow

1. In the Console, click the **Oracle Cloud** icon at the top of the page to go to the Console home page.

The page has a **Quick Actions** section to take you directly to common tasks.

2. Click the quick action for **Networking Solutions: Create an IPSec VPN connection**.

The workflow starts.

### Related Topics

- [VPN Connect Overview](#)
- [Setting Up VPN Connect](#)
- [Supported IPSec Parameters](#)
- [CPE Configuration](#)
- [Verified CPE Devices](#)
- [Working with VPN Connect](#)
- [VPN Connect FAQ](#)
- [Using the API for VPN Connect](#)

## Supported IPSec Parameters

This topic lists the supported phase 1 (ISAKMP) and phase 2 (IPSec) configuration parameters for VPN Connect. Oracle chose these values to maximize security and to cover a wide range of CPE devices. If your CPE device is not on the [list of verified devices](#), use the information here to configure your device.



### Important

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

## Supported Encryption Domain or Proxy ID

The values for the encryption domain (also known as a proxy ID, security parameter index (SPI), or traffic selector) depend on whether your CPE supports route-based tunnels or policy-based tunnels. For more information about the correct encryption domain values to use, see [Supported Encryption Domain or Proxy ID](#).

## Supported Parameters for the Commercial Cloud

This section lists the supported parameters if your VPN Connect is for the commercial cloud. For a list of the commercial cloud regions, see [Regions and Availability Domains](#).

For some parameters, Oracle supports multiple values, and the recommended one is highlighted in *red italics*.

Oracle supports the following parameters for IKEv1 or IKEv2. Check the documentation for your particular CPE to confirm which parameters the CPE supports for IKEv1 or IKEv2.

**Phase 1 (ISAKMP)**

Parameter	Options
ISAKMP Protocol	Version 1
Exchange type	Main mode
Authentication method	Pre-shared keys
Encryption algorithm	<i>AES-256-cbc</i> AES-192-cbc AES-128-cbc
Authentication algorithm	<i>SHA-2 384</i> SHA-2 256 SHA-1 (also called SHA or SHA1-96)
Diffie-Hellman group	group 1 (MODP 768) group 2 (MODP 1024) group 5 (MODP 1536) group 14 (MODP 2048) group 19 (ECP 256) <i>group 20 (ECP 384) *</i>
IKE session key lifetime	28800 seconds (8 hours)
* Group 20 will be supported in all Oracle Cloud Infrastructure regions very soon.	

**Phase 2 (IPSec)**

Parameter	Options
IPSec Protocol	ESP, tunnel mode
Encryption algorithm	<i>AES-256-gcm</i> AES-192-gcm AES-128-gcm AES-256-cbc AES-192-cbc AES-128-cbc
Authentication algorithm	If using GCM (Galois/Counter Mode), no authentication algorithm is required because authentication is included with GCM encryption.  If not using GCM, these are supported:  <i>HMAC-SHA-256-128</i>  HMAC-SHA-196
IPSec session key lifetime	3600 seconds (1 hour)
Perfect Forward Secrecy (PFS)	enabled, group 5

**Supported Parameters for the Government Cloud**

This section lists the supported parameters if your VPN Connect is for the Government Cloud. For more information, see [For All Government Cloud Customers](#).

For some parameters, Oracle supports multiple values, and the recommended one is highlighted in *red italics*.

Oracle supports the following parameters for IKEv1 or IKEv2. Check the documentation for your particular CPE to confirm which parameters the CPE supports for IKEv1 or IKEv2.

### Phase 1 (ISAKMP)

Parameter	Options
ISAKMP protocol	Version 1
Exchange type	Main mode
Authentication method	Pre-shared keys
Encryption algorithm	<i>AES-256-cbc</i> AES-192-cbc AES-128-cbc
Authentication algorithm	<i>SHA-2 384</i> SHA-2 256 SHA-1 (also called SHA or SHA1-96)
Diffie-Hellman group	group 14 (MODP 2048) group 19 (ECP 256) <i>group 20 (ECP 384) *</i>
IKE session key lifetime	28800 seconds (8 hours)
* Group 20 will be supported in all Oracle Cloud Infrastructure regions very soon.	

**Phase 2 (IPSec)**

Parameter	Options
IPSec protocol	ESP, tunnel mode
Encryption algorithm	<i>AES-256-gcm</i> AES-192-gcm AES-128-gcm AES-256-cbc AES-192-cbc AES-128-cbc
Authentication algorithm	If using GCM (Galois/Counter Mode), no authentication algorithm is required because authentication is included with GCM encryption.  If not using GCM, use HMAC-SHA-256-128.
IPSec session key lifetime	3600 seconds (1 hour)
Perfect Forward Secrecy (PFS)	enabled, group 14

## Supported Encryption Domain or Proxy ID

The IPSec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an IPsec tunnel, the packet is encrypted and sent to the other end of the tunnel.
- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

## Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPSec SAs. This pair is referred to as an *encryption domain*.



#### **Important**

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

If you use policy-based IPSec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

## Setting Up VPN Connect

This topic gives instructions for setting up VPN Connect for your VCN. For general information about IPsec VPNs, see [VPN Connect Overview](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Before You Get Started

To prepare, do these things first:

- Read this section: [Routing for the Oracle IPsec VPN](#)
- Answer these questions:

Question	Answer
What is your VCN's CIDR?	
<p>What is the public IP address of your CPE device? If you have multiple devices for redundancy, get the IP address for each.</p> <p><b>Note:</b> If your CPE device is behind a NAT device, see <a href="#">If Your CPE Is Behind a NAT Device</a> and also <a href="#">Information About Your CPE Device</a>.</p>	
Will you be doing port address translation (PAT) between each CPE device and your VCN?	
<p>What type of routing do you plan to use? If it's BGP dynamic routing, what are the BGP session IP addresses to use and the ASN of your network? The IP addresses must be part of the IPsec VPN's encryption domain.</p> <p>If static routing, what are the static routes for your on-premises network? See <a href="#">Routing for the Oracle IPsec VPN</a>.</p>	
Do you want to provide each tunnel's shared secret or let Oracle assign them? See <a href="#">About the Tunnel Shared Secret</a> .	

- Draw a diagram of your network layout (for examples, see the first task in [Example: Setting Up a Proof of Concept IPsec VPN](#)). Think about which parts of your on-premises network need to communicate with your VCN, and the reverse. Map out the [routing](#) and [security rules](#) that you need for your VCN.



**Tip**

If you have an existing Oracle IPsec VPN that uses static routing, you can [change the tunnels to instead use BGP dynamic routing](#).

### Overall Process

Here's the overall process for setting up an IPSec VPN:

1. **Complete the tasks listed in [Before You Get Started](#).**
2. **Set up the IPSec VPN components** (instructions in [Example: Setting Up a Proof of Concept IPSec VPN](#)):
  - a. Create your VCN.
  - b. Create a DRG.
  - c. Attach the DRG to your VCN.
  - d. Create a route table and route rule for the DRG.
  - e. Create a security list and required rules.
  - f. Create a subnet in the VCN.
  - g. Create a CPE object and provide your CPE device's public IP address.
  - h. Create an IPSec connection to the CPE object and provide required routing information.
3. **Have your network engineer configure your CPE device:** Your network engineer must configure your CPE device with information that Oracle provides during the previous steps. There is general information about the VCN, and specific information for each IPSec tunnel. This is the only part of the setup that you can't execute by using the Console or API. Without this configuration, traffic will not flow between your VCN and on-premises network. For more information, see [CPE Configuration](#).
4. **Validate connectivity.**

## Example: Setting Up a Proof of Concept IPsec VPN

**Tip**

Oracle offers a quickstart workflow to make it easier to set up VPN Connect. For more information, see [VPN Connect Quickstart](#).

This example scenario shows how to set up a single IPsec VPN with a simple layout that you might use for a proof of concept (POC). It follows tasks 1 and 2 in [Overall Process](#) and shows each component in the layout being created. For each task, there's a corresponding screenshot from the Console to help you understand what information is needed. For more complex layouts, see [Example Layout with Multiple Geographic Areas](#) or [Example Layout with PAT](#).

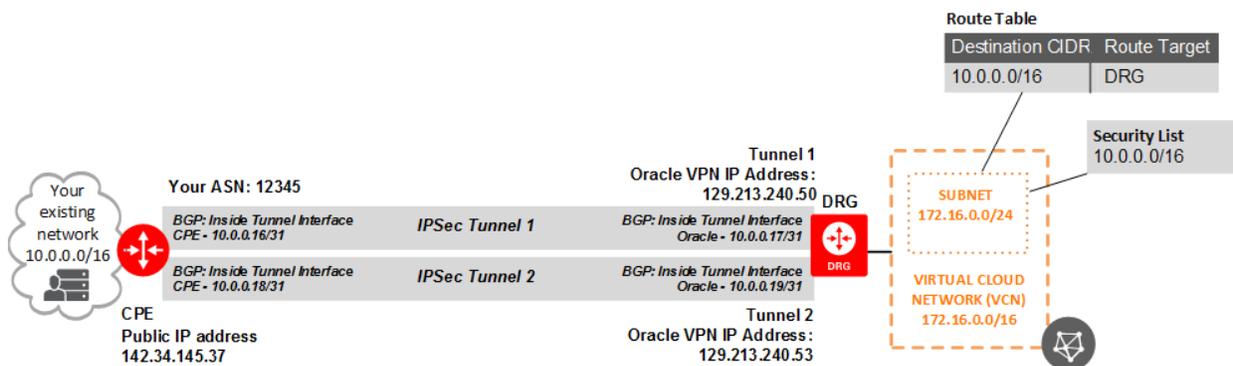
## Task 1: Gather information

Question	Answer
What is your VCN's CIDR?	172.16.0.0/16
What is the public IP address of your CPE device? If you have multiple devices for redundancy, get the IP address for each. <b>Note:</b> If your CPE device is behind a NAT device, see <a href="#">If Your CPE Is Behind a NAT Device</a> and also <a href="#">Information About Your CPE Device</a> .	142.34.145.37
Will you be doing port address translation (PAT) between each CPE device and your VCN?	No

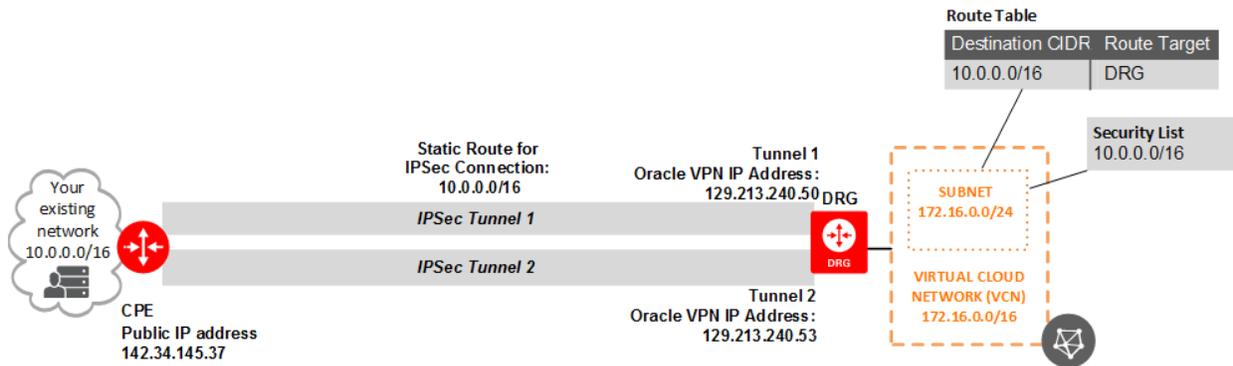
Question	Answer
<p>What type of routing do you plan to use? If it's BGP dynamic routing, what are the BGP session IP addresses to use and the ASN of your network? If static routing, what are the static routes for your on-premises network? See <a href="#">Routing for the Oracle IPSec VPN</a>.</p>	<p><b>BGP dynamic routing example:</b></p> <p>Tunnel 1:</p> <ul style="list-style-type: none"> <li>• BGP Inside tunnel interface - CPE: 10.0.0.16/31</li> <li>• BGP Inside tunnel interface - Oracle: 10.0.0.17/31</li> </ul> <p>Tunnel 2:</p> <ul style="list-style-type: none"> <li>• BGP Inside tunnel interface - CPE: 10.0.0.18/31</li> <li>• BGP Inside tunnel interface - Oracle: 10.0.0.19/31</li> </ul> <p>Network ASN:</p>

Question	Answer
	12345  <b>Static routing example:</b>  Use 10.0.0.0/16 for the static route for a simple POC.
Do you want to provide each tunnel's shared secret or let Oracle assign them? See <a href="#">About the Tunnel Shared Secret</a> .	Let Oracle assign

Here's an example diagram for task 1 if you plan to use BGP dynamic routing:



Here's an example diagram for task 1 if you plan to use static routing:



### Task 2a: Create the VCN

If you already have a VCN, skip to the next task.



#### Tip

When you use the Console to create a VCN, you can create only the VCN, or you can create the VCN with several related resources. This task creates only the VCN, and the subsequent tasks create the other required resources.

## CHAPTER 23 Networking

Create Virtual Cloud Network [help](#) [cancel](#)

CREATE IN COMPARTMENT  
Sandbox

NAME (OPTIONAL)  
MyExampleVCN

CREATE VIRTUAL CLOUD NETWORK ONLY  
 CREATE VIRTUAL CLOUD NETWORK PLUS RELATED RESOURCES

Creates a Virtual Cloud Network only. You'll still need to set up at least one Subnet, Gateway, and Route Rule to have a working Virtual Cloud Network.

CIDR BLOCK  
172.16.0.0/16  
Specified IP addresses: 172.16.0.0-172.16.255.255 (65,536 IP addresses)

DNS RESOLUTION  
 USE DNS HOSTNAMES IN THIS VCN ?  
Allows assignment of DNS hostname when launching an Instance

DNS LABEL  
myexamplevcn  
Only letters and numbers, starting with a letter. 15 characters max.

DNS DOMAIN NAME (READ-ONLY)  
myexamplevcn.oraclevcn.com

TAGS  
Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.  
[Learn more about tagging](#)

TAG NAMESPACE	TAG KEY	VALUE
None (apply a free-form tag)		

View detail page after this resource is created

**Create Virtual Cloud Network**

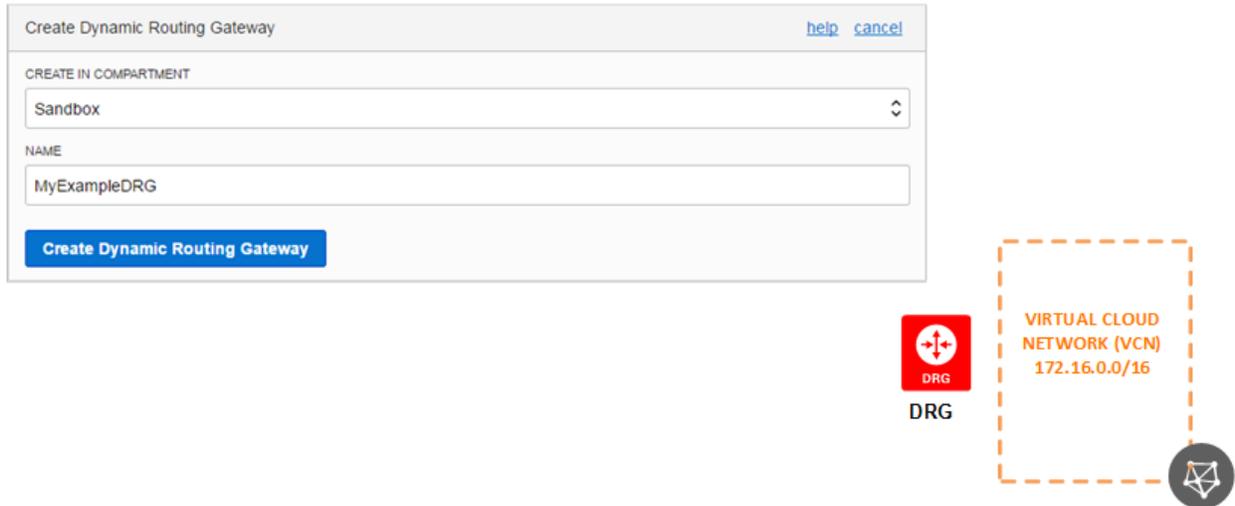


1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator. For more information, see [Access Control](#).
3. Click **Create Virtual Cloud Network**.
4. Enter the following values:

- **Create in Compartment:** Leave as is.
  - **Name:** A descriptive name for the cloud network. It doesn't have to be unique, and it can't be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create Virtual Cloud Network Only:** Select this option.
  - **CIDR Block:** A single, contiguous CIDR block for the cloud network (for example, 172.16.0.0/16). You *can't* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, use a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
5. You can provide values for the rest of the options, or you can ignore them:
- **DNS Resolution:** Required for assignment of DNS hostnames to hosts in the VCN, and required if you plan to use the VCN's default DNS feature (called the *Internet and VCN Resolver*). If the check box is selected, you can specify a DNS label for the VCN, or the Console will generate one for you. The dialog box automatically displays the corresponding **DNS Domain Name** for the VCN (*<VCN DNS label>.oraclevcn.com*). For more information, see [DNS in Your Virtual Cloud Network](#).
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click **Create Virtual Cloud Network**.

The VCN is created and displayed on the page. Ensure that it's done being provisioned before continuing.

### Task 2b: Create the DRG



The screenshot shows the 'Create Dynamic Routing Gateway' console form. The form has a title bar with 'Create Dynamic Routing Gateway' and links for 'help' and 'cancel'. Below the title bar, there are two input fields: 'CREATE IN COMPARTMENT' with a dropdown menu showing 'Sandbox', and 'NAME' with a text input field containing 'MyExampleDRG'. A blue button labeled 'Create Dynamic Routing Gateway' is positioned below the name field. To the right of the form, there is a diagram showing a red square icon with a white cross and the text 'DRG' below it. This icon is connected by a dashed orange line to a larger dashed orange rectangle representing a 'VIRTUAL CLOUD NETWORK (VCN)' with the IP address '172.16.0.0/16'. A small circular icon with a network diagram is at the bottom right of the VCN box.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click **Create Dynamic Routing Gateway**.
3. Enter the following values:
  - **Create in Compartment:** Leave as is (the VCN's compartment).
  - **Name:** A descriptive name for the DRG. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
4. Click **Create Dynamic Routing Gateway**.

The DRG is created and displayed on the page. Ensure that it's done being provisioned before continuing.



### Tip

You could also use this DRG as the gateway for [Oracle Cloud Infrastructure FastConnect](#), which is an alternative way to connect your on-premises network to your VCN.

### Task 2c: Attach the DRG to the VCN

Attach to Virtual Cloud Network [help](#) [close](#)

VIRTUAL CLOUD NETWORK in **Sandbox** ([Change Compartment](#))

My Example VCN

[Show Advanced Options](#)

**Attach to Virtual Cloud Network**



1. Click the name of the DRG that you just created.
2. Under **Resources**, click **Virtual Cloud Networks**.
3. Click **Attach to Virtual Cloud Network**.
4. Select the VCN. Ignore the section for advanced options, which is only for an advanced

routing scenario called transit routing, which is not relevant here.

5. Click **Attach**.

The attachment will be in the Attaching state for a short period before it's ready.

### Task 2d: Create a route table and route rule for the DRG

Although the VCN comes with a default route table (without any rules), in this task you create a custom route table with a route rule for the DRG. In this example, your on-premises network is 10.0.0.0/16. You create a route rule that takes any traffic destined for 10.0.0.0/16 and routes it to the DRG. When you create a subnet in task 2f, you associate this custom route table with the subnet.



#### Tip

If you already have an existing VCN with a subnet, you don't need to create a route table or subnet. Instead you can update the existing subnet's route table to include the route rule for the DRG.

Create Route Table [help](#) [cancel](#)

CREATE IN COMPARTMENT  
Sandbox

NAME  
MyExampleRouteTable

**Route Rules**

**Important:** For a route rule that targets a Private IP, you must first enable "Skip Source/Destination Check" on the VNIC that the Private IP is assigned to.

TARGET TYPE: Dynamic Routing Gateway  
DESTINATION CIDR BLOCK: 10.0.0/16  
COMPARTMENT: Sandbox  
TARGET DYNAMIC ROUTING GATEWAY: MyExampleDRG

Specified IP addresses: 10.0.0.0-10.0.255.255 (65,536 IP addresses)

+ Another Route Rule

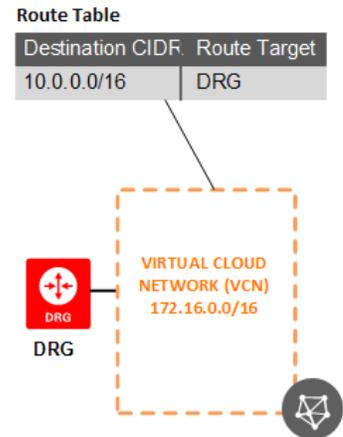
TAGS

Tagging is a metadata system that allows you to organize and track resources within your tenancy. Tags are composed of keys and values which can be attached to resources.

[Learn more about tagging](#)

TAG NAMESPACE: None (apply a free-form tag)  
TAG KEY:   
VALUE:

Create Route Table



1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click your VCN.
3. Click **Route Tables** to see your VCN's route tables.
4. Click **Create Route Table**.
5. Enter the following values:
  - **Name:** A descriptive name for the route table (for example, MyExampleRouteTable). The name doesn't have to be unique, and it can't be changed later in the Console (but you can change it in the API). Avoid entering confidential information.
  - **Create in compartment:** Leave as is.
  - Click **+ Additional Route Rule**, and enter the following:

- **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
- **Destination CIDR Block:** The CIDR for your on-premises network (10.0.0.0/16 in this example).
- **Tags:** Leave the tag information as is.

6. Click **Create Route Table**.

The route table is created and displayed on the page. However, the route table doesn't do anything unless you associate it with a subnet during subnet creation (see task 2f).

### Task 2e: Create a security list

By default, incoming traffic to the instances in your VCN is set to DENY on all ports and all protocols. In this task, you set up two ingress rules and one egress rule to allow basic required network traffic. Your VCN comes with a default security list with a set of default rules. However, in this task you create a separate security list with a more restrictive set of rules focused on traffic with your on-premises network. When you create a subnet in task 2f, you associate this security list with the subnet.



#### Tip

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

## CHAPTER 23 Networking

Create Security List help cancel

CREATE IN COMPARTMENT  
Sandbox  
dnx2 inst/Sandbox

SECURITY LIST NAME  
MyExampleSecurityList

### Allow Rules for Ingress

**Ingress Rule 1**

Allows TCP traffic for ports 22 SSH Remote Login Protocol

STATELESS [View Information](#)

SOURCE TYPE: CIDR  
SOURCE CIDR: 10.0.0.0/16  
IP PROTOCOL: TCP

SOURCE PORT RANGE (OPTIONAL): All  
DESTINATION PORT RANGE (OPTIONAL): 22

**Ingress Rule 2**

Allows ICMP traffic for: 3,4 Destination Unreachable, Fragmentation Needed and Don't Fragment was Set

STATELESS [View Information](#)

SOURCE TYPE: CIDR  
SOURCE CIDR: 10.0.0.0/16  
IP PROTOCOL: ICMP

TYPE AND CODE (OPTIONAL): 3,4

### Allow Rules for Egress

**Egress Rule 1**

Allows TCP traffic for ports all

STATELESS [View Information](#)

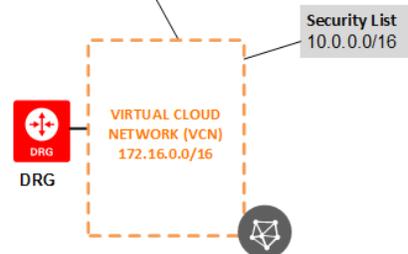
DESTINATION TYPE: CIDR  
DESTINATION CIDR: 10.0.0.0/16  
IP PROTOCOL: TCP

SOURCE PORT RANGE (OPTIONAL): All  
DESTINATION PORT RANGE (OPTIONAL): All

[+ Another Egress Rule](#)

### Route Table

Destination CIDR	Route Target
10.0.0.0/16	DRG



### Important

In the following procedure, ensure that the on-premises CIDR that you specify in the security list rules is the same (or smaller) than the CIDR that you specified in the route rule in the preceding task. Otherwise, traffic will be blocked by the security lists.

1. While still viewing your VCN, click **Security Lists** on the left side of the page.
2. Click **Create Security List**.
3. Enter the following values:
  - **Name:** A descriptive name for the security list. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it in the API). Avoid entering confidential information.
  - **Create in compartment:** Leave as is.
4. In the **Allow Rules for Ingress** section, click **Add Ingress Rule** and enter the following values for the ingress rule, which allows incoming SSH on TCP port 22 from your on-premises network:
  - **Source Type:** CIDR
  - **Source CIDR:** The CIDR for your on-premises network (10.0.0.0/16 in this example)
  - **IP Protocol:** TCP.
  - **Source Port Range:** Leave as is (the default All).
  - **Destination Port Range:** 22 (for SSH traffic).
5. In the **Allow Rules for Egress** section, click **Add Egress Rule** enter the following values for the egress rule, which allows outgoing TCP traffic on all ports to your on-premises network:
  - **Destination Type:** CIDR
  - **Destination CIDR:** The CIDR for your on-premises network (10.0.0.0/16 in this example).
  - **IP Protocol:** TCP.
  - **Source Port Range:** Leave as is (the default All).
  - **Destination Port Range:** Leave as is (the default All).
6. Leave the tag information as is.
7. Click **Create Security List**.

## CHAPTER 23 Networking

The security list is created and displayed on the page. However, the security list doesn't do anything unless you associate it with a subnet during subnet creation (see task 2f).

### Task 2f: Create a subnet

In this task, you create a subnet in the VCN. Typically a subnet has a CIDR block smaller than the VCN's CIDR. Any instances that you create in this subnet have access to your on-premises network. Oracle recommends using [regional subnets](#). Here you create a regional private subnet.

Create Subnet [help](#) [cancel](#)

If the Route Table, DHCP Options, or Security Lists are in a different Compartment than the Subnet, enable Compartment selection for those resources: [Click here](#)

NAME

SUBNET TYPE

REGIONAL (RECOMMENDED)  
Instances in the subnet can be created in any availability domain in the region. Useful for high availability.

AVAILABILITY DOMAIN-SPECIFIC  
Instances in the subnet can only be created in one availability domain in the region.

CIDR BLOCK  
  
Specified IP addresses: 172.16.0.0-172.16.0.255 (256 IP addresses)

ROUTE TABLE

SUBNET ACCESS

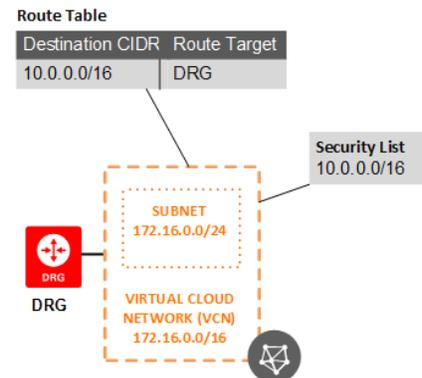
PRIVATE SUBNET  
Prohibit public IP addresses for instances in this Subnet

PUBLIC SUBNET  
Allow public IP addresses for instances in this Subnet

**Security Lists**

SECURITY LIST

[+ Additional Security List](#)



1. While still viewing your VCN, click **Subnets** on the left side of the page.
2. Click **Create Subnet**.
3. Enter the following values:
  - **Name:** A descriptive name for the subnet. It doesn't have to be unique, and it can't be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Regional or AD-specific subnet:** Select the radio button for **Regional**. Oracle recommends using [regional subnets](#).
  - **CIDR Block:** A single, contiguous CIDR block for the subnet (for example, 172.16.0.0/24). It must be within the cloud network's CIDR block and can't overlap with any other subnets. You *can't* change this value later. See [Allowed VCN Size and Address Ranges](#). For reference, use a [CIDR calculator](#).
  - **Enable IPv6 Address Assignment:** This option is available only if the VCN is in the US Government Cloud. For more information, see [IPv6 Addresses](#).
  - **Route Table:** The route table that you created earlier.
  - **Private Subnet:** Select this option. For more information, see [Access to the Internet](#).
  - **Use DNS Hostnames in this Subnet:** Leave as is (selected).
  - **DHCP Options:** The set of DHCP options to associate with the subnet. Select the default set of DHCP options for the VCN.
  - **Security Lists:** The security list that you created earlier.
  - **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
4. Click **Create Subnet**.

The subnet is created and displayed on the page. The basic VCN in this example is now set up, and you're ready to create the remaining components for the IPSec VPN.

Task 2g: Create a CPE object and provide your CPE device's public IP address

In this task, you create the CPE object, which is a virtual representation of your CPE device.

**Create Customer-Premises Equipment** [help](#) [cancel](#)

CREATE IN COMPARTMENT  
Sandbox

NAME  
MyExampleCPE

IP ADDRESS  
142.34.45.37

**Create**

**Route Table**

Destination CIDR	Route Target
10.0.0.0/16	DRG

**Security List**  
10.0.0.0/16

Your existing network  
10.0.0.0/16

CPE  
142.34.145.37

DRG

SUBNET  
172.16.0.0/24

VIRTUAL CLOUD NETWORK (VCN)  
172.16.0.0/16

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.
2. Click **Create Customer-Premises Equipment**.
3. Enter the following values:
  - **Create in Compartment:** Leave as is (the VCN's compartment).
  - **Name:** A descriptive name for the CPE object. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **IP Address:** The public IP address of the CPE device at your end of the VPN (see the list of information to gather in [Before You Get Started](#)).

- **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).

4. Click **Create**.

The CPE object is created and displayed on the page.

### Task 2h: Create an IPSec connection to the CPE object

In this task, you create the IPSec tunnels and configure the type of routing for them (BGP dynamic routing or static routing).



#### Tip

If you have an existing Oracle IPSec VPN that uses static routing, you can [change the tunnels to instead use BGP dynamic routing](#).

### For BGP dynamic routing

In this example, you configure both tunnels to use BGP dynamic routing.

## CHAPTER 23 Networking

Create IPsec Connection [help](#) [cancel](#)

CREATE IN COMPARTMENT  
Sandbox

NAME OPTIONAL  
MyExampleIPsecConnection

CUSTOMER-PREMISES EQUIPMENT COMPARTMENT  
Sandbox

CUSTOMER-PREMISES EQUIPMENT  
MyExampleCPE

DYNAMIC ROUTING GATEWAY COMPARTMENT  
Sandbox

DYNAMIC ROUTING GATEWAY  
MyExampleDRG

Show Advanced Options →

Create IPsec Connection

**CPE IKE Identifier** | Tunnel 1 | Tunnel 2 | Tags

NAME OPTIONAL  
Tunnel 1

PROVIDE CUSTOM SHARED SECRET ⓘ

SHARED SECRET

Only numbers, letters, and spaces are allowed.

ROUTING TYPE ⓘ

STATIC ROUTING  
Use static routes specified for the IPsec connection (above).

BGP DYNAMIC ROUTING  
Use BGP and the following settings for this tunnel:

BGP ASN  
12345

INSIDE TUNNEL INTERFACE - CPE ⓘ  
10.0.0.16/31

INSIDE TUNNEL INTERFACE - ORACLE ⓘ  
10.0.0.17/31

Your existing network  
10.0.0.0/16

**Your ASN: 12345**

**CPE**  
Public IP address  
142.34.145.37

BGP: Inside Tunnel Interface	IPsec Tunnel	BGP: Inside Tunnel Interface
CPE - 10.0.0.16/31	IPsec Tunnel 1	Oracle - 10.0.0.17/31
CPE - 10.0.0.18/31	IPsec Tunnel 2	Oracle - 10.0.0.19/31

**Tunnel 1**  
Oracle VPN IP Address:  
129.213.240.50

**Tunnel 2**  
Oracle VPN IP Address:  
129.213.240.53

Destination CIDR	Route Target
10.0.0.0/16	DRG

**SUBNET**  
172.16.0.0/24

**VIRTUAL CLOUD NETWORK (VCN)**  
172.16.0.0/16

**Security List**  
10.0.0.0/16

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPsec Connections**.
2. Click **Create IPsec Connection**.
3. Enter the following values:

- **Create in Compartment:** Leave as is (the VCN's compartment).
  - **Name:** Enter a descriptive name for the IPSec connection. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
  - **Customer-Premises Equipment Compartment:** Leave as is (the VCN's compartment).
  - **Customer-Premises Equipment:** Select the CPE object that you created earlier.
  - **Dynamic Routing Gateway Compartment:** Leave as is (the VCN's compartment).
  - **Dynamic Routing Gateway:** Select the DRG that you created earlier.
  - **Static Route CIDR:** Leave empty because this IPSec connection uses BGP dynamic routing. You configure the two tunnels to use BGP in subsequent steps.
4. Click **Show Advanced Options**.
  5. On the **CPE IKE Identifier** tab (optional): Oracle defaults to using the public IP address of the CPE. But if your [CPE is behind a NAT device](#), you might need to enter a different value. You can either enter the new value here, or [change the value](#) later.
  6. On the **Tunnel 1** tab (required):
    - **Name:** Enter a descriptive name for the tunnel. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
    - **Provide custom shared secret** (optional): By default, Oracle provides the shared secret for the tunnel. If you want to provide it instead, select this check box and enter the shared secret. You can [change the shared secret later](#).
    - **IKE Version:** The Internet Key Exchange (IKE) version to use for this tunnel. Only select [IKEv2](#) if your CPE supports it. You must also then configure the CPE to use only IKEv2 for this tunnel.
    - **Routing Type:** Select the radio button for **BGP Dynamic Routing**.

- **BGP ASN:** Enter your network's ASN.
  - **Inside Tunnel Interface - CPE:** Enter the BGP IP address with subnet mask (either /30 or /31) for the CPE end of the tunnel. For example: 10.0.0.16/31. The IP address must be part of the IPsec VPN's encryption domain.
  - **Inside Tunnel Interface - Oracle:** Enter the BGP IP address with subnet mask (either /30 or /31) for the Oracle end of the tunnel. For example: 10.0.0.17/31. The IP address must be part of the IPsec VPN's encryption domain.
7. On the **Tunnel 2** tab (required):
- **Name:** Enter a descriptive name for the tunnel. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
  - **Provide custom shared secret** (optional): By default, Oracle provides the shared secret for the tunnel. If you want to provide it instead, select this check box and enter the shared secret. You can [change the shared secret later](#).
  - **IKE Version:** The Internet Key Exchange (IKE) version to use for this tunnel. Only select [IKEv2](#) if your CPE supports it. You must also then configure the CPE to use only IKEv2 for this tunnel.
  - **Routing Type:** Select the radio button for **BGP Dynamic Routing**.
  - **BGP ASN:** Enter your network's ASN.
  - **Inside Tunnel Interface - CPE:** Enter the BGP IP address with subnet mask (either /30 or /31) for the CPE end of the tunnel. Use a different IP address than for tunnel 1. For example: 10.0.0.18/31. The IP address must be part of the IPsec VPN's encryption domain.
  - **Inside Tunnel Interface - Oracle:** Enter the BGP IP address with subnet mask (either /30 or /31) for the Oracle end of the tunnel. Use a different IP address than for tunnel 1. For example: 10.0.0.19/31. The IP address must be part of the IPsec VPN's encryption domain.
8. On the **Tags** tab (optional): Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).

9. Click **Create IPSec Connection**.

The IPSec connection is created and displayed on the page. It will be in the Provisioning state for a short period.

The displayed tunnel information includes:

- The Oracle VPN IP address (for the Oracle VPN headend).
- The tunnel's IPSec status (possible values are Up, Down, and Down for Maintenance). At this point, the status is Down. Your network engineer still must configure your CPE device.
- The tunnel's BGP status. At this point, the status is Down. Your network engineer still must configure your CPE device.

To view the tunnel's shared secret, click the tunnel to view its details, and then click **Show** next to **Shared Secret**.

10. Copy the Oracle VPN IP address and shared secret for each of the tunnels to an email or other location so you can deliver it to the network engineer who will configure your CPE device.

You can view this tunnel information here in the Console at any time.

You have now created all the components required for the IPSec VPN. But your network engineer must configure your CPE device before network traffic can flow between your on-premises network and VCN.

For static routing

The screenshot shows the 'Create IPSec Connection' console page. A red arrow points from the 'Show Advanced Options' link to the 'Tunnel 1' tab in the 'Advanced Options' panel. The 'Advanced Options' panel shows 'ROUTING TYPE' set to 'STATIC ROUTING' and 'INSIDE TUNNEL INTERFACE - CPE' set to '10.0.0.16/31'. Below the screenshot is a network diagram illustrating the setup:

- Your existing network:** 10.0.0.0/16, connected to a CPE with public IP address 142.34.145.37.
- Static Route for IPSec Connection:** 10.0.0.0/16.
- Tunnel 1:** Oracle VPN IP Address: 129.213.240.50, DRG.
- Tunnel 2:** Oracle VPN IP Address: 129.213.240.53, DRG.
- Virtual Cloud Network (VCN):** 172.16.0.0/16, containing a **SUBNET** 172.16.0.0/24.
- Security List:** 10.0.0.0/16.
- Route Table:**

Destination CIDR	Route Target
10.0.0.0/16	DRG

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

2. Click **Create IPSec Connection**.
3. Enter the following values:
  - **Create in Compartment:** Leave as is (the VCN's compartment).
  - **Name:** Enter a descriptive name for the IPSec connection. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
  - **Customer-Premises Equipment Compartment:** Leave as is (the VCN's compartment).
  - **Customer-Premises Equipment:** Select the CPE object that you created earlier.
  - **Dynamic Routing Gateway Compartment:** Leave as is (the VCN's compartment).
  - **Dynamic Routing Gateway:** Select the DRG that you created earlier.
  - **Static Route CIDR:** Enter at least one static route CIDR (see the list of information to gather in [Before You Get Started](#)). For this example, enter 10.0.0.0/16. You can enter up to 10 static routes, and you can [change the static routes](#) later if you like.
4. Click **Show Advanced Options**.
5. On the **CPE IKE Identifier** tab (optional): Oracle defaults to using the public IP address of the CPE. But if your [CPE is behind a NAT device](#), you might need to enter a different value. You can either enter the new value here, or [change the value](#) later.
6. On the **Tunnel 1** tab (optional):
  - **Tunnel Name:** Enter a descriptive name for the tunnel. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
  - **Provide custom shared secret** (optional): By default, Oracle provides the shared secret for the tunnel. If you want to provide it instead, select this check box and enter the shared secret. You can [change the shared secret later](#).

- **IKE Version:** The Internet Key Exchange (IKE) version to use for this tunnel. Only select [IKEv2](#) if your CPE supports it. You must also then configure the CPE to use only IKEv2 for this tunnel.
  - **Routing Type:** Leave the radio button for **Static Routing** selected.
  - **Inside Tunnel Interface - CPE** (optional): You can provide an IP address with subnet mask (either /30 or /31) for the CPE end of the tunnel for the purposes of tunnel troubleshooting or monitoring. For example: 10.0.0.16/31. The IP address must be part of the IPSec VPN's encryption domain.
  - **Inside Tunnel Interface - Oracle** (optional): You can provide an IP address with subnet mask (either /30 or /31) for the Oracle end of the tunnel for the purposes of tunnel troubleshooting or monitoring. For example: 10.0.0.17/31. The IP address must be part of the IPSec VPN's encryption domain.
7. On the **Tunnel 2** tab (optional):
- **Tunnel Name:** Enter a descriptive name for the tunnel. It doesn't have to be unique, and you can change it later if you like. Avoid entering confidential information.
  - **Provide custom shared secret** (optional): By default, Oracle provides the shared secret for the tunnel. If you want to provide it instead, select this check box and enter the shared secret. You can [change the shared secret later](#).
  - **IKE Version:** The Internet Key Exchange (IKE) version to use for this tunnel. Only select [IKEv2](#) if your CPE supports it. You must also then configure the CPE to use only IKEv2 for this tunnel.
  - **Routing Type:** Leave the radio button for **Static Routing** selected.
  - **Inside Tunnel Interface - CPE** (optional): You can provide an IP address with subnet mask (either /30 or /31) for the CPE end of the tunnel for the purposes of tunnel troubleshooting or monitoring. Use a different IP address than for tunnel 1. For example: 10.0.0.18/31. The IP address must be part of the IPSec VPN's encryption domain.

- **Inside Tunnel Interface - Oracle** (optional): You can provide an IP address with subnet mask (either /30 or /31) for the Oracle end of the tunnel for the purposes of tunnel troubleshooting or monitoring. Use a different IP address than for tunnel 1. For example: 10.0.0.19/31. The IP address must be part of the IPSec VPN's encryption domain.
8. **Tags:** Leave as is. You can add tags later if you want. For more information, see [Resource Tags](#).
  9. Click **Create IPSec Connection**.

The IPSec connection is created and displayed on the page. It will be in the Provisioning state for a short period.

The displayed tunnel information includes:

    - The Oracle VPN IP address (for the Oracle VPN headend).
    - The tunnel's IPSec status (possible values are Up, Down, and Down for Maintenance). At this point, the status is Down. Your network engineer still must configure your CPE device.

To view the tunnel's shared secret, click the tunnel to view its details, and then click **Show** next to **Shared Secret**.
  10. Copy the Oracle VPN IP address and shared secret for each of the tunnels to an email or other location so you can deliver it to the network engineer who will configure the CPE device.
- You can view this tunnel information here in the Console at any time.

You have now created all the components required for the IPSec VPN. But your network engineer must configure the CPE device before network traffic can flow between your on-premises network and VCN.

For more information, see [CPE Configuration](#).

### Task 3: Have your network engineer configure your CPE

Provide your network engineer with the following information:

- The VCN's OCID.
- The VCN's CIDR and subnet mask.
- The shared secret and Oracle VPN IP address for each tunnel (from task 2h).
- If using BGP: the BGP session information for each tunnel, which includes the IP addresses and Oracle BGP ASN. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- The information required to configure your particular CPE device, which is listed in the [topic for each verified type of CPE device](#).
- The [general IPSec parameters](#) that Oracle supports.
- A link to this topic: [CPE Configuration](#).



### Important

Be sure to have your network engineer configure your CPE device to support both of the tunnels in case one fails or Oracle takes one down for maintenance. If you're using BGP, see [Important Routing Details for an Oracle IPSec VPN](#).

### Task 4: Validate connectivity

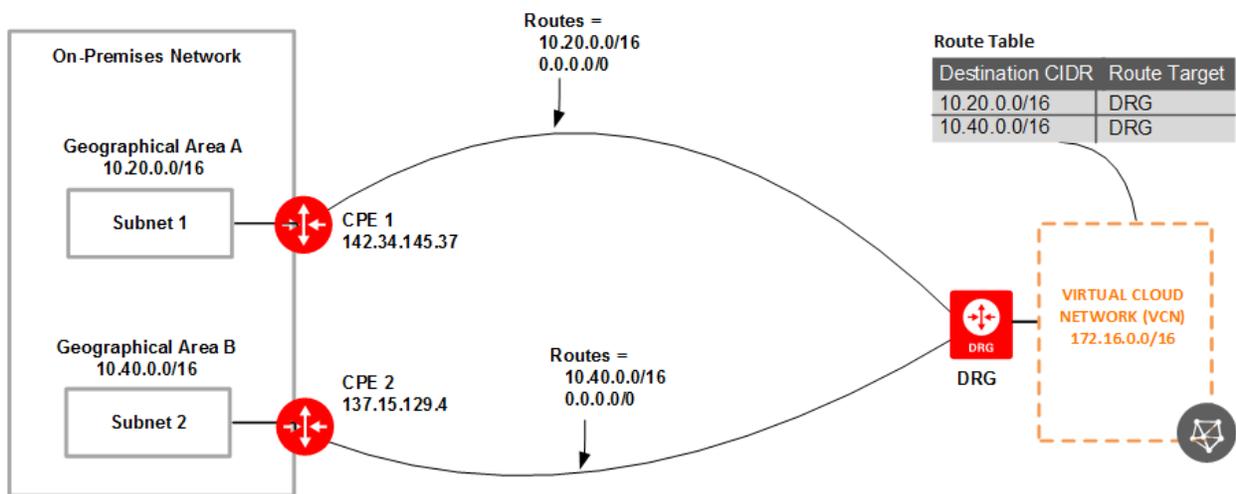
After the network engineer configures your CPE device, you can confirm that the tunnel's IPSec status is Up and green. Next, you can create a Linux instance in the subnet in your VCN. You should then be able to use SSH to connect to the instance's private IP address from a host in your on-premises network. For more information, see [Creating an Instance](#).

## Example Layout with Multiple Geographic Areas

The following diagram shows an example with this configuration:

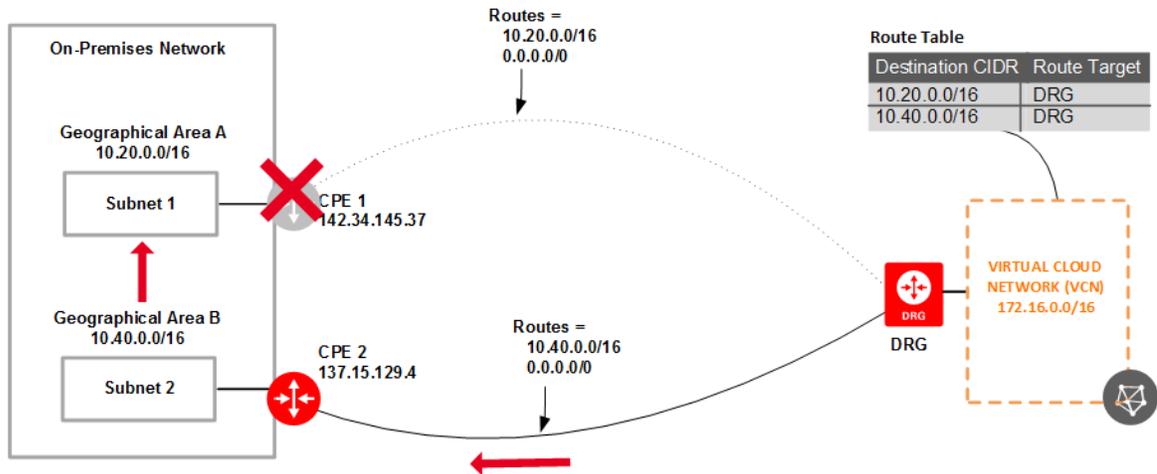
- Two networks in separate geographical areas that each connect to your VCN
- A single CPE device in each area
- Two IPsec VPNs (one for each CPE device)

Notice that each IPsec VPN has two routes associated with it: one for the particular geographical area's subnet, and a default 0.0.0.0/0 route. Oracle learns about the available routes for each tunnel either through BGP (if the tunnels use BGP), or because you've set them as static routes for the IPsec connection (if the tunnels use static routing).

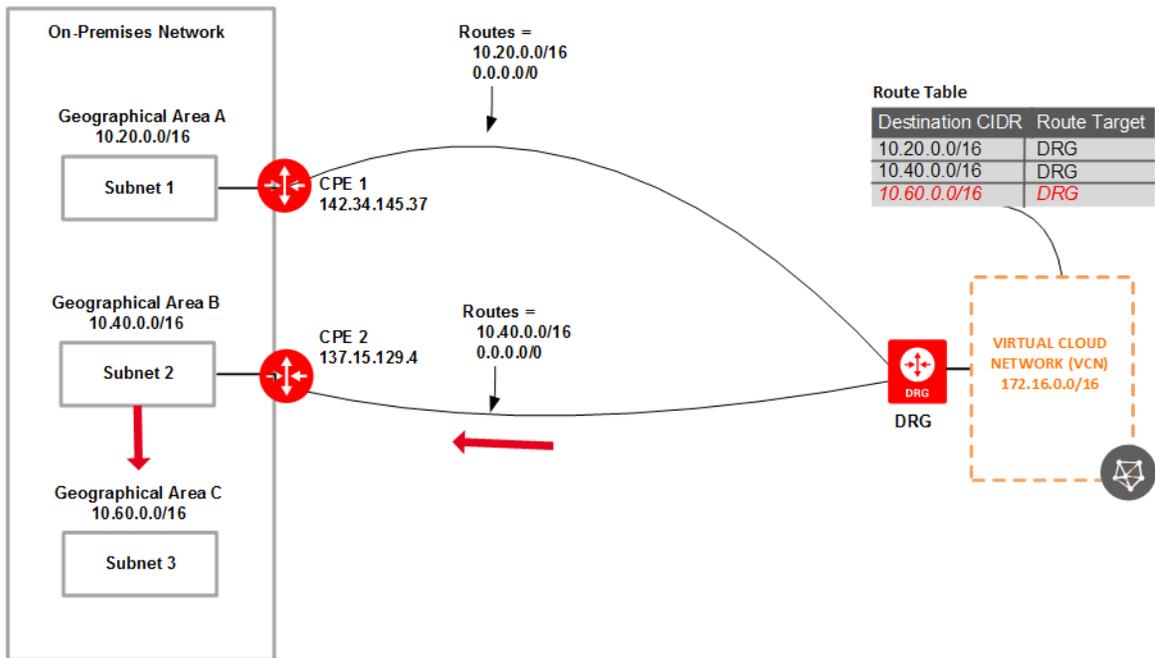


Following are some examples of situations in which the 0.0.0.0/0 route can provide flexibility:

- Assume that the CPE 1 device goes down (see the next diagram). If Subnet 1 and Subnet 2 can communicate with each other, your VCN could still reach the systems in Subnet 1 because of the 0.0.0.0/0 route that goes to CPE 2.



- Assume that your organization adds a new geographical area with Subnet 3 and initially just connects it to Subnet 2 (see the next diagram). If you added a route rule to your VCN's route table for Subnet 3, the VCN could reach systems in Subnet 3 because of the 0.0.0.0/0 route that goes to CPE 2.



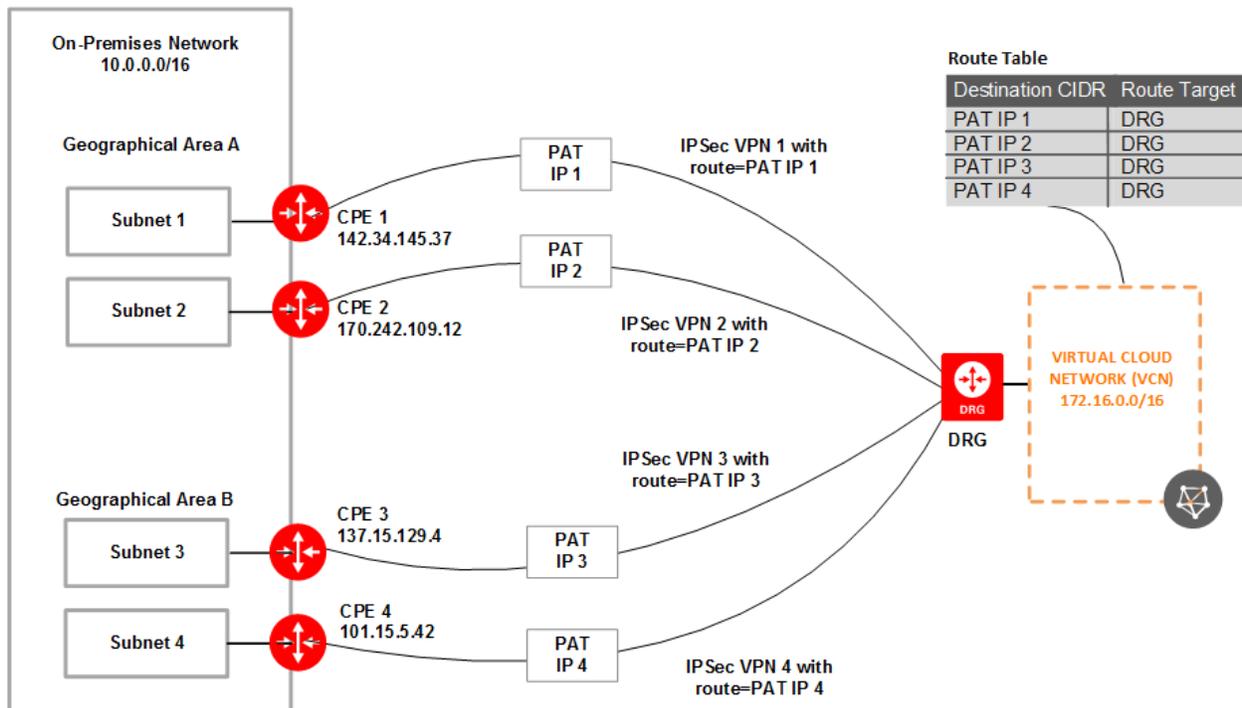
### Example Layout with PAT

The following diagram shows an example with this configuration:

- Two networks in separate geographical areas that each connect to your VCN
- Redundant CPE devices (two in each geographical area)
- Four IPsec VPNs (one for each CPE device)
- Port address translation (PAT) for each CPE device

For each of the four IPsec connections, the route that Oracle needs to know about is the PAT IP address for the specific CPE device. Oracle learns about the PAT IP address route for each tunnel either through BGP (if the tunnels use BGP), or because you've set the relevant address as a static route for the IPsec connection (if the tunnels use static routing).

When you set up the route rules for the VCN, you specify a rule for each PAT IP address (or an aggregate CIDR that covers them all) with your DRG as the rule's target.



### What's Next?

See these related topics and procedures:

- [VPN Connect Quickstart](#)
- [CPE Configuration](#)

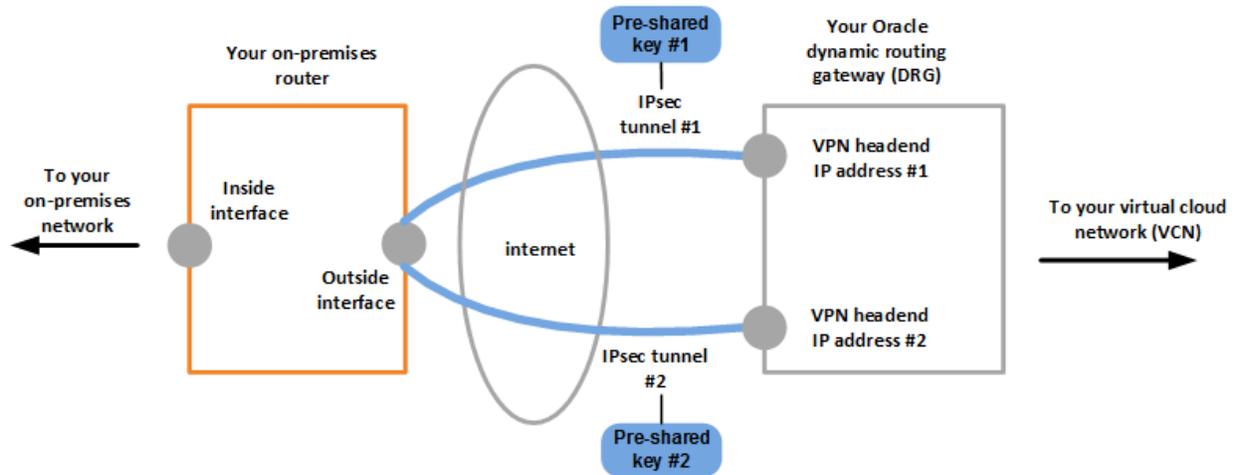
- [Verified CPE Devices](#)
- [Changing from Static Routing to BGP Dynamic Routing](#)
- [Working with VPN Connect](#)
- [VPN Connect FAQ](#)
- [Using the API for VPN Connect](#)
- [VPN Connect Metrics](#)
- [VPN Connect Troubleshooting](#)

### CPE Configuration

This topic is for network engineers. It explains how to configure the on-premises device (the customer-premises equipment, or CPE) at your end of the IPsec VPN so traffic can flow between your on-premises network and virtual cloud network (VCN). See these related topics:

- [Overview of Networking](#): For general information about the parts of a VCN
- [VPN Connect](#): For various topics about IPsec VPNs
- [Verified CPE Devices](#): For a list of CPE devices Oracle has verified

The following figure shows the basic layout of the IPsec VPN connection.



### Requirements and Prerequisites

There are several requirements and prerequisites to be aware of before moving forward.

#### Routing Considerations

For important details about routing for your IPsec VPN see [Routing for the Oracle IPsec VPN](#).

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

If you use BGP dynamic routing with your IPsec VPN, you can configure routing so that Oracle prefers one tunnel over the other.

Note that the [Cisco ASA policy-based configuration](#) uses a single tunnel.

### Creation of Cloud Network Components

You or someone in your organization must have already used the Oracle Console to create a VCN and an IPSec connection, which consists of multiple IPSec tunnels for redundancy. You must gather the following information about those components:

- VCN OCID: The VCN [OCID](#) is a unique Oracle Cloud Infrastructure identifier that has a UUID at the end. You can use this UUID or any other string that helps you identify this VCN in the device configuration and doesn't conflict with other object-group or access-list names.
- VCN CIDR
- VCN CIDR subnet mask
- For each IPSec tunnel:
  - The IP address of the Oracle IPSec tunnel endpoint (the VPN headend)
  - The shared secret

### Information About Your CPE Device

You also need some basic information about the inside and outside interfaces of your on-premises device (your CPE). For a list of the required information for your particular CPE, see the links in this list: [Verified CPE Devices](#).

Oracle recommends that you disable NAT-T at your CPE when establishing IPSec tunnels with Oracle Cloud Infrastructure. Unless you have multiple CPEs sharing the same NAT IP, NAT-T is not required.

If your CPE is behind a NAT device, you can provide Oracle with your CPE's IKE identifier. For more information, see [If Your CPE Is Behind a NAT Device](#).

### Route-Based Versus Policy-Based IPSec

The Oracle VPN headends use route-based tunnels, but can work with policy-based tunnels with some caveats listed in the following section.

Oracle supports only a single encryption domain (also known as a proxy ID, security parameter index (SPI), or traffic selector).

For more information, see [Supported Encryption Domain or Proxy ID](#).

### IPSec VPN Best Practices

- **Configure all tunnels for every IPSec connection:** Oracle deploys multiple IPSec headends for all your connections to provide high availability for your mission-critical workloads. Configuring all the available tunnels is a key part of the "Design for Failure" philosophy. (Exception: [Cisco ASA policy-based configuration](#), which uses a single tunnel.)
- **Have redundant CPEs in your on-premises locations:** Each of your sites that connects with IPSec to Oracle Cloud Infrastructure should have redundant CPE devices. You add each CPE to the Oracle Cloud Infrastructure Console and create a separate IPSec connection between your dynamic routing gateway (DRG) and each CPE. For each IPSec connection, Oracle provisions two tunnels on geographically redundant IPSec headends. Oracle may use any tunnel that is "up" to send traffic back to your on-premises network. For more information, see [Routing for the Oracle IPSec VPN](#).
- **Consider backup aggregate routes:** If you have multiple sites connected via IPSec VPNs to Oracle Cloud Infrastructure, and those sites are connected to your on-premises backbone routers, consider configuring your IPSec connection routes with both the local site aggregate route as well as a default route.

Note that the DRG routes learned from the IPSec connections are only used by traffic you route from your VCN to your DRG. The default route will only be used by traffic sent to your DRG whose destination IP address does not match the more specific routes of any of your tunnels.

### Confirming the Status of the Connection

After you configure the IPSec connection, you can test the connection by launching an instance into the VCN and then pinging it from your on-premises network. For information about launching an instance, see [Launching an Instance](#). To ping the instance, the VCN's security rules must [allow ping traffic](#).

You can get the status of the IPSec tunnels in the API or Console. For instructions, see [To view the status and configuration information for the IPSec tunnels](#).

### Device Configurations

For links to the specific configuration information for each verified CPE device, see [Verified CPE Devices](#).

### Verified CPE Devices

The following devices or software have been verified for use with VPN Connect.



#### Note

Oracle provides configuration instructions for the vendors and devices in the following table. Make sure to use the configuration instructions for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the following table, or if you're already familiar with configuring your device for IPSec, see the list of [supported IPSec parameters](#) and consult your vendor's documentation for assistance.

Vendor	Device	Minimum Software Version	Configuration
Check Point	2200 or Open Server	R80.20	<a href="#">Check Point Configuration Options</a>
Cisco	ASA	9.7.1 ( <a href="#">recommended</a> )	<a href="#">Cisco ASA Configuration Options</a>
Cisco	2921	IOS version 15.4(3)M3	<a href="#">Cisco IOS</a>
FortiGate	FortiGate-VM	6.0.4	<a href="#">FortiGate</a>
Juniper	MX 240	JunOS 15.1	<a href="#">Juniper MX</a>
Juniper	SRX 240	JunOS 11.0	<a href="#">Juniper SRX</a>
Libreswan (or Openswan)		3.18	<a href="#">Libreswan</a>
NEC	IX3315	10.1.16	<a href="#">NEC IX Series</a>
NEC	IX2106	10.1.16	<a href="#">NEC IX Series</a>
Palo Alto	PA-500	PanOS version 8.0.0	<a href="#">Palo Alto</a>
WatchGuard	Firebox	Fireware v12	<a href="#">WatchGuard</a>
Yamaha	RTX1210	Firmware Rev.14.01.28	<a href="#">Yamaha RTX Series</a>
Yamaha	RTX830	Firmware Rev.15.02.03	<a href="#">Yamaha RTX Series</a>

## Check Point Configuration Options

Choose the configuration that suits your situation:

- [Check Point: Route-Based](#)
- [Check Point: Policy-Based](#)

### Check Point: Route-Based

This topic provides a route-based configuration for Check Point CloudGuard. The instructions were validated with Check Point CloudGuard version R80.20.

This topic is for route-based (VTI-based) configuration. If you instead want policy-based configuration, see [Check Point: Policy-Based](#).

Check Point experience is required. This topic does not include how to add Check Point CloudGuard Security Gateway to Check Point CloudGuard Security Manager. For more information about using Check Point products, see the Check Point documentation.



#### **Important**

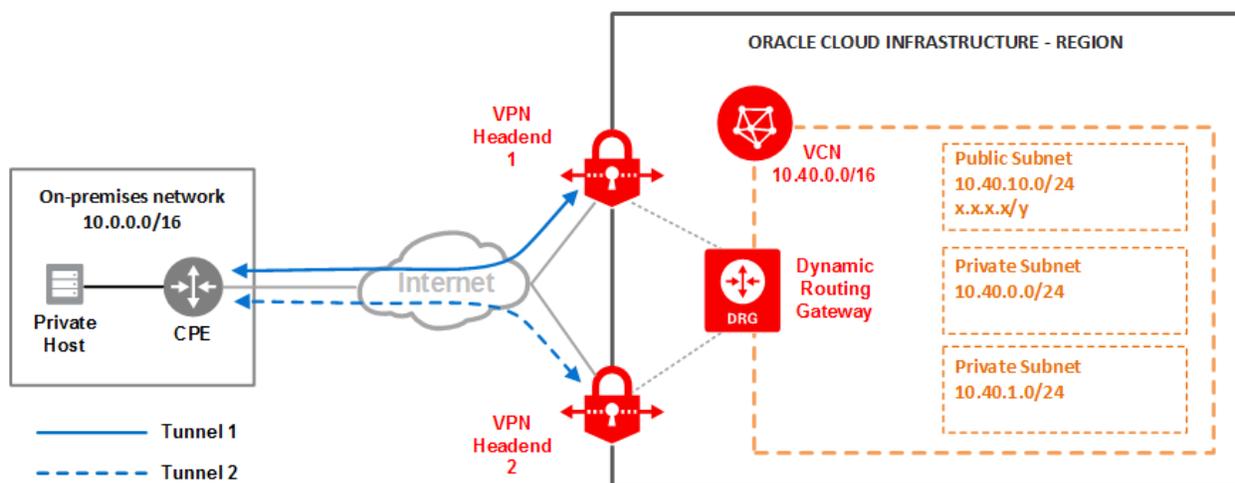
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPSec, see the list of [supported IPSec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPSec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



### Best Practices

This section covers general best practices and considerations for using VPN Connect.

#### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

#### HAVE REDUNDANT CPES IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two

tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

## Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

## Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

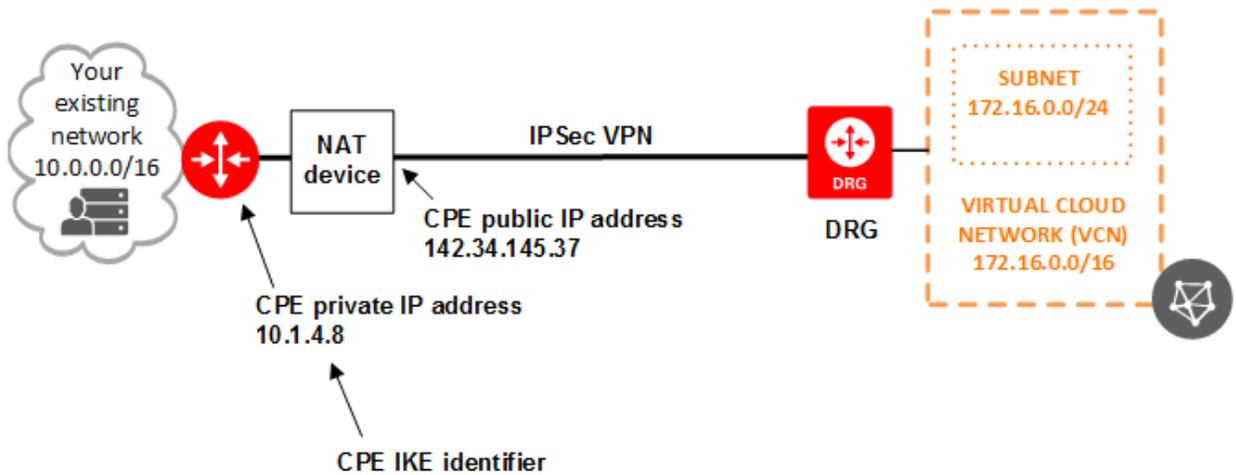
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

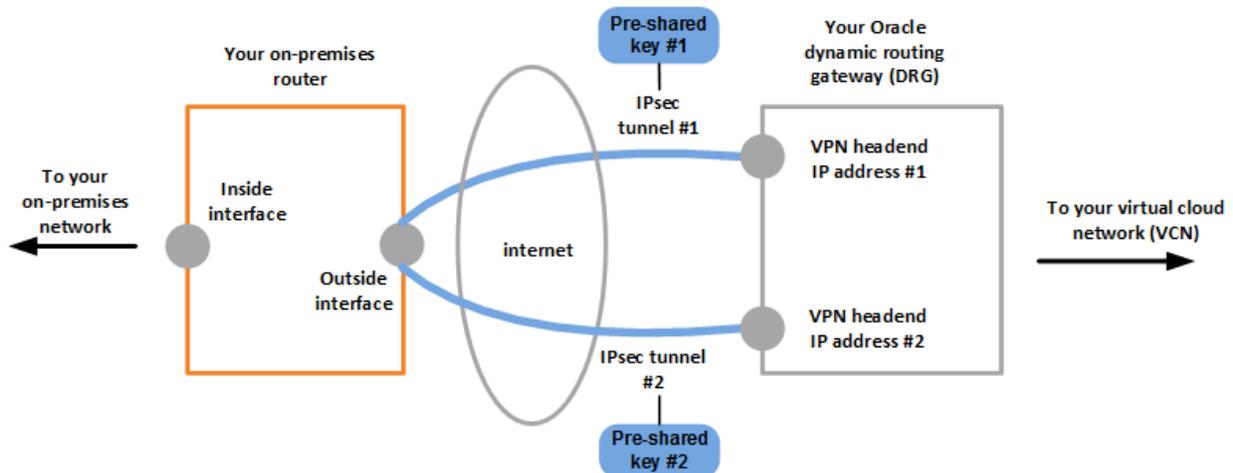
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration (Route-Based)

**Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



#### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to

## CHAPTER 23 Networking

use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

If you want to use IKEv2, there's a variation on one of the tasks presented in the next section. Specifically, in [task 4](#), when configuring encryption, select **IKEv2 only** for the encryption method.

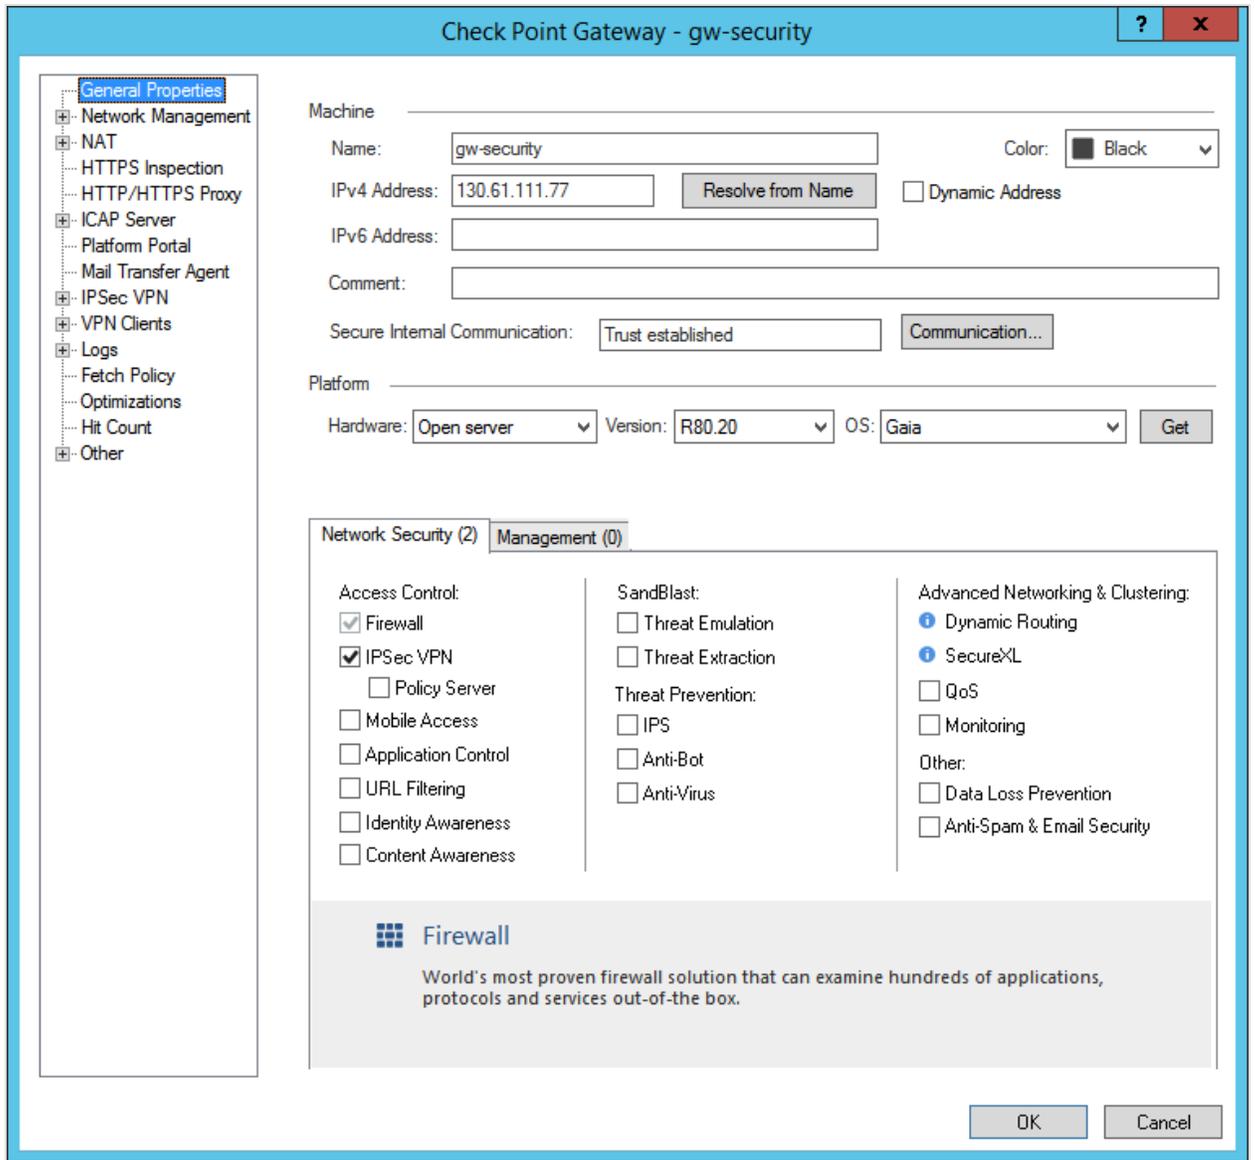
### CONFIGURATION PROCESS

#### Task 1: Install IPSec VPN on Check Point CloudGuard Security Gateway

**Prerequisite:** Before starting, add Check Point CloudGuard Security Gateway to Check Point CloudGuard Security Manager. Also establish the Secure Internal Communication (SIC) so you can configure the IPSec tunnel by using the Check Point Smart Console. For instructions to add the Security Gateway to CloudGuard or to establish the SIC, see the Check Point documentation.

Status	Name	IP	Version	Active Blades	Hardware	CPU Usage
✓	 gw-manager		R80.20	 	Open server	 4%
✓	 gw-security	130.61.111.77	R80.20	 	Open server	 0%

1. Install the IPSec VPN module. Oracle recommends that you also install the Monitoring module for traffic analysis.

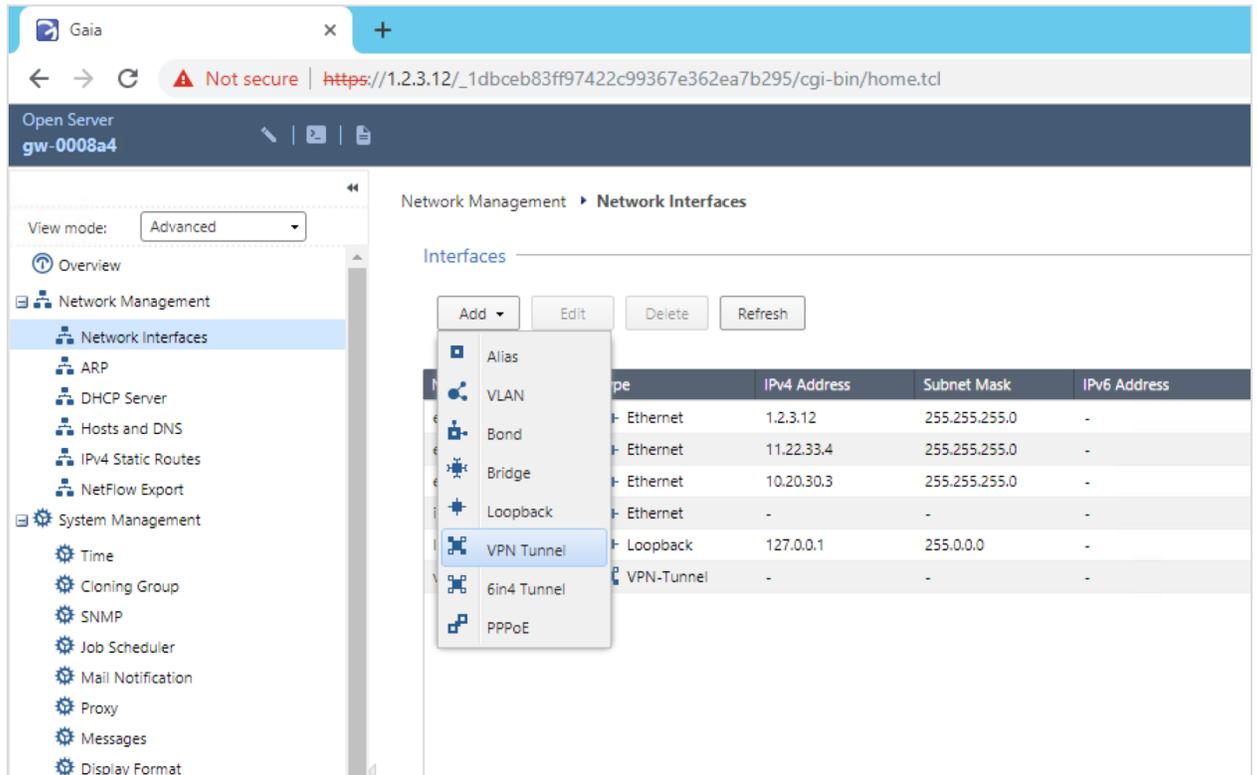


2. Click **OK** to save your changes.

## Task 2: Create the VTI interface from GAIA

In this task, you configure a VTI interface that passes traffic by using routing rules from the VTI interface to the newly created IPsec tunnel.

1. Log in to the GAIA portal using the Check Point CloudGuard Security Gateway public or private IP address.
2. On the GAIA portal, select the **Advanced** view.
3. Under **Network Management**, go to **Network Interfaces**.
4. Click **Add**, and then click **VPN Tunnel**.



5. Specify the following items:

- **VPN Tunnel ID:** A number that will be added to the VTI interface called vpnt\*, where the asterisk is the VPN tunnel ID number specified . For VPN tunnel ID = 1, the interface is labeled vpnt1.
- **Peer:** The name of the interoperable device that you created earlier for the IPsec tunnel. In this case, the name is OCI-VPN\_BGP1. **Important:** If the name you specify here does not match the name of the interoperable device, traffic does not flow through the IPsec tunnel.
- **Numbered:** Select **Numbered** to create a numbered interface.
- **Local Address:** The local IP address that was specified in the Oracle Console as the **Inside Tunnel Interface - CPE**.
- **Remote Address:** The remote IP address that was specified in the Oracle Console as the **Inside Tunnel Interface - Oracle**.

**Add VPN Tunnel**

Type: VPN-Tunnel

Enable:

Comment:

**VPN Tunnel**

VPN Tunnel ID:

Peer:

**VPN Tunnel Type**

Numbered  Unnumbered

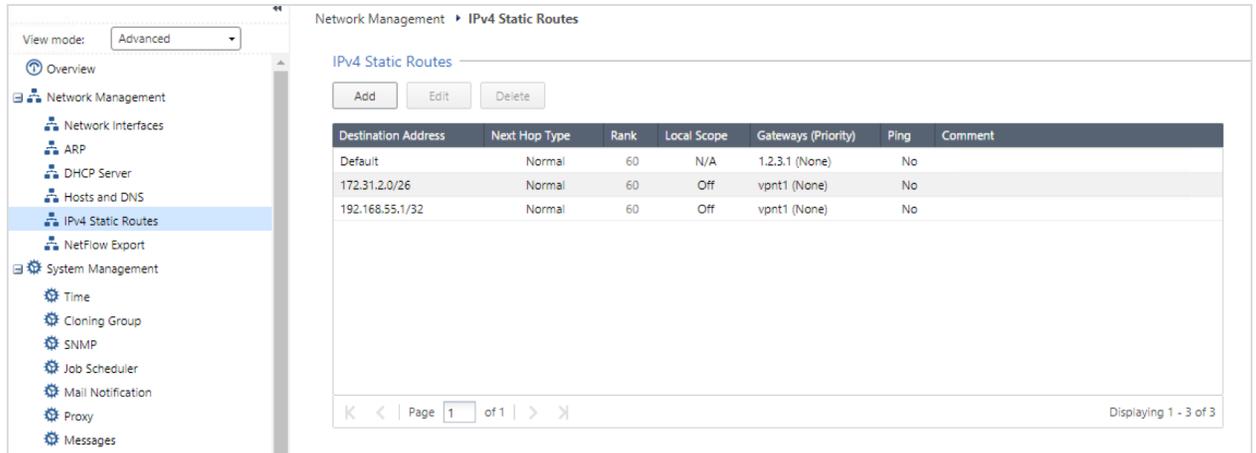
Local Address:  Physical device:

Remote Address:

OK Cancel

6. Click **OK**.
7. Under **Network Management**, go to **IPv4 Static Routes**.
8. Specify the following items:
  - **Static route for the Oracle IP address:** Add an IP address with /32 mask for the remote IP address that was specified in the Oracle Console as the **Inside Tunnel Interface - Oracle**.
  - **Static routes to the VCN subnets:** If you're using static routing for this IPsec connection to Oracle, add at least one subnet for the Oracle VCN to be reached

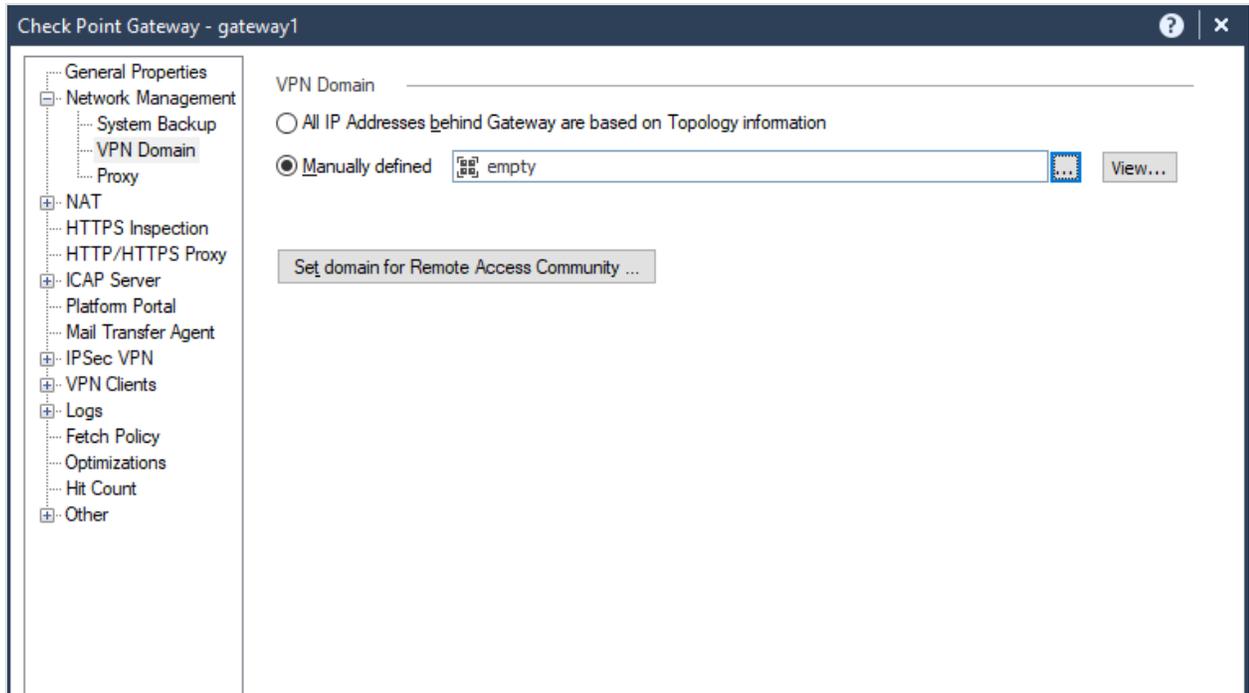
through the IPsec tunnel. The following screenshot shows a static route to 172.31.2.0/26. If you're using BGP for this IPsec connection to Oracle, skip this item because those routes are learned through BGP (see the next section).



Now all traffic with a specific destination learned from a static route will pass through the newly created IPsec tunnel.

9. Get the interfaces and verify that the VPN tunnel is in the list:
  - a. In the Smart Console, go to **Gateways & Servers**.
  - b. Select the **Check Point Security Gateways**, and double-click.
  - c. Under **General Properties**, on the **Network Management** page, select **Get Interfaces**.  
The VPN tunnel interface should appear in the list.
10. To force a route-based VPN to take priority, create an empty group and assign it to the VPN domain:
  - a. On the **VPN Domain** page, select **Manually defined**, and then select **Create empty group**.
  - b. Click **New**, select **Group**, and then select **Simple Group**.

- c. Enter an **Object Name**, and then click **OK**. Do not assign any objects to this empty group.

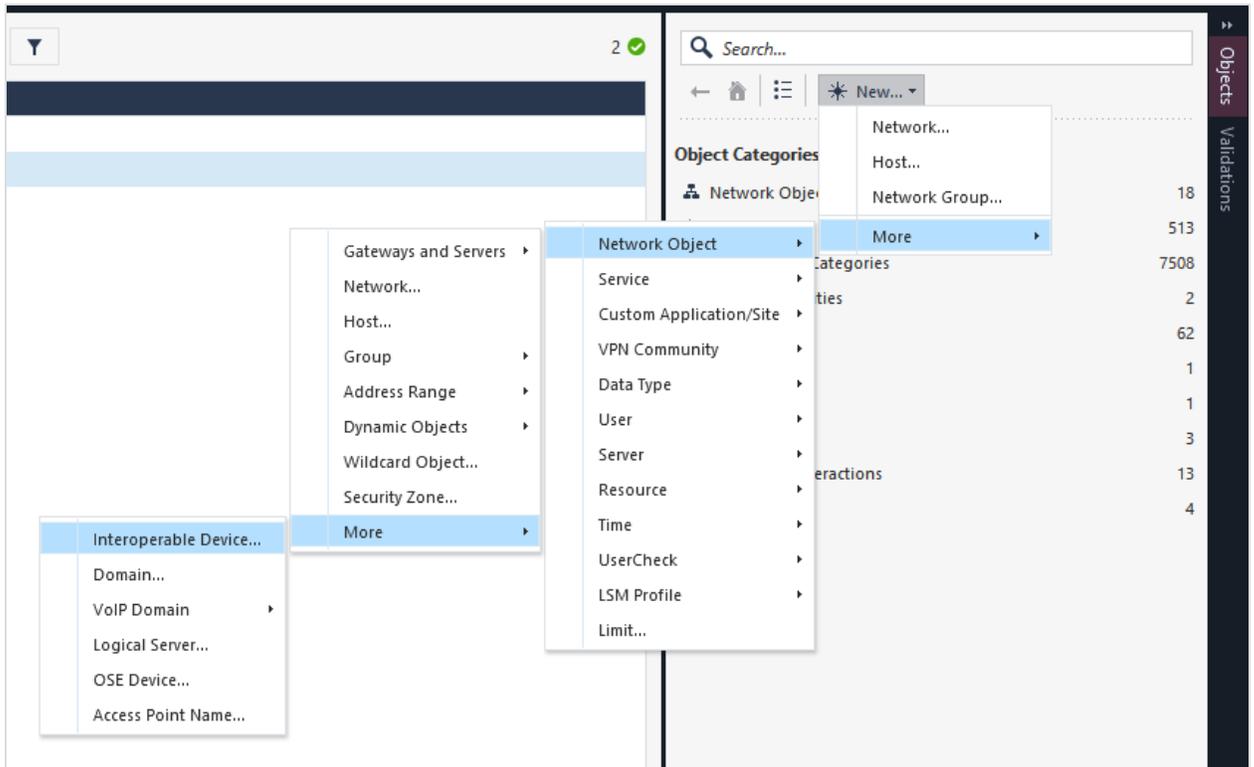


### Task 3: Create an interoperable device

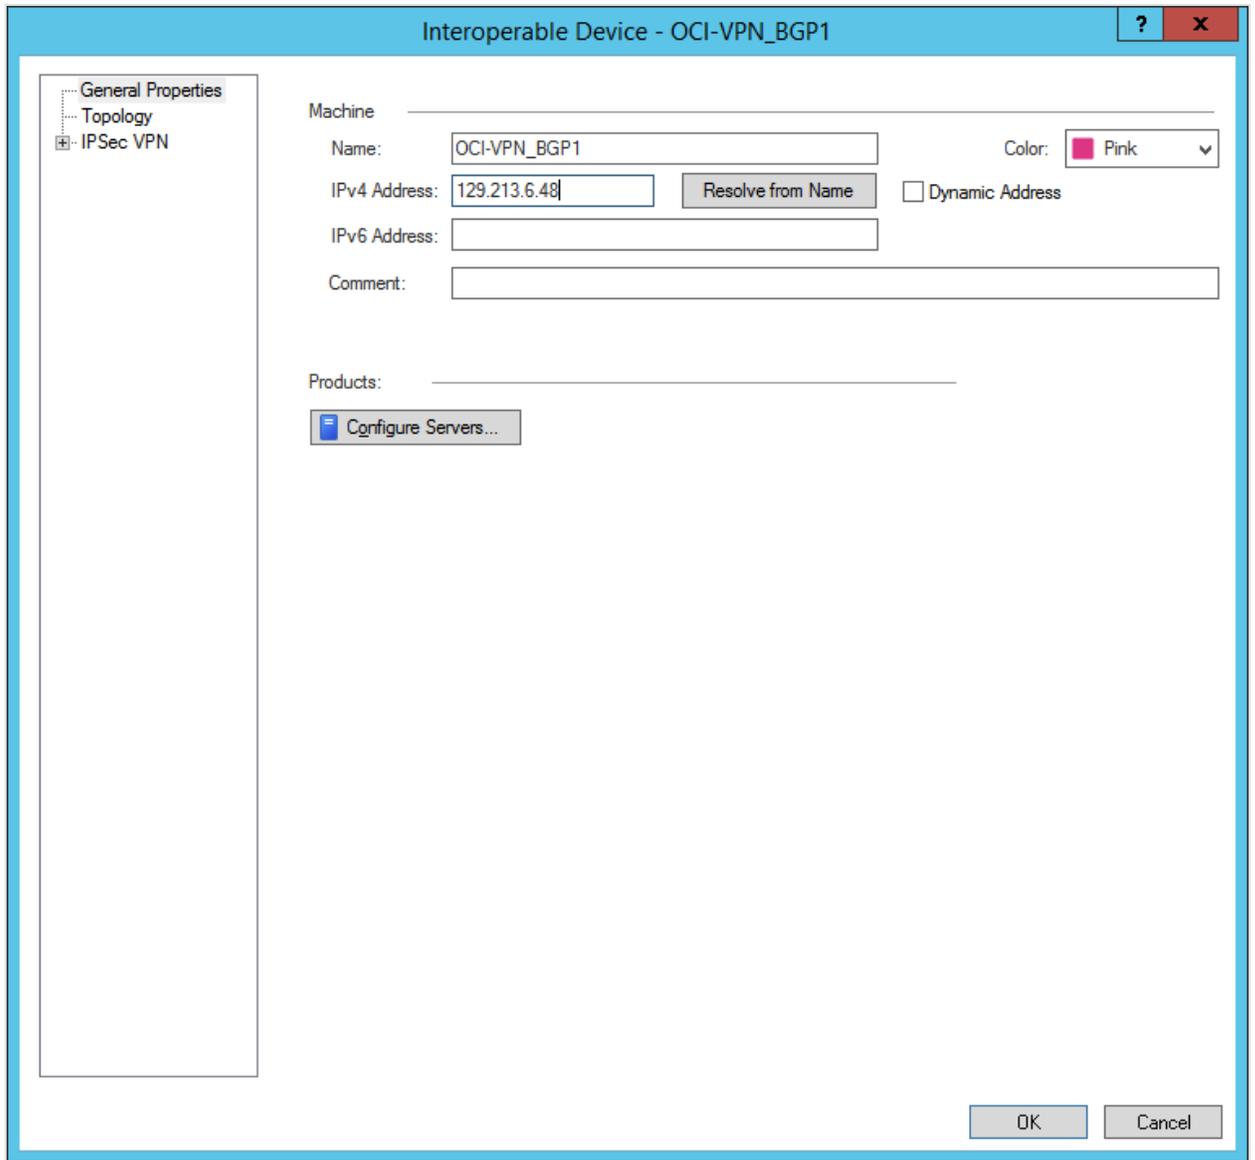
Later, you will create a VPN Community. Before you can, you must create an **Interoperable Device** that will be used in Check Point CloudGuard Security Gateway to define the Oracle DRG.

## CHAPTER 23 Networking

1. Create the new interoperable device.



2. On the **General Properties** page of the new interoperable device, add a name to identify the IPsec tunnel. Enter the IP address that Oracle assigned for the Oracle end of the tunnel when creating the IPsec connection.

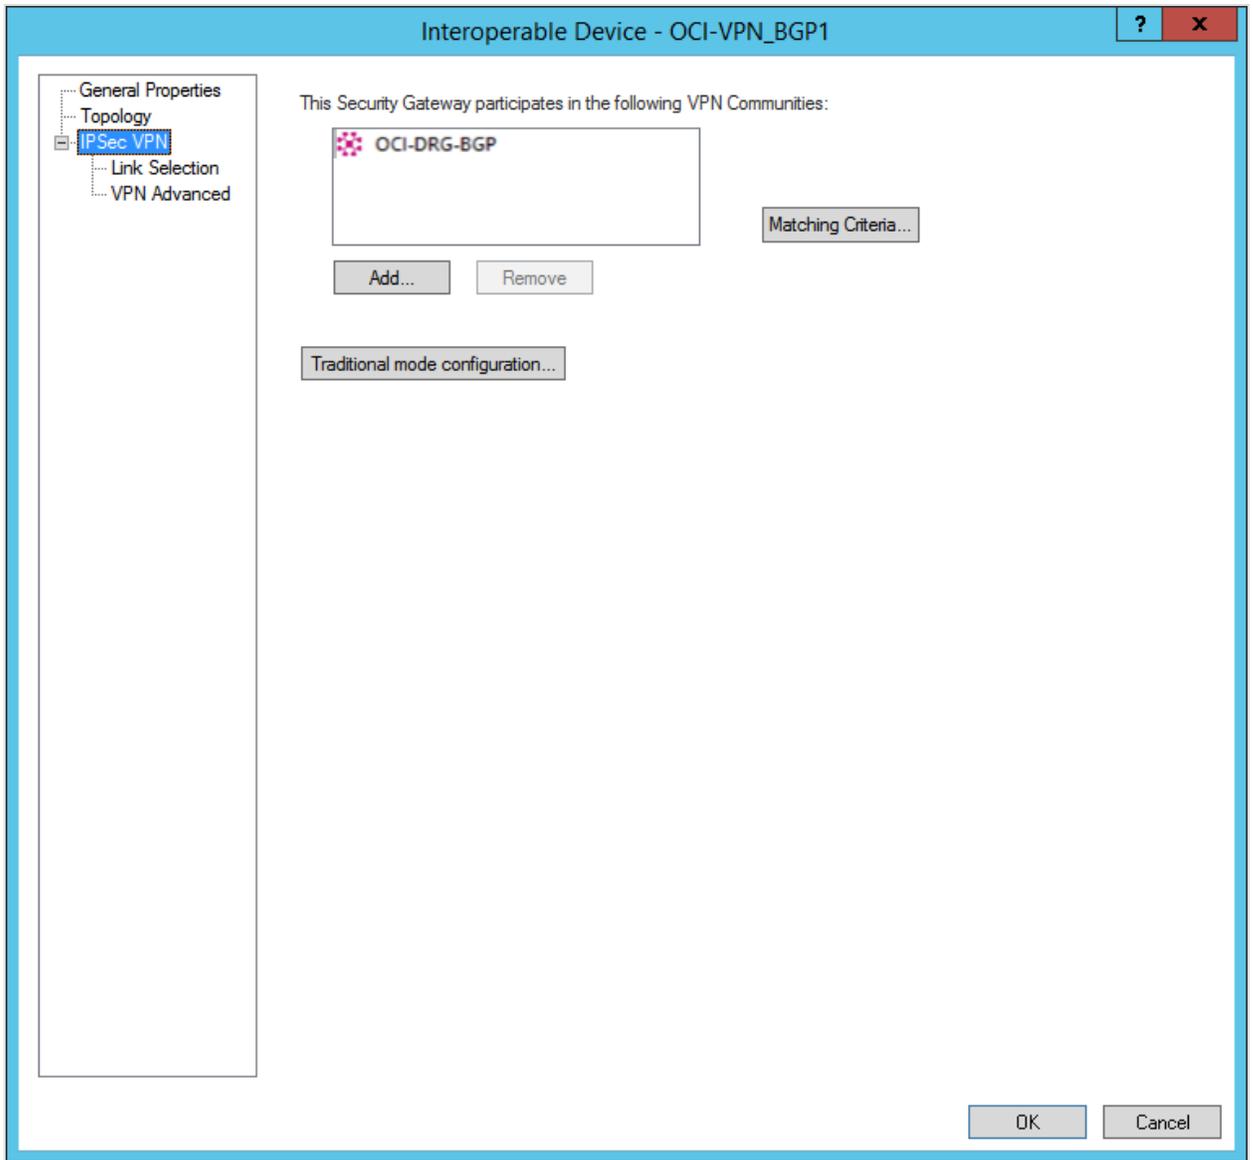


3. To force the route-based VPN to take priority, you must create an empty group and assign it to the VPN domain. To do that, on the **Topology** page, in the **VPN Domain**

section, select **Manually defined**, and select the empty group.

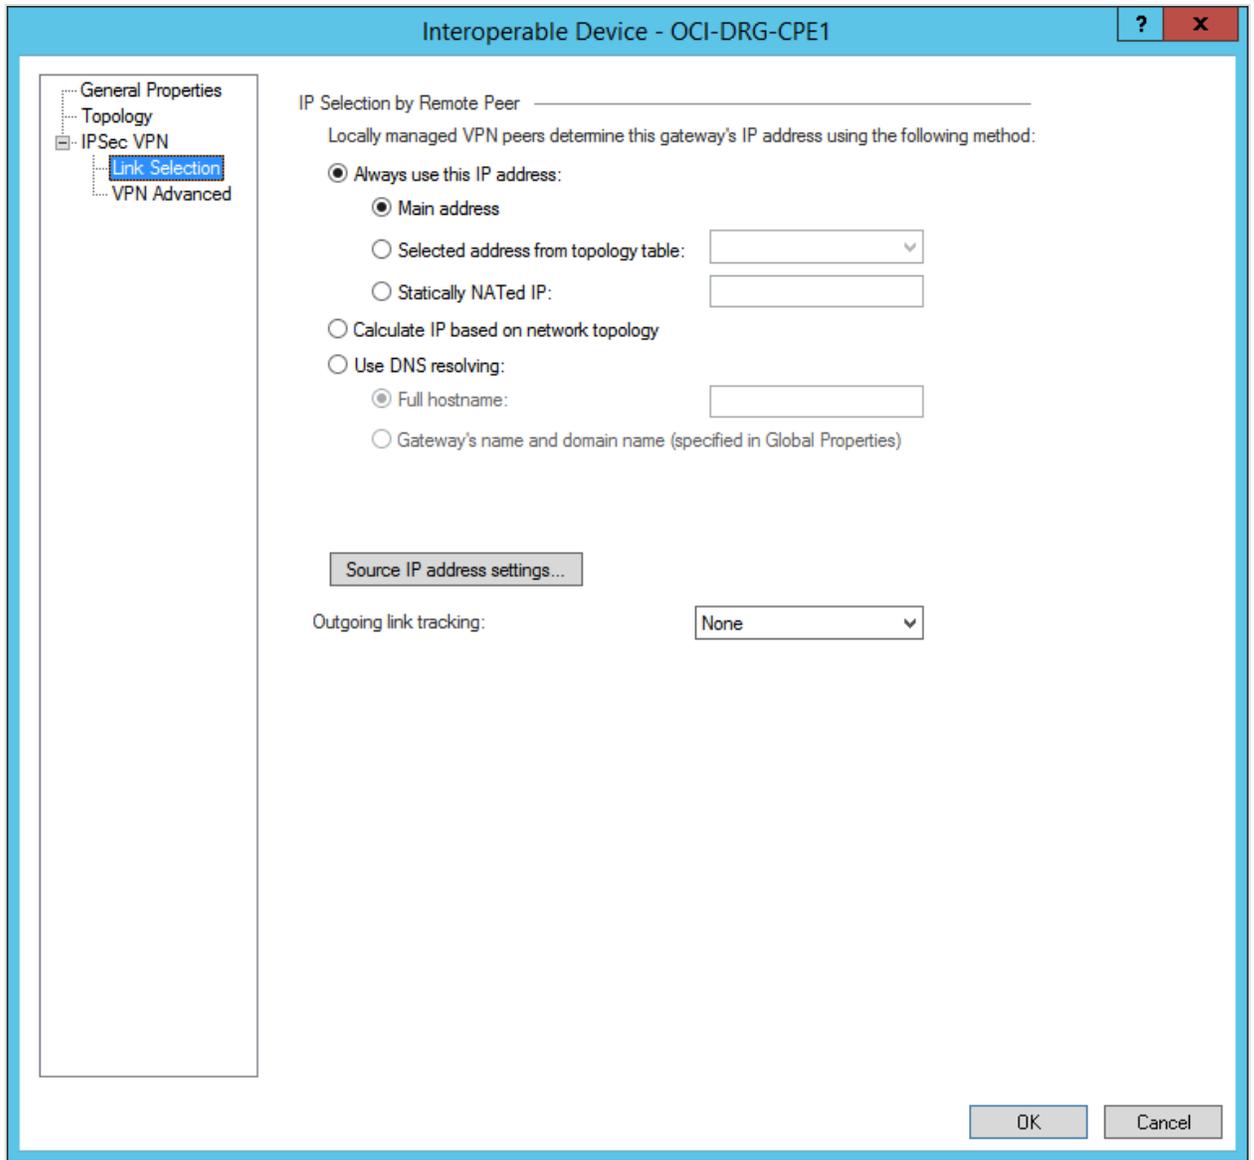
4. On the **IPSec VPN** page, you can optionally add the new interoperable device to an existing VPN Community. You can skip this step if you don't yet have any VPN Communities created.

Notice that you skip the **Traditional mode configuration**, because you will define all the Phase 1 and Phase 2 parameters in the VPN Community in a later step. The VPN Community applies those parameters to all interoperable devices that belong to the VPN Community.



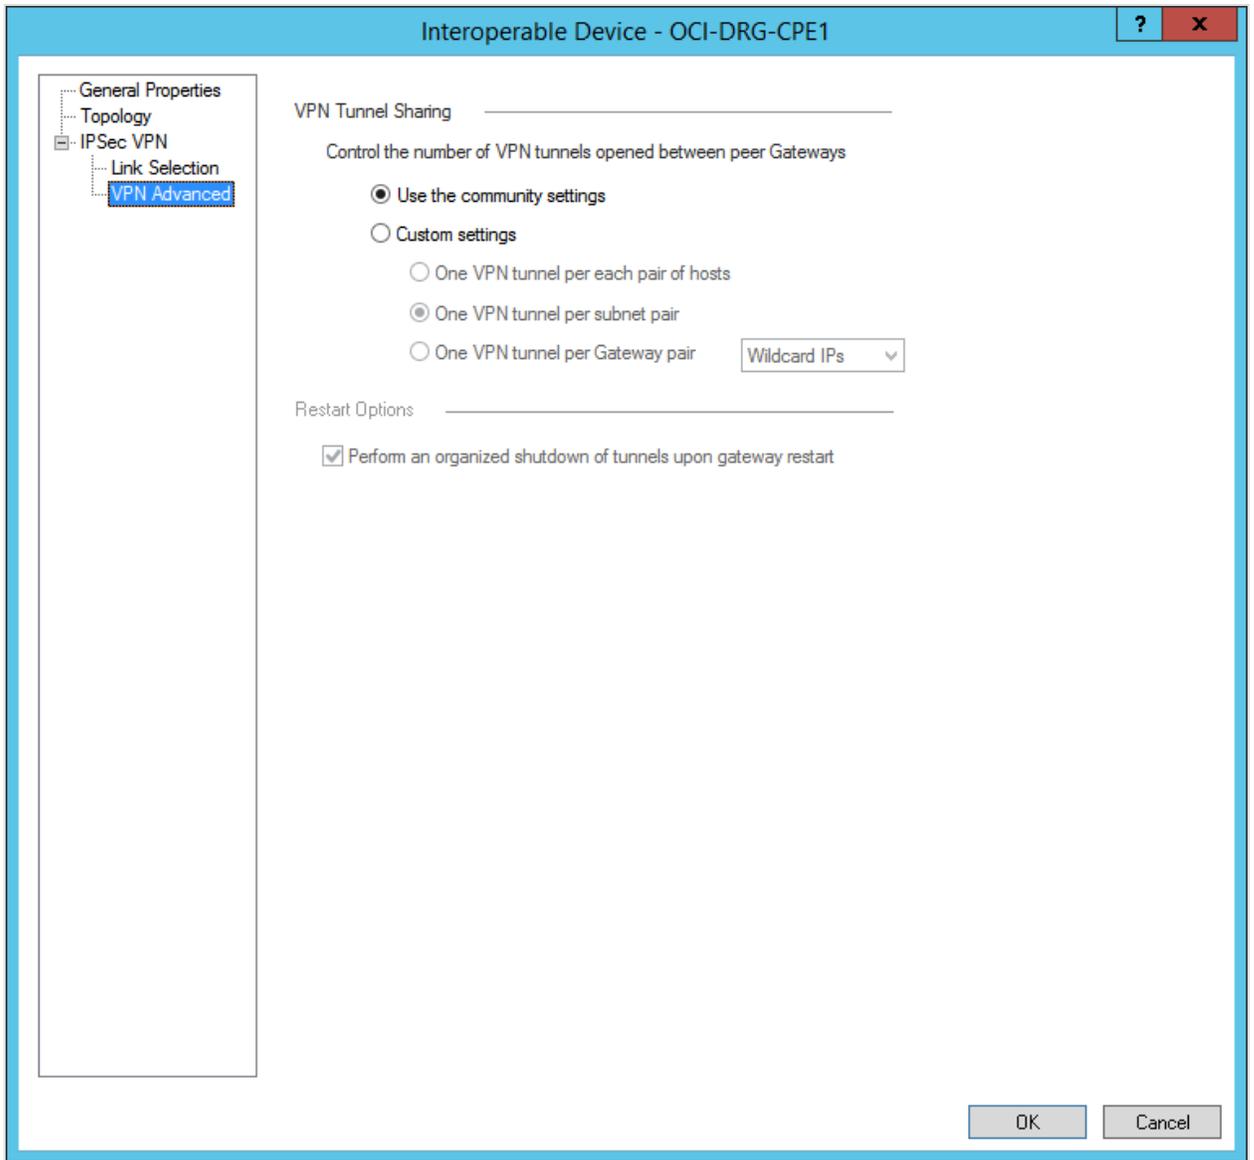
5. On the **Link Selection** page, under **Always use this IP address**, select **Main address**, which was the address you specified when creating the interoperable device.

If necessary, you can use a specific IP address that will be used as the IKE ID.



6. On the **VPN Advanced** page, select **Use the community settings**, which applies all the options and values in the VPN Community, including the Phase 1 and Phase 2

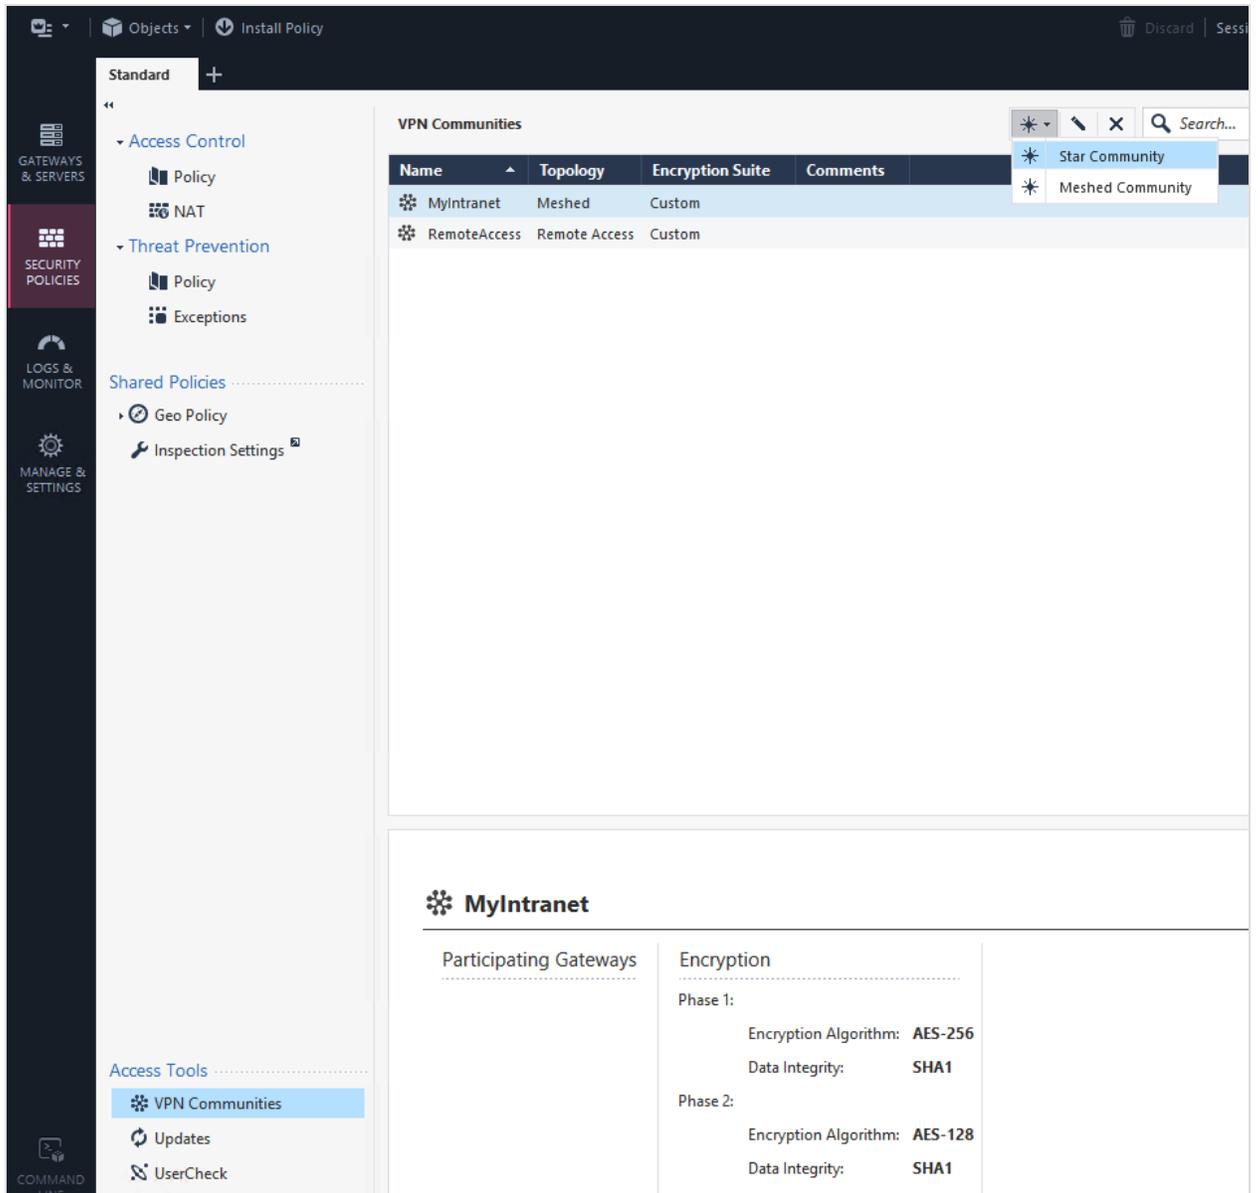
parameters.



7. Click **OK** to save your changes.

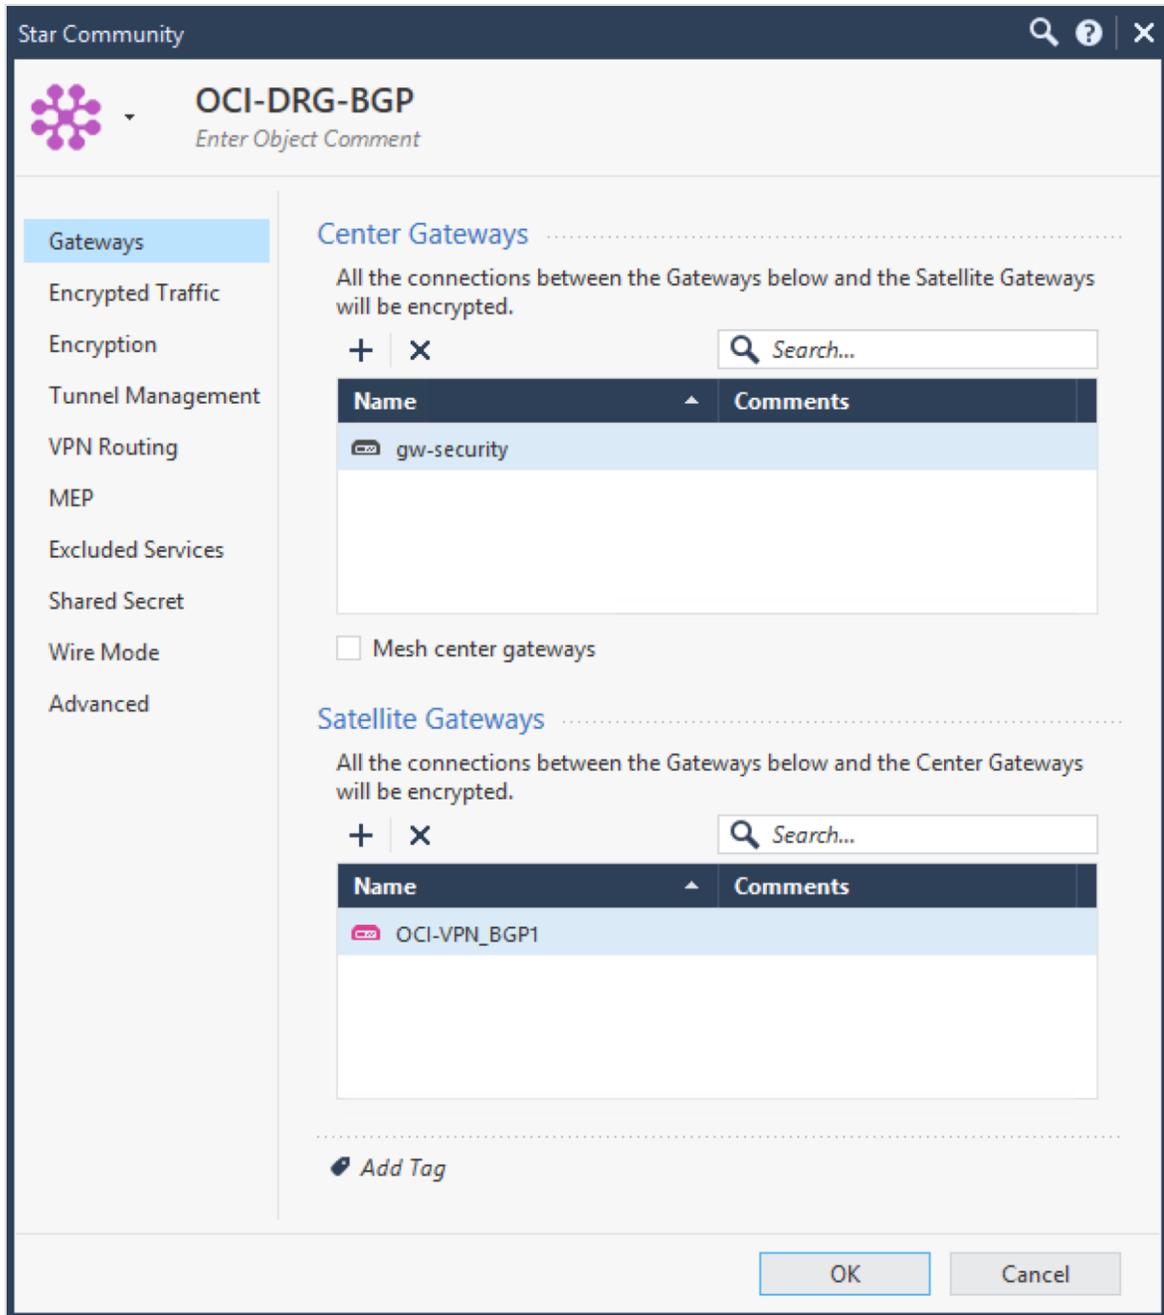
#### Task 4: Create a VPN community

1. Go to **Security Policies**, and then from **Access Tools**, select **VPN Communities**.
2. Create a **Star Community**.

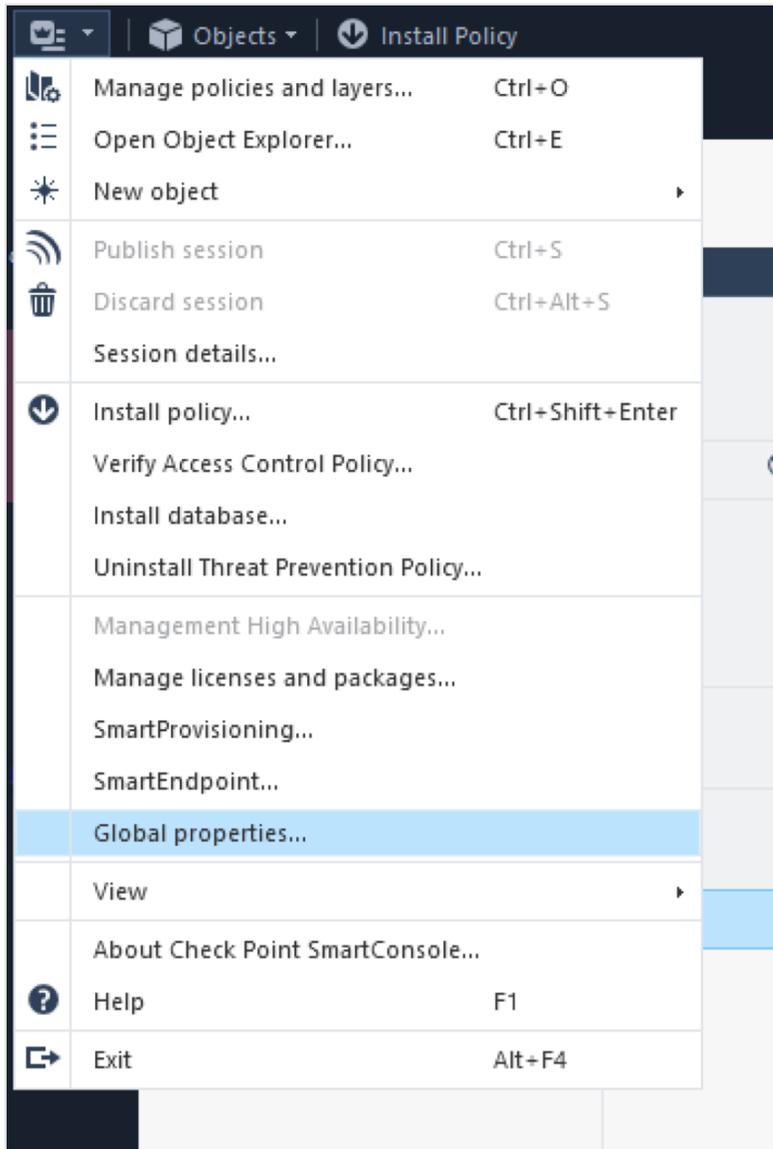


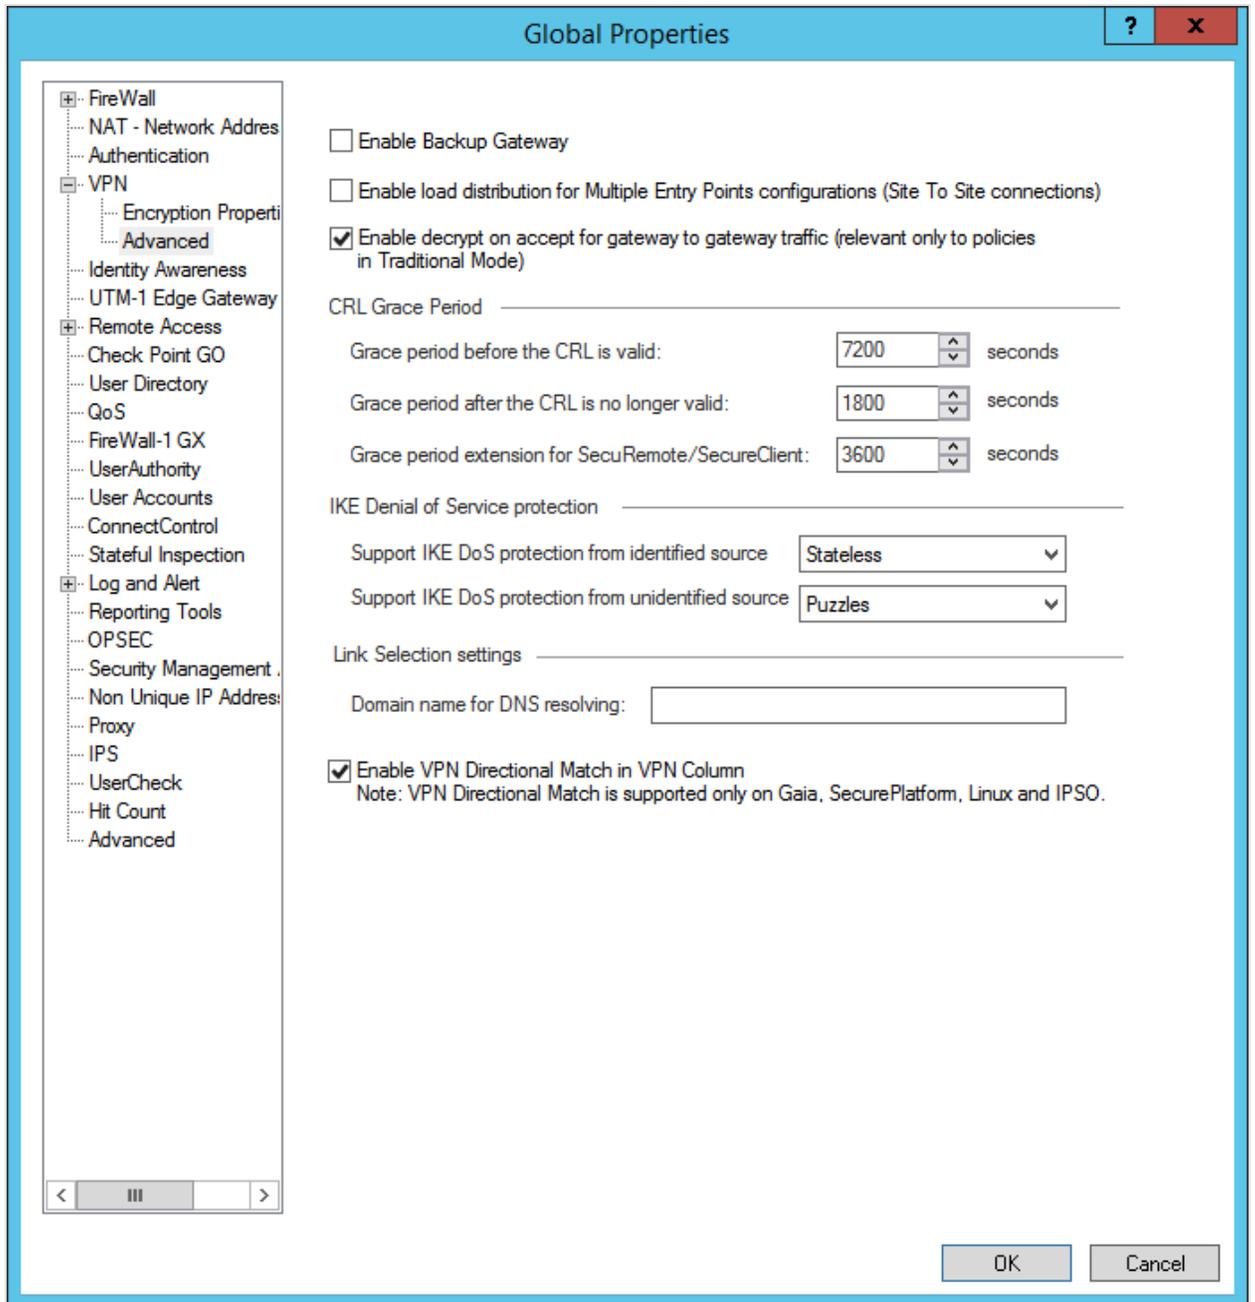
3. For the star community, add a name.

4. On the **Gateways** page, select the values for **Center Gateways** and **Satellite Gateways**. This star community acts as a settings template for the interoperable devices you specify in **Center Gateways** and **Satellite Gateways**.
  - **Center Gateways**: For the Check Point CloudGuard Security Gateway.
  - **Satellite Gateways**: For the CPE that connects to the Oracle DRG for each IPSec tunnel.

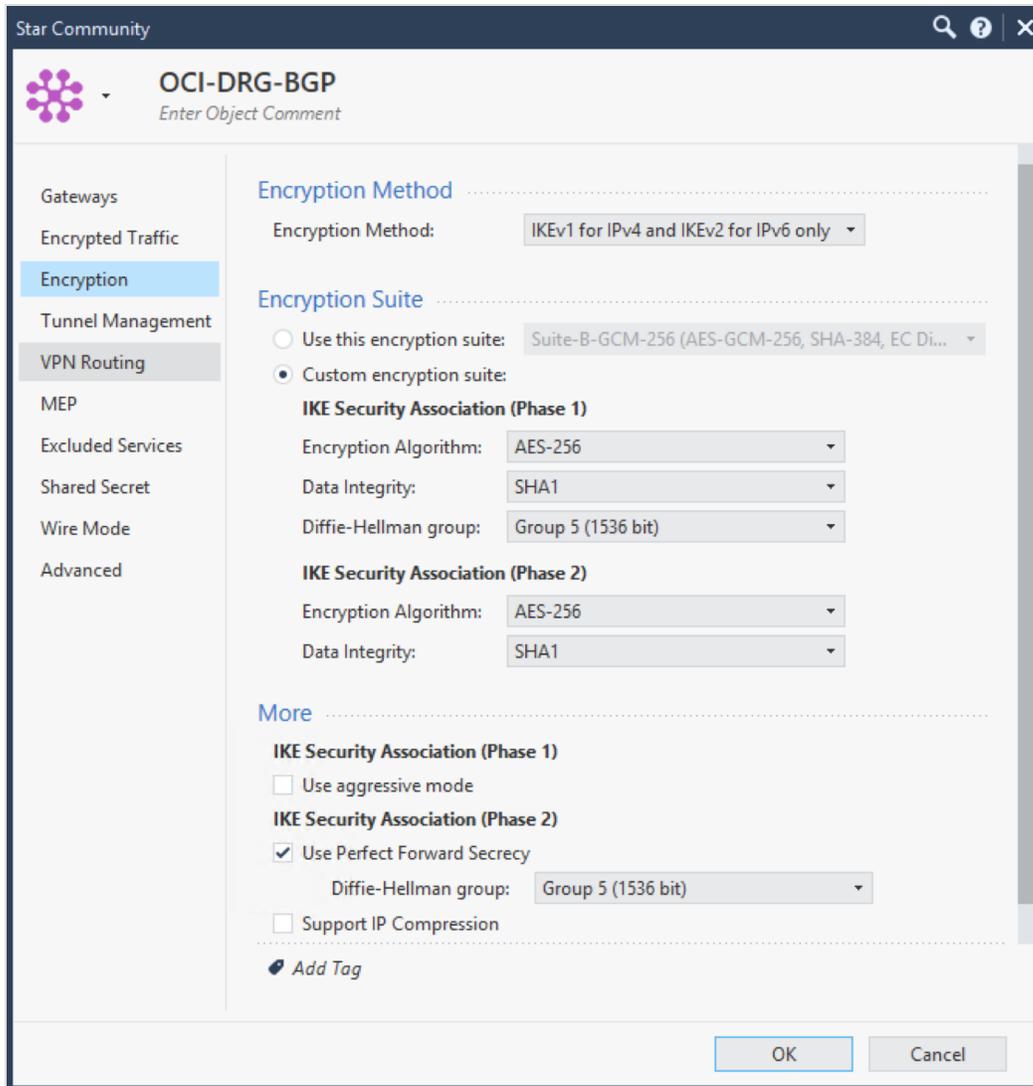


5. To allow traffic, go to **Global Properties**, and then **VPN**, and then **Advanced**.





6. Select the check box for **Enable VPN Directional Match in VPN Column**. Later you will create a security policy that uses a directional match condition to allow traffic to pass based on routing rules.
7. Click **OK**.
8. On the **Encryption** page, configure the Phase 1 and Phase 2 parameters that Oracle supports. For a list of those values, see [Supported IPSec Parameters](#).  
If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#).  
Notice that if you want to use IKEv2, for the **Encryption Method**, instead select **IKEv2 only**.



9. On the **Tunnel Management** page, select **Set Permanent Tunnels**. Oracle recommends that you:

- Select **On all tunnels in the community** to keep all the Oracle IPSec tunnels up all the time.
- In the **VPN Tunnel Sharing** section, select **One VPN tunnel per Gateway pair**.

The latter option generates only one pair of IPSec security associations (SAs), and each SA with only one security parameter index (SPI) (unidirectional).

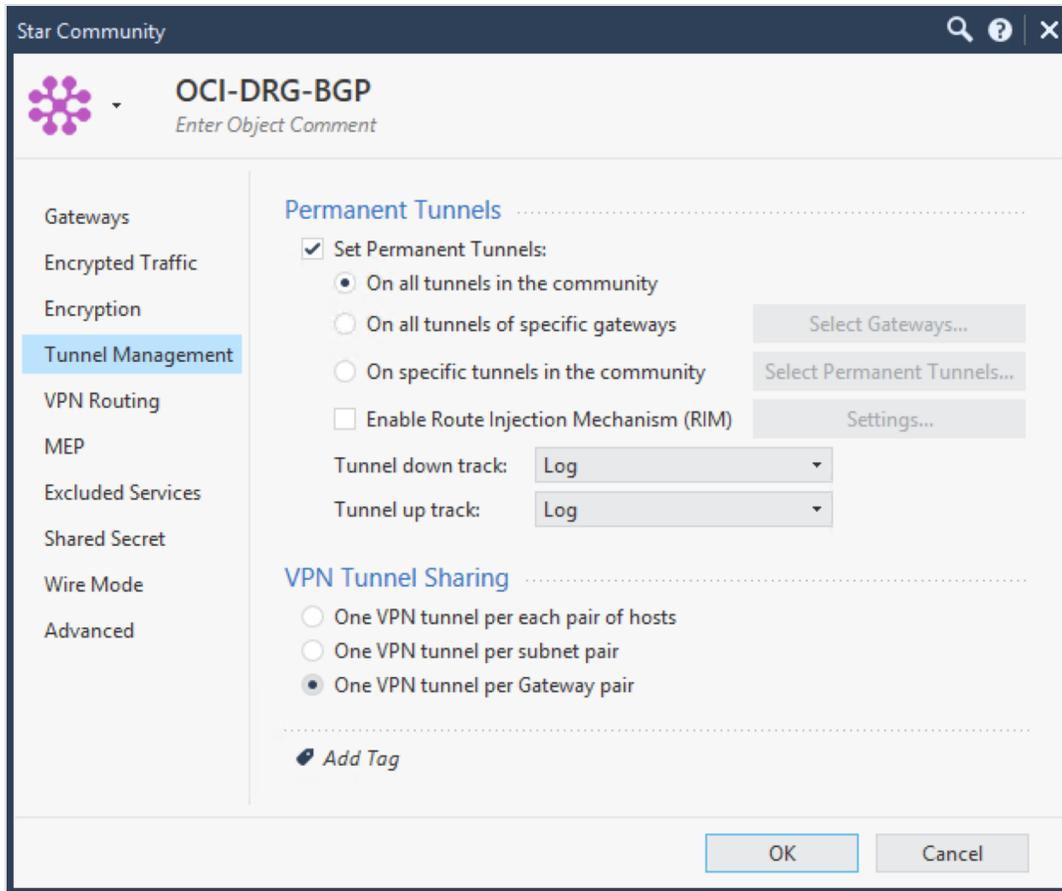
When you use policy-based tunnels, every policy entry generates a pair of IPSec SAs, (also referred to as an *encryption domain*).



### Important

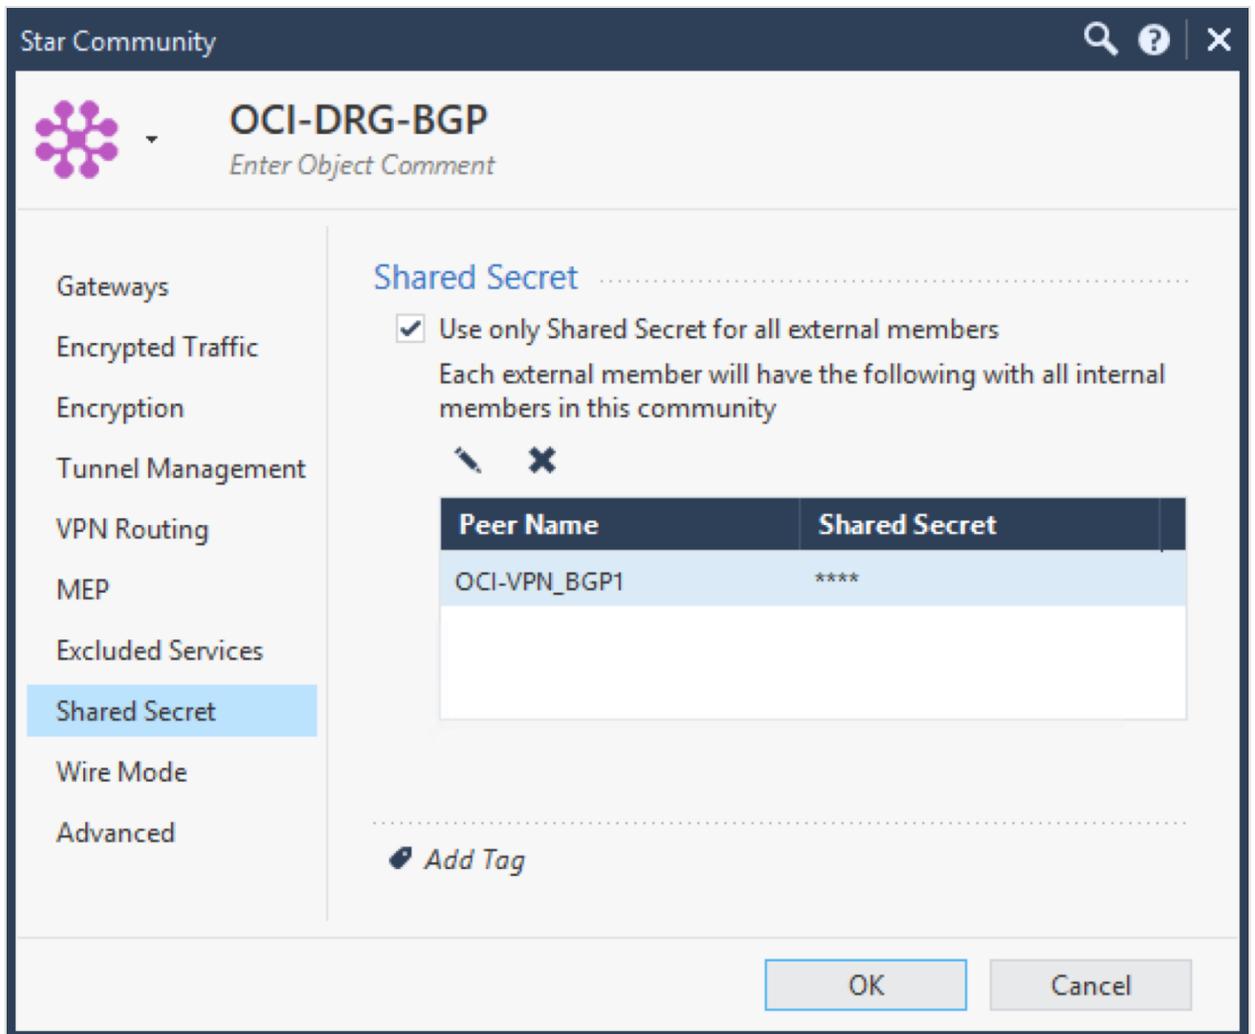
The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

Oracle creates a route-based IPSec connection, which means that everything is routed through an encryption domain that has 0.0.0.0/0 (any) for local traffic and 0.0.0.0/0 (any) for remote traffic. For more information, see [Supported Encryption Domain or Proxy ID](#).



10. On the **Shared Secret** page, select **Use only Shared Secret for all external members**, and add the shared secret that Oracle generated for the tunnel when creating the IPSec connection.

Currently Oracle supports only shared secret keys. Note that you can [change the shared secret](#) in the Oracle Console.



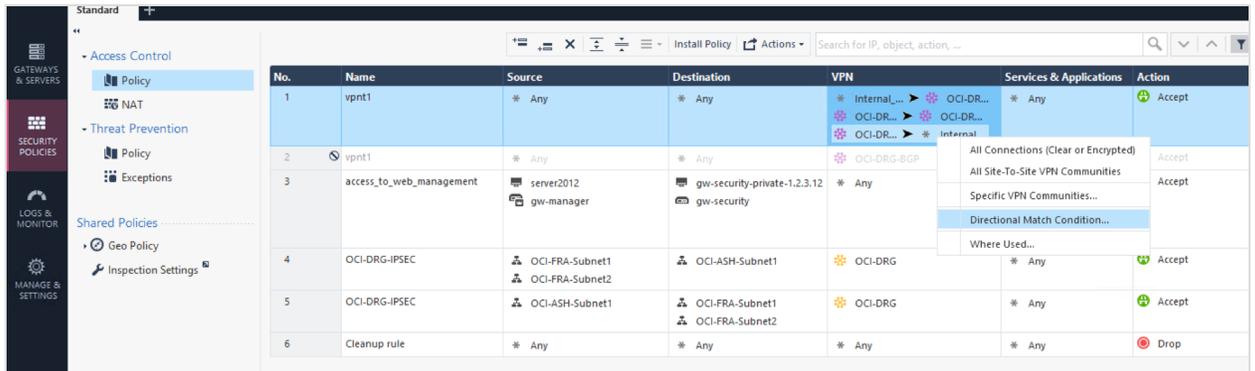
11. Click **OK** to save your changes.

### Task 5: Create a security policy

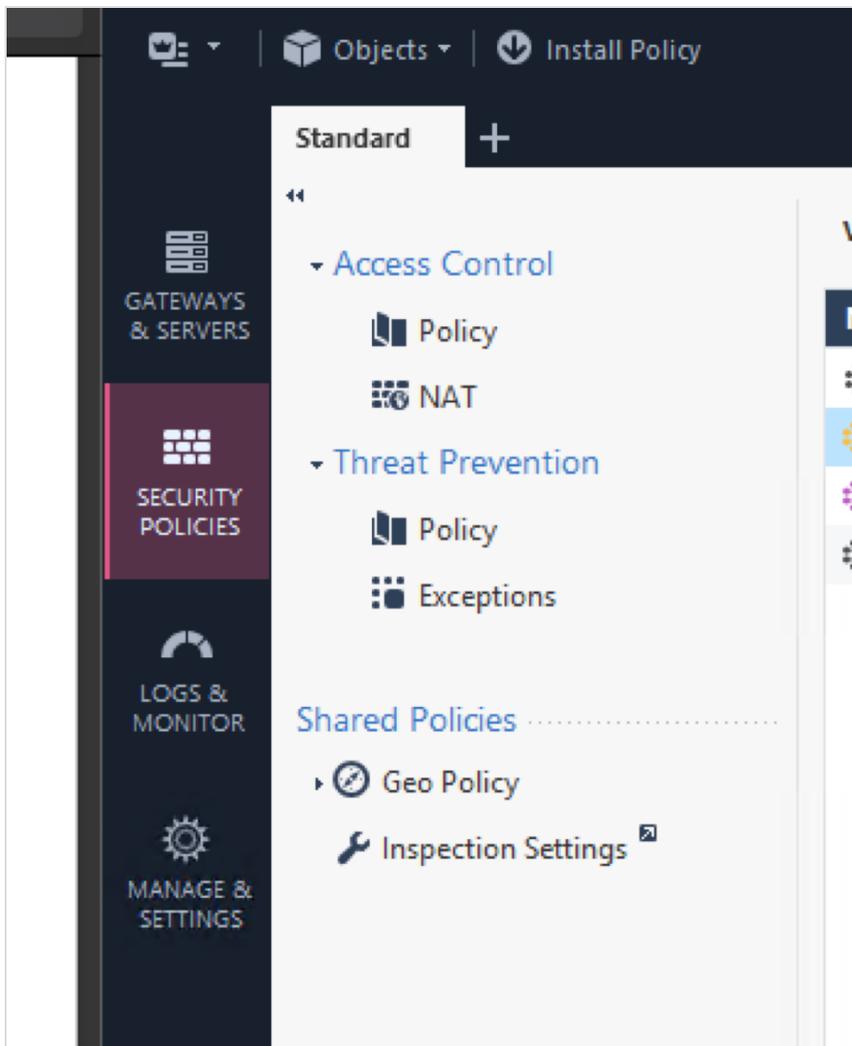
1. Go to **Access Control**, and then the **Policy** tab. Create specific security policies by using **Directional Match Condition**, which allows traffic to pass based on routing tables. Set up the condition with these settings:

- **Internal\_Clear** > *VPN Community* created
- *VPN Community* created > *VPN Community* created
- *VPN Community* created > **Internal\_Clear**

In this case, the *VPN Community* is **OCI-DRG-BGP** and the **Internal\_Clear** is predefined by Check Point.



2. Click **OK** to save your changes.
3. Click **Install Policy** to apply the configuration.



**REDUNDANCY WITH BGP OVER IPSEC**

For redundancy, Oracle recommends using BGP over IPsec. By default, if you have two connections of the same type (for example, two IPsec VPNs that both use BGP), and you advertise the same routes across both connections, Oracle prefers the oldest established route when responding to requests or initiating connections. If you want to force routing to be

symmetric, Oracle recommends using BGP and AS path prepending with your routes to influence which path Oracle uses when responding to and initiating connections. For more information, see [Routing Details for Connections to Your On-Premises Network](#).

The Oracle DRG uses /30 or /31 as subnets for configuring IP addresses on the interface tunnels. Remember that the IP address must be part of the IPSec VPN's encryption domain and must be allowed in the firewall policy to reach the peer VPN through the interface tunnel. You might need to implement a static route through the tunnel interface for the peer IP address.

Oracle's BGP ASN in commercial regions is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

For your side, you can use a private ASN. Private ASNs are in the range 64512–65534.

### Task 1: Enable BGP

Perform the following steps for each tunnel.

1. Go to **Advanced Routing**, and then **BGP**.
2. Under **BGP Global Settings**, click **Change Global Settings**, and then add a router ID and local ASN.

Change Global Settings

Router ID and Cluster ID

Router ID: 1 . 2 . 3 . 12

The Router ID cannot be changed while BGP is configured and active.

Cluster ID for Route Reflectors: . . .

Autonomous System

Unconfigured

Local Autonomous System Number: 62113

Confederation Identifier:

Loops Permitted in AS Path: 1

Routing Domain Identifier:

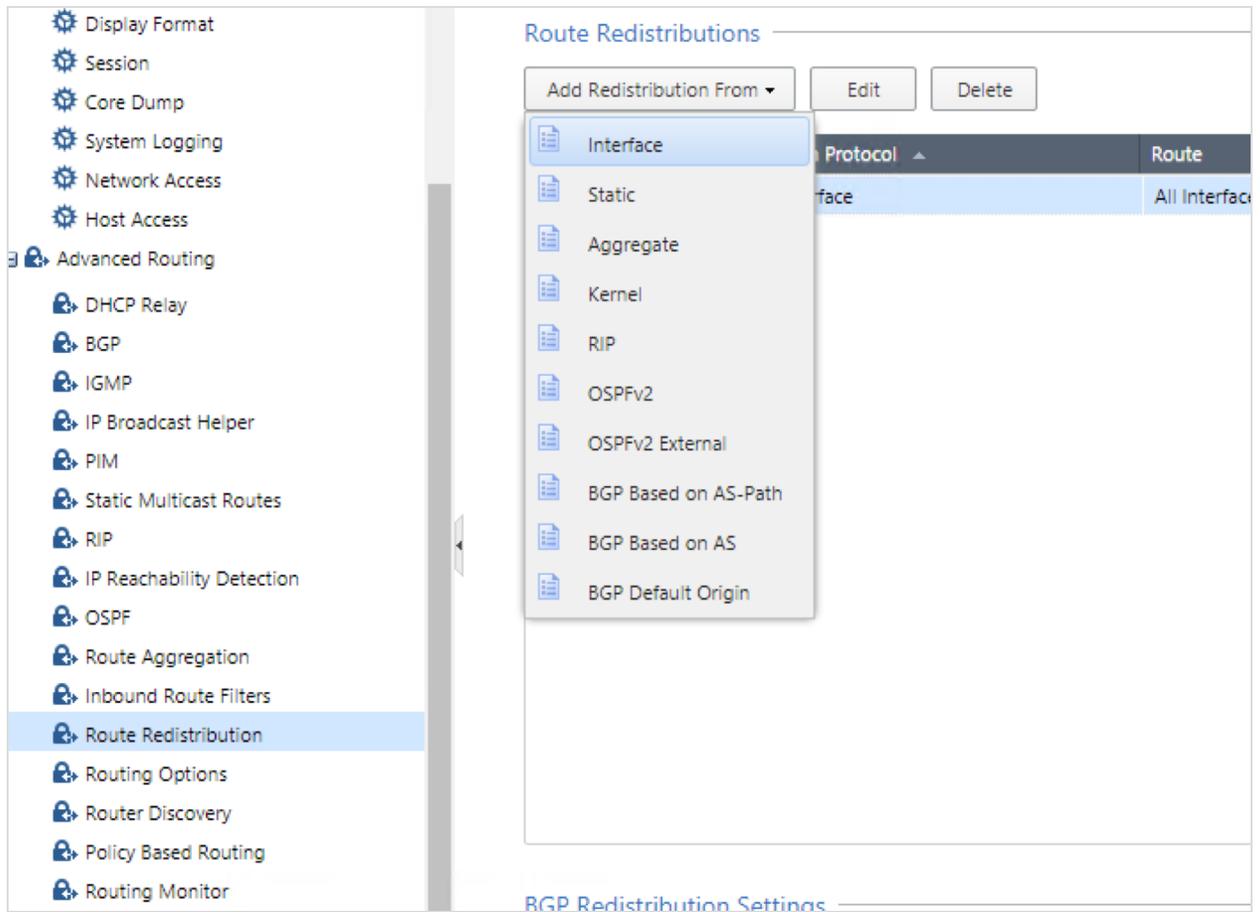
Loops Permitted in AS Path: 1

Save Cancel

3. Click **Save**.

### Task 2: Redistribute routes into BGP

1. Go to **Advanced Routing**, and then **Route Distribution**.
2. Click **Add Redistribution From**, and then select **Interface**, which is for adding all connected subnets.



3. In the **Add Redistribution from Interface** dialog, configure the following items:
  - **To Protocol:** Select **BGP AS 31898**, which is the Oracle ASN for commercial regions. If you're configuring VPN Connect for the Government Cloud, see [Oracle's BGP ASN](#).
  - **Interface:** Select **all** to advertise all connected subnets.

The screenshot shows a dialog box titled "Add Redistribution from Interface". It contains the following fields and options:

- To Protocol:** BGP AS 31898
- From Protocol:** Interface
- Interface:** Select one (dropdown menu is open, showing options: all, eth0, eth1, eth2, ip\_vti0, vpnt1)
- Metric:** (empty field)

Buttons: Save, Cancel

4. Click **Save**.

Now the BGP session should be up and advertising and receiving subnets.

### Verification

The following CLI command verifies BGP peers and routing.

```
show bgp peers
```

The following command verifies that you're receiving BGP routes.

```
show route bgp
```

The following command verifies the routes that are being advertised. In this example, replace *<remote\_IP\_address>* with the remote IP address that was specified in the Oracle Console as the **Inside Tunnel Interface - Oracle**

```
show bgp peer <remote_IP_address> advertise
```

The following command verifies the routes that are being received.

```
show bgp peer <remote_IP_address> received
```

Use options 2 and 4 in the following command to verify security associations (SAs).

```
vpn tunnelutil

***** Select Option *****

(1) List all IKE SAs
(2) * List all IPsec SAs
(3) List all IKE SAs for a given peer (GW) or user (Client)
(4) * List all IPsec SAs for a given peer (GW) or user (Client)
(5) Delete all IPsec SAs for a given peer (GW)
(6) Delete all IPsec SAs for a given User (Client)
(7) Delete all IPsec+IKE SAs for a given peer (GW)
(8) Delete all IPsec+IKE SAs for a given User (Client)
(9) Delete all IPsec SAs for ALL peers and users
(0) Delete all IPsec+IKE SAs for ALL peers and users

* To list data for a specific CoreXL instance, append "-i <instance number>" to your selection.

(Q) Quit

```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Check Point: Policy-Based

This topic provides a policy-based configuration for Check Point CloudGuard. The instructions were validated with Check Point CloudGuard version R80.20.

This topic is for policy-based configuration. If you instead want route-based (VTI-based) configuration, see [Check Point: Route-Based](#).

Check Point experience is required. This topic does not include how to add Check Point CloudGuard Security Gateway to Check Point CloudGuard Security Manager. For more information about using Check Point products, see the Check Point documentation.



#### **Important**

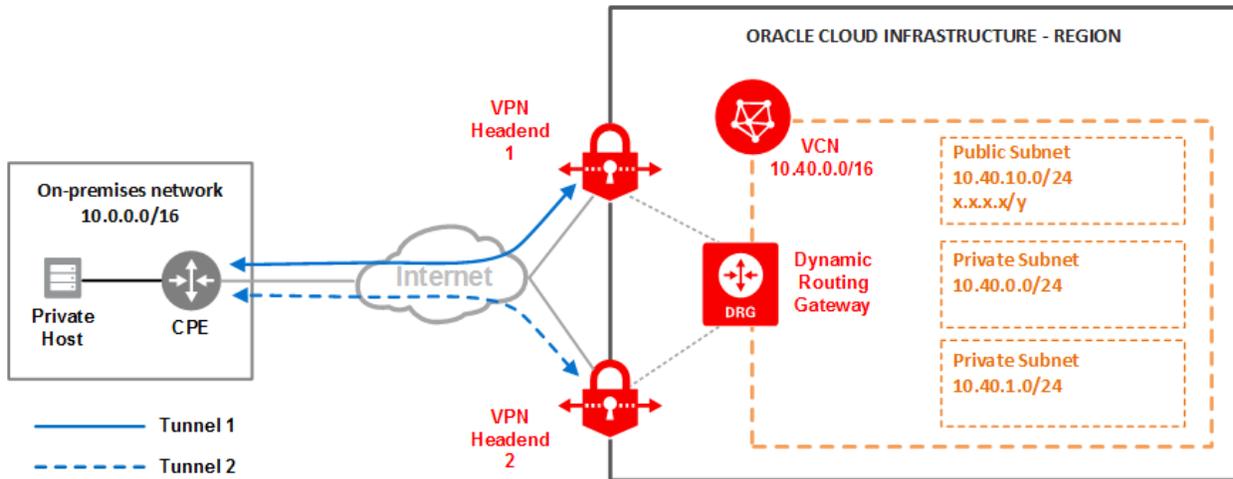
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPSec, see the list of [supported IPSec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPSec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPSec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

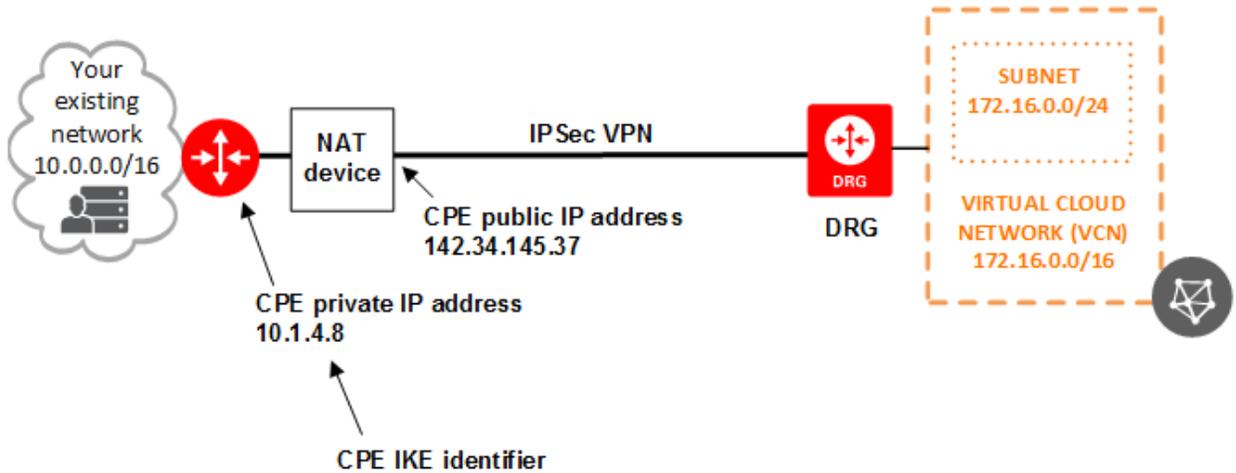
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

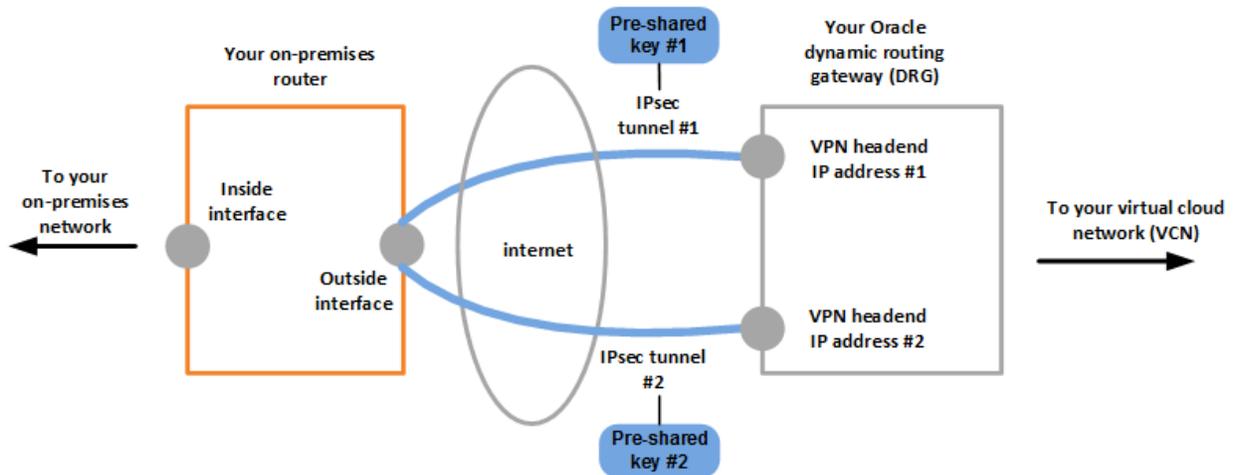
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration (Policy-Based)

 **Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to

## CHAPTER 23 Networking

use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

If you want to use IKEv2, there's a variation on one of the tasks presented in the next section. Specifically, in [task 4](#), when configuring encryption, select **IKEv2 only** for the encryption method.

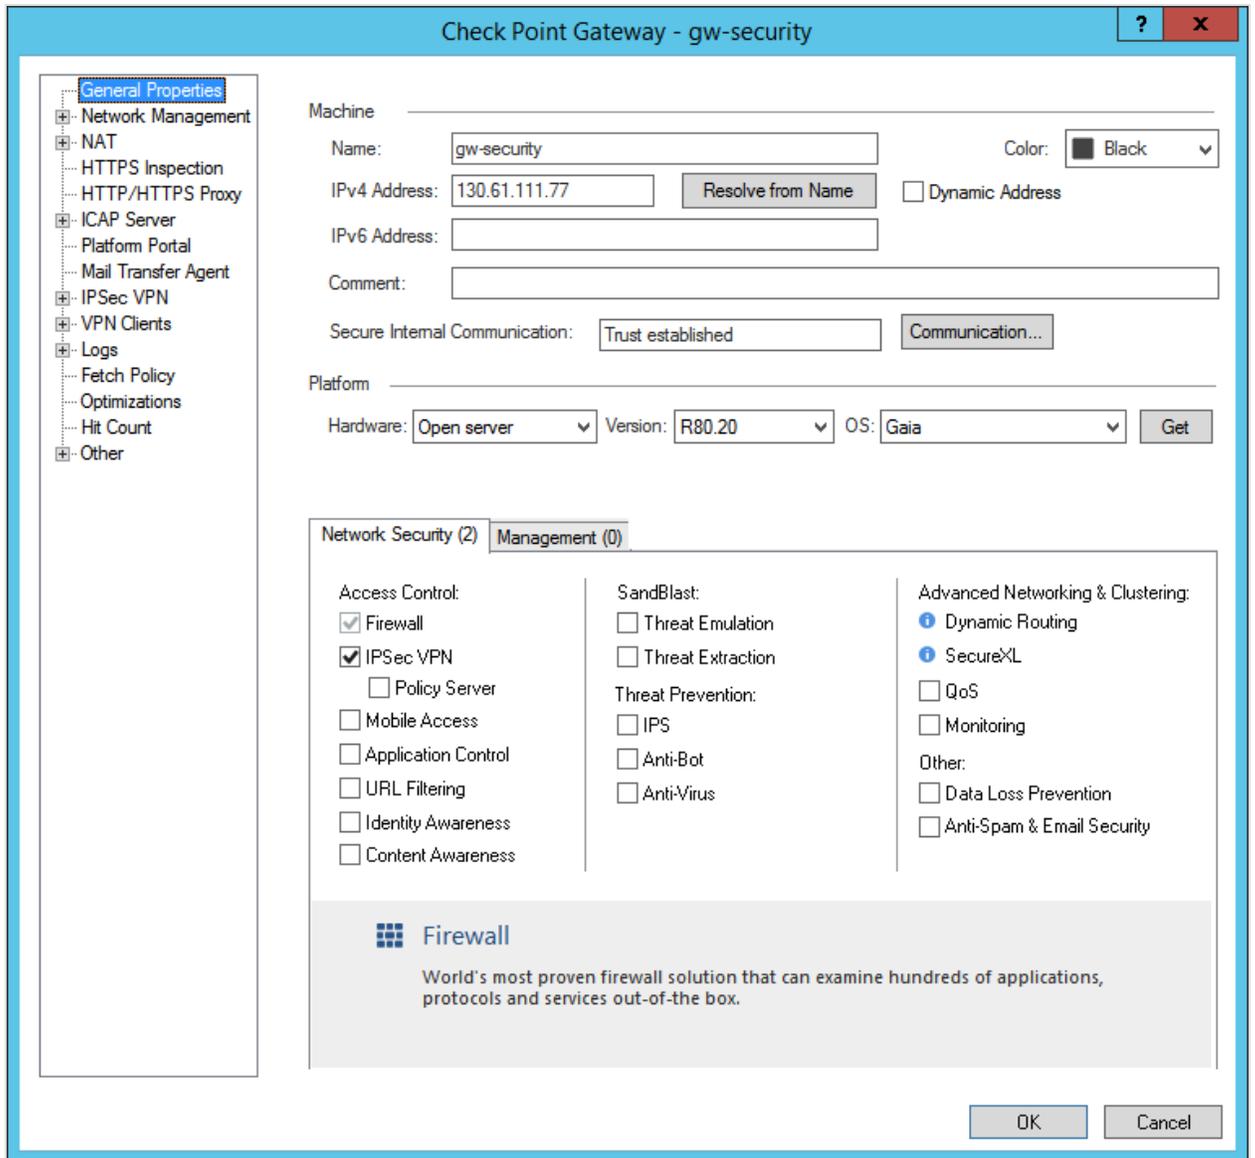
### CONFIGURATION PROCESS

#### Task 1: Install IPSec VPN on Check Point CloudGuard Security Gateway

**Prerequisite:** Before starting, add Check Point CloudGuard Security Gateway to Check Point CloudGuard Security Manager. Also establish the Secure Internal Communication (SIC) so you can configure the IPSec tunnel by using the Check Point Smart Console. For instructions to add the Security Gateway to CloudGuard or to establish the SIC, see the Check Point documentation.

Status	Name	IP	Version	Active Blades	Hardware	CPU Usage
✓	 gw-manager		R80.20	 	Open server	 4%
✓	 gw-security	130.61.111.77	R80.20	 	Open server	 0%

1. Install the IPSec VPN module. Oracle recommends that you also install the Monitoring module for traffic analysis.



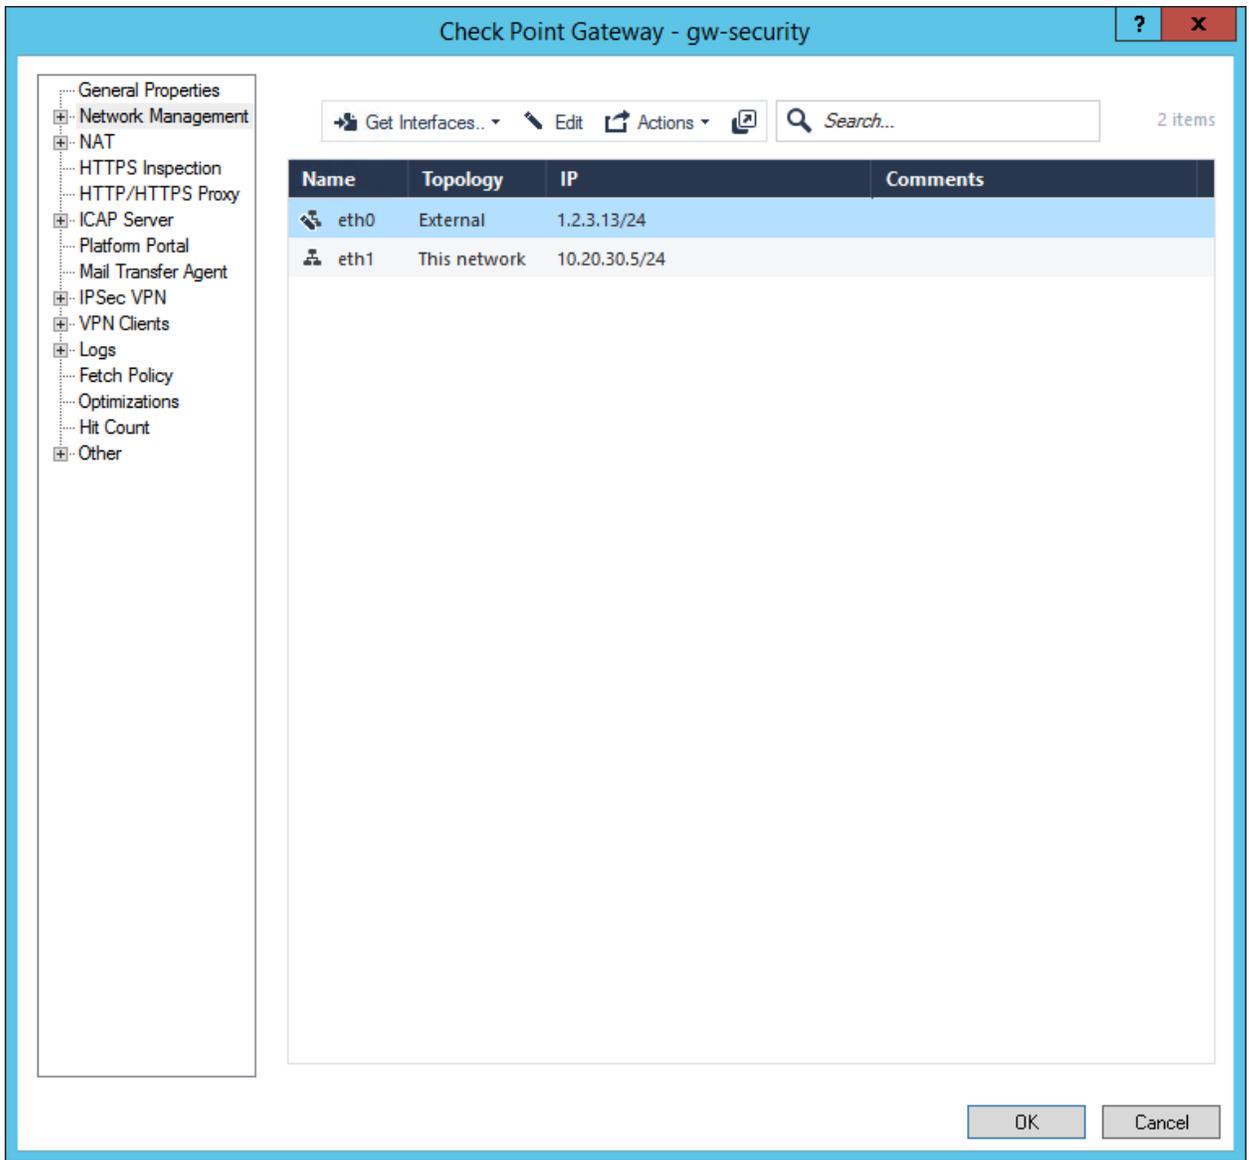
2. Click **OK** to save your changes.

### Task 2: Configure IPSec settings for Check Point CloudGuard Security Gateway

This task covers the most important options used for an IPSec tunnel with Oracle Cloud Infrastructure.

1. On the **Network Management** page, import all the interfaces. You can do this by clicking **Get Interfaces**, which contains options for **Get Interfaces With Topology** and **Get Interfaces Without Topology**. This example uses **Get Interfaces Without Topology** so that you can define the purpose of each interface as an external or internal network.

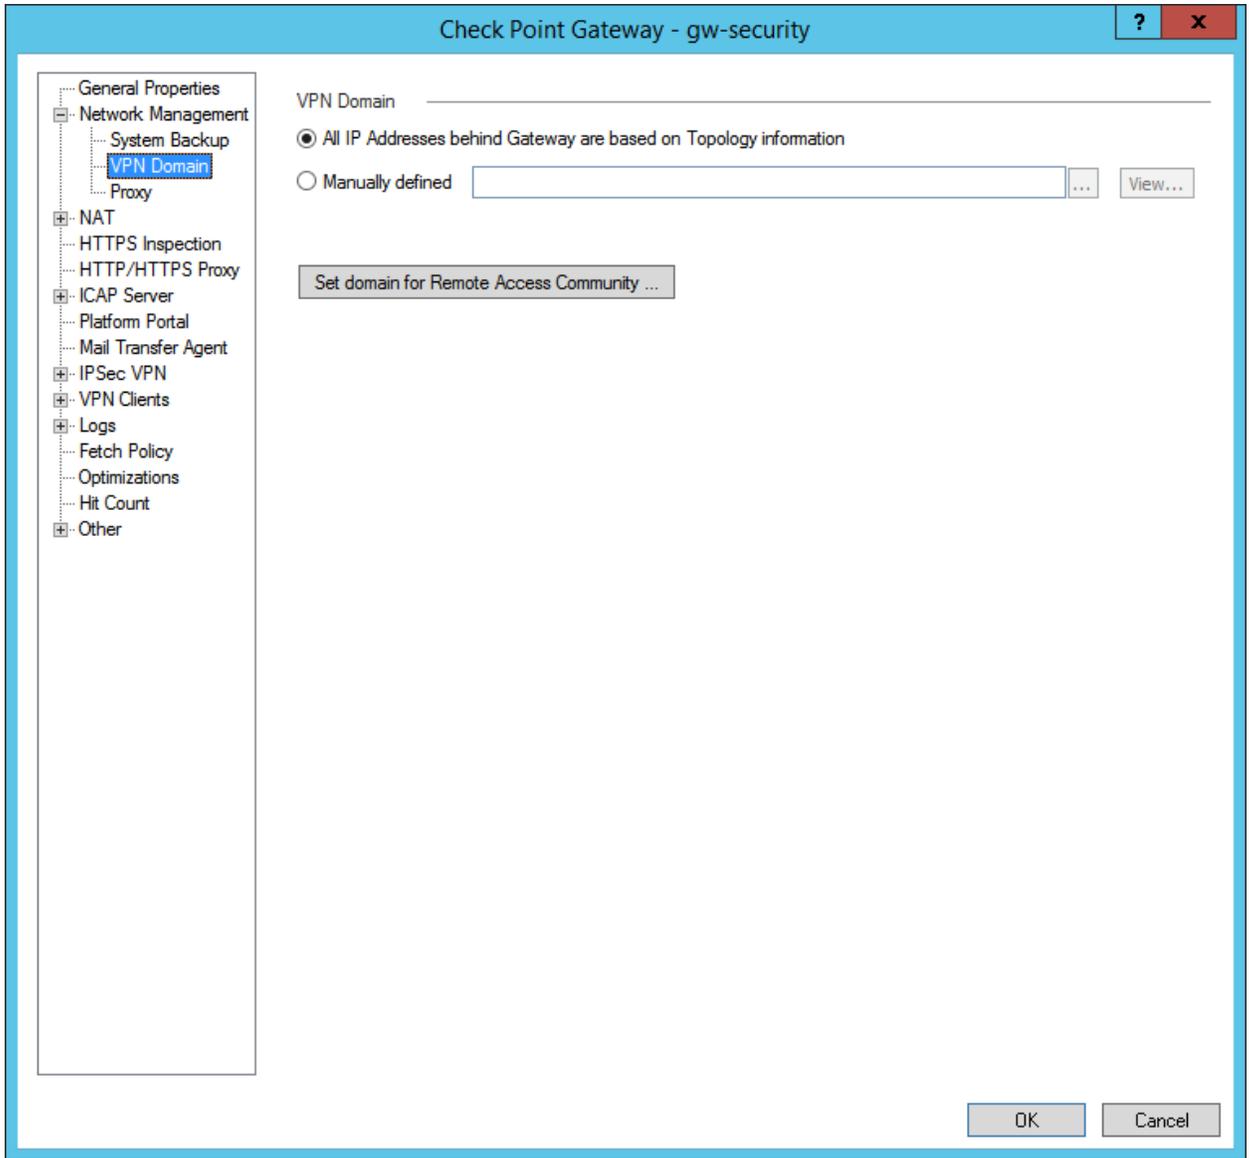
All of these interfaces will be used in the **VPN Domain** as subnets advertised by Check Point CloudGuard Security Gateway in the IPSec encryption domain.



2. On the **VPN Domain** page, Oracle recommends that you select the option **for All IP Addresses behind Gateway are based on Topology information**. This option

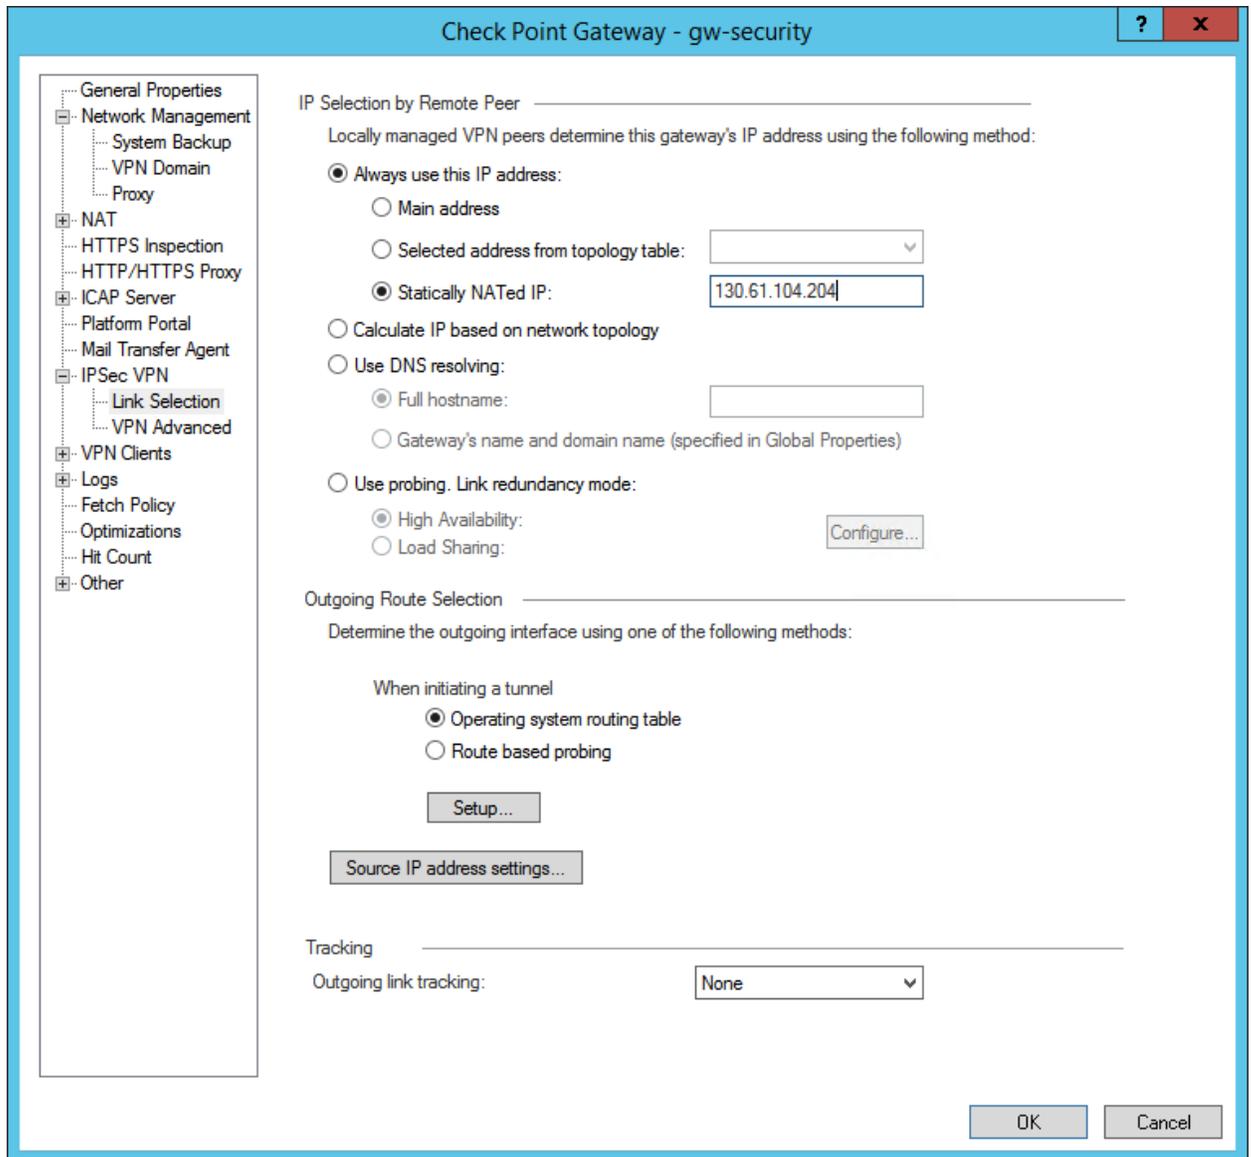
adds all the subnets discovered in **Network Management** to the IPSec Encryption Domain.

You can instead select the option for **Manually defined**. However, that requires a **Network Object** with all subnets to include in the IPSec encryption domain.



3. If your Check Point CloudGuard Security Gateway uses 1:1 NAT to map private IP addresses to public IP addresses: On the **Link Selection** page, under **Always use this**

**IP address**, select **Statically NATed IP** and specify the IP address that you want to use as your IKE ID.



If you don't want to use a public IP address as the local IKE ID, you can use another value (such as a private IP address), but the value will not match the one expected on

the Oracle DRG. To resolve this, you can change the value that Oracle uses in the Oracle Console (see the instructions that follow).

### To change the CPE IKE identifier that Oracle uses (Oracle Console)

- a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

A list of the IPSec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

- b. For the IPSec connection you're interested in, click the Actions icon (three dots), and then click **Edit**.

The current CPE IKE identifier that Oracle is using is displayed at the bottom of the dialog.

- c. Enter your new values for **CPE IKE Identifier Type** and **CPE IKE Identifier**, and then click **Save Changes**.

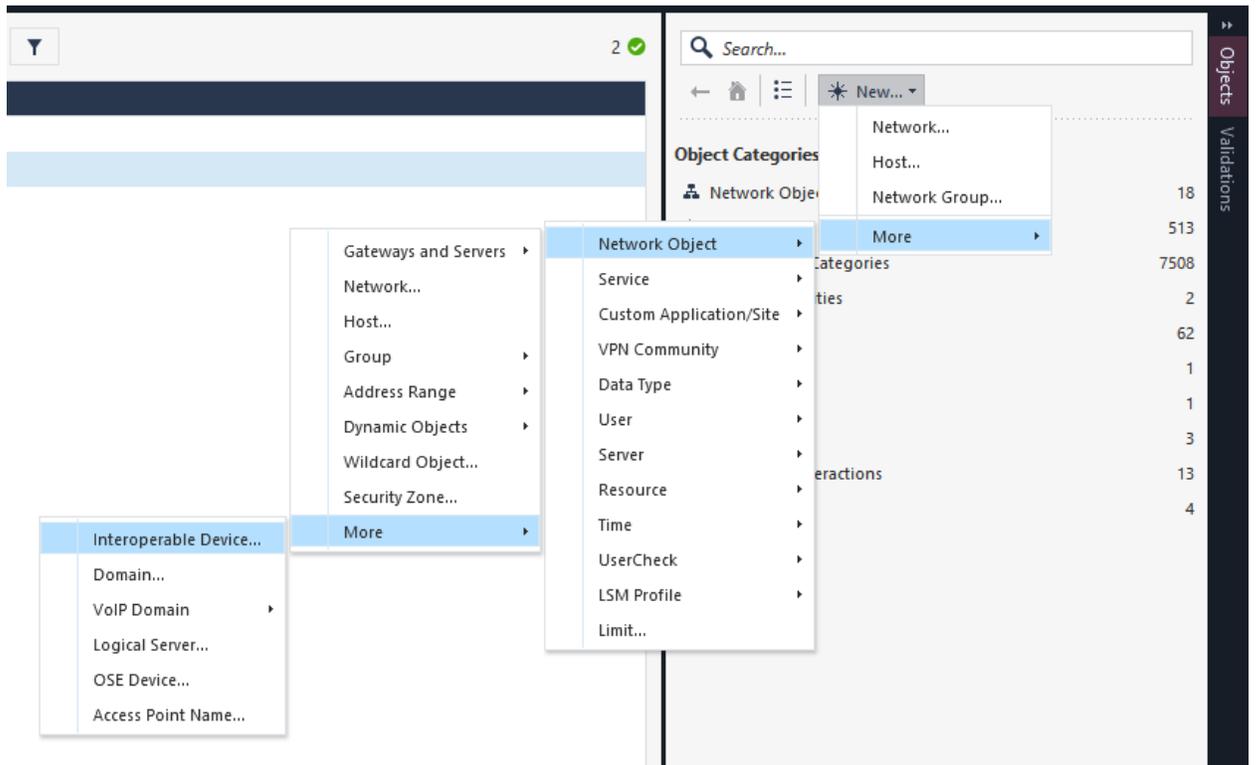
4. Click **OK** to save your changes.

### Task 3: Create an interoperable device

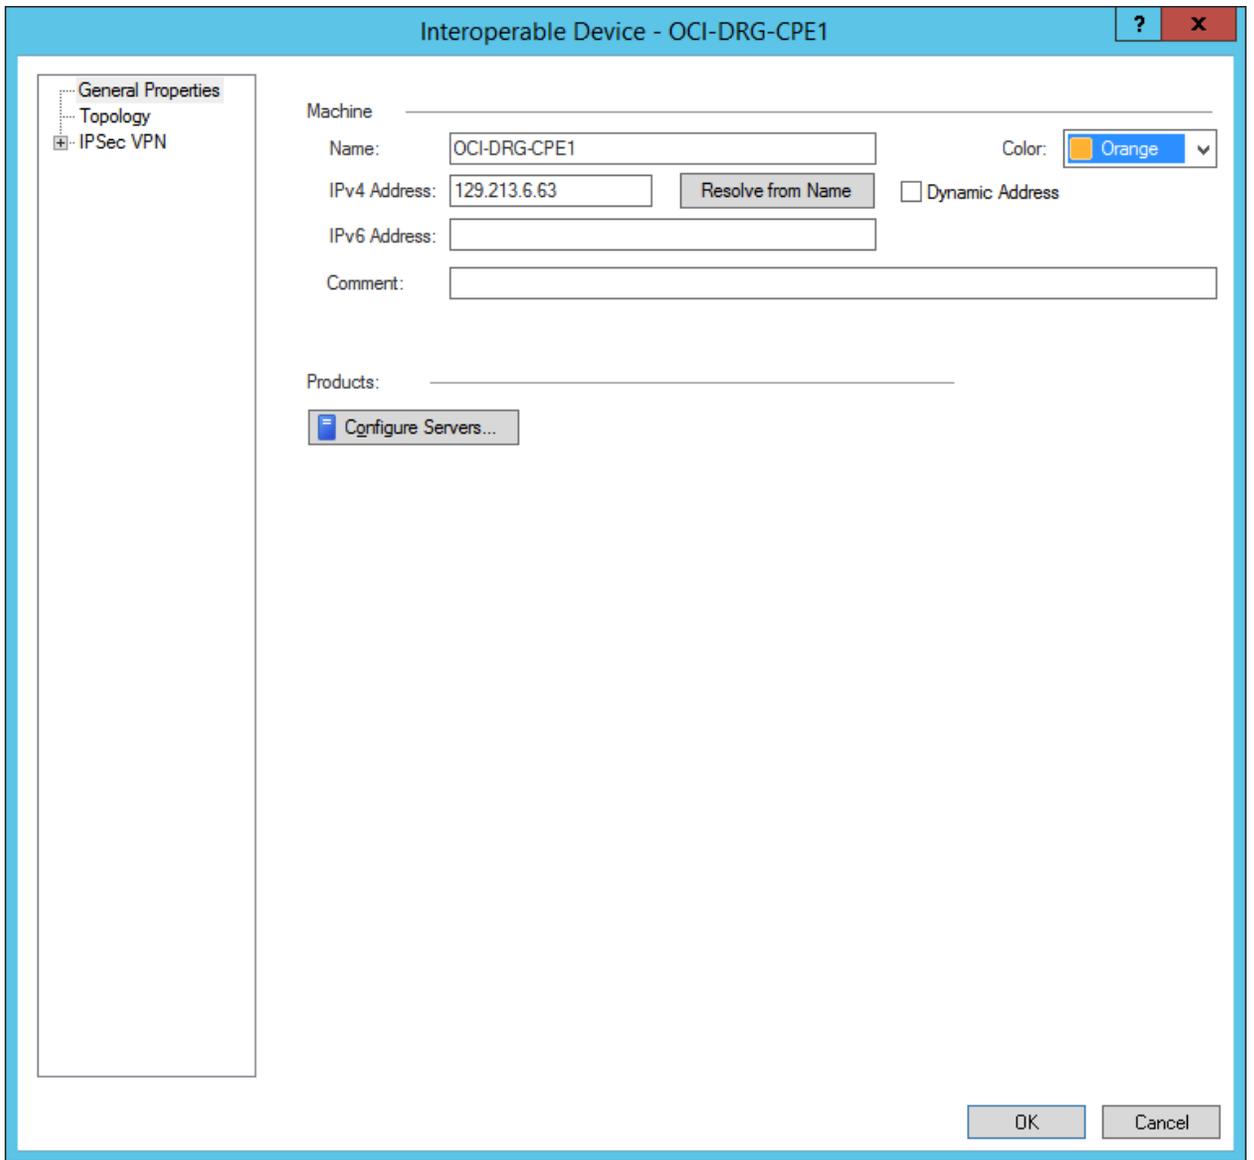
Later, you will create a VPN Community. Before you can, you must create an **Interoperable Device** that will be used in Check Point CloudGuard Security Gateway to define the Oracle DRG.

## CHAPTER 23 Networking

1. Create the new interoperable device.

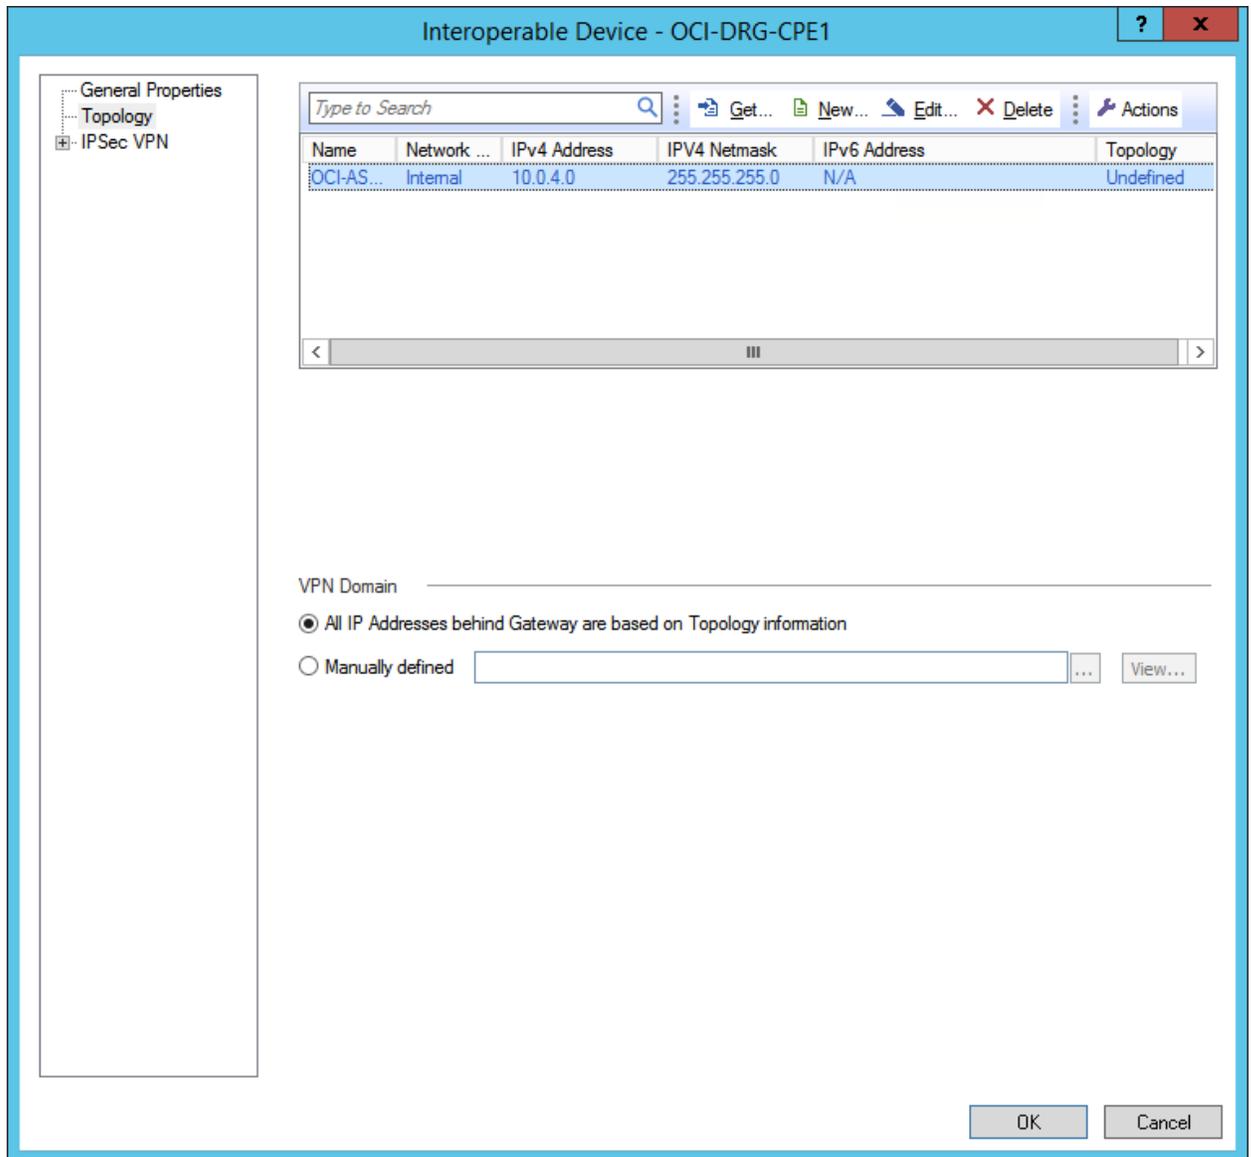


2. On the **General Properties** page of the new interoperable device, add a name to identify the IPsec tunnel. Enter the IP address that Oracle assigned for the Oracle end of the tunnel when creating the IPsec connection.



3. On the **Topology** page, Oracle recommends that you create a new topology by clicking **New** and then adding the Oracle VCN subnets to be used for the tunnel.

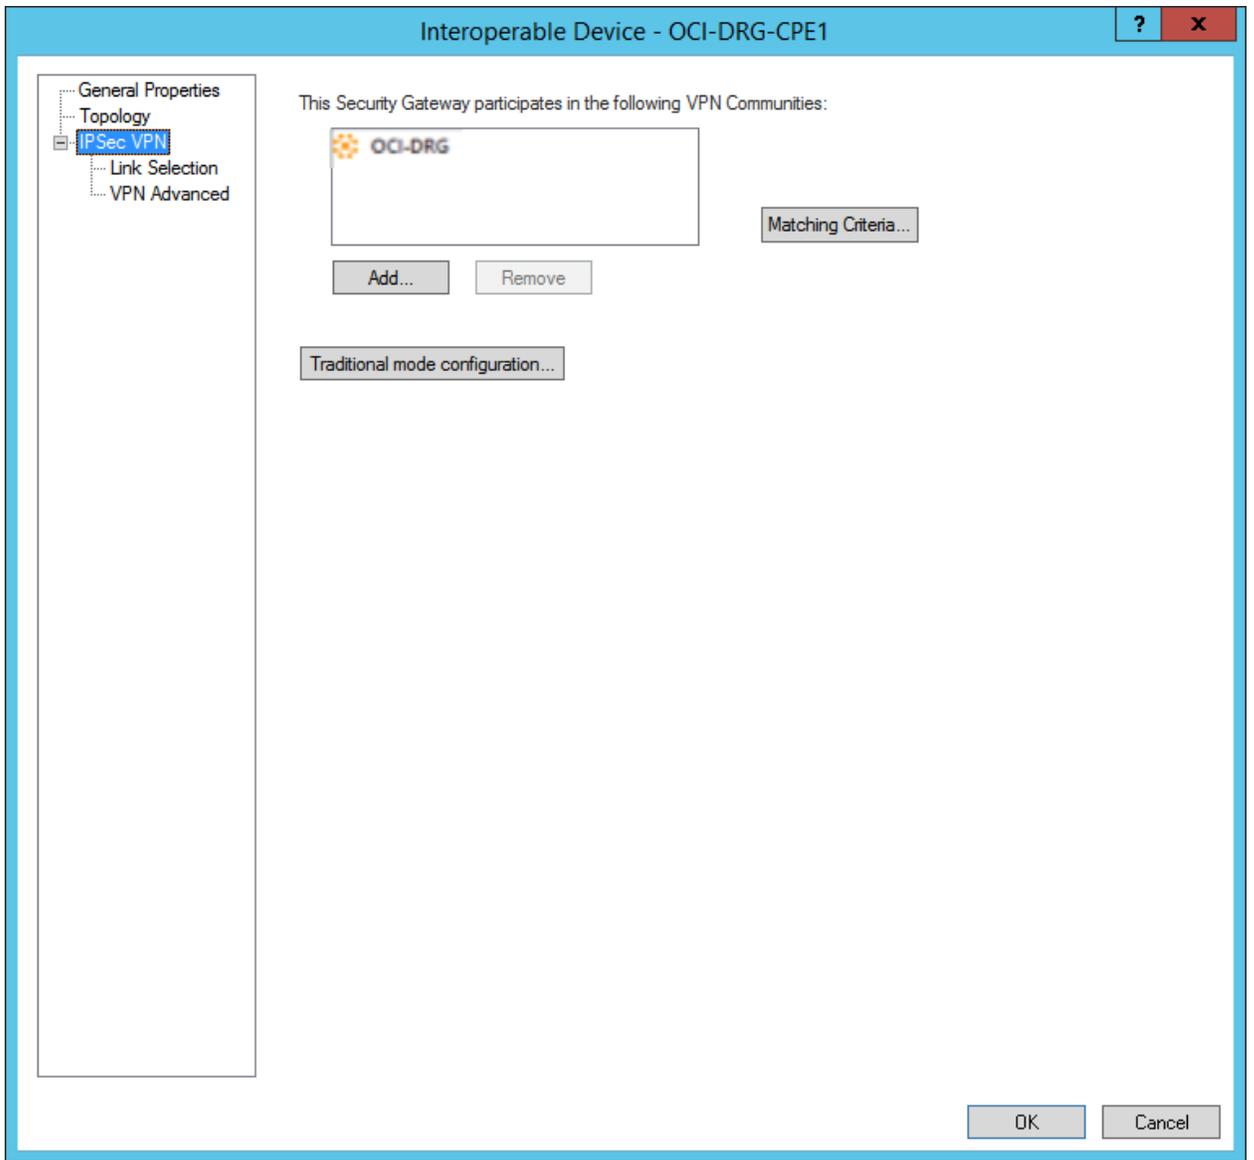
You can instead select the option for **Manually defined**. However, that requires a **Network Object** with all subnets to include in the IPSec Encryption Domain.



4. On the **IPSec VPN** page, you can optionally add the new interoperable device to an existing VPN Community. You can skip this step if you don't yet have any VPN

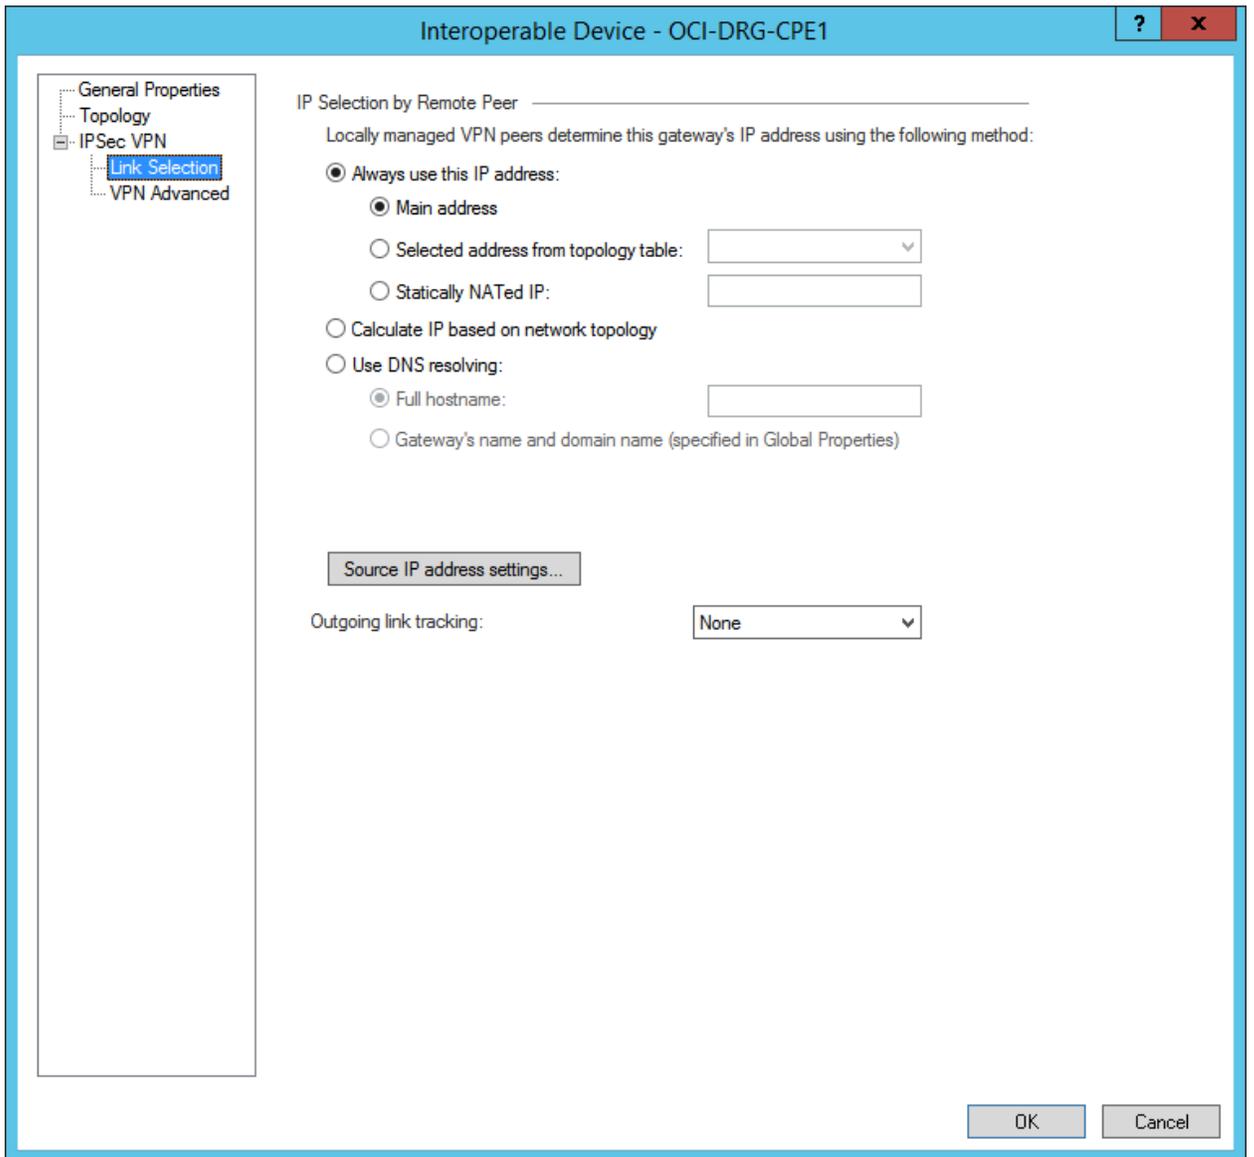
Communities created.

Notice that you skip the **Traditional mode configuration**, because you will define all the Phase 1 and Phase 2 parameters in the VPN Community in a later step. The VPN Community applies those parameters to all interoperable devices that belong to the VPN Community.

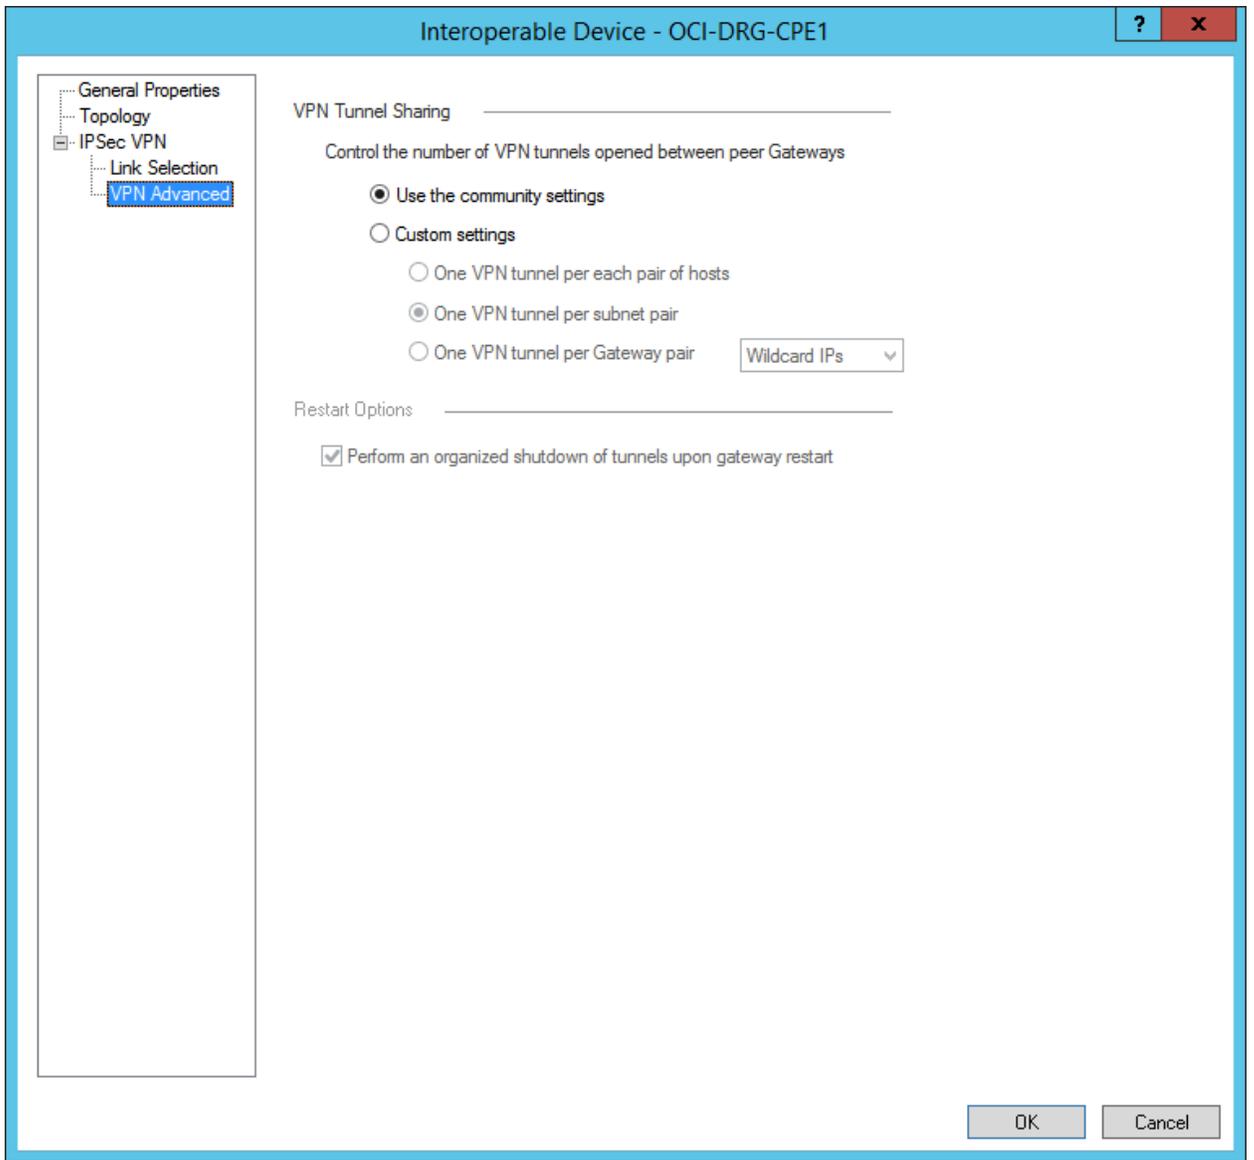


5. On the **Link Selection** page, under **Always use this IP address**, select **Main address**, which is the address that you specified when creating the interoperable

device. If necessary, you can use a specific IP address that will be used as the IKE ID.



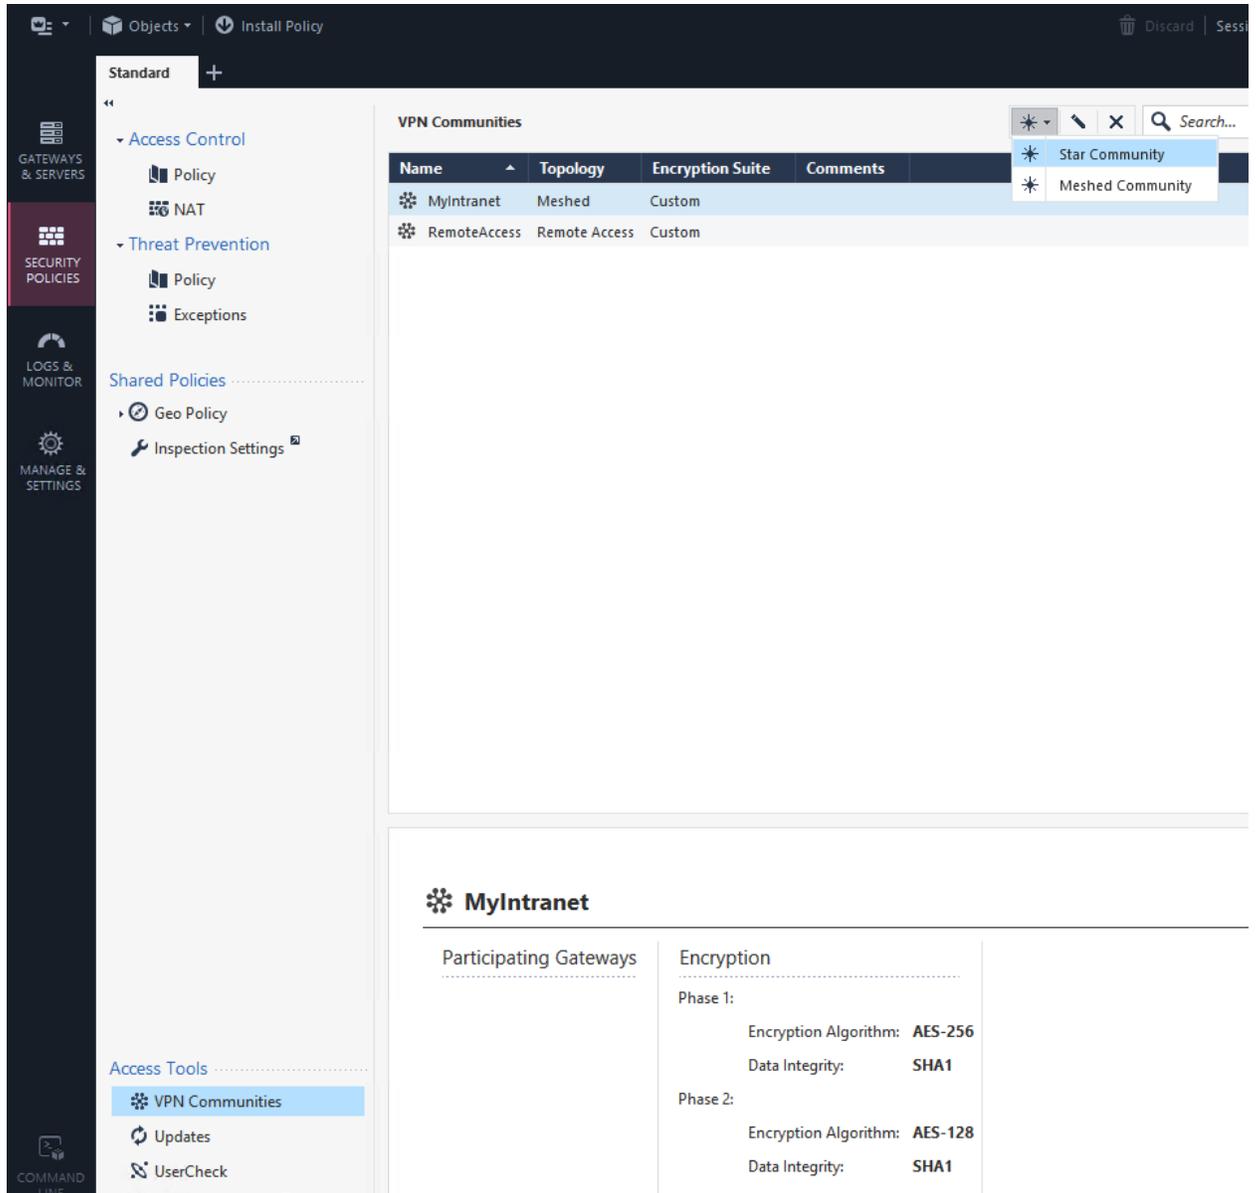
6. On the **VPN Advanced** page, select **Use the community settings**, which applies all the options and values in the VPN Community, including the Phase 1 and Phase 2 parameters.



7. Click **OK** to save your changes.

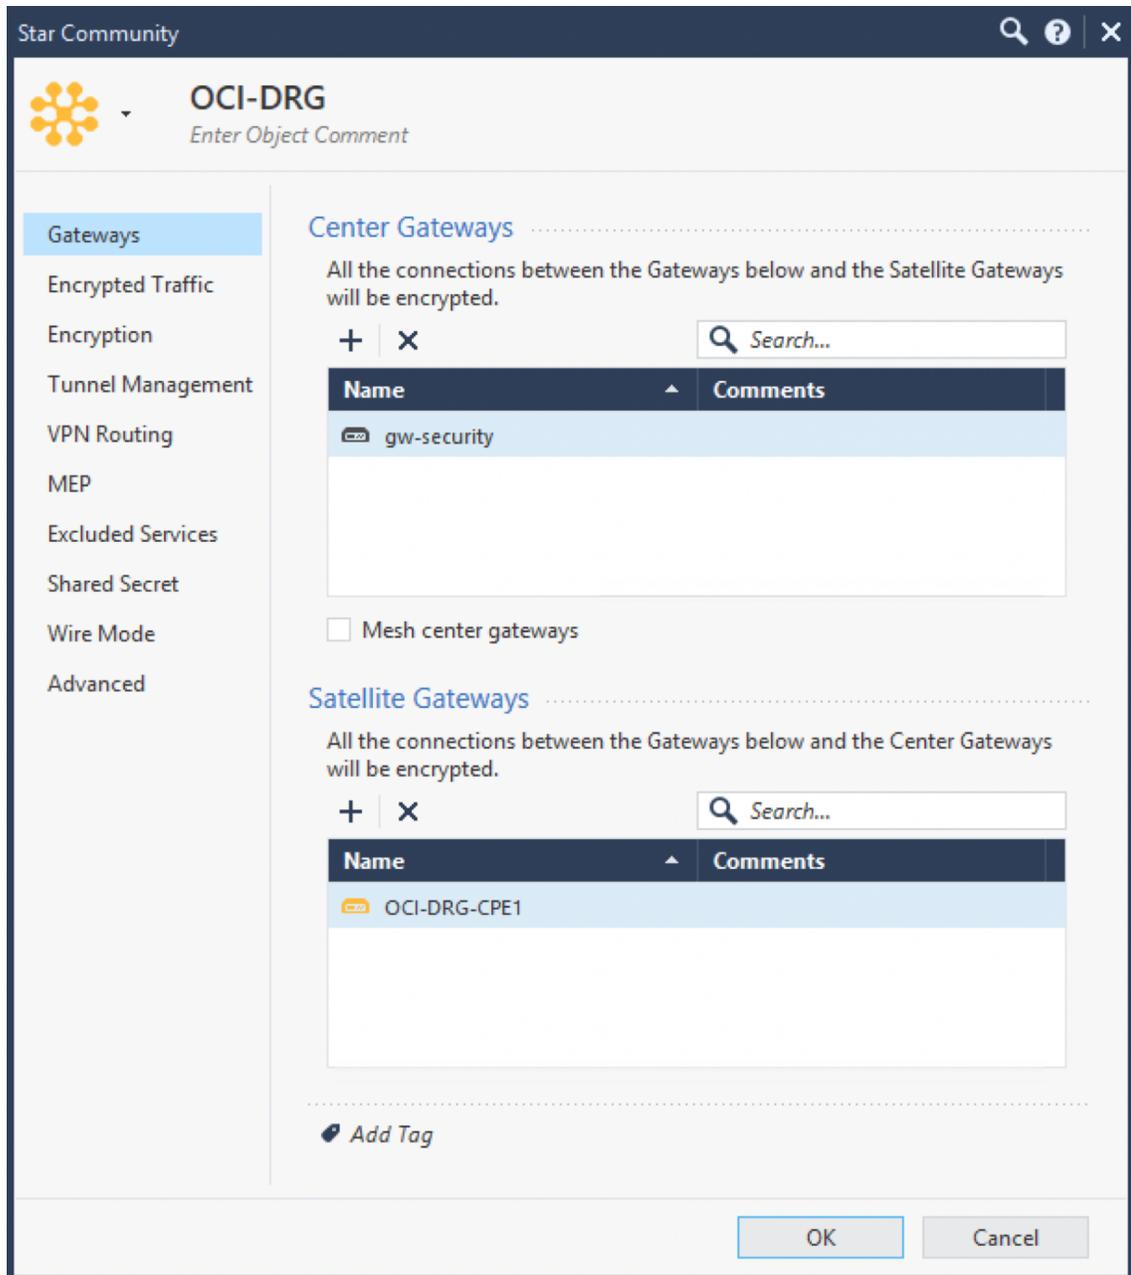
Task 4: Create a VPN community

1. Go to **Security Policies**, and then from **Access Tools**, select **VPN Communities**.
2. Create a **Star Community**.

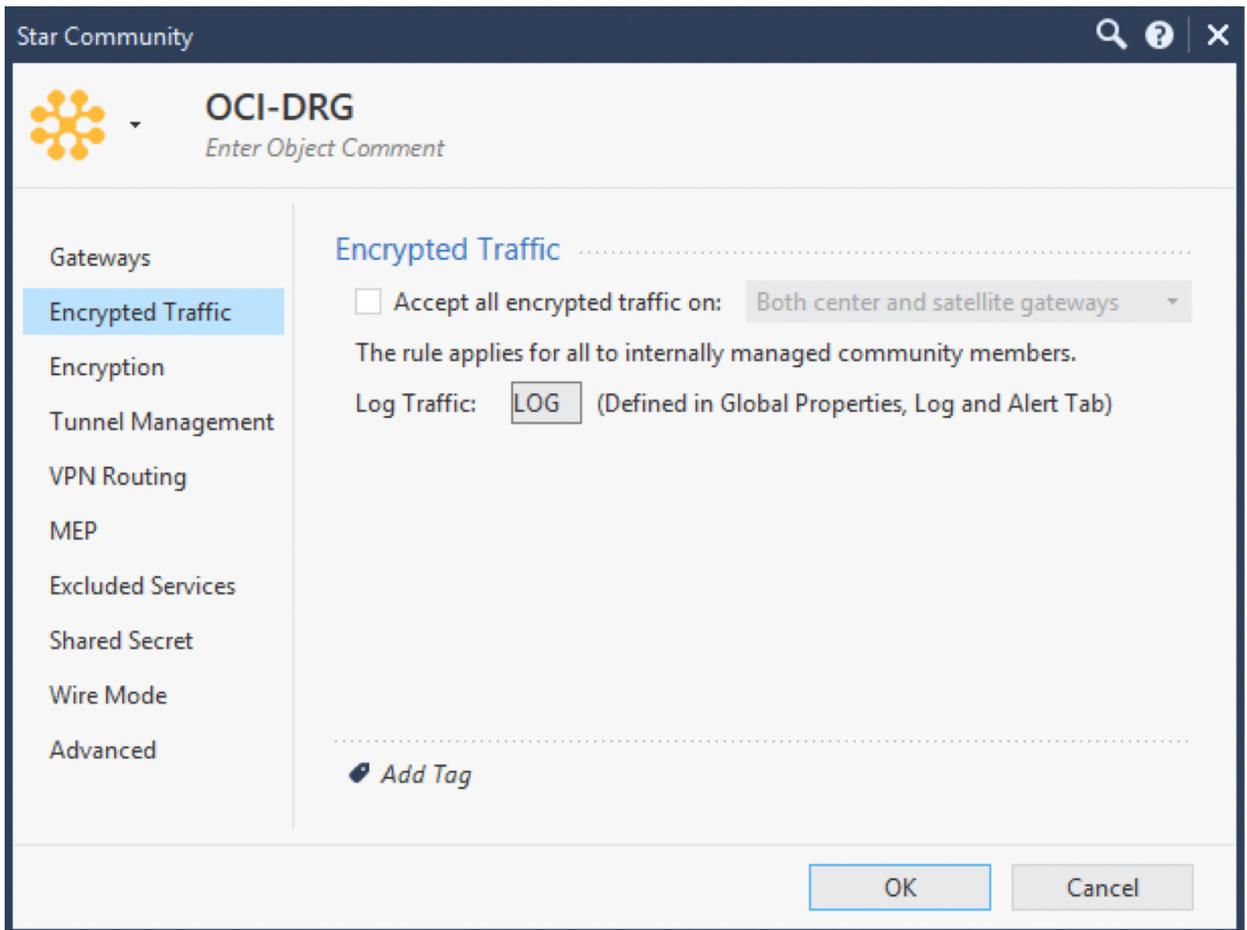


3. For the star community, add a name.

4. On the **Gateways** page, select the values for **Center Gateways** and **Satellite Gateways**. This star community acts as a settings template for the interoperable devices you specify in **Center Gateways** and **Satellite Gateways**.
  - **Center Gateways**: For the Check Point CloudGuard Security Gateway.
  - **Satellite Gateways**: For the CPE that connects to the Oracle DRG for each IPSec tunnel.



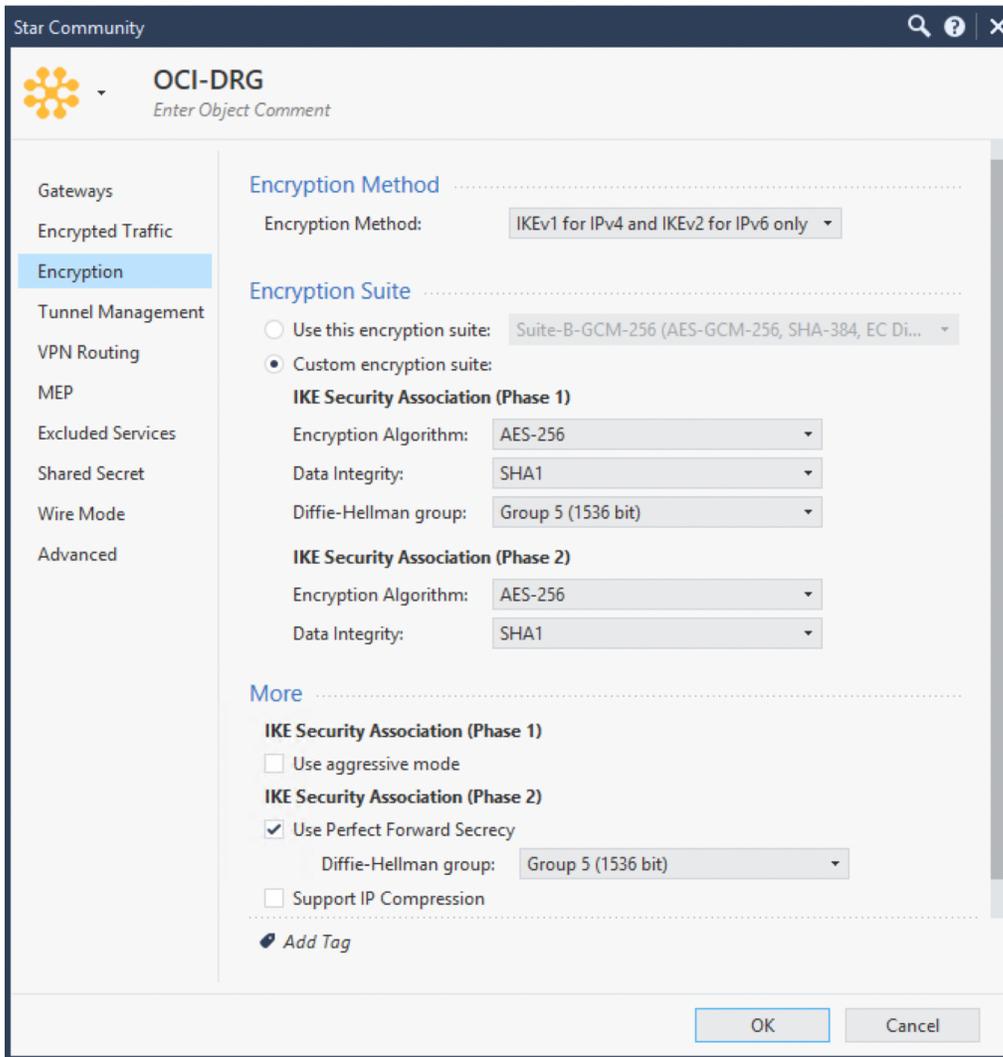
5. If this is a proof of concept (POC) scenario: On the **Encrypted Traffic** page, select the check box for **Accept all encrypted traffic on**. The default value for this setting allows the traffic between both center and satellite gateways. This setting is appropriate for a POC scenario. However, for a production scenario, Oracle recommends that you instead create specific security policies under **Access Control** and on the **Policy** tab. That is covered in the final task in this process.



6. On the **Encryption** page, configure the Phase 1 and Phase 2 parameters that Oracle supports. For a list of those values, see [Supported IPSec Parameters](#).

If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#).

Notice that if you want to use IKEv2, for the **Encryption Method**, instead select **IKEv2 only**.



7. On the **Tunnel Management** page, select **Set Permanent Tunnels**. Oracle recommends that you:
  - Select **On all tunnels in the community** to keep all the Oracle IPsec tunnels up all the time.

- In the **VPN Tunnel Sharing** section, select **One VPN tunnel per Gateway pair**.

The latter option generates only one pair of IPSec security associations (SAs), and each SA with only one security parameter index (SPI) (unidirectional).

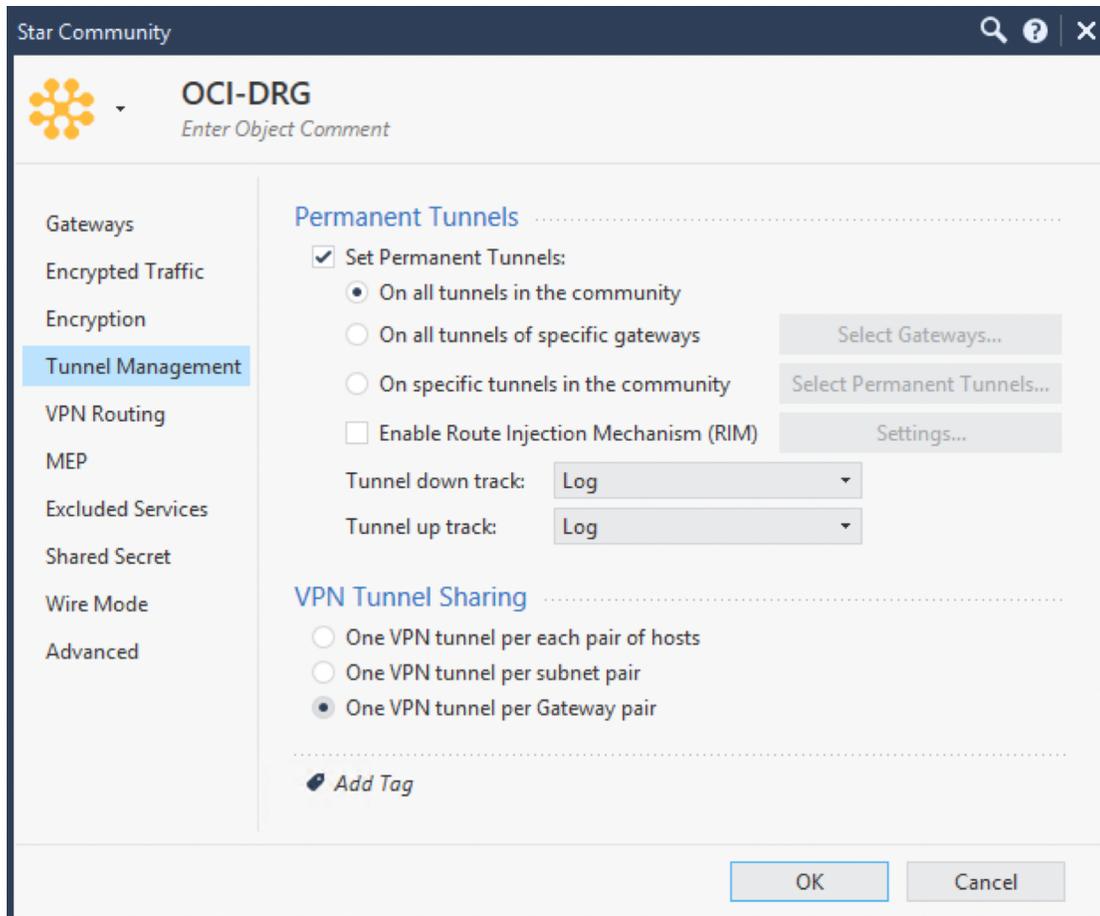
When you use policy-based tunnels, every policy entry generates a pair of IPSec SAs, (also referred to as an *encryption domain*).



### Important

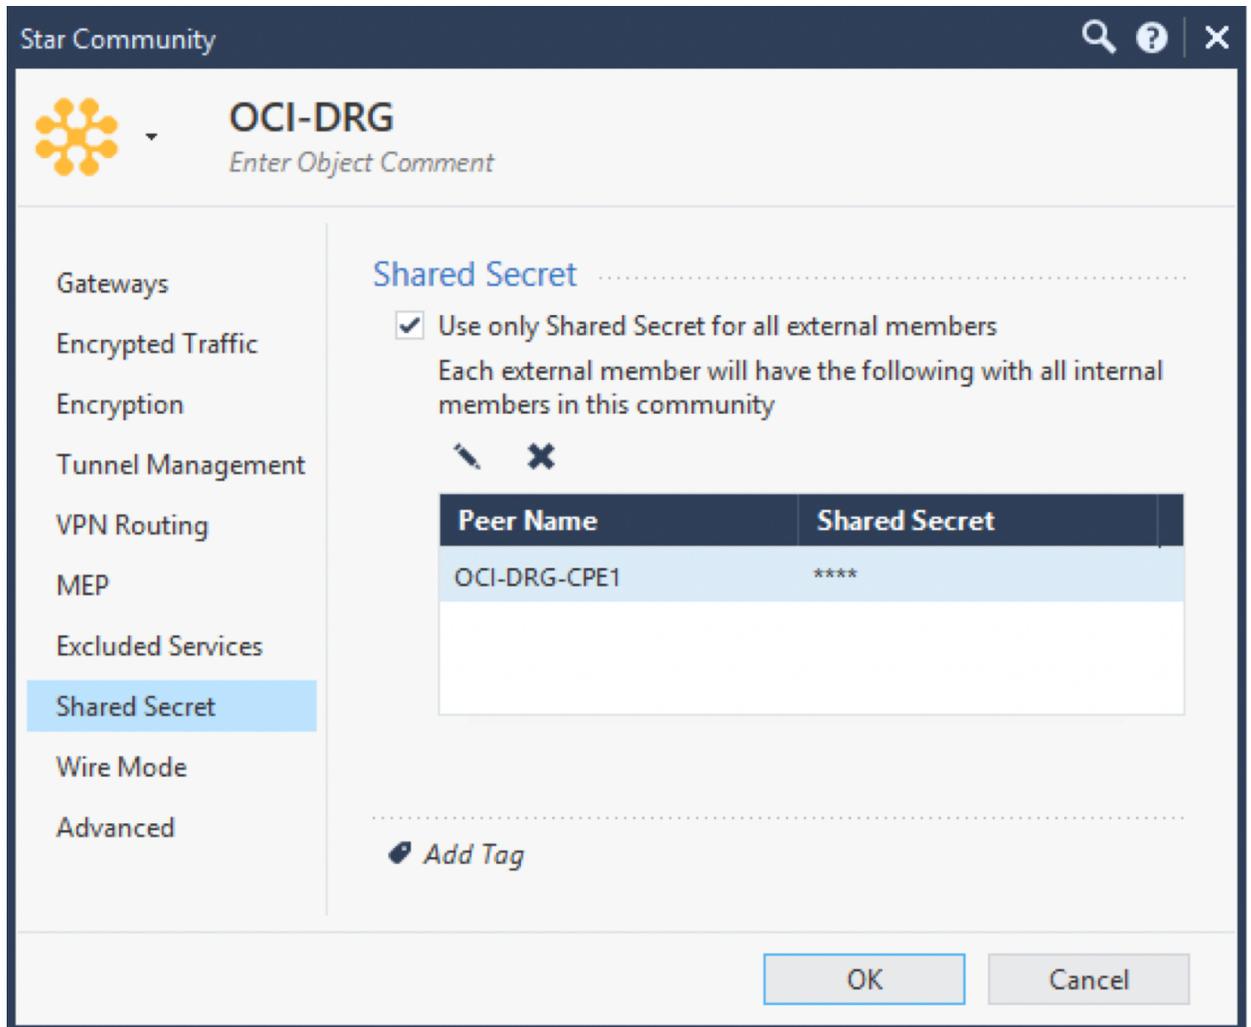
The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

Oracle creates a route-based IPSec connection, which means that everything is routed through an encryption domain that has 0.0.0.0/0 (any) for local traffic and 0.0.0.0/0 (any) for remote traffic. For more information, see [Supported Encryption Domain or Proxy ID](#).



8. On the **Shared Secret** page, select **Use only Shared Secret for all external members**, and add the shared secret that Oracle generated for the tunnel when creating the IPsec connection.

Currently Oracle supports only shared secret keys. Note that you can [change the shared secret](#) in the Oracle Console.



9. Click **OK** to save your changes.

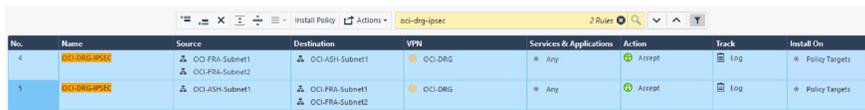
### Task 5: Create a security policy (recommended for a production scenario)

If this is a proof of concept (POC) scenario, earlier you selected **Accept all encrypted**

## CHAPTER 23 Networking

**traffic** on the **Encrypted Traffic** page. If this is instead a production scenario, Oracle recommends creating security policies.

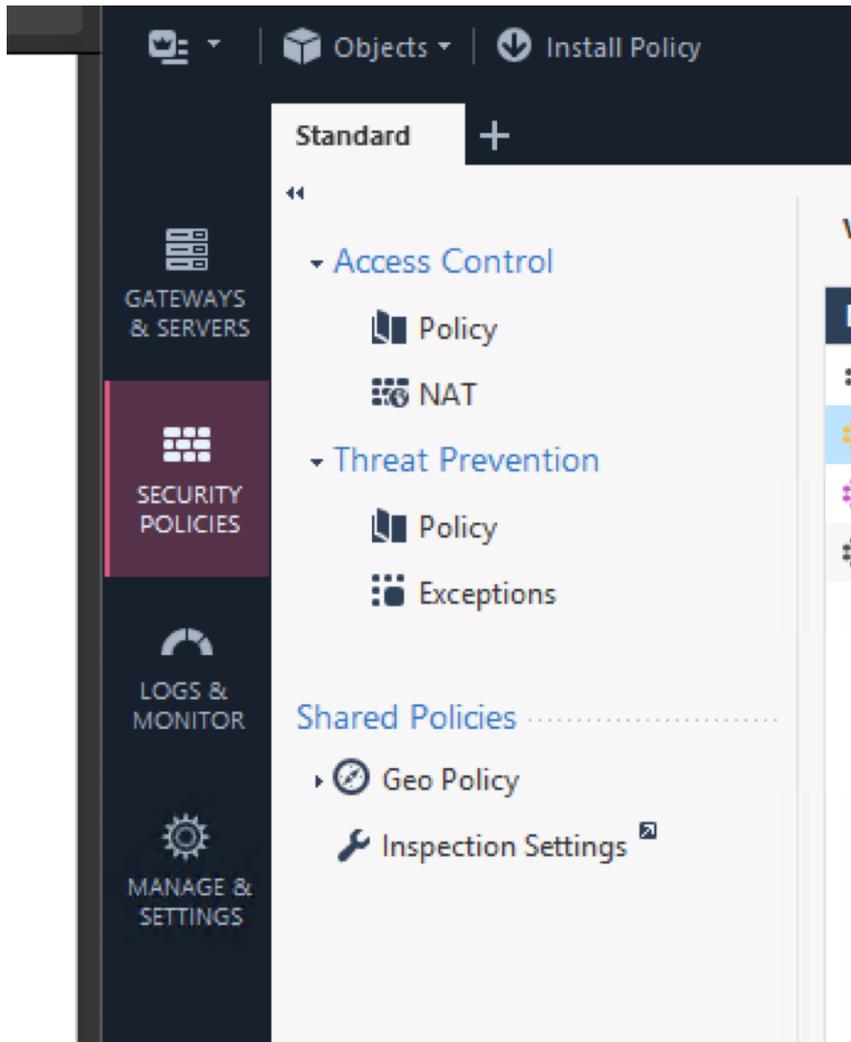
1. Under Security Policies, click Access Control, and then select the Policy tab.
2. Configure the required values.



No.	Name	Source	Destination	VPN	Services & Applications	Action	Track	Install On
4	oci-ash-ash1	OCI-FRA-Subnet1 OCI-FRA-Subnet2	OCI-ASH-Subnet1	OCI-DRG	* Any	Accept	Log	* Policy Targets
5	oci-fra-fra1	OCI-ASH-Subnet1	OCI-FRA-Subnet1 OCI-FRA-Subnet2	OCI-DRG	* Any	Accept	Log	* Policy Targets

3. Click **OK** to save your changes.

4. Click **Install Policy** to apply the configuration.



The IPSec tunnel should now be up.

### Verification

Use options 2 and 4 in the following command to verify security associations (SAs).

```
vpn tunnelutil

***** Select Option *****

(1) List all IKE SAs
(2) * List all IPsec SAs
(3) List all IKE SAs for a given peer (GW) or user (Client)
(4) * List all IPsec SAs for a given peer (GW) or user (Client)
(5) Delete all IPsec SAs for a given peer (GW)
(6) Delete all IPsec SAs for a given User (Client)
(7) Delete all IPsec+IKE SAs for a given peer (GW)
(8) Delete all IPsec+IKE SAs for a given User (Client)
(9) Delete all IPsec SAs for ALL peers and users
(0) Delete all IPsec+IKE SAs for ALL peers and users

* To list data for a specific CoreXL instance, append "-i <instance number>" to your selection.

(Q) Quit

```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Cisco ASA Configuration Options

Choose the configuration based on the ASA software version:

- **9.7.1 or newer:** [Route-based configuration](#)
- **8.5 to 9.7.0:** [Policy-based configuration](#)
- **Older than 8.5:** Not supported by the Oracle configuration instructions. Consider upgrading to a newer version.



### Important

Oracle recommends using a [route-based configuration](#) to avoid interoperability issues and to achieve tunnel redundancy with a single Cisco ASA device.

The Cisco ASA does not support route-based configuration for software versions older than 9.7.1. For the best results, if your device allows it, Oracle recommends that you upgrade to a software version that supports route-based configuration.

With policy-based configuration, you can configure only a single tunnel between your Cisco ASA and your dynamic routing gateway (DRG).

## Cisco ASA: Route-Based

This topic provides a route-based configuration for a Cisco ASA that is running software version 9.7.1 (or newer).

As a reminder, Oracle provides different configurations based on the ASA software:

- **9.7.1 or newer:** Route-based configuration (this topic)
- **8.5 to 9.7.0:** [Policy-based configuration](#)
- **Older than 8.5:** Not supported by the Oracle configuration instructions. Consider upgrading to a newer version.



### Important

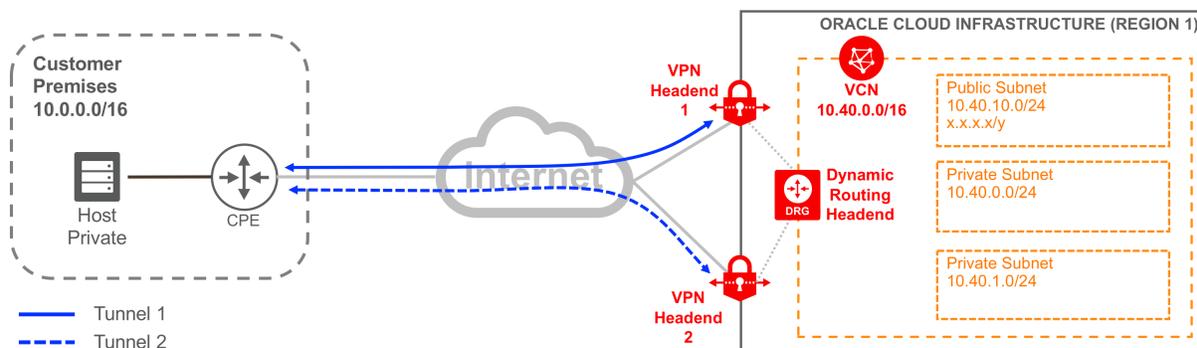
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. The IP addresses in this diagram are examples only and not for literal use.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle

created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### **OTHER IMPORTANT CPE CONFIGURATIONS**

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### **Specific to Cisco ASA: Caveats and Limitations**

This section covers important characteristics and limitations that are specific to Cisco ASA.

#### **TUNNEL MTU AND PATH MTU DISCOVERY**

You have two options for addressing tunnel MTU and path MTU discovery with Cisco ASA:

- [Option 1: TCP MSS adjustment](#)
- [Option 2: Clear/set the Don't Fragment bit](#)

### OPTION 1: TCP MSS ADJUSTMENT

The maximum transmission unit (packet size) through the IPsec tunnel is less than 1500 bytes. You can fragment packets that are too large to fit through the tunnel. Or, you can signal back to the hosts that are communicating through the tunnel that they need to send smaller packets.

You can configure the Cisco ASA to change the maximum segment size (MSS) for any new TCP flows through the tunnel. The ASA looks at any TCP packets where the SYN flag is set and changes the MSS value to the configured value. This configuration might help new TCP flows avoid using path maximum transmission unit discovery (PMTUD).

Use the following command to change the MSS. This command is not part of the sample configuration in the [CPE Configuration](#) section of this topic. Apply the TCP MSS adjustment command manually, if needed.

```
sysopt connection tcpmss 1387
```

### OPTION 2: CLEAR/SET THE DON'T FRAGMENT BIT

Path MTU discovery requires that all TCP packets have the **Don't Fragment** (DF) bit set. If the DF bit is set and a packet is too large to go through the tunnel, the ASA drops the packet when it arrives. The ASA sends an ICMP packet back to the sender indicating that the received packet was too large for the tunnel. The ASA offers three options for handling the DF bit. Choose one of the options and apply it to the configuration:

- **Set the DF bit (recommended):** Packets have the DF bit set in their IP header. The ASA may still fragment the packet if the original received packet cleared the DF bit.

```
crypto ipsec df-bit set-df ${outsideInterface}
```

- **Clear the DF bit:** The DF bit is cleared in the packet's IP header. Allows the packet to be fragmented and sent to the end host in Oracle Cloud Infrastructure for reassembly.

```
crypto ipsec df-bit clear-df ${outsideInterface}
```

- **Ignore (copy) the DF bit:** The ASA looks at the original packet's IP header information and copies the DF bit setting.

```
crypto ipsec df-bit copy-df ${outsideInterface}
```

### **VPN TRAFFIC MIGHT ENTER ONE TUNNEL AND EXIT ANOTHER**

If VPN traffic enters an interface with the same security level as an interface toward the packet's next hop, you must allow that traffic. By default, the packets between interfaces that have identical security levels on your ASA are dropped.

Add the following command manually if you need to permit traffic between interfaces with the same security levels. This command is not part of the sample configuration in the [CPE Configuration](#) section.

```
same-security-traffic permit inter-interface
```

### **General Caveats and Limitations**

This section covers general characteristics and limitations of VPN Connect.

#### **ASYMMETRIC ROUTING**

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

#### **ROUTE-BASED OR POLICY-BASED IPSEC VPN**

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an IPsec tunnel, the packet is encrypted and sent to the other end of the tunnel.
- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

## Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPsec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



#### **Important**

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

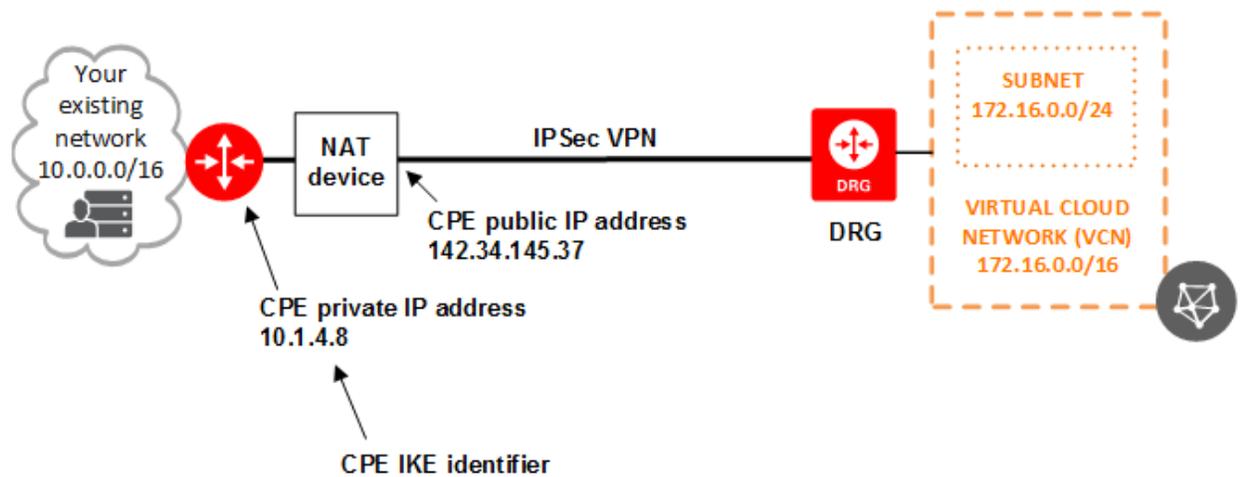
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

**If Your CPE Is Behind a NAT Device**

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.





### Note

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

### Supported IPsec Parameters

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

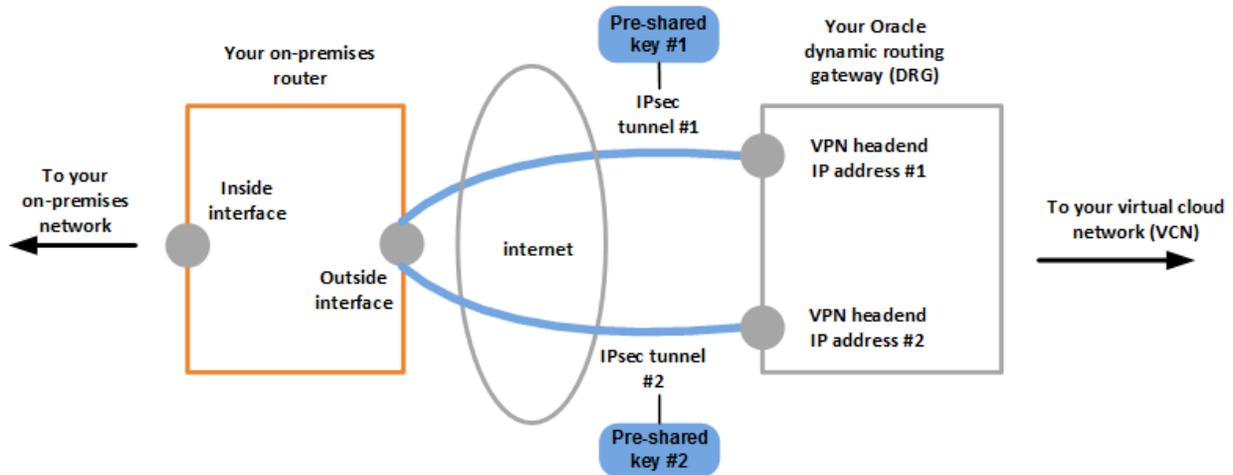
### CPE Configuration



### Important

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPsec connection.



The configuration template provided is for a Cisco router running Cisco ASA 9.7.1 software (or later). The template provides information for each tunnel that you must configure. Oracle recommends setting up all configured tunnels for maximum redundancy.

The configuration template refers to these items that you must provide:

- **CPE public IP address:** The internet-routable IP address that is assigned to the external interface on the CPE. You or your Oracle administrator provides this value to Oracle when creating the CPE object in the Oracle Console.
- **Inside tunnel interface (required if using BGP):** The IP addresses for the CPE and Oracle ends of the inside tunnel interface. You provide these values when creating the IPsec connection in the Oracle Console.
- **BGP ASN (required if using BGP):** Your BGP ASN.

In addition, you must:

- Configure internal routing that routes traffic between the CPE and your local network.
- Ensure that you permit traffic between your ASA and your Oracle VCN.

- Identify the IPSec profile used (the following configuration template references this group policy as `oracle-vcn-vpn-policy`).
- Identify the transform set used for your crypto map (the following configuration template references this transform set as `oracle-vcn-transform`).
- Identify the virtual tunnel interface names used (the following configuration template references these as variables `${tunnelInterfaceName1}` and `${tunnelInterfaceName2}`).



### Important

This following configuration template from Oracle Cloud Infrastructure **is a starting point for what you need to apply to your CPE**. Some of the parameters referenced in the template must be unique on the CPE, and the uniqueness can only be determined by accessing the CPE. Ensure that the parameters are valid on your CPE and do not overwrite any previously configured values. In particular, ensure these values are unique:

- Policy names or numbers
- Interface names or numbers
- Access list numbers (if applicable)

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

There's a separate configuration template for IKEv1 versus IKEv2.

### IKEv1 Configuration Template

```
!-----
!-----
! IKEv1 Configuration Template
! The configuration consists of two IPsec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.
!-----
!-----
! The configuration template involves setting up the following:
! ISAKMP Policy
! IPsec Configuration
! IPsec Tunnel Group Configuration
! VTI Interface Configuration
! IP Routing (BGP or Static)
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
provide these values when creating the IPsec connection in the Oracle Console.
! ${OracleInsideTunnelIpAddress2} = Inside tunnel IP address of Oracle-side for the second tunnel. You
provide these values when creating the IPsec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${oracleHeadend2} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPsec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${sharedSecret2} = You provide when you set up the IPsec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
! ${tunnelInterfaceName1} = The name of the first VTI used on your ASA.
! ${tunnelInterfaceName2} = The name of the second VTI used on your ASA.
! ${cpeInsideTunnelIpAddress1} = The CPE's inside tunnel IP for the first tunnel.
```

## CHAPTER 23 Networking

---

```
! ${cpeInsideTunnelIpAddress2} = The CPE's inside tunnel IP for the second tunnel.
! ${cpeInsideTunnelNetmask1} = The CPE's inside tunnel netmask for the first tunnel.
! ${cpeInsideTunnelNetmask2} = The CPE's inside tunnel netmask for the second tunnel.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
!-----

! ISAKMP Policy

! ISAKMP Phase 1 configuration.
! IKEv1 is enabled on the outside interface.
! IKEv1 policy is created for Phase 1 which specifies to use a Pre-Shared Key, AES256, SHA1, Diffie-
Hellman Group 5, and a Phase 1 lifetime of 28800 seconds (8 hours).
! If different parameters are required, modify this template before applying the configuration.
! WARNING: The IKEv1 group policy is created with a priority of 10. Make sure this doesn't conflict with
any pre-existing configuration on your ASA.

crypto ikev1 enable $(outsideInterface)

crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
 lifetime 28800

! IPSec Configuration

! Create an IKEv1 transform set named 'oracle-vcn-transform' which defines a combination of IPSec (Phase
2) policy options. Specifically, AES256 for encryption and SHA1 for authentication.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec ikev1 transform-set oracle-vcn-transform esp-aes-256 esp-sha-hmac

! A IPSec profile named 'oracle-vcn-vpn-policy' is created.
! The previously created transform set is added to this policy along with settings for enabling PFS
Group 5 and the security association lifetime to 3600 seconds (1 hour).
! If different parameters are required, modify this template before applying the configuration.
```

## CHAPTER 23 Networking

---

```
crypto ipsec profile oracle-vcn-vpn-policy
 set ikev1 transform-set oracle-vcn-transform
 set pfs group5
 set security-association lifetime seconds 3600

! IPsec Tunnel Group Configuration

! A tunnel group is created for each Oracle VPN Headend. Each tunnel group defines the pre-shared key
used for each respective tunnel.

tunnel-group ${oracleHeadend1} type ipsec-l2l
tunnel-group ${oracleHeadend1} ipsec-attributes
 ikev1 pre-shared-key ${sharedSecret1}

tunnel-group ${oracleHeadend2} type ipsec-l2l
tunnel-group ${oracleHeadend2} ipsec-attributes
 ikev1 pre-shared-key ${sharedSecret2}

! VTI Interface Configuration

! A virtual tunnel interface (VTI) is a logical interface representing the local end of a VPN tunnel to
a remote VPN peer. Two VTIs are created representing two tunnels, one to each Oracle VPN Headend. The IP
address of each VPN headend is provided when you create your IPsec connection in Oracle Console.
! All traffic routed to a VTI will be encrypted and sent across the tunnel towards Oracle Cloud
Infrastructure.
! Each VTI configuration also references the previously created IPsec profile 'oracle-vcn-vpn-policy'
for its IPsec parameters.

interface ${tunnelInterfaceName1}
 nameif ORACLE-VPN1
 ip address ${cpeInsideTunnelIpAddress1} ${cpeInsideTunnelNetmask1}
 tunnel source interface ${outsideInterface}
 tunnel destination ${oracleHeadend1}
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile oracle-vcn-vpn-policy

interface ${tunnelInterfaceName2}
 nameif ORACLE-VPN2
 ip address ${cpeInsideTunnelIpAddress2} ${cpeInsideTunnelNetmask2}
 tunnel source interface ${outsideInterface}
```

## CHAPTER 23 Networking

---

```
tunnel destination #{oracleHeadend2}
tunnel mode ipsec ipv4
tunnel protection ipsec profile oracle-vcn-vpn-policy

! IP Routing
! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying
configuration.

! Border Gateway Protocol (BGP) Configuration
! Uncomment below lines if you want to use BGP.

! router bgp #{bgpASN}
! address-family ipv4 unicast
! neighbor #{OracleInsideTunnelIpAddress1} remote-as 31898
! neighbor #{OracleInsideTunnelIpAddress1} activate
! neighbor #{OracleInsideTunnelIpAddress2} remote-as 31898
! neighbor #{OracleInsideTunnelIpAddress2} activate
! network #{onPremCidrNetwork} mask #{onPremCidrNetmask}
! no auto-summary
! no synchronization
! exit-address-family

! Static Route Configuration
! Each static route references the other VTI by its nameif value.
! Uncomment below line if you want to use static routing.

! route ORACLE-VPN1 #{VcnCidrNetwork} #{VcnCidrNetmask} #{OracleInsideTunnelIpAddress1} 1 track
! route ORACLE-VPN2 #{VcnCidrNetwork} #{VcnCidrNetmask} #{OracleInsideTunnelIpAddress2} 100

! Configuration for Tunnel Failover

! Uncomment the below IP SLA lines if using static routing.
! Use this IP SLA configuration for static route failover This IP SLA configuration is used for static
route failover between the two tunnels.
! Make sure that the SLA monitor and tracking numbers used do not conflict with any existing
configuration on your ASA.

! sla monitor 10
! type echo protocol ipIcmpEcho #{oracleHeadend1} interface outside
! frequency 5
! sla monitor schedule 10 life forever start-time now
```

```
! track 1 rtr 10 reachability
```

### IKEv2 Configuration Template

```
!-----
!-----
! IKEv2 Configuration Template
! The configuration consists of two IPSec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.
!-----
!-----
! The configuration template involves setting up the following:
! IKEv2 Policy
! IPSec Configuration
! IPSec Tunnel Group Configuration
! VTI Interface Configuration
! IP Routing (BGP or Static)
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${OracleInsideTunnelIpAddress2} = Inside tunnel IP address of Oracle-side for the second tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${oracleHeadend2} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${sharedSecret2} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
```

## CHAPTER 23 Networking

---

```
! ${tunnelInterfaceName1} = The name of the first VTI used on your ASA.
! ${tunnelInterfaceName2} = The name of the second VTI used on your ASA.
! ${cpeInsideTunnelIpAddress1} = The CPE's inside tunnel IP for the first tunnel.
! ${cpeInsideTunnelIpAddress2} = The CPE's inside tunnel IP for the second tunnel.
! ${cpeInsideTunnelNetmask1} = The CPE's inside tunnel netmask for the first tunnel.
! ${cpeInsideTunnelNetmask2} = The CPE's inside tunnel netmask for the second tunnel.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
!-----

! IKEv2 Policy

! IKEv2 is enabled on the outside interface.
! IKEv2 policy is created and specifies use of a Pre-Shared Key, AES256, SHA1, Diffie-Hellman Group 5,
and a lifetime of 28800 seconds (8 hours).
! If different parameters are required, modify this template before applying the configuration.
! WARNING: The IKEv2 group policy is created with a priority of 10. Make sure this doesn't conflict with
any pre-existing configuration on your ASA.

crypto ikev2 enable ${outsideInterface}

crypto ikev2 policy 10
 encryption aes-256
 integrity sha
 group 5
 prf sha
 lifetime seconds 28800

! IPSec Configuration

! Create an IKEv2 IPSec proposal named 'oracle_v2_ipsec_proposal' which defines AES256 for encryption
and SHA1 for authentication.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec ikev2 ipsec-proposal oracle_v2_ipsec_proposal
 protocol esp encryption aes-256
 protocol esp integrity sha-1
```

## CHAPTER 23 Networking

---

```
! An IPSec profile named 'oracle-vcn-vpn-policy' is created.
! The previously created IPSec proposal is added to this policy along with settings for enabling PFS
Group 5 and the security association lifetime to 3600 seconds (1 hour).
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec profile oracle-vcn-vpn-policy
 set ikev2 ipsec-proposal oracle_v2_ipsec_proposal
 set pfs group5
 set security-association lifetime seconds 3600

! IPSec Tunnel Group Configuration

group-policy oracle_v2_group_policy internal
group-policy oracle_v2_group_policy attributes
 vpn-tunnel-protocol ikev2

! A tunnel group is created for each Oracle VPN Headend. Each tunnel group defines the pre-shared key
used for each respective tunnel.

tunnel-group ${oracleHeadend1} type ipsec-l2l
tunnel-group ${oracleHeadend1} general-attributes
 default-group-policy oracle_v2_group_policy
tunnel-group ${oracleHeadend1} ipsec-attributes
 ikev2 local-authentication pre-shared-key ${sharedSecret1}
 ikev2 remote-authentication pre-shared-key ${sharedSecret1}

tunnel-group ${oracleHeadend2} type ipsec-l2l
tunnel-group ${oracleHeadend2} general-attributes
 default-group-policy oracle_v2_group_policy
tunnel-group ${oracleHeadend2} ipsec-attributes
 ikev2 local-authentication pre-shared-key ${sharedSecret2}
 ikev2 remote-authentication pre-shared-key ${sharedSecret2}

! VTI Interface Configuration

! A virtual tunnel interface (VTI) is a logical interface representing the local end of a VPN tunnel to
a remote VPN peer. Two VTIs are created representing two tunnels, one to each Oracle VPN Headend. The IP
address of each VPN headend is provided when you create your IPSec connection in Oracle Console.
! All traffic routed to a VTI will be encrypted and sent across the tunnel towards Oracle Cloud
Infrastructure.
! Each VTI configuration also references the previously created IPSec profile 'oracle-vcn-vpn-policy'
```

## CHAPTER 23 Networking

---

```
for its IPsec parameters.

interface ${tunnelInterfaceName1}
 nameif ORACLE-VPN1
 ip address ${cpeInsideTunnelIpAddress1} ${cpeInsideTunnelNetmask1}
 tunnel source interface ${outsideInterface}
 tunnel destination ${oracleHeadend1}
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile oracle-vcn-vpn-policy

interface ${tunnelInterfaceName2}
 nameif ORACLE-VPN2
 ip address ${cpeInsideTunnelIpAddress2} ${cpeInsideTunnelNetmask2}
 tunnel source interface ${outsideInterface}
 tunnel destination ${oracleHeadend2}
 tunnel mode ipsec ipv4
 tunnel protection ipsec profile oracle-vcn-vpn-policy

! IP Routing
! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying
configuration.

! Border Gateway Protocol (BGP) Configuration
! Uncomment below lines if you want to use BGP.

! router bgp ${bgpASN}
! address-family ipv4 unicast
! neighbor ${OracleInsideTunnelIpAddress1} remote-as 31898
! neighbor ${OracleInsideTunnelIpAddress1} activate
! neighbor ${OracleInsideTunnelIpAddress2} remote-as 31898
! neighbor ${OracleInsideTunnelIpAddress2} activate
! network ${onPremCidrNetwork} mask ${onPremCidrNetmask}
! no auto-summary
! no synchronization
! exit-address-family

! Static Route Configuration
! Each static route references the other VTI by its nameif value.
! Uncomment below line if you want to use static routing.

! route ORACLE-VPN1 ${VcnCidrNetwork} ${VcnCidrNetmask} ${OracleInsideTunnelIpAddress1} 1 track
```

## CHAPTER 23 Networking

---

```
! route ORACLE-VPN2 ${VcnCidrNetwork} ${VcnCidrNetmask} ${OracleInsideTunnelIpAddress2} 100

! Configuration for Tunnel Failover

! Uncomment the below IP SLA lines if using static routing.
! Use this IP SLA configuration for static route failover This IP SLA configuration is used for static
route failover between the two tunnels.
! Make sure that the SLA monitor and tracking numbers used do not conflict with any existing
configuration on your ASA.

! sla monitor 10
! type echo protocol ipIcmpEcho ${oracleHeadend1} interface outside
! frequency 5
! sla monitor schedule 10 life forever start-time now

! track 1 rtr 10 reachability
```

### Verification

The following ASA commands are included for basic troubleshooting. For more exhaustive information, refer to Cisco's [IPSec Troubleshooting](#) document.

Use the following command to verify that ISAKMP security associations are being built between the two peers.

```
show crypto isakmp sa
```

Use the following command to verify the status of all your BGP connections.

```
show bgp summary
```

Use the following command to verify the ASA's route table.

```
show route
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Cisco ASA: Policy-Based

This topic provides a policy-based configuration for a Cisco ASA that is running software version 8.5 to 9.7.0.

As a reminder, Oracle provides different configurations based on the ASA software:

- **9.7.1 or newer:** [Route-based configuration](#)
- **8.5 to 9.7.0:** Policy-based configuration (this topic)
- **Older than 8.5:** Not supported by the Oracle configuration instructions. Consider upgrading to a newer version.



#### Important

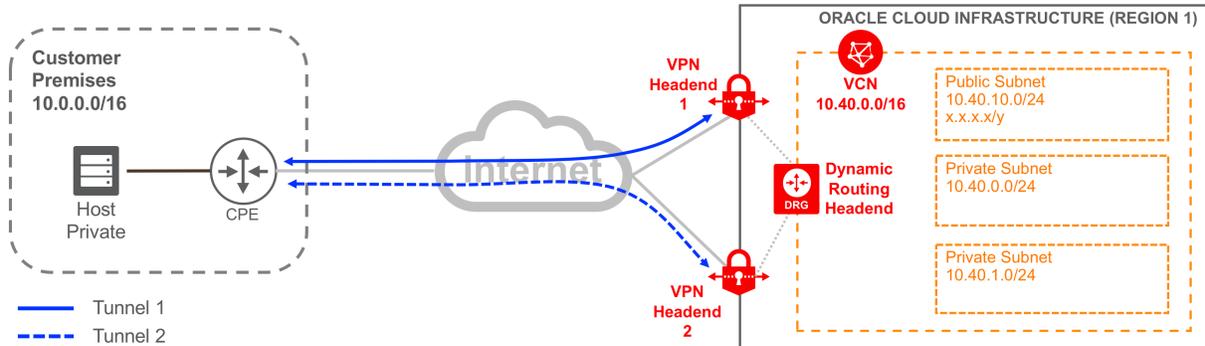
Oracle recommends using a [route-based configuration](#) to avoid interoperability issues and to achieve tunnel redundancy with a single Cisco ASA device.

The Cisco ASA does not support route-based configuration for software versions older than 9.7.1. For the best results, if your device allows it, Oracle recommends that you upgrade to a software version that supports route-based configuration.

With policy-based configuration, you can configure only a single tunnel between your Cisco ASA and your dynamic routing gateway (DRG).

VPN Connect is the IPSec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPSec connection to Oracle Cloud Infrastructure with redundant tunnels. The IP addresses in this diagram are examples only and not for literal use.



### Important

Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPSec, see the list of [supported IPSec parameters](#) and consult your vendor's documentation for assistance.

### Best Practices

This section covers best practices and considerations for using VPN Connect.

### SPECIFIC TO CISCO ASA: VPN FILTERS

VPN filters let you further filter traffic either before it enters or after it exits a tunnel. Use VPN filters if you need additional granularity for filtering different traffic types or source/destination flows. For more information, see Cisco's [VPN Filter documentation](#).

VPN filter configuration is not included in the configuration template that appears in the [CPE Configuration](#) section. To use VPN filters, add the following configuration items manually.

- **Access control list (ACL):** Create an ACL that the VPN filter can use to restrict the traffic permitted through the tunnels. If you have an ACL already used for a VPN filter, do not also use it for an interface access group.

```
access-list #{vpnFilterAclName} extended permit ip #{VcnCidrNetwork} #{VcnCidrNetmask}
#{onPremCidrNetwork} #{onPremCidrNetmask}
```

- **Group policy:** Apply the VPN filter to your group policy.

```
group-policy oracle-vcn-vcn-vcn-policy attributes
 vpn-filter value #{vpnFilterAclName}
```

- **Tunnel group:** Apply the group policy to your tunnel group.

```
tunnel-group #{oracleHeadend1} general-attributes
 default-group-policy oracle-vcn-vcn-vcn-policy
```

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPES IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two

tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Specific to Cisco ASA: Caveats and Limitations

This section covers important characteristics and limitations that are specific to Cisco ASA.

### TUNNEL MTU AND PATH MTU DISCOVERY

You have two options for addressing tunnel MTU and path MTU discovery with Cisco ASA:

- [Option 1: TCP MSS adjustment](#)
- [Option 2: Clear/set the Don't Fragment bit](#)

#### *OPTION 1: TCP MSS ADJUSTMENT*

The maximum transmission unit (packet size) through the IPsec tunnel is less than 1500 bytes. You can fragment packets that are too large to fit through the tunnel. Or, you can signal back to the hosts that are communicating through the tunnel that they need to send smaller packets.

You can configure the Cisco ASA to change the maximum segment size (MSS) for any new TCP flows through the tunnel. The ASA looks at any TCP packets where the SYN flag is set and changes the MSS value to the configured value. This configuration might help new TCP flows avoid using path maximum transmission unit discovery (PMTUD).

Use the following command to change the MSS. This command is not part of the sample configuration in the [CPE Configuration](#) section of this topic. Apply the TCP MSS adjustment command manually, if needed.

```
sysopt connection tcpmss 1387
```

#### *OPTION 2: CLEAR/SET THE DON'T FRAGMENT BIT*

Path MTU discovery requires that all TCP packets have the **Don't Fragment** (DF) bit set. If the DF bit is set and a packet is too large to go through the tunnel, the ASA drops the packet when it arrives. The ASA sends an ICMP packet back to the sender indicating that the received packet was too large for the tunnel. The ASA offers three options for handling the DF bit. Choose one of the options and apply it to the configuration:

- **Set the DF bit (recommended):** Packets have the DF bit set in their IP header. The ASA may still fragment the packet if the original received packet cleared the DF bit.

```
crypto ipsec df-bit set-df ${outsideInterface}
```

- **Clear the DF bit:** The DF bit is cleared in the packet's IP header. Allows the packet to

be fragmented and sent to the end host in Oracle Cloud Infrastructure for reassembly.

```
crypto ipsec df-bit clear-df ${outsideInterface}
```

- **Ignore (copy) the DF bit:** The ASA looks at the original packet's IP header information and copies the DF bit setting.

```
crypto ipsec df-bit copy-df ${outsideInterface}
```

### VPN TRAFFIC MIGHT ENTER ONE TUNNEL AND EXIT ANOTHER

If VPN traffic enters an interface with the same security level as an interface toward the packet's next hop, you must allow that traffic. By default, the packets between interfaces that have identical security levels on your ASA are dropped.

Add the following command manually if you need to permit traffic between interfaces with the same security levels. This command is not part of the sample configuration in the [CPE Configuration](#) section.

```
same-security-traffic permit inter-interface
```

### General Caveats and Limitations

This section covers general characteristics and limitations of VPN Connect.

#### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPSec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPSec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPSec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.
- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPSec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

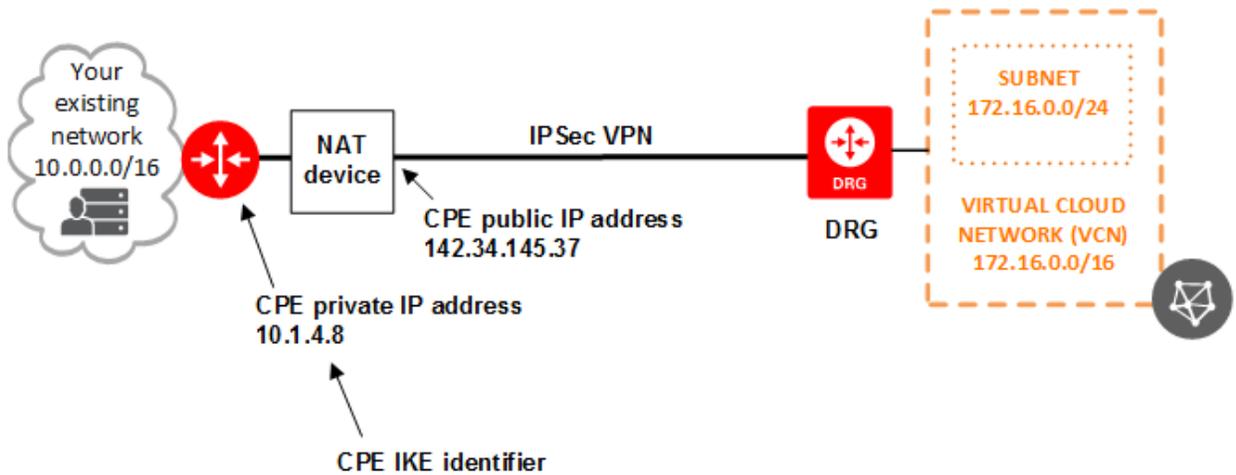
If you use policy-based IPSec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPSec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

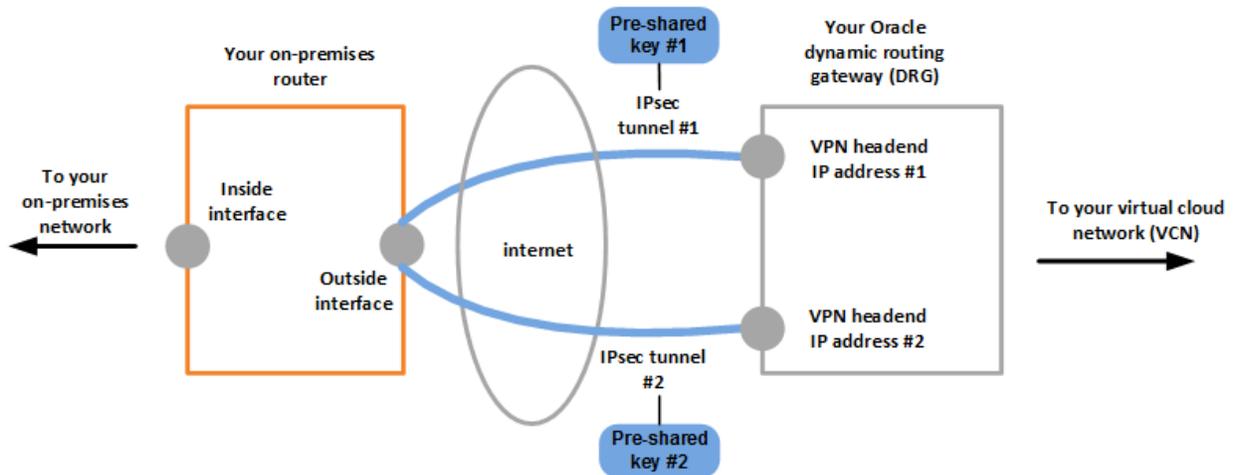
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

**Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



The configuration template provided is for a Cisco ASA running version 8.5 software (or later).



### Note

Cisco ASA versions 9.7.1 and newer support [route-based configuration](#), which is the recommended method to avoid interoperability issues.

If you want tunnel redundancy with a single Cisco ASA device, you must use the route-based configuration. With policy-based configuration, you can configure only a single tunnel between your Cisco ASA and your dynamic routing gateway (DRG).

The configuration template refers to these items that you must provide:

- **CPE public IP address:** The internet-routable IP address that is assigned to the external interface on the CPE. You or your Oracle administrator provides this value to Oracle when creating the CPE object in the Oracle Console.
- **Inside tunnel interface (required if using BGP):** The IP addresses for the CPE and Oracle ends of the inside tunnel interface. You provide these values when creating the IPsec connection in the Oracle Console.
- **BGP ASN (required if using BGP):** Your BGP ASN.

In addition, you must:

- Configure internal routing that routes traffic between the CPE and your local network.
- Ensure that you permit traffic between your ASA and your Oracle VCN (the following configuration template references this access list with the variable `${outboundAclName}`).
- Identify the internal VPN group policy (the following configuration template references this group policy as `oracle-vcn-vpn-policy`).
- Identify the transform set used for your crypto map (the following configuration template references this transform set as `oracle-vcn-transform`).

- Identify the crypto map name and sequence number (the following configuration template references the map name as `oracle-vpn-map-v1` and sequence number 1).
- Identify the operation number for IP SLA continuous ping (the following configuration template uses operation number 1).



### Important

This following configuration template from Oracle Cloud Infrastructure **is a starting point for what you need to apply to your CPE**. The syntax for each CPE device configuration may be different and depends on the model and software versions. Be sure to compare your CPE model and version to the appropriate configuration template.

Some of the parameters referenced in the template must be unique on the CPE, and the uniqueness can only be determined by accessing the CPE. Ensure that the parameters are valid on your CPE and do not overwrite any previously configured values. In particular, ensure that the following values are unique:

- Policy names or numbers
- Crypto map names and sequence numbers
- Interface names
- Access list names or numbers (if applicable)

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

There's a separate configuration template for IKEv1 versus IKEv2.

### IKEv1 Configuration Template

```
!-----
!-----
! IKEv1 Configuration Template
! The configuration consists of a single IPsec tunnel.
!-----
!-----
! The configuration template involves setting up the following:
! Access Lists
! ISAKMP Policy
! Base VPN Policy
! IPsec Configuration
! IPsec Tunnel Group Configuration
! IP Routing (BGP or Static)
! Optional: Disable NAT for VPN Traffic
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
provide these values when creating the IPsec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${outboundAclName} = ACL used to control traffic out of your inside and outside interfaces
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPsec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
! ${cryptoMapAclName} = Name of ACL which will be associated with the IPsec security association.
! ${vcnHostIp} = IP address of a VCN host. Used for IP SLA continuous ping to maintain tunnel UP state.
```

## CHAPTER 23 Networking

---

```
!-----
!-----

! Access Lists

! Permit Traffic Between Your ASA and Your Oracle VCN
! Assuming there is an access-list controlling traffic in and out of your Internet facing interface, you
will need to permit traffic between your CPE and the Oracle VPN Headend
! WARNING: The new ACL entry you add to permit the traffic between your ASA and VPN headend needs to be
above any deny statements you might have in your existing access-list

access-list ${outboundAclName} extended permit ip host ${oracleHeadend1} host ${cpePublicIpAddress}

! Crypto ACL
! Create an ACL named ${cryptoMapAclName} which will later be associated with the IPSec security
association using the 'crypto-map' command. This will define which source/destination traffic needs to
be encrypted and sent across the VPN tunnel.
! Keep this ACL to a single entry. In a policy based configuration each ACL line will establish a
separate encryption domain.
! The single encryption domain used in this configuration sample will have a source/destination of
any/VCN CIDR. Refer to the 'Encryption domain for policy-based tunnels' subsection for supported
alternatives.

access-list ${cryptoMapAclName} extended permit ip any ${vcnCidrNetwork} ${vcnCidrNetmask}

! ISAKMP Policy

! ISAKMP Phase 1 configuration.
! IKEv1 is enabled on the outside interface.
! IKEv1 policy is created for Phase 1 which specifies to use a Pre-Shared Key, AES256, SHA1, Diffie-
Hellman Group 5, and a Phase 1 lifetime of 28800 seconds (8 hours).
! If different parameters are required, modify this template before applying the configuration.
! WARNING: The IKEv1 group policy is created with a priority of 10. Make sure this doesn't conflict with
any pre-existing configuration on your ASA.

crypto ikev1 enable outside
crypto ikev1 policy 10
 authentication pre-share
 encryption aes-256
 hash sha
 group 5
```

## CHAPTER 23 Networking

---

```
lifetime 28800

! Base VPN Policy

! An internal VPN group policy named 'oracle-vcn-vpn-policy' is created to define some basic VPN tunnel
settings
! Idle and session timeouts are disabled to maintain the tunnel UP state and tunnel protocol is set to
IKEv1

group-policy oracle-vcn-vpn-policy internal
group-policy oracle-vcn-vpn-policy attributes
 vpn-idle-timeout none
 vpn-session-timeout none
 vpn-tunnel-protocol ikev1

! IPSec Configuration

! Create an IKEv1 transform set named 'oracle-vcn-transform' which defines a combination of IPSec (Phase
2) policy options. Specifically, AES256 for encryption and SHA1 for authentication.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec ikev1 transform-set oracle-vcn-transform esp-aes-256 esp-sha-hmac

! A crypto map is used to tie together the important traffic that needs encryption (via crypto map ACL)
with defined security policies (from the transform set along with other crypto map statements), and the
destination of the traffic to a specific crypto peer.
! In this configuration example, a single crypto map is created named 'oracle-vpn-map-v1' This crypto
map references the previously created crypto map ACL, the 'oracle-vcn-transform' transform set and
further defines PFS Group 5 and the security association lifetime to 3600 seconds (1 hour).
! WARNING: Make sure your crypto map name and sequence numbers do not overlap with existing crypto maps.
! WARNING: DO NOT use the 'originate-only' option with an Oracle IPSec VPN tunnel. It causes the
tunnel's traffic to be inconsistently blackholed. The command is only for tunnels between two Cisco
devices. Here's an example of the command that you should NOT use for the Oracle IPSec VPN tunnels:
crypto map <map name> <sequence number> set connection-type originate-only

crypto map oracle-vpn-map-v1 1 match address #{cryptoMapAclName}
crypto map oracle-vpn-map-v1 1 set pfs group5
crypto map oracle-vpn-map-v1 1 set peer #{oracleHeadend1}
crypto map oracle-vpn-map-v1 1 set ikev1 transform-set oracle-vcn-transform
crypto map oracle-vpn-map-v1 1 set security-association lifetime seconds 3600
```

## CHAPTER 23 Networking

---

! WARNING: The below command will apply the 'oracle-vpn-map-v1' crypto map to the outside interface. The Cisco ASA supports a single crypto map per interface. Make sure no other crypto map is applied to the outside interface before using this command.

```
crypto map oracle-vpn-map-v1 interface outside
```

```
! IPsec Tunnel Group Configuration
```

```
! This configuration matches the group policy 'oracle-vcn-vpn-policy' with an Oracle VPN headend endpoint.
```

```
! The pre-shared key for each Oracle VPN headend is defined in the corresponding tunnel group.
```

```
tunnel-group #{oracleHeadend1} type ipsec-l2l
tunnel-group #{oracleHeadend1} general-attributes
 default-group-policy oracle-vcn-vpn-policy
tunnel-group #{oracleHeadend1} ipsec-attributes
 ikev1 pre-shared-key #{sharedSecret1}
```

```
! IP SLA Configuration
```

```
! The Cisco ASA doesn't establish a tunnel if there's no interesting traffic trying to pass through the tunnel.
```

```
! You must configure IP SLA on your device for a continuous ping so that the tunnel remains up at all times.
```

```
! You must allow ICMP on the outside interface.
```

```
! Make sure that the SLA monitor number used is unique.
```

```
sla monitor 1
 type echo protocol ipIcmpEcho #{vcnHostIp} interface outside
 frequency 5
sla monitor schedule 1 life forever start-time now
```

```
icmp permit any #{outsideInterface}
```

```
! IP Routing
```

```
! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying configuration.
```

```
! Border Gateway Protocol (BGP) Configuration
! Uncomment below lines if you want to use BGP.
```

## CHAPTER 23 Networking

---

```
! router bgp $(bgpASN)
! address-family ipv4 unicast
! neighbor $(OracleInsideTunnelIpAddress1) remote-as 31898
! neighbor $(OracleInsideTunnelIpAddress1) activate
! network $(onPremCidrNetwork) mask $(onPremCidrNetmask)
! no auto-summary
! no synchronization
! exit-address-family

! Static Route Configuration
! Uncomment below line if you want to use static routing.

! route outside $(VcnCidrNetwork) $(VcnCidrNetmask) $(OracleInsideTunnelIpAddress1)

! Disable NAT for VPN Traffic

! If you are using NAT for traffic between your inside and outside interfaces, you might need to disable
NAT for traffic between your on-premises network and the Oracle VCN.
! Two objects are created for this NAT exemption. 'obj-OnPrem' represents the on-premises network as a
default route, and 'obj-oracle-vcn-1' represents the VCN CIDR block used in Oracle Cloud Infrastructure.
! If different address ranges are required, modify this template before applying the configuration.

! object network obj-onprem
! subnet 0.0.0.0 0.0.0.0
! object network obj-oracle-vcn-1
! subnet $(vcnCidrNetwork) $(vcnCidrNetmask)
! nat (inside,outside) source static obj-onprem obj-onprem destination static obj-oracle-vcn-1 obj-
oracle-vcn-1
```

### IKEv2 Configuration Template

```
!-----
!-----
! IKEv2 Configuration Template
! The configuration consists of a single IPsec tunnel.
!-----
!-----
! The configuration template involves setting up the following:
! Access Lists
```

## CHAPTER 23 Networking

---

```
! IKEv2 Policy
! Base VPN Policy
! IPSec Configuration
! IPSec Tunnel Group Configuration
! IP Routing (BGP or Static)
! Optional: Disable NAT for VPN Traffic
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${outboundAclName} = ACL used to control traffic out of your inside and outside interfaces
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
! ${cryptoMapAclName} = Name of ACL which will be associated with the IPSec security association.
! ${vcnHostIp} = IP address of a VCN host. Used for IP SLA continuous ping to maintain tunnel UP state.
!-----
!-----
! Access Lists

! Permit Traffic Between Your ASA and Your Oracle VCN
! Assuming there is an access-list controlling traffic in and out of your Internet facing interface, you
will need to permit traffic between your CPE and the Oracle VPN Headend
! WARNING: The new ACL entry you add to permit the traffic between your ASA and VPN headend needs to be
above any deny statements you might have in your existing access-list
```

## CHAPTER 23 Networking

---

```
access-list $(outboundAclName) extended permit ip host $(oracleHeadend1) host $(cpePublicIpAddress)

! Crypto ACL
! Create an ACL named $(cryptoMapAclName) which will later be associated with the IPSec security
association using the 'crypto-map' command. This will define which source/destination traffic needs to
be encrypted and sent across the VPN tunnel.
! Keep this ACL to a single entry. In a policy based configuration each ACL line will establish a
separate encryption domain.
! The single encryption domain used in this configuration sample will have a source/destination of
any/VCN CIDR. Refer to the 'Encryption domain for policy-based tunnels' subsection for supported
alternatives.

access-list $(cryptoMapAclName) extended permit ip any $(vcnCidrNetwork) $(vcnCidrNetmask)

! IKEv2 Policy

! IKEv2 is enabled on the outside interface.
! IKEv2 policy is created and specifies use of a Pre-Shared Key, AES256, SHA1, Diffie-Hellman Group 5,
and a lifetime of 28800 seconds (8 hours).
! If different parameters are required, modify this template before applying the configuration.
! WARNING: The IKEv2 group policy is created with a priority of 10. Make sure this doesn't conflict with
any pre-existing configuration on your ASA.

crypto ikev2 enable outside
crypto ikev2 policy 10
 encryption aes-256
 integrity sha384
 group 5
 prf sha
 lifetime seconds 28800

! Base VPN Policy

! An internal VPN group policy named 'oracle-vcn-vpn-policy' is created to define some basic VPN tunnel
settings
! Idle and session timeouts are disabled to maintain the tunnel UP state and tunnel protocol is set to
IKEv2

group-policy oracle-vcn-vpn-policy internal
group-policy oracle-vcn-vpn-policy attributes
 vpn-idle-timeout none
```

## CHAPTER 23 Networking

```
vpn-session-timeout none
vpn-tunnel-protocol ikev2

! IPsec Configuration

! Create an IKEv2 IPsec proposal named 'oracle_v2_ipsec_proposal' which defines AES256 for encryption
and SHA1 for authentication.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec ikev2 ipsec-proposal oracle_v2_ipsec_proposal
 protocol esp encryption aes-256
 protocol esp integrity sha-1

! A crypto map is used to tie together the important traffic that needs encryption (via crypto map ACL)
with defined security policies (from the IPsec proposal along with other crypto map statements), and the
destination of the traffic to a specific crypto peer.
! In this configuration example, a single crypto map is created named 'oracle-vpn-map-v2' This crypto
map references the previously created crypto map ACL, the 'oracle_v2_ipsec_proposal' IPsec proposal and
further defines PFS Group 5 and the security association lifetime to 3600 seconds (1 hour).
! WARNING: Make sure your crypto map name and sequence numbers do not overlap with existing crypto maps.
! WARNING: DO NOT use the 'originate-only' option with an Oracle IPsec VPN tunnel. It causes the
tunnel's traffic to be inconsistently blackholed. The command is only for tunnels between two Cisco
devices. Here's an example of the command that you should NOT use for the Oracle IPsec VPN tunnels:
crypto map <map name> <sequence number> set connection-type originate-only

crypto map oracle-vpn-map-v2 1 match address ${cryptoMapAclName}
crypto map oracle-vpn-map-v2 1 set pfs group5
crypto map oracle-vpn-map-v2 1 set peer ${oracleHeadend1}
crypto map oracle-vpn-map-v2 1 set ikev2 ipsec-proposal oracle_v2_ipsec_proposal
crypto map oracle-vpn-map-v2 1 set security-association lifetime seconds 3600

! WARNING: The below command will apply the 'oracle-vpn-map-v2' crypto map to the outside interface. The
Cisco ASA supports a single crypto map per interface. Make sure no other crypto map is applied to the
outside interface before using this command.

crypto map oracle-vpn-map-v2 interface outside

! IPsec Tunnel Group Configuration

! This configuration matches the group policy 'oracle-vcn-vpn-policy' with an Oracle VPN headend
endpoint.
```

## CHAPTER 23 Networking

---

```
! The pre-shared key for each Oracle VPN headend is defined in the corresponding tunnel group.

tunnel-group ${oracleHeadend1} type ipsec-l2l
tunnel-group ${oracleHeadend1} general-attributes
 default-group-policy oracle-vcn-vpn-policy
tunnel-group ${oracleHeadend1} ipsec-attributes
 ikev2 local-authentication pre-shared-key ${sharedSecret1}
 ikev2 remote-authentication pre-shared-key ${sharedSecret1}

! IP SLA Configuration

! The Cisco ASA doesn't establish a tunnel if there's no interesting traffic trying to pass through the
tunnel.
! You must configure IP SLA on your device for a continuous ping so that the tunnel remains up at all
times.
! You must allow ICMP on the outside interface.
! Make sure that the SLA monitor number used is unique.

sla monitor 1
type echo protocol ipIcmpEcho ${vcnHostIp} interface outside
frequency 5
sla monitor schedule 1 life forever start-time now

icmp permit any ${outsideInterface}

! IP Routing
! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying
configuration.

! Border Gateway Protocol (BGP) Configuration
! Uncomment below lines if you want to use BGP.

! router bgp ${bgpASN}
! address-family ipv4 unicast
! neighbor ${OracleInsideTunnelIpAddress1} remote-as 31898
! neighbor ${OracleInsideTunnelIpAddress1} activate
! network ${onPremCidrNetwork} mask ${onPremCidrNetmask}
! no auto-summary
! no synchronization
! exit-address-family
```

## CHAPTER 23 Networking

---

```
! Static Route Configuration
! Uncomment below line if you want to use static routing.

! route outside ${VcnCidrNetwork} ${VcnCidrNetmask} ${OracleInsideTunnelIpAddress1}

! Disable NAT for VPN Traffic

! If you are using NAT for traffic between your inside and outside interfaces, you might need to disable
NAT for traffic between your on-premises network and the Oracle VCN.
! Two objects are created for this NAT exemption. 'obj-OnPrem' represents the on-premises network as a
default route, and 'obj-oracle-vcn-1' represents the VCN CIDR block used in Oracle Cloud Infrastructure.
! If different address ranges are required, modify this template before applying the configuration.

! object network obj-onprem
! subnet 0.0.0.0 0.0.0.0
! object network obj-oracle-vcn-1
! subnet ${vcnCidrNetwork} ${vcnCidrNetmask}
! nat (inside,outside) source static obj-onprem obj-onprem destination static obj-oracle-vcn-1 obj-
oracle-vcn-1
```

### Verification

The following ASA commands are included for basic troubleshooting. For more exhaustive information, refer to Cisco's [IPSec Troubleshooting](#) document.

Use the following command to verify that ISAKMP security associations are being built between the two peers.

```
show crypto isakmp sa
```

Use the following command to verify the status of all your BGP connections.

```
show bgp summary
```

Use the following command to verify the ASA's route table.

```
show route
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Cisco IOS

This topic provides a route-based configuration for a Cisco IOS device. The configuration was validated using a Cisco 2921 running IOS version 15.4(3)M3.



#### **Important**

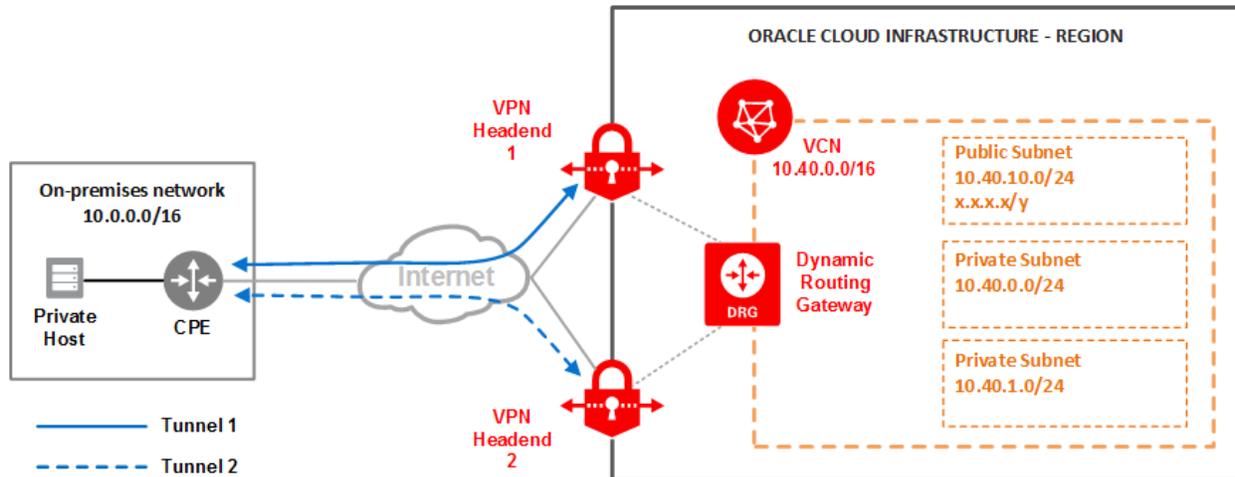
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

## Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

## Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

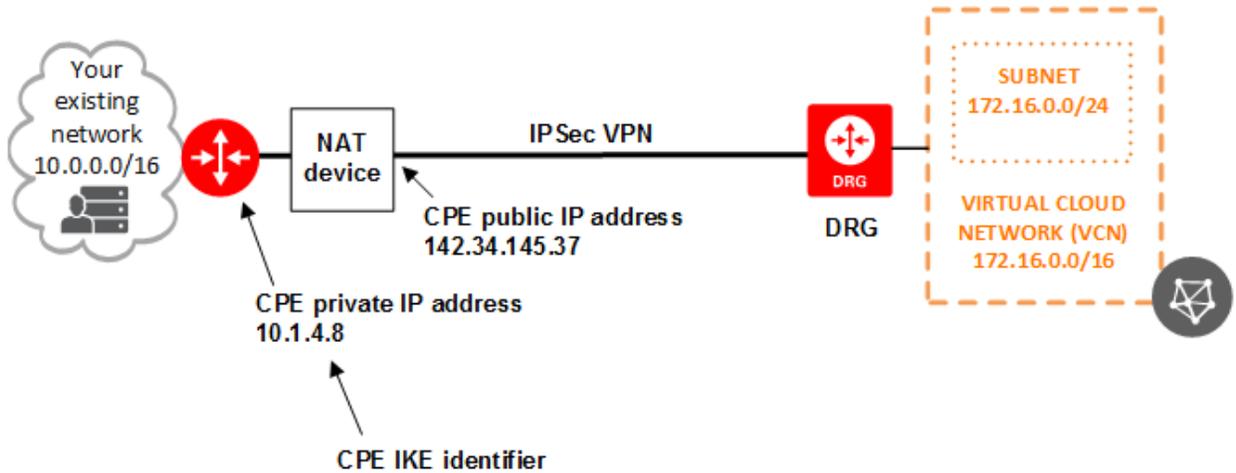
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

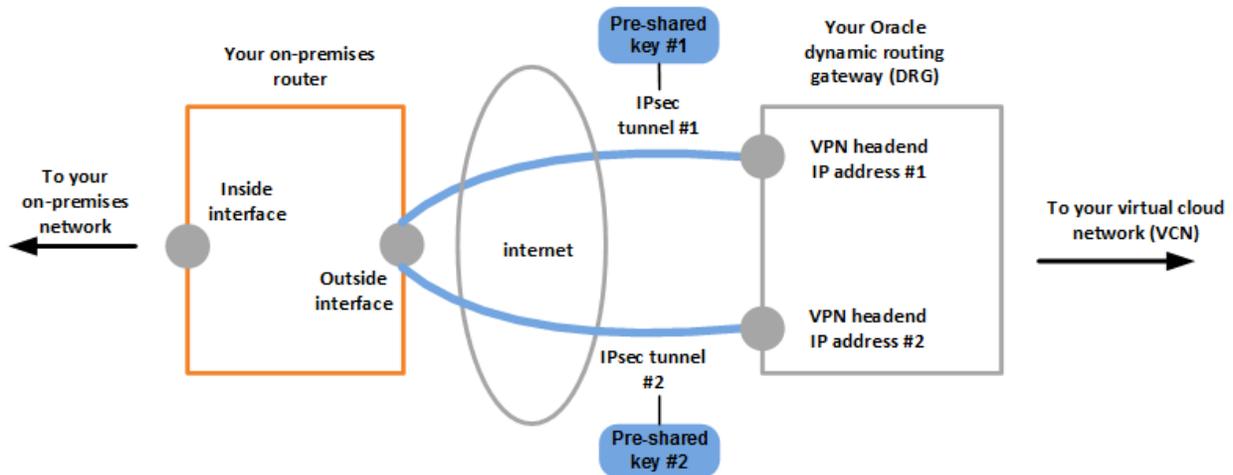
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

**Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



The configuration template was validated using a Cisco 2921 running IOS version 15.4(3)M3. The template provides information for each tunnel that you must configure. Oracle recommends setting up all configured tunnels for maximum redundancy.

The configuration template refers to these items that you must provide:

- **CPE public IP address:** The internet-routable IP address that is assigned to the external interface on the CPE. You or your Oracle administrator provides this value to Oracle when creating the CPE object in the Oracle Console.
- **Inside tunnel interface (required if using BGP):** The IP addresses for the CPE and Oracle ends of the inside tunnel interface. You provide these values when creating the IPsec connection in the Oracle Console.
- **BGP ASN (required if using BGP):** Your BGP ASN.

In addition, you must:

- Configure internal routing that routes traffic between the CPE and your local network.
- Ensure that you permit traffic between your CPE and your Oracle VCN.
- Identify the IPsec profile used (the following configuration template references this group policy as `oracle-vpn`).
- Identify the transform set used for your crypto map (the following configuration template references this transform set as `oracle-vpn-transform`).



### Important

This following configuration template from Oracle Cloud Infrastructure **is a starting point for what you need to apply to your CPE**. Some of the parameters referenced in the template must be unique on the CPE, and the uniqueness can only be determined by accessing the CPE. Ensure the parameters are valid on your CPE and do not overwrite any previously configured values. In particular, ensure these values are unique:

- Policy names or numbers
- Interface names
- Keyrings
- Access list numbers (if applicable)

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

There's a separate configuration template for IKEv1 versus IKEv2.

### IKEv1 Configuration Template

```
!-----
!-----
! IKEv1 Configuration Template
! The configuration consists of two IPSec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.
!-----
!-----
! The configuration template involves setting up the following:
```

## CHAPTER 23 Networking

---

```
! Keyring (Pre-Shared Key)
! Basic ISAKMP Options
! ISAKMP and IPSec Policy Configuration
! IPSec Peers
! Virtual Tunnel Interfaces
! IP Routing (BGP or Static)
! Update Any Internet Facing Access List to Allow IPSec and ISAKMP Packets
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${OracleInsideTunnelIpAddress2} = Inside tunnel IP address of Oracle-side for the second tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${oracleHeadend2} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${sharedSecret2} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
!-----
!-----
! Keyring (Pre-Shared Key)

! For authentication during IKE a separate keyring is defined for each Oracle VPN Headend peer.
! Add the pre-shared key for each Oracle VPN headend under the corresponding keyring.
```

## CHAPTER 23 Networking

---

```
crypto keyring oracle-vpn-{oracleHeadend1}
 local-address {cpePublicIpAddress}
 pre-shared-key address {oracleHeadend1} key {sharedSecret1}
crypto keyring oracle-vpn-{oracleHeadend2}
 local-address {cpePublicIpAddress}
 pre-shared-key address {oracleHeadend2} key {sharedSecret2}

! Basic ISAKMP Options

! Optional IPSec settings are included here.
! All optional settings included are recommended by Oracle. Remove or comment out any unneeded commands
prior to applying this configuration.
! WARNING: These settings are global and may impact other IPSec connections

! Enables fragmentation of IKE packets prior to encryption.
crypto isakmp fragmentation

! Enables Dead Peer Detection (DPD)
crypto isakmp keepalive 10 10

! The Router will clear the DF-bit in the IP header. Allows the packet to be fragmented and sent to the
end host in Oracle Cloud Infrastructure for reassembly.
crypto ipsec df-bit clear

! Increases security association anti-replay window. An increased window size is helpful for scenarios
where packets are regularly being dropped due to delays.
crypto ipsec security-association replay window-size 128

! ISAKMP and IPSec Policy Configuration

! An ISAKMP policy is created for Phase 1 which specifies to use a Pre-Shared Key, AES256, SHA384,
Diffie-Hellman Group 5, and a Phase 1 lifetime of 28800 seconds (8 hours).
! If different parameters are required, modify this template before applying the configuration.
! WARNING: The ISAKMP group policy is created with a priority of 10. Make sure this doesn't conflict
with any pre-existing configuration before applying.

crypto isakmp policy 10
 encr aes 256
 hash sha384
 authentication pre-share
 group 5
```

## CHAPTER 23 Networking

---

```
lifetime 28800

! Create an IPSec transform set named 'oracle-vpn-transform' which defines a combination of IPSec (Phase
2) policy options. Specifically, AES256 for encryption and SHA1 for authentication. This is also where
tunnel mode is set for IPSec.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec transform-set oracle-vpn-transform esp-aes 256 esp-sha-hmac
mode tunnel

! A IPSec profile named 'oracle-vpn' is created.
! The previously created transform set is added to this policy along with settings for enabling PFS
Group 5 and the security association lifetime to 3600 seconds (1 hour).
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec profile oracle-vpn
set pfs group5
set security-association lifetime seconds 3600
set transform-set oracle-vpn-transform

! IPSec Peers

! Two ISAKMP profiles are created for each Oracle VPN Headend.
! An ISAKMP profile is used as a repository for various Phase 1 commands tied to a specific IPSec peer.
In this case, we match the previously created keyrings to an Oracle VPN headend.

crypto isakmp profile oracle-vpn-${oracleHeadend1}
keyring oracle-vpn-${oracleHeadend1}
self-identity address
match identity address ${oracleHeadend1} 255.255.255.255
crypto isakmp profile oracle-vpn-${oracleHeadend2}
keyring oracle-vpn-${oracleHeadend2}
self-identity address
match identity address ${oracleHeadend2} 255.255.255.255

! Virtual Tunnel Interfaces

! Each tunnel interface is a logical interface representing the local end of a VPN tunnel to a remote
VPN peer. Each tunnel interface represents a single tunnel to a different Oracle VPN Headend. The IP
address of each VPN headend is provided when you create your IPSec connection in Oracle Console.
! All traffic routed to a tunnel interface will be encrypted and sent across the tunnel towards Oracle
```

## CHAPTER 23 Networking

---

Cloud Infrastructure.

! Each tunnel interface configuration also references the previously created IPsec profile 'oracle-vpn' for its IPsec parameters.

! WARNING: When doing static routing you do NOT have to set IPs on the tunnel interfaces unless you have pre-configured inside tunnel interfaces in Oracle Console when creating your IPsec connection. Inside tunnel interfaces are required if using BGP.

```
interface Tunnel${tunnelNumber1}
ip address ${cpeInsideTunnelIpAddress1} ${cpeInsideTunnelNetmask1}
tunnel source ${cpePublicIpAddress}
tunnel mode ipsec ipv4
tunnel destination ${oracleHeadend1}
tunnel protection ipsec profile oracle-vpn
```

```
interface Tunnel${tunnelNumber2}
ip address ${cpeInsideTunnelIpAddress2} ${cpeInsideTunnelNetmask2}
tunnel source ${cpePublicIpAddress}
tunnel mode ipsec ipv4
tunnel destination ${oracleHeadend2}
tunnel protection ipsec profile oracle-vpn
```

! IP Routing

! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying configuration.

! Border Gateway Protocol (BGP) Configuration

! Uncomment below lines if you want to use BGP.

```
! router bgp ${bgpASN}
! neighbor ${OracleInsideTunnelIpAddress1} remote-as 31898
! neighbor ${OracleInsideTunnelIpAddress2} remote-as 31898
! network ${onPremCidrNetwork} mask ${onPremCidrNetmask}
```

! Static Route Configuration

! Uncomment below lines if you want to use static routing.

```
! ip route ${vcnCidrNetwork} ${vcnCidrNetmask} Tunnel${tunnelNumber1}
! ip route ${vcnCidrNetwork} ${vcnCidrNetmask} Tunnel${tunnelNumber2}
```

! Update Any Internet Facing Access List to Allow IPsec and ISAKMP Packets

! You may need to allow IPsec and ISAKMP packets out your internet facing interface.

## CHAPTER 23 Networking

---

```
! Uncomment below lines to create a new ACL allowing IPSec and ISAKMP traffic and apply it to the
outside interface.

! ip access-list extended INTERNET-INGRESS
! permit udp host ${oracleHeadend1} host ${cpePublicIpAddress} eq isakmp
! permit esp host ${oracleHeadend1} host ${cpePublicIpAddress}
! permit udp host ${oracleHeadend2} host ${cpePublicIpAddress} eq isakmp
! permit esp host ${oracleHeadend2} host ${cpePublicIpAddress}
! permit icmp any any echo
! permit icmp any any echo-reply
! permit icmp any any unreachable

! interface ${outsideInterface}
! ip address ${cpePublicIpAddress} ${netmask}
! ip access-group INTERNET-INGRESS in
```

## IKEv2 Configuration Template

```
!-----
!-----
! IKEv2 Configuration Template
! The configuration consists of two IPSec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.
!-----
!-----
! The configuration template involves setting up the following:
! Keyring (Pre-Shared Key)
! IKEv2 and IPSec Policy Configuration
! IPSec Peers
! Virtual Tunnel Interfaces
! IP Routing (BGP or Static)
! Update Any Internet Facing Access List to Allow IPSec and ISAKMP Packets
!-----
!-----
! The configuration template has various parameters that you must define before applying the
configuration.
!-----
!-----
! PARAMETERS REFERENCED:
! ${OracleInsideTunnelIpAddress1} = Inside tunnel IP address of Oracle-side for the first tunnel. You
```

## CHAPTER 23 Networking

---

```
provide these values when creating the IPSec connection in the Oracle Console.
! ${OracleInsideTunnelIpAddress2} = Inside tunnel IP address of Oracle-side for the second tunnel. You
provide these values when creating the IPSec connection in the Oracle Console.
! ${bgpASN} = Your BGP ASN
! ${cpePublicIpAddress} = The public IP address for the CPE. This is the IP address of your outside
interface
! ${oracleHeadend1} = Oracle public IP endpoint obtained from the Oracle Console.
! ${oracleHeadend2} = Oracle public IP endpoint obtained from the Oracle Console.
! ${sharedSecret1} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${sharedSecret2} = You provide when you set up the IPSec connection in the Oracle Console, or you can
use the default Oracle-provided value.
! ${outsideInterface} = The public interface or outside of tunnel interface which is configured with the
CPE public IP address.
! ${vcnCidrNetwork} = VCN IP range
! ${vcnCidrNetmask} = Subnet mask for VCN
! ${onPremCidrNetwork} = On-premises IP range
! ${onPremCidrNetmask} = ON-premises subnet mask
!-----

! Keyring (Pre-Shared Key)

! For authentication during IKE a separate keyring is defined for each Oracle VPN Headend peer.
! Add the pre-shared key for each Oracle VPN headend under the corresponding keyring.

crypto ikev2 keyring oracle-vpn-${oracleHeadend1}
peer oracle_vpn
 address ${oracleHeadend1}
 pre-shared-key local ${sharedSecret1}
 pre-shared-key remote ${sharedSecret1}

crypto ikev2 keyring oracle-vpn-${oracleHeadend2}
peer oracle_vpn
 address ${oracleHeadend2}
 pre-shared-key local ${sharedSecret2}
 pre-shared-key remote ${sharedSecret2}

! Optional IPSec settings are included here.
! All optional settings included are recommended by Oracle. Remove or comment out any unneeded commands
```

## CHAPTER 23 Networking

---

```
prior to applying this configuration.
! WARNING: These settings are global and may impact other IPSec connections

! The Router will clear the DF-bit in the IP header. Allows the packet to be fragmented and sent to the
end host in Oracle Cloud Infrastructure for reassembly.
crypto ipsec df-bit clear

! Increases security association anti-replay window. An increased window size is helpful for scenarios
where packets are regularly being dropped due to delays.
crypto ipsec security-association replay window-size 128

! IKEv2 and IPSec Policy Configuration

! An IKEv2 proposal is created and specifies use of a Pre-Shared Key, AES256, SHA384, and Diffie-Hellman
Group 5.
! If different parameters are required, modify this template before applying the configuration.

crypto ikev2 proposal oracle_v2_proposal
 encryption aes-cbc-256
 integrity sha384
 group 5

crypto ikev2 policy oracle_v2_policy
 proposal oracle_v2_proposal

! Create an IPSec transform set named 'oracle-vpn-transform' which defines a combination of IPSec (Phase
2) policy options. Specifically, AES256 for encryption and SHA1 for authentication. This is also where
tunnel mode is set for IPSec.
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec transform-set oracle-vpn-transform esp-aes 256 esp-sha-hmac
 mode tunnel

! An IPSec profile named 'oracle_v2_ipsec_profile_tunnel#' is created for each tunnel.
! The previously created transform set is added to this policy along with settings for enabling PFS
Group 5 and the security association lifetime to 3600 seconds (1 hour).
! If different parameters are required, modify this template before applying the configuration.

crypto ipsec profile oracle_v2_ipsec_profile_tunnel1
 set ikev2-profile oracle_v2_profile_tunnel1
```

## CHAPTER 23 Networking

---

```
set pfs group5
set security-association lifetime seconds 3600
set transform-set oracle-vpn-transform

crypto ipsec profile oracle_v2_ipsec_profile_tunnel2
set ikev2-profile oracle_v2_profile_tunnel2
set pfs group5
set security-association lifetime seconds 3600
set transform-set oracle-vpn-transform

! IPSec Peers

! Two IKEv2 profiles are created for each Oracle VPN Headend.

crypto ikev2 profile oracle-vpn-${oracleHeadend1}
keyring oracle-vpn-${oracleHeadend1}
identity local address ${cpePublicIpAddress}
match identity remote address ${oracleHeadend1} 255.255.255.255
authentication remote pre-share
authentication local pre-share

crypto ikev2 profile oracle-vpn-${oracleHeadend2}
keyring oracle-vpn-${oracleHeadend2}
identity local address ${cpePublicIpAddress}
match identity remote address ${oracleHeadend2} 255.255.255.255
authentication remote pre-share
authentication local pre-share

! Virtual Tunnel Interfaces

! Each tunnel interface is a logical interface representing the local end of a VPN tunnel to a remote
VPN peer. Each tunnel interface represents a single tunnel to a different Oracle VPN Headend. The IP
address of each VPN headend is provided when you create your IPSec connection in Oracle Console.
! All traffic routed to a tunnel interface will be encrypted and sent across the tunnel towards Oracle
Cloud Infrastructure.
! Each tunnel interface configuration also references the previously created IPSec profile 'oracle-vpn'
for its IPSec parameters.
! WARNING: When doing static routing you do NOT have to set IPs on the tunnel interfaces unless you have
pre-configured inside tunnel interfaces in Oracle Console when creating your IPSec connection. Inside
tunnel interfaces are required if using BGP.
```

## CHAPTER 23 Networking

```
interface Tunnel${tunnelNumber1}
ip address ${cpeInsideTunnelIpAddress1} ${cpeInsideTunnelNetmask1}
tunnel source ${cpePublicIpAddress}
tunnel mode ipsec ipv4
tunnel destination ${oracleHeadend1}
tunnel protection ipsec profile oracle_v2_ipsec_profile_tunnel1

interface Tunnel${tunnelNumber2}
ip address ${cpeInsideTunnelIpAddress2} ${cpeInsideTunnelNetmask2}
tunnel source ${cpePublicIpAddress}
tunnel mode ipsec ipv4
tunnel destination ${oracleHeadend2}
tunnel protection ipsec profile oracle_v2_ipsec_profile_tunnel2

! IP Routing
! Pick either dynamic (BGP) or static routing. Uncomment the corresponding commands prior to applying
configuration.

! Border Gateway Protocol (BGP) Configuration
! Uncomment below lines if you want to use BGP.

! router bgp ${bgpASN}
! neighbor ${OracleInsideTunnelIpAddress1} remote-as 31898
! neighbor ${OracleInsideTunnelIpAddress2} remote-as 31898
! network ${onPremCidrNetwork} mask ${onPremCidrNetmask}

! Static Route Configuration
! Uncomment below lines if you want to use static routing.
! ip route ${vcnCidrNetwork} ${vcnCidrNetmask} Tunnel${tunnelNumber1}
! ip route ${vcnCidrNetwork} ${vcnCidrNetmask} Tunnel${tunnelNumber2}

! Update Any Internet Facing Access List to Allow IPSec and ISAKMP Packets

! You may need to allow IPSec and ISAKMP packets out your internet facing interface.
! Uncomment below lines to create a new ACL allowing IPSec and ISAKMP traffic and apply it to the
outside interface.

! ip access-list extended INTERNET-INGRESS
! permit udp host ${oracleHeadend1} host ${cpePublicIpAddress} eq isakmp
! permit esp host ${oracleHeadend1} host ${cpePublicIpAddress}
! permit udp host ${oracleHeadend2} host ${cpePublicIpAddress} eq isakmp
```

## CHAPTER 23 Networking

---

```
! permit esp host ${oracleHeadend2} host ${cpePublicIpAddress}
! permit icmp any any echo
! permit icmp any any echo-reply
! permit icmp any any unreachable

! interface ${outsideInterface}
! ip address ${cpePublicIpAddress} ${netmask}
! ip access-group INTERNET-INGRESS in
```

### Verification

The following IOS commands are included for basic troubleshooting.

Use the following command to verify that ISAKMP security associations are being built between the two peers.

```
show crypto isakmp sa
```

Use the following command to verify the status of all your BGP connections or neighbors.

```
show ip bgp summary
show ip bgp neighbors
```

Use the following command to verify the route table.

```
show ip route
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### FortiGate

This topic provides configuration for a FortiGate that is running software version 6.0.4.

FortiGate experience is recommended. For more details on how to use FortiGate products, visit their official site. For FortiGate documentation for high availability (HA) or manual deployment, see the [Fortinet Document Library](#).



### Important

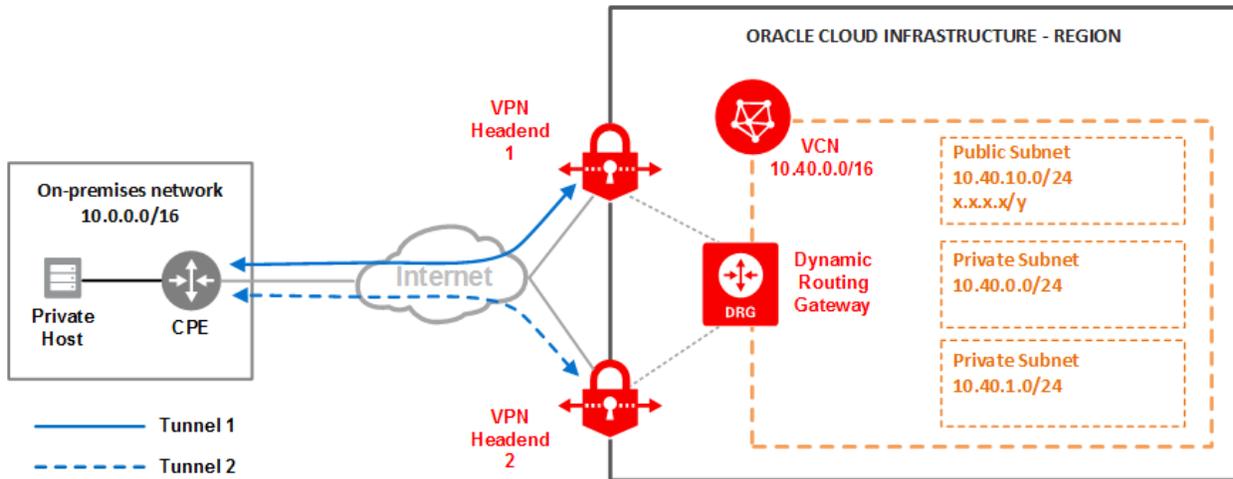
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

## Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

## Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

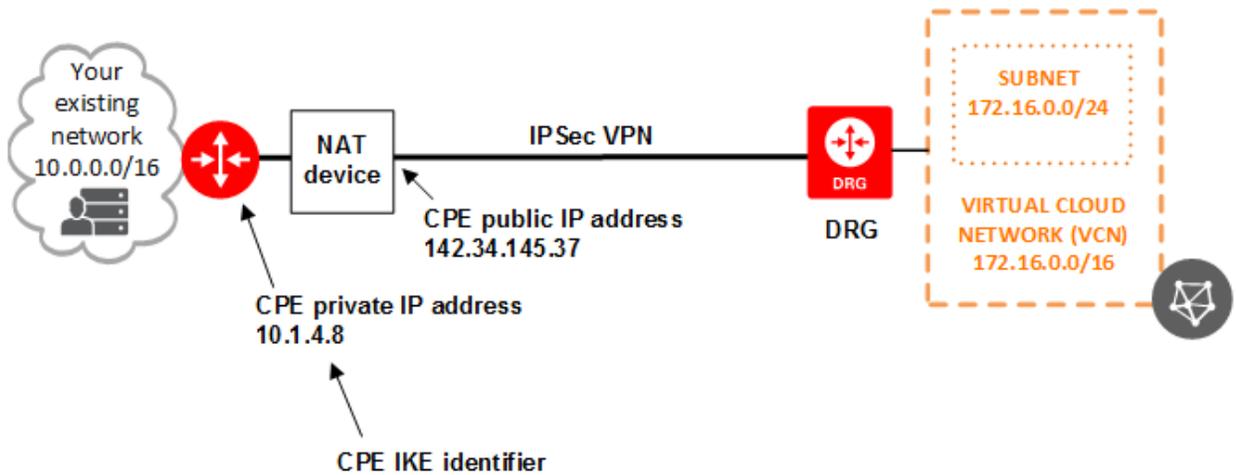
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

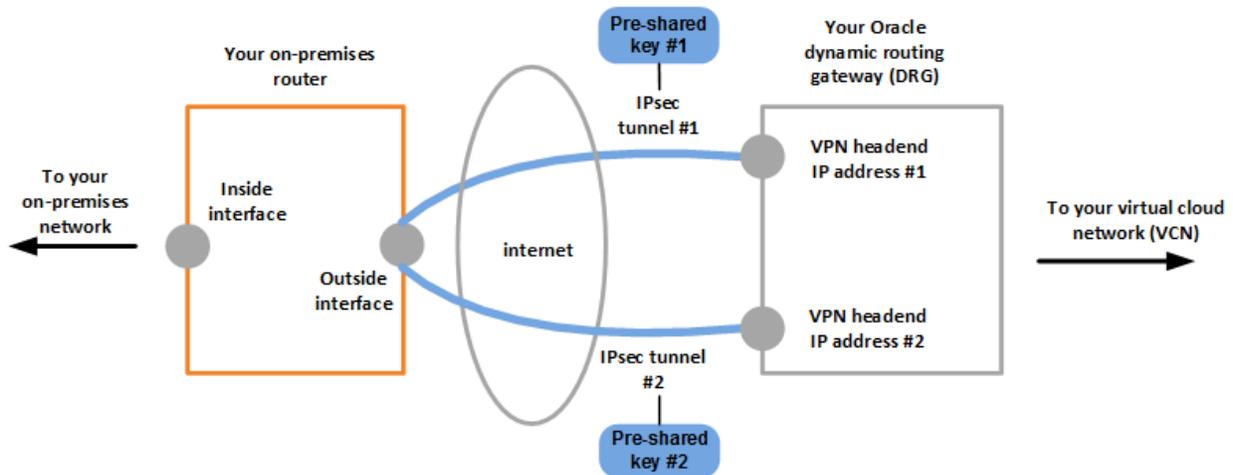
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

 **Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

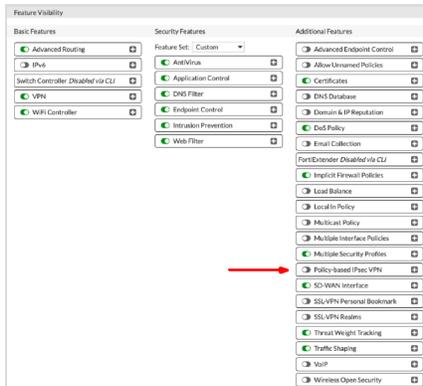
The following figure shows the basic layout of the IPSec connection.



By default, FortiGate provisions the IPSec tunnel in route-based mode. This topic focuses on FortiGate with a route-based VPN configuration.

## CHAPTER 23 Networking

If necessary, you can have FortiGate provision the IPsec tunnel in policy-based mode. To enable the feature, go to **System**, and then to **Feature Visibility**. Under **Additional Features**, enable the **Policy-based IPsec VPN** feature.



### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPsec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPsec Parameters](#).

If you want to use IKEv2, there's a variation on one of the tasks presented in the next section. Specifically, in [task 2](#), when configuring authentication, select IKE version 2.

### CONFIGURATION PROCESS



#### Important

Before starting, ensure you have a valid license or trial license to configure FortiGate.

### Task 1: Use the wizard to create the VPN

1. Go to **VPN**, and then to **IPsec Wizard** to create a new VPN tunnel.
2. On the **VPN Creation Wizard** page, specify the following items:
  - **Name:** Description used to identify the IPsec tunnel.
  - **Template Type:** Site to Site
  - **Remote Device Type:** Cisco
  - **NAT Configuration:** No NAT between sites



3. Click **Next**.
4. On the **Authentication** page, specify the following items:
  - **Remote Device:** IP Address
  - **IP Address:** IP address for the Oracle VPN headend. Oracle generated this value when creating the IPsec tunnel.
  - **Outgoing Interface:** The WAN interface configured for external traffic.
  - **Authentication Method:** Pre-shared Key. Oracle supports only shared secret keys.
  - **Pre-shared Key:** The shared secret. Oracle generated this value when creating the IPsec tunnel.

## CHAPTER 23 Networking

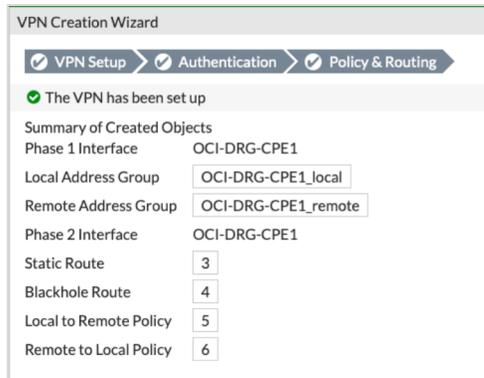


5. Click **Next**.
6. On the **Policy & Routing** page, specify the following items:
  - **Local Interface:** The LAN interface configured for internal traffic.
  - **Local Subnets:** The subnet used for internal traffic.
  - **Remote Subnets:** The Oracle VCN subnets that will be used for the IPsec tunnel.
  - **Internet Access:** None



7. Click **Create**.

A summary message is shown with details about the configuration. Notice that the wizard automatically creates security policies with the subnets that you specified and adds the required static routes.



### Task 2: Add Phase 1 and Phase 2 parameters to each IPsec tunnel

You must convert each newly created IPsec tunnel into a custom tunnel to add the recommended parameters for Phase 1 and Phase 2.

Perform the following steps for each tunnel.

1. Go to **VPN**, and then click **IPsec Tunnels**.
2. Select the tunnel and click **Edit** to view the **Edit VPN Tunnel** page.
3. Click **Convert to Custom Tunnel**.

## CHAPTER 23 Networking

**Edit VPN Tunnel**

Name: OCI-DRG-CPE1 [Convert To Custom Tunnel](#)

Comments: VPN: OCI-DRG-CPE1 (Created by VPN wizard) 41/255

**Network** [Edit](#)

Remote Gateway: Static IP Address (129.213.6.50), Interface: port1

**Authentication** [Edit](#)

Authentication Method: Pre-shared Key  
IKE Version: 1, Mode: Main (ID protection)

**Phase 1 Proposal** [Edit](#)

Algorithms: AES256-SHA384  
Diffie-Hellman Group: 5

**XAUTH** [Edit](#)

Type: Disabled

**Phase 2 Selectors**

Name	Local Address	Remote Address	<a href="#">Add</a>
OCI-DRG-CPE1	OCI-DRG-CPE1_local	OCI-DRG-CPE1_remote	<a href="#">Edit</a>

4. Edit the relevant sections to match the required settings shown in the following screenshots. Remember to click the check mark icon in the top-right corner of each section after making your changes.

The IP address shown in the first screenshot is an example address.

Notice that if you want to use IKEv2, on the **Authentication** screen, instead select **IKE Version 2**.

**Network** [Check](#) [Refresh](#)

IP Version: IPv4

Remote Gateway: Static IP Address

IP Address: 129.213.6.50

Interface: port1

Local Gateway:

Mode Config:

NAT Traversal:  Enable  Disable  Forced

Dead Peer Detection:  Disable  On Idle  On Demand

## CHAPTER 23 Networking

---

**Authentication** 🔍 ↻

Method

Pre-shared Key

IKE

Version

Mode

**Phase 1 Proposal** ➕ Add 🔍 ↻

Encryption  Authentication

Diffie-Hellman Group  31  30  29  28  27  21  
 20  19  18  17  16  15  
 14  5  2  1

Key Lifetime (seconds)

Local ID

Phase 2 Selectors		
Name	Local Address	Remote Address
OCI-DRG-CPE1	OCI-DRG-CPE1_local	OCI-DRG-CPE1_remote

**Edit Phase 2**

Name: OCI-DRG-CPE1

Comments: VPN: OCI-DRG-CPE1 (Created by VPN wizard)

Local Address: Named Address | OCI-DRG-CPE1\_local

Remote Address: Named Address | OCI-DRG-CPE1\_remo

**Phase 2 Proposal**

Encryption: AES256 | Authentication: SHA1

Enable Replay Detection:

Enable Perfect Forward Secrecy (PFS):

Diffie-Hellman Group:  31  30  29  28  27  21  20  19  18  17  16  15  5  2  1

Local Port: All

Remote Port: All

Protocol: All

Auto-negotiate:

Autokey Keep Alive:

Key Lifetime: Seconds

Seconds: 3600

5. After configuring all sections, click **OK** to save and close the dialogs.

### Task 3: Verify the IPsec connection

At this point, the IPsec tunnel will not be established by default because FortiGate uses the IP address assigned on the WAN interface. In this case, this IP address is a private IP address because Oracle does 1:1 NAT. This private IP address will be used as the local IKE ID and will not match the one expected on the Oracle DRG. To resolve this, you can manually change the local IKE ID on your FortiGate by using the CPE's CLI, or you can change the value that Oracle uses in the Oracle Console (see the instructions that follow). Either way, this fixes the incompatibility and brings up the IPsec tunnel.

### To change the CPE IKE identifier that Oracle uses (Oracle Console)

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

A list of the IPSec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. For the IPSec connection you're interested in, click the Actions icon (three dots), and then click **Edit**.

The current CPE IKE identifier that Oracle is using is displayed at the bottom of the dialog.

3. Enter your new values for **CPE IKE Identifier Type** and **CPE IKE Identifier**, and then click **Save Changes**.

#### REDUNDANCY WITH BGP OVER IPSEC

For redundancy, Oracle recommends using BGP over IPSec. By default, if you have two connections of the same type (for example, two IPSec VPNs that both use BGP), and you advertise the same routes across both connections, Oracle prefers the oldest established route when responding to requests or initiating connections. If you want to force routing to be symmetric, Oracle recommends using BGP and AS path prepending with your routes to influence which path Oracle uses when responding to and initiating connections. For more information, see [Routing Details for Connections to Your On-Premises Network](#).

The Oracle DRG uses /30 or /31 as subnets for configuring IP addresses on the interface tunnels. Remember that the IP address must be part of the IPSec VPN's encryption domain and must be allowed in the firewall policy to reach the peer VPN through the interface tunnel. You might need to implement a static route through the tunnel interface for the peer IP address.

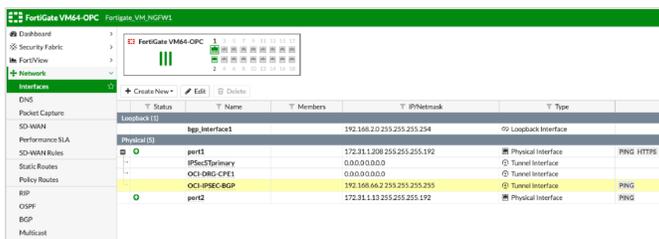
Oracle's BGP ASN in commercial regions is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

For your side, you can use a private ASN. Private ASNs are in the range 64512–65534.

### Task 1: Edit the tunnel interface

In the first task, you add the BGP IP address to the newly created FortiGate tunnel interface. Perform the following steps for each tunnel.

1. Go to **Network**, and then **Interface**.
2. Select the interface you're interested in and click **Edit**.



3. Configure the following items:

- **IP:** Enter the BGP IP address that you assigned to the FortiGate end of the tunnel interface. The following screenshot shows an example value of 192.168.66.2.
- **Remote IP/Network Mask:** Add the BGP IP address that you assigned to the Oracle end of the tunnel interface. Include either a /30 or /31 mask, depending on how you specified the addresses in the Oracle Console. In the following screenshot, 192.168.66.0/30 was used, where 192.168.66.2 is assigned to the FortiGate end, and 192.168.66.1 is assigned to the Oracle end.
- **Ping access** (recommended): In the **Administrative Access** section, enable

ping access.

The screenshot shows the 'Edit Interface' configuration page for 'OCI-IPSEC-BGP'. The interface is a Tunnel Interface connected to 'port1'. The role is set to 'Undefined'. The addressing mode is 'Manual' with an IP of '192.168.66.2' and a network mask of '255.255.255.252'. The remote IP/network mask is '192.168.66.1/255.255.255.252'. Under 'Administrative Access', the 'PING' checkbox is checked, while others like 'HTTPS', 'SSH', and 'SNMP' are unchecked. A 'DHCP Server' checkbox is at the bottom.

4. Click **OK**.

### Task 2: Add a static route for the Oracle IP address

For each tunnel, add a /32 static route towards the Oracle IP address through the tunnel, as shown in the following screenshot.

The screenshot shows the 'Static Routes' configuration page in FortiGate VM. The table below represents the data visible in the configuration:

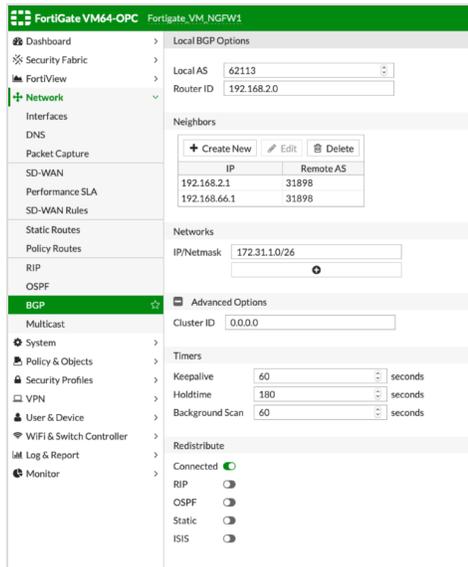
Destination	Gateway	Interface	Priority
192.168.66.1	192.168.66.1	port1	1

### Task 3: Configure BGP

Perform the following steps for each tunnel.

1. Go to **Network**, and then **BGP**.
2. Enter the following items:
  - **Local AS:** Your BGP ASN. You can use a private ASN. Private ASNs are in the range 64512–65534.
  - **Router ID:** A value to provide a unique identity for this BGP router among its peers.
  - **Neighbors:** Click **Create New** and enter the BGP IP address for the Oracle end of the tunnel, and the Oracle BGP ASN (31898 for commercial regions). If you're configuring VPN Connect for connecting to the Government Cloud, see [Oracle's BGP ASN](#).
  - **Networks:** Optionally use this field to advertise a specific subnet over BGP. You can also advertise subnets by using the **Redistribute** section in the **Advanced**

**Options section.**



3. Click **OK**.

**Verification**

The following CLI command is useful for gathering statistical data such as the number of packets encrypted versus decrypted, the number of bytes sent versus received, the encryption domain (SPI) identifier, and so on. This kind of information can be critical for determining an issue with the VPN.

```
diagnose vpn tunnel list
```

The following command indicates a lack of firewall policy, a lack of forwarding route, and policy ordering issues. If there are no communication issues, this command returns blank output.

```
diagnose debug flow
```

## CHAPTER 23 Networking

---

The following command verifies BGP neighbor status information. Remember that an "Active" state doesn't mean that the BGP session is up. "Active" refers to a BGP state message. For more information, see [BGP Background and Concepts](#) in the FortiGate documentation.

```
get router info bgp summary
```

The following command provides more detailed information about a BGP neighbor.

```
get router info bgp neighbors
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Juniper MX

This topic provides configuration for a Juniper MX that is running software version JunOS 15.0 (or newer).



### Important

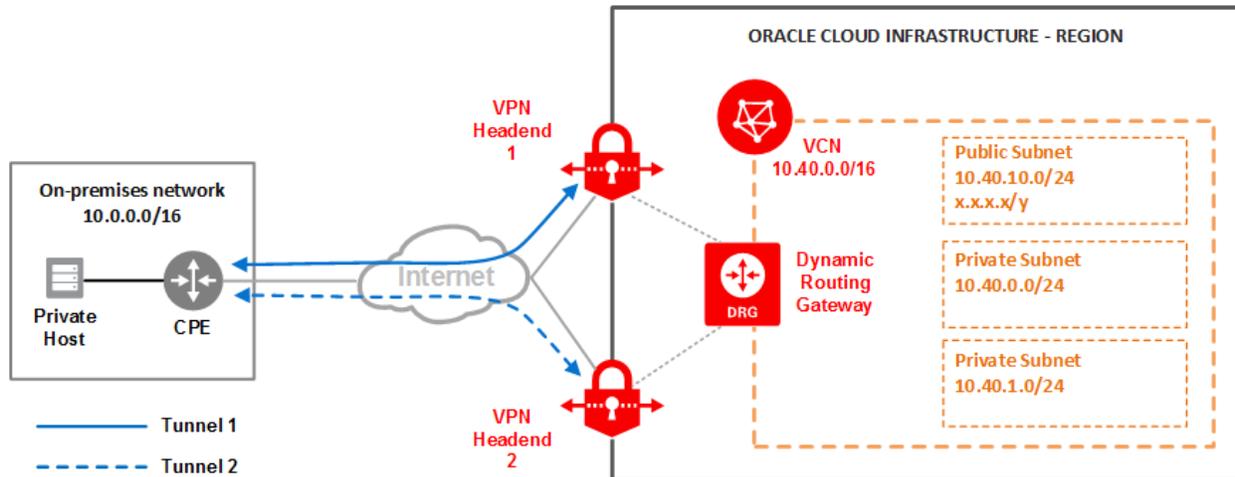
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

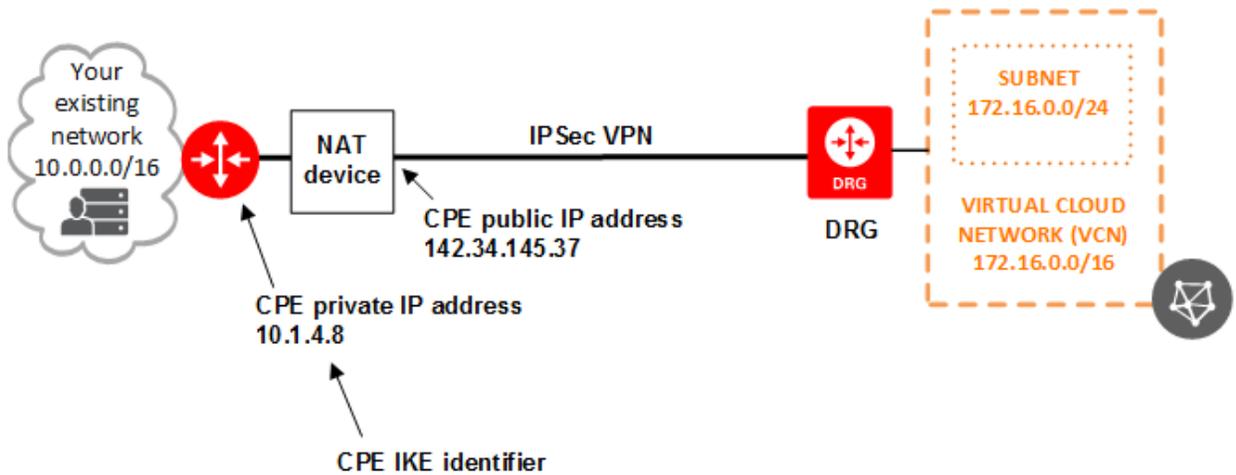
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

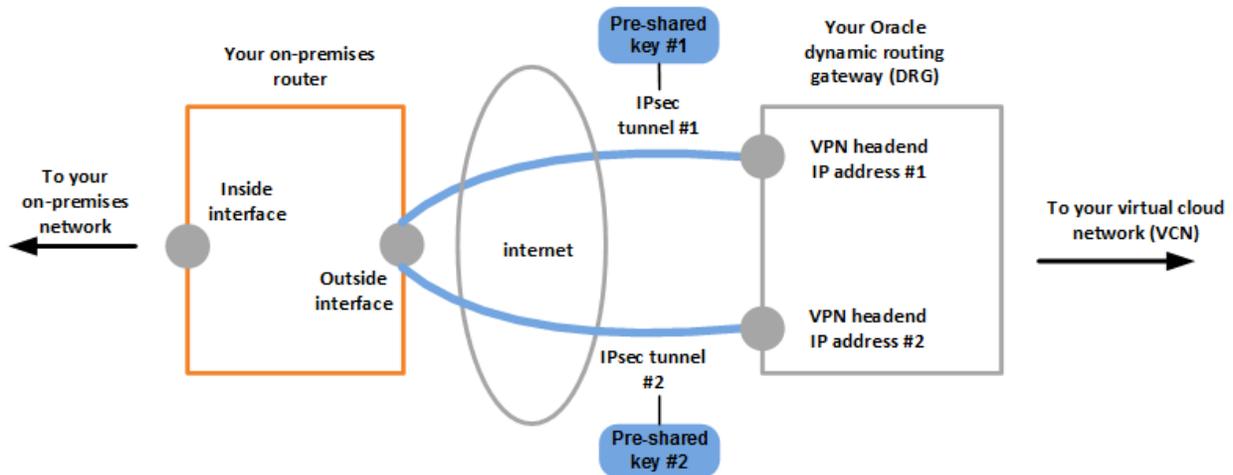
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

 **Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPsec connection.



The configuration template provided is for a Juniper MX router running JunOS 15.0 (or newer). The template provides information for each tunnel that you must configure. Oracle recommends setting up all configured tunnels for maximum redundancy.

The configuration template refers to these items that you must provide:

- **CPE public IP address:** The internet-routable IP address that is assigned to the external interface on the CPE. You or your Oracle administrator provides this value to Oracle when creating the CPE object in the Oracle Console.
- **Inside tunnel interface (required if using BGP):** The IP addresses for the CPE and Oracle ends of the inside tunnel interface. You provide these values when creating the IPSec connection in the Oracle Console.
- **BGP ASN (required if using BGP):** Your BGP ASN.

In addition, you must:

- Configure the Juniper MX public interface (the CPE public IP address is bound to this interface).
- Configure internal routing that routes traffic between the CPE and your local network.
- Configure the tunnel interfaces. See the next section for more information.

### ABOUT THE TUNNEL INTERFACES

In the following configuration template, the tunnel interfaces are referred to with the following variables:

- `msInterface#` - one per tunnel
  - These interfaces correspond to one of the four encryption ASICs on the MS-MPC card.
  - You can distribute load across the ASICs by spreading your tunnels across them.
  - Example values: `ms-2/3/0`, `ms-2/3/1`
- `insideMsUnit#` and `outsideMsUnit#` - one pair per tunnel
  - For every tunnel, you need an ms-mpc interface pair of units.
  - One represents the outside of the IPSec tunnel. The other represents the inside of the tunnel.
  - The router forwards packets from your on-premises network to your VCN into the inside unit.

- The encryption ASIC then encrypts the packets based on the rules and policies.
- Then the encrypted packet egresses out the outside unit as an ESP packet, ready to be forwarded to the Oracle VPN headend routers.
- There are over 16,000 possible values for unit numbers.
  - One way to allocate the units is to offset them by 8,000.
  - You can pick values between 0 - 7999 for `insideMsUnit#` and 8000-15999 for `outsideMsUnit#`.



### Important

This following configuration template from Oracle Cloud Infrastructure **is a starting point for what you need to apply to your CPE**. Some of the parameters referenced in the template must be unique on the CPE, and the uniqueness can only be determined by accessing the CPE. Ensure the parameters are valid on your CPE and do not overwrite any previously configured values. In particular, ensure these values are unique:

- Policy names or numbers
- Interface names
- Access list numbers (if applicable)

To find parameters that you must define before applying the configuration, search for the keyword `USER_DEFINED` in the template.

### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

You specify the IKE version when defining the IKE policy. In the following configuration, there's a comment showing how to configure the IKE policy for IKEv1 versus IKEv2.

### CONFIGURATION TEMPLATE

```


Configuration Template
The configuration consists of two IPSec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.

The configuration template involves setting up the following:
PHASE 1
PHASE 2
SETTING THE TUNNEL INTERFACES FOR ORACLE
SETTING THE SERVICES FOR ORACLE.
SETTING BGP/STATIC ROUTING
SETTING ROUTING-INSTANCES FOR ORACLE (OPTIONAL).

The configuration template has various parameters that you must define before applying the
configuration.
Search in the template for the keyword "USER_DEFINED" to find those parameters.

PARAMETERS REFERENCED:
oracle_headend_1 = Oracle public IP endpoint obtained from the Oracle Console.
oracle_headend_2 = Oracle public IP endpoint obtained from the Oracle Console.
connection_presharedkey_1 = You provide when you set up the IPSec connection in the Oracle Console, or
you can use the default Oracle-provided value.
connection_presharedkey_2 = You provide when you set up the IPSec connection in the Oracle Console, or
you can use the default Oracle-provided value.
cpe_public_ip_address = The internet-routable IP address that is assigned to the public interface on
the CPE. You provide this when creating the CPE object in the Oracle Console.
```

## CHAPTER 23 Networking

---

```
cpe_public_interface = The name of the Juniper interface where the CPE IP address is configured. Eg:
ge-0/0/1.0
msInterface1 = The interface correspond to one of the four encryption ASICs on the MS-MPC card. Eg:
ms-2/3/0, ms-2/3/1
msInterface2 = Second tunnel interface that needs to be configured. Eg: ms-2/3/0, ms-2/3/1
insideMsUnit1 = The inside interface of the MS-MPC interface pair for tunnel_1
insideMsUnit2 = The inside interface of the MS-MPC interface pair for tunnel_2
outsideMsUnit1 = The outside interface of the MS-MPC interface pair for tunnel_1
outsideMsUnit2 = The outside interface of the MS-MPC interface pair for tunnel_2
inside_tunnel_interface_ip_address = The IP addresses for the CPE and Oracle ends of the inside tunnel
interface. You provide these when creating the IPSec connection in the Oracle Console.
inside_tunnel_interface_ip_address_neighbor = The neighbor IP address between the MX and Oracle end
points of the inside tunnel interface.
bgp_asn = Your ASN
vcn_range = VCN IP Range

OPTIONAL PARAMETERS:
customer_on-prem_to_oracle = Name of the routing instance to be defined on the CPE for the tunnel
interfaces connecting to the Oracle headends.
internet_routing_instance = Name of the routing instance to be defined on the CPE for the tunnel
interfaces that are connected to the Internet.

IPSec Tunnel 1

#1: Internet Key Exchange (IKE) Configuration (Phase 1)
Defining the IKE Proposal for Oracle
This IKE (Phase 1) configuration template uses AES256, SHA384, Diffie-Hellman Group 5, and 28800
second (8 hours) IKE session key lifetime.
If different parameters are required, modify this template before applying the configuration.

set services ipsec-vpn ike proposal oracle-ike-proposal authentication-method pre-shared-keys
set services ipsec-vpn ike proposal oracle-ike-proposal authentication-algorithm sha-384
set services ipsec-vpn ike proposal oracle-ike-proposal encryption-algorithm aes-256-cbc
set services ipsec-vpn ike proposal oracle-ike-proposal lifetime-seconds 28800
set services ipsec-vpn ike proposal oracle-ike-proposal dh-group group5

Defining the IKE Policy for Oracle
USER_DEFINED: Replace the parameters in the section below as needed

If using IKEv1, uncomment the following two lines, and comment out the line after (the line with
```

## CHAPTER 23 Networking

```
"version 2" at the end)
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 mode main
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 version 1

set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 version 2
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 proposals oracle-ike-proposal
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 local-id ipv4_addr <cpe_public_ip_address>
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 remote-id ipv4_addr <oracle_headend_1>
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_1 pre-shared-key ascii-text <connection_
presharedkey_1>

Setting up Public Interface with the CPE Public IP.
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces <cpe_public_interface> unit 0 family inet address <cpe_public_ip_address>

#2: IPSec Configuration

Defining the IPSec (Phase 2) Proposal for Oracle
The IPSec proposal defines the protocol, authentication, encryption, and lifetime parameters for the
IPsec security association.
The configuration template sets AES256 for encryption, SHA256 for authentication, enables PFS group
14, and sets the IPSec session key lifetime to 3600 seconds (1 hour).
The IPsec policy incorporates the Diffie-Hellman group and the IPsec proposal.
If different parameters are required, modify this template before applying the configuration.

set services ipsec-vpn ipsec proposal oracle-ipsec-proposal
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal protocol esp
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal authentication-algorithm hmac-sha-256-128
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal encryption-algorithm aes-256-cbc
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal lifetime-seconds 3600

Defining the IPSec (PHASE 2) policy for Oracle

set services ipsec-vpn ipsec policy oracle-ipsec-policy perfect-forward-secrecy keys group14
set services ipsec-vpn ipsec policy oracle-ipsec-policy proposals oracle-ipsec-proposal

Defining Security Association for Oracle
USER_DEFINED: Replace the parameters in the section below as needed.
The IKE and IPSEC policies are associated with the tunnel interface. Eg: ms-2/3/0.101
The IPsec Dead Peer Detection option causes periodic messages to be sent to ensure a Security
```

## CHAPTER 23 Networking

Association remains operational.

```
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 from ipsec-inside-interface
<msInterface1>.<insideMsUnit1>
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then remote-gateway <oracle_headend_1>
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then dynamic ike-policy oracle-ike-policy-tunnel_
1
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then dynamic ipsec-policy oracle-ipsec-policy
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then tunnel-mtu 1430
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then initiate-dead-peer-detection
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then dead-peer-detection
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then dead-peer-detection interval 5
set services ipsec-vpn rule oracle-vpn-tunnel_1 term 1 then dead-peer-detection threshold 4
set services ipsec-vpn rule oracle-vpn-tunnel_1 match-direction input

#3: Tunnel Interface Configuration

Defining the Tunnel Interfaces
USER_DEFINED: Replace the parameters in the section below as needed.

set interfaces <msInterface1> unit <insideMsUnit1> description oracle-vpn-tunnel-1-INSIDE
set interfaces <msInterface1> unit <insideMsUnit1> family inet address <inside_tunnel_interface_ip_
address>
set interfaces <msInterface1> unit <insideMsUnit1> service-domain inside

set interfaces <msInterface1> unit <outsideMsUnit1> description oracle-vpn-tunnel-1-OUTSIDE
set interfaces <msInterface1> unit <outsideMsUnit1> family inet
set interfaces <msInterface1> unit <outsideMsUnit1> service-domain outside

#4: Service Set Configuration

USER_DEFINED: Replace the parameters in the section below as needed
Service set configuration to direct traffic to the tunnel interfaces and associating the appropriate
IPSec-VPN-Rule.

set services service-set oracle-vpn-tunnel_1 next-hop-service inside-service-interface
<msInterface1>.<insideMsUnit1>
set services service-set oracle-vpn-tunnel_1 next-hop-service outside-service-interface
<msInterface1>.<outsideMsUnit1>
set services service-set oracle-vpn-tunnel_1 ipsec-vpn-options local-gateway <cpe_public_ip_address>
set services service-set oracle-vpn-tunnel_1 ipsec-vpn-rules oracle-vpn-tunnel-tunnel_1
```

## CHAPTER 23 Networking

```
This option causes the router to reduce the Maximum Segment Size of TCP packets to prevent packet
fragmentation.

set services service-set oracle-vpn-tunnel_1 tcp-mss 1387

#5a: Border Gateway Protocol (BGP) Configuration

USER_DEFINED: Replace the parameters in the section below as needed

BGP is used within the tunnel to exchange prefixes between the Dynamic Routing Gateway and your CPE.
The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG
advertises the VCN's subnets.
The configuration template uses a basic route policy to advertise a default route to the DRG.
To advertise additional prefixes to the Oracle VCN, add additional prefixes to the term ORACLE-DEFAULT
policy. Make sure the prefix is present in the routing table of the device with a valid next-hop.

You configure the local BGP Autonomous System Number (BGP ASN) when you set up the IPsec connection in
the Oracle Console. If you later need to change the ASN, you must recreate the CPE object and IPsec
connection in the Oracle Console.

set policy-options policy-statement ORACLE-DEFAULT term default from route-filter 0.0.0.0/0 exact

set policy-options policy-statement ORACLE-DEFAULT term default then accept
set policy-options policy-statement ORACLE-DEFAULT term reject then reject

set protocols bgp group ebgp type external
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> export ORACLE-
DEFAULT
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> peer-as 31898
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> local-as <bgp_asn>

#5b: Static Route Configuration

USER_DEFINED: Replace the parameters in the section below as needed
In case you plan to use static routing to get traffic through the IPsec tunnels, you can point the
routes down to the tunnel interfaces. You should redistribute these routes into your on-premises
network. Configuration for CPE to VCN static routes:

set routing-options static route <vcn_range> next-hop <msInterface1>.<insideMsUnit1>
```

## CHAPTER 23 Networking

---

```
##6: Routing Instances Configuration (Optional)
USER_DEFINED: Replace the parameters in the section below as needed.
If you are using routing-instances on your CPE, you need to make sure you account for them in your
configuration. Merge the following configuration into the template provided above.

set routing-instances <customer_on-prem_to_oracle> interface <msInterfacel>.<insideMsUnit1>
set routing-instances <internet_routing_instance> interface <msInterfacel>.<outsideMsUnit1>
set services service-set oracle-vpn-tunnel-tunnel_1 ipsec-vpn-options local-gateway <cpe_public_ip_
address> routing-instance <internet_routing_instance>

IPSec Tunnel 2

#1: Internet Key Exchange (IKE) Configuration (Phase 1)

Defining the IKE Proposal for Oracle
This IKE (Phase 1) configuration template uses AES256, SHA384, Diffie-Hellman Group 5, and 28800
second (8 hours) IKE session key lifetime.
If different parameters are required, modify this template before applying the configuration.

set services ipsec-vpn ike proposal oracle-ike-proposal authentication-method pre-shared-keys
set services ipsec-vpn ike proposal oracle-ike-proposal authentication-algorithm sha-384
set services ipsec-vpn ike proposal oracle-ike-proposal encryption-algorithm aes-256-cbc
set services ipsec-vpn ike proposal oracle-ike-proposal lifetime-seconds 28800
set services ipsec-vpn ike proposal oracle-ike-proposal dh-group group5

Defining the IKE Policy for Oracle
USER_DEFINED: Replace the parameters in the section below as needed

If using IKEv1, uncomment the following two lines, and comment out the line after (the line with
"version 2" at the end)
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 mode main
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 version 1

set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 version 2
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 proposals oracle-ike-proposal
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 local-id ipv4_addr <cpe_public_ip_address>
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 remote-id ipv4_addr <oracle_headend_2>
set services ipsec-vpn ike policy oracle-ike-policy-tunnel_2 pre-shared-key ascii-text <connection_
presharedkey_2>
```

## CHAPTER 23 Networking

```
Setting up Public Interface with the CPE Public IP.
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces <cpe_public_interface> unit 0 family inet address <cpe_public_ip_address>

#2: IPSec Configuration

Defining the IPSec (Phase 2) Proposal for Oracle
The IPSec proposal defines the protocol, authentication, encryption, and lifetime parameters for the
IPsec security association.
The configuration template sets AES256 for encryption, SHA256 for authentication, enables PFS group
14, and sets the IPSec session key lifetime to 3600 seconds (1 hour).
The IPsec policy incorporates the Diffie-Hellman group and the IPsec proposal.
If different parameters are required, modify this template before applying the configuration.

set services ipsec-vpn ipsec proposal oracle-ipsec-proposal
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal protocol esp
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal authentication-algorithm hmac-sha-256-128
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal encryption-algorithm aes-256-cbc
set services ipsec-vpn ipsec proposal oracle-ipsec-proposal lifetime-seconds 3600

Defining the IPSec (PHASE 2) policy for Oracle

set services ipsec-vpn ipsec policy oracle-ipsec-policy perfect-forward-secrecy keys group14
set services ipsec-vpn ipsec policy oracle-ipsec-policy proposals oracle-ipsec-proposal

Defining Security Association for Oracle
USER_DEFINED: Replace the parameters in the section below as needed
The IKE and IPSEC policies are associated with the tunnel interface. Eg: ms-2/3/0.101
The IPsec Dead Peer Detection option causes periodic messages to be sent to ensure a Security
Association remains operational.

set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 from ipsec-inside-interface
<msInterface2>.<insideMsUnit2>
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then remote-gateway <oracle_headend_2>
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then dynamic ike-policy oracle-ike-policy-tunnel_
2
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then dynamic ipsec-policy oracle-ipsec-policy
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then tunnel-mtu 1420
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then initiate-dead-peer-detection
```

## CHAPTER 23 Networking

```
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then dead-peer-detection
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then dead-peer-detection interval 5
set services ipsec-vpn rule oracle-vpn-tunnel_2 term 1 then dead-peer-detection threshold 4
set services ipsec-vpn rule oracle-vpn-tunnel_2 match-direction input

#3: Tunnel Interface Configuration

Defining the Tunnel Interfaces
USER_DEFINED: Replace the parameters in the section below as needed.

set interfaces <msInterface2> unit <insideMsUnit2> description oracle-vpn-tunnel-2-INSIDE
set interfaces <msInterface2> unit <insideMsUnit2> family inet address <inside_tunnel_interface_ip_
address>
set interfaces <msInterface2> unit <insideMsUnit2> service-domain inside

set interfaces <msInterface2> unit <outsideMsUnit2> description oracle-vpn-tunnel-2-OUTSIDE
set interfaces <msInterface2> unit <outsideMsUnit2> family inet
set interfaces <msInterface2> unit <outsideMsUnit2> service-domain outside

#4: Service Set Configuration

USER_DEFINED: Replace the parameters in the section below as needed
Service set configuration to direct traffic to the tunnel interfaces and associating the appropriate
IPSec-VPN-Rule.

set services service-set oracle-vpn-tunnel_2 next-hop-service inside-service-interface
<msInterface2>.<insideMsUnit2>
set services service-set oracle-vpn-tunnel_2 next-hop-service outside-service-interface
<msInterface2>.<outsideMsUnit2>
set services service-set oracle-vpn-tunnel_2 ipsec-vpn-options local-gateway <cpe_public_ip_address>
set services service-set oracle-vpn-tunnel_2 ipsec-vpn-rules oracle-vpn-tunnel-tunnel_2

This option causes the router to reduce the Maximum Segment Size of TCP packets to prevent packet
fragmentation.

set services service-set oracle-vpn_1 tcp-mss 1387

#5a: Border Gateway Protocol (BGP) Configuration

USER_DEFINED: Replace the parameters in the section below as needed
```

## CHAPTER 23 Networking

```
BGP is used within the tunnel to exchange prefixes between the dynamic routing gateway and your CPE.
The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG
advertises the VCN's subnets.
The configuration templates uses a basic route policy to advertise a default route to the DRG.
To advertise additional prefixes to the Oracle VCN, add additional prefixes to the term ORACLE-DEFAULT
policy. Make sure the prefix is present in the routing table of the device with a valid next-hop.

You configure the local BGP Autonomous System Number (BGP ASN) when you set up the IPsec connection in
the Oracle Console. If you later need to change the ASN, you must recreate the CPE object and IPsec
connection in the Oracle Console.

set policy-options policy-statement ORACLE-DEFAULT term default from route-filter 0.0.0.0/0 exact

set policy-options policy-statement ORACLE-DEFAULT term default then accept
set policy-options policy-statement ORACLE-DEFAULT term reject then reject

set protocols bgp group ebgp type external
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> export ORACLE-
DEFAULT
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> peer-as 31898
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> local-as <bgp_asn>

#5b: Static Route Configuration

USER_DEFINED: Replace the parameters in the section below as needed
In case you plan to use static routing to get traffic through the IPsec tunnels, you can point the
routes down to the tunnel interfaces. You should redistribute these routes into your on-premises
network. Configuration for CPE to VCN static routes:

set routing-options static route <vcn_range> next-hop <msInterface2>.<insideMsUnit2>

##6: Routing Instances Configuration (Optional)
USER_DEFINED: Replace the parameters in the section below as needed.
If you are using routing-instances on your CPE, you need to make sure you account for them in your
configuration. Merge the following configuration into the template provided above.

set routing-instances <customer_on-prem_to_oracle> interface <msInterface2>.<insideMsUnit2>
set routing-instances <internet_routing_instance> interface <msInterface2>.<outsideMsUnit2>
set services service-set oracle-vpn-tunnel-tunnel_2 ipsec-vpn-options local-gateway <cpe_public_ip_
address> routing-instance <internet_routing_instance>
```

### Verification

Use the following command to verify security associations (SAs).

```
show services ipsec-vpn ipsec security-associations detail
```

Use the following command to check the BGP status.

```
show bgp summary
```

Use the following commands to check the routes advertised to and received from Oracle Cloud Infrastructure. If you've configured the CPE to use routing instances, use the commands with table *<table-name>* at the end.

```
show route advertising-protocol bgp <neighbor-address>
```

```
show route receive-protocol bgp <neighbor-address>
```

```
show route advertising-protocol bgp <neighbor-address> table <table-name>
```

```
show route receive-protocol bgp <neighbor-address> table <table-name>
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Juniper SRX

This topic provides configuration for a Juniper SRX that is running software version JunOS 11.0 (or newer).



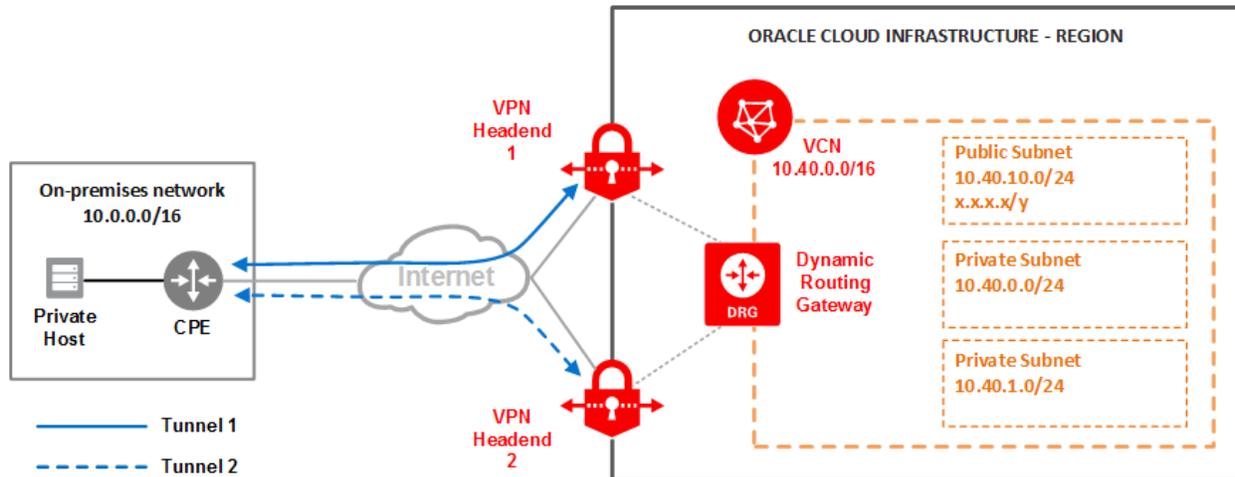
### Important

Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPSec, see the list of [supported IPSec parameters](#) and consult your vendor's documentation for assistance.

The following diagram shows a basic IPSec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

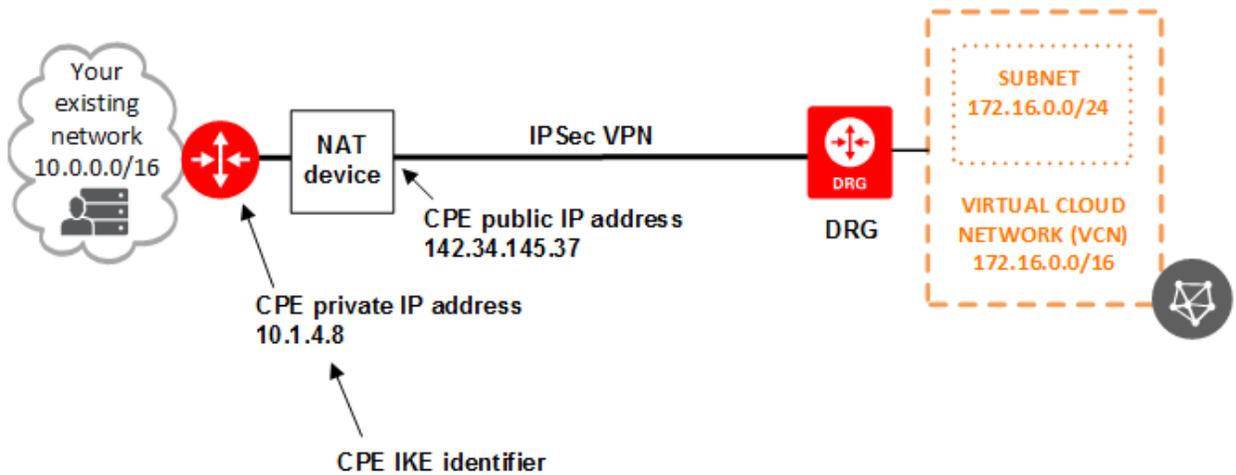
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

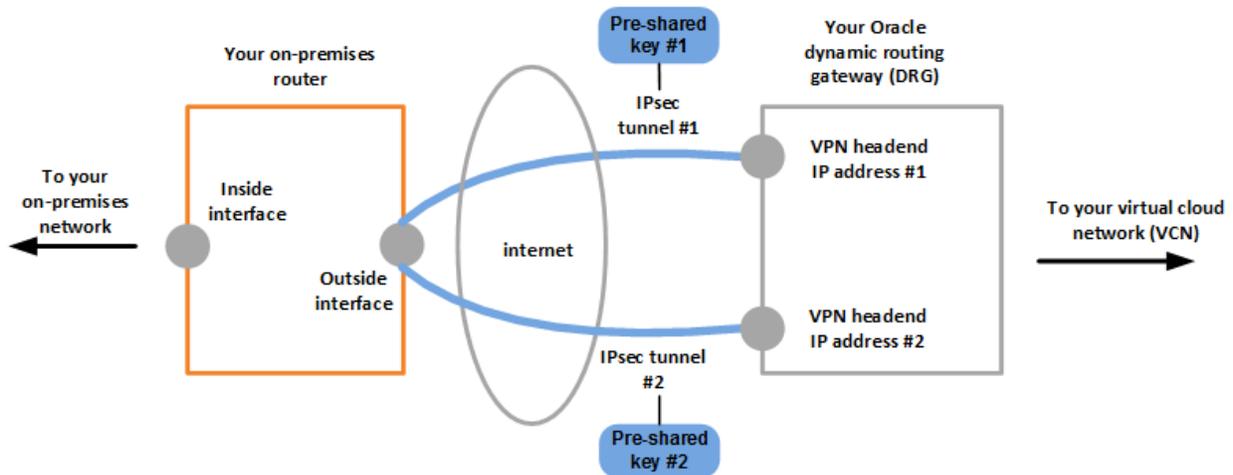
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

**✓ Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



The configuration template provided is for a Juniper SRX router running JunOS 11.0 software (or later). The template provides information for each tunnel that you must configure. Oracle recommends setting up all configured tunnels for maximum redundancy.

The configuration template refers to these items that you must provide:

- **CPE public IP address:** The internet-routable IP address that is assigned to the external interface on the CPE. You or your Oracle administrator provides this value to Oracle when creating the CPE object in the Oracle Console.
- **Inside tunnel interface (required if using BGP):** The IP addresses for the CPE and Oracle ends of the inside tunnel interface. You provide these values when creating the IPSec connection in the Oracle Console.
- **BGP ASN (required if using BGP):** Your BGP ASN.

In addition, you must:

- Configure the outside tunnel interface (the CPE public IP address is bound to this interface).
- Configure the tunnel interface IDs (referred to as st0.1 and st0.2 in the following configuration template). You need multiple tunnel unit numbers per IPSec connection.
- Configure internal routing that routes traffic between the CPE and your local network.
- Identify the security zone for the outside interface (the following configuration template references this zone as `internet_untrust`).
- Identify the security zone for the inside interface (the following configuration template references this zone as `oracle_trust`).
- Identify the security zone for the tunnel interface (the following configuration template references this zone as `oracle_vpn`).



### Important

This following configuration template from Oracle Cloud Infrastructure **is a starting point for what you need to apply to your CPE**. Some of the parameters referenced in the template must be unique on the CPE, and the uniqueness can only be determined by accessing the CPE. Ensure the parameters are valid on your CPE and do not overwrite any previously configured values. In particular, ensure these values are unique:

- Policy names or numbers
- Interface names
- Access list numbers (if applicable)

To find parameters that you must define before applying the configuration, search for the keyword `USER_DEFINED` in the template.

### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPsec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPsec Parameters](#).

You specify the IKE version when defining the IKE gateway. In the following configuration, there's a comment showing how to configure the IKE gateway for IKEv1 versus IKEv2.

### CONFIGURATION TEMPLATE

```


Configuration Template
```

## CHAPTER 23 Networking

---

```
The configuration consists of two IPSec tunnels. Oracle highly recommends that you configure both
tunnels for maximum redundancy.

The configuration template involves setting up the following:
PHASE 1
PHASE 2
SETTING THE SECURITY ZONES FOR ORACLE
SETTING THE SECURITY POLICIES FOR ORACLE
SETTING THE SECURITY SETTING FOR ORACLE
SETTING BGP/STATIC ROUTING

The configuration template has various parameters that you must define before applying the
configuration.
Search in the template for the keyword "USER_DEFINED" to find those parameters.

PARAMETERS REFERENCED:
oracle_headend_1 = Oracle public IP endpoint obtained from the Oracle Console.
oracle_headend_2 = Oracle public IP endpoint obtained from the Oracle Console.
connection_presharedkey_1 = You provide when you set up the IPSec connection in the Oracle Console, or
you can use the default Oracle-provided value.
connection_presharedkey_2 = You provide when you set up the IPSec connection in the Oracle Console, or
you can use the default Oracle-provided value.
outside_public_interface = The public interface or outside of tunnel interface which is configured
with the CPE public IP address. Example: ge-0/0/1.0
cpe_public_ip_address = The internet-routable IP address that is assigned to the public interface on
the CPE. You provide this when creating the CPE object in the Oracle Console.
inside_tunnel_interface = The internal-facing interface for the on-premises network behind the SRX
that needs to reach the Oracle VCN. Example: ge-0/0/0.0
inside_tunnel_interface_ip_address = The IP addresses for the CPE and Oracle ends of the inside tunnel
interface. You provide these when creating the IPSec connection in the Oracle Console.
inside_tunnel_interface_ip_address_neighbor = The neighbor IP address between the SRX and Oracle end
points of the inside tunnel interface.
internal_network_ip_range = Internal on-premise network behind the SRX that needs to reach resources
in the Oracle VCN.
bgp_asn = Your ASN
vcn_range = VCN IP Range

```

## CHAPTER 23 Networking

```
IPsec Tunnel 1

#1: Internet Key Exchange (IKE) Configuration (Phase 1)
Defining the IKE Proposal for Oracle
This IKE (Phase 1) configuration template uses AES256, SHA384, Diffie-Hellman Group 5, and 28800
second (8 hours) IKE session key lifetime.
If different parameters are required, modify this template before applying the configuration.

set security ike proposal oracle-ike-proposal authentication-method pre-shared-keys
set security ike proposal oracle-ike-proposal authentication-algorithm sha-384
set security ike proposal oracle-ike-proposal encryption-algorithm aes-256-cbc
set security ike proposal oracle-ike-proposal lifetime-seconds 28800
set security ike proposal oracle-ike-proposal dh-group group5

Defining the IKE Policy for Oracle
USER_DEFINED: Replace the parameters in the section below as needed

set security ike policy ike_pol_oracle-vpn-<oracle_headend_1> mode main
set security ike policy ike_pol_oracle-vpn-<oracle_headend_1> proposals oracle-ike-proposal
set security ike policy ike_pol_oracle-vpn-<oracle_headend_1> pre-shared-key ascii-text <connection_
presharedkey_1>

Setting up Public Interface with the CPE Public IP.
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces <outside_public_interface> unit 0 family inet address <cpe_public_ip_address>

Defining the IKE Gateway for Oracle
USER_DEFINED: Replace the parameters in the section below as needed.
This option enables IPsec Dead Peer Detection, which causes periodic messages to be sent to ensure a
Security Association remains operational.
If you want to use IKEv1 instead, comment out the line below that ends with "version v2-only".

set security ike gateway gw_oracle-<oracle_headend_1> ike-policy ike_pol_oracle-vpn-<oracle_headend_1>
set security ike gateway gw_oracle-<oracle_headend_1> external-interface <outside_public_interface>
set security ike gateway gw_oracle-<oracle_headend_1> address <oracle_headend_1>
set security ike gateway gw_oracle-<oracle_headend_1> dead-peer-detection
set security ike gateway gw_oracle-<oracle_headend_1> version v2-only
set security ike gateway gw_oracle-<oracle_headend_1> local-identity inet <cpe_public_ip_address>

#2: IPsec Configuration
```

## CHAPTER 23 Networking

```
Defining the IPSec (Phase 2) Proposal for Oracle
The IPSec proposal defines the protocol, authentication, encryption, and lifetime parameters for the
IPsec security association.
The configuration template sets AES256 for encryption, SHA1 for authentication, enables PFS group 5,
and sets the IPSec session key lifetime to 3600 seconds (1 hour).
The IPsec policy incorporates the Diffie-Hellman group and the IPsec proposal.
If different parameters are required, modify this template before applying the configuration.

set security ipsec vpn-monitor-options
set security ipsec proposal oracle-ipsec-proposal protocol esp
set security ipsec proposal oracle-ipsec-proposal authentication-algorithm hmac-sha1-96;
set security ipsec proposal oracle-ipsec-proposal encryption-algorithm aes-256-cbc
set security ipsec proposal oracle-ipsec-proposal lifetime-seconds 3600

Defining the IPSec (PHASE 2) policy for Oracle
set security ipsec policy ipsec_pol_oracle-vpn perfect-forward-secrecy keys group5
set security ipsec policy ipsec_pol_oracle-vpn proposals oracle-ipsec-proposal

Defining Security Association for Oracle
USER_DEFINED: Replace the parameters in the section below as needed
The IPsec Policy and IKE gateways are associated with a tunnel interface (st0.1). If other tunnels are
defined on your router, you must specify a unique interface name (for example, st0.2).
The df-bit clear option allows the SRX to fragment the packet and send it to the end host in Oracle
Cloud Infrastructure to reassemble the packet.

set security ipsec vpn oracle-vpn-<oracle_headend_1> bind-interface st0.1
set security ipsec vpn oracle-vpn-<oracle_headend_1> vpn-monitor
set security ipsec vpn oracle-vpn-<oracle_headend_1> ike gateway gw_oracle-<oracle_headend_1>
set security ipsec vpn oracle-vpn-<oracle_headend_1> ike ipsec-policy ipsec_pol_oracle-vpn
set security ipsec vpn oracle-vpn-<oracle_headend_1> df-bit clear
set security ipsec vpn establish-tunnels immediately

#3: Tunnel Interface Configuration

Defining the Tunnel Interface
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces st0.1 family inet address <inside_tunnel_interface_ip_address>
set interfaces st0.1 family inet mtu 1430
set interfaces <inside_tunnel_interface> unit 0 family inet address <internal_network_ip_range>

Setting the Security Zones for Oracle
```

## CHAPTER 23 Networking

---

```
USER_DEFINED: Replace the parameters in the section below as needed
Tunnel interface st0.1, inside_tunnel_interface and outside_public_interface are each defined in its
own security zones.

set security zones security-zone oracle_vpn interfaces st0.1
set security zones security-zone oracle_trust interfaces <inside_tunnel_interface>
set security zones security-zone internet_untrust interfaces <outside_public_interface>

The security zone protecting outside interface of the router must be configured to allow IKE and ping
inbound traffic.

set security zones security-zone internet_untrust interfaces <outside_public_interface> host-inbound-
traffic system-services ike
set security zones security-zone internet_untrust interfaces <outside_public_interface> host-inbound-
traffic system-services ping

The security zone protecting the logical tunnel interface must be configured to allow BGP inbound
traffic.

set security zones security-zone oracle_vpn interfaces st0.1 host-inbound-traffic protocols bgp

This option causes the router to reduce the Maximum Segment Size of TCP packets to prevent packet
fragmentation.

set security flow tcp-mss ipsec-vpn mss 1387

#4: Policies

Setting the Security Policies for Oracle
Policies basically define the permitted flow of traffic between defined security zones.
The configuration template permits any ipv4 traffic sourced and destined between security zones
oracle_trust and oracle_vpn.

set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match source-address any-
ipv4
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match destination-address
any-ipv4
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match application any
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match source-identity any
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out then permit
```

## CHAPTER 23 Networking

---

```
#5a: Border Gateway Protocol (BGP) Configuration

USER_DEFINED: Replace the parameters in the section below as needed

BGP is used within the tunnel to exchange prefixes between the dynamic routing gateway and your CPE.
The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG
advertises the VCN's subnets.
The configuration template uses a basic route policy to advertise a default route to the DRG.
To advertise additional prefixes to the Oracle VCN, add additional prefixes to the term ORACLE-DEFAULT
policy. Make sure the prefix is present in the routing table of the device with a valid next-hop.

You configure the local BGP Autonomous System Number (BGP ASN) when you set up the IPSec connection in
the Oracle Console. If you later need to change the ASN, you must recreate the CPE object and IPSec
connection in the Oracle Console.

set policy-options policy-statement ORACLE-DEFAULT term default from route-filter 0.0.0.0/0 exact

set policy-options policy-statement ORACLE-DEFAULT term default then accept
set policy-options policy-statement ORACLE-DEFAULT term reject then reject

set protocols bgp group ebgp type external
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> export ORACLE-
DEFAULT
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> peer-as 31898
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> local-as <bgp_asn>

#5b: Static Route Configuration

USER_DEFINED: Replace the parameters in the section below as needed
In case you plan to use static routing to get traffic through the IPSec tunnels, you can point the
routes down to the tunnel interfaces. You should redistribute these routes into your on-premises
network. Configuration for CPE to VCN static routes:

set routing-options static route <vcn_range> next-hop st0.1

IPSec Tunnel 2

#1: Internet Key Exchange (IKE) Configuration (Phase 1)
```

## CHAPTER 23 Networking

---

```
Defining the IKE Proposal for Oracle
This IKE (Phase 1) configuration template uses AES256, SHA384, Diffie-Hellman Group 5, and 28800
second (8 hours) IKE session key lifetime.
If different parameters are required, modify this template before applying the configuration.

set security ike proposal oracle-ike-proposal authentication-method pre-shared-keys
set security ike proposal oracle-ike-proposal authentication-algorithm sha-384
set security ike proposal oracle-ike-proposal encryption-algorithm aes-256-cbc
set security ike proposal oracle-ike-proposal lifetime-seconds 28800
set security ike proposal oracle-ike-proposal dh-group group5

Defining the IKE Policy for Oracle
USER_DEFINED: Replace the parameters in the section below as needed

set security ike policy ike_pol_oracle-vpn-<oracle_headend_2> mode main
set security ike policy ike_pol_oracle-vpn-<oracle_headend_2> proposals oracle-ike-proposal
set security ike policy ike_pol_oracle-vpn-<oracle_headend_2> pre-shared-key ascii-text <connection_
prsharedkey_2>

Setting up Public Interface with the CPE Public IP.
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces <outside_public_interface> unit 0 family inet address <cpe_public_ip_address>

Defining the IKE Gateway for Oracle
USER_DEFINED: Replace the parameters in the section below as needed.
This option enables IPsec Dead Peer Detection, which causes periodic messages to be sent to ensure a
Security Association remains operational.
If you want to use IKEv1 instead, comment out the line below that ends with "version v2-only".

set security ike gateway gw_oracle-<oracle_headend_2> ike-policy ike_pol_oracle-vpn-<oracle_headend_2>
set security ike gateway gw_oracle-<oracle_headend_2> external-interface <outside_public_interface>
set security ike gateway gw_oracle-<oracle_headend_2> address <oracle_headend_2>
set security ike gateway gw_oracle-<oracle_headend_2> dead-peer-detection
set security ike gateway gw_oracle-<oracle_headend_2> version v2-only
set security ike gateway gw_oracle-<oracle_headend_2> local-identity inet <cpe_public_ip_address>

#2: IPSec Configuration

Defining the IPSec (Phase 2) Proposal for Oracle
The IPSec proposal defines the protocol, authentication, encryption, and lifetime parameters for our
IPsec security association.
```

## CHAPTER 23 Networking

```
The configuration template sets AES256 for encryption, SHA1 for authentication, enables PFS group 5,
and sets the IPsec session key lifetime to 3600 seconds (1 hour).
The IPsec policy incorporates the Diffie-Hellman group and the IPsec proposal.
If different parameters are required, modify this template before applying the configuration.

set security ipsec vpn-monitor-options
set security ipsec proposal oracle-ipsec-proposal protocol esp
set security ipsec proposal oracle-ipsec-proposal authentication-algorithm hmac-sha1-96;
set security ipsec proposal oracle-ipsec-proposal encryption-algorithm aes-256-cbc
set security ipsec proposal oracle-ipsec-proposal lifetime-seconds 3600

Defining the IPsec (PHASE 2) policy for Oracle
set security ipsec policy ipsec_pol_oracle-vpn perfect-forward-secrecy keys group5
set security ipsec policy ipsec_pol_oracle-vpn proposals oracle-ipsec-proposal

Defining Security Association for Oracle
USER_DEFINED: Replace the parameters in the section below as needed
The IPsec Policy and IKE gateways are associated with a tunnel interface (st0.2). If other tunnels are
defined on your router, you must specify a unique interface name.
The df-bit clear option allows the SRX to fragment the packet and send it to the end host in Oracle
Cloud Infrastructure to reassemble the packet.

set security ipsec vpn oracle-vpn-<oracle_headend_2> bind-interface st0.2
set security ipsec vpn oracle-vpn-<oracle_headend_2> vpn-monitor
set security ipsec vpn oracle-vpn-<oracle_headend_2> ike gateway gw_oracle-<oracle_headend_2>
set security ipsec vpn oracle-vpn-<oracle_headend_2> ike ipsec-policy ipsec_pol_oracle-vpn
set security ipsec vpn oracle-vpn-<oracle_headend_2> df-bit clear
set security ipsec vpn establish-tunnels immediately

#3: Tunnel Interface Configuration

Defining the Tunnel Interface
USER_DEFINED: Replace the parameters in the section below as needed

set interfaces st0.2 family inet address <inside_tunnel_interface_ip_address>
set interfaces st0.2 family inet mtu 1430
set interfaces <inside_tunnel_interface> unit 0 family inet address <internal_network_ip_range>

Setting the Security Zones for Oracle
USER_DEFINED: Replace the parameters in the section below as needed
Tunnel interface st0.2, inside_tunnel_interface and outside_public_interface are each defined in its
own security zones.
```

## CHAPTER 23 Networking

---

```
set security zones security-zone oracle_vpn interfaces st0.2
set security zones security-zone oracle_trust interfaces <inside_tunnel_interface>
set security zones security-zone internet_untrust interfaces <outside_public_interface>

The security zone protecting outside interface of the router must be configured to allow IKE and ping
inbound traffic.

set security zones security-zone internet_untrust interfaces <outside_public_interface> host-inbound-
traffic system-services ike
set security zones security-zone internet_untrust interfaces <outside_public_interface> host-inbound-
traffic system-services ping

The security zone protecting the logical tunnel interface must be configured to allow BGP inbound
traffic.

set security zones security-zone oracle_vpn interfaces st0.2 host-inbound-traffic protocols bgp

This option causes the router to reduce the Maximum Segment Size of TCP packets to prevent packet
fragmentation.

set security flow tcp-mss ipsec-vpn mss 1387

#4: Policies

Setting the Security Policies for Oracle
Policies basically define the permitted flow of traffic between defined security zones.
The configuration template permits any IPv4 traffic sourced and destined between security zones
oracle_trust and oracle_vpn.

set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match source-address any-
ipv4
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match destination-address
any-ipv4
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match application any
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out match source-identity any
set security policies from-zone oracle_trust to-zone oracle_vpn policy vpn-out then permit

#5a: Border Gateway Protocol (BGP) Configuration

USER_DEFINED: Replace the parameters in the section below as needed
```

## CHAPTER 23 Networking

---

```
BGP is used within the tunnel to exchange prefixes between the dynamic routing gateway and your CPE.
The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG
advertises the VCN's subnets.
The configuration template uses a basic route policy to advertise a default route to the DRG.
To advertise additional prefixes to the Oracle VCN, add additional prefixes to the term ORACLE-DEFAULT
policy. Make sure the prefix is present in the routing table of the device with a valid next-hop.

You configure the local BGP Autonomous System Number (BGP ASN) when you set up the IPSec connection in
the Oracle Console. If you later need to change the ASN, you must recreate the CPE object and IPSec
connection in the Oracle Console.

set policy-options policy-statement ORACLE-DEFAULT term default from route-filter 0.0.0.0/0 exact

set policy-options policy-statement ORACLE-DEFAULT term default then accept
set policy-options policy-statement ORACLE-DEFAULT term reject then reject

set protocols bgp group ebgp type external
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> export ORACLE-
DEFAULT
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> peer-as 31898
set protocols bgp group ebgp neighbor <inside_tunnel_interface_ip_address_neighbor> local-as <bgp_asn>

#5b: Static Route Configuration

USER_DEFINED: Replace the parameters in the section below as needed
In case you plan to use static routing to get traffic through the IPSec tunnels, you can point the
routes down to the tunnel interfaces. You should redistribute these routes into your on-premises
network. Configuration for CPE to VCN static routes:

set routing-options static route <vcn_range> next-hop st0.2
```

### Verification

Use the following command to verify security associations (SAs).

```
show security ipsec security-associations
```

Use the following command to check the BGP status.

```
show bgp summary
```

Use the following commands to check the routes advertised to and received from Oracle Cloud Infrastructure.

## CHAPTER 23 Networking

---

```
show route advertising-protocol bgp <neighbor-address>
show route receive-protocol bgp <neighbor-address>
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### Libreswan

[Libreswan](#) is an open source IPsec implementation that is based on FreeS/WAN and Openswan. Most Linux distributions include Libreswan or make it easy to install. You can install it on hosts in either your on-premises network or a cloud provider network. For an example of setting up a Libreswan host in another cloud provider to connect to your Oracle Cloud Infrastructure VCN, see [Access to Other Clouds with Libreswan](#).

This topic provides configuration for CPE that is running Libreswan. Virtual tunnel interface (VTI) support for this route-based configuration requires minimum Libreswan version 3.18 and a recent Linux 3.x or 4.x kernel. This configuration was validated using Libreswan version 3.29.



### Important

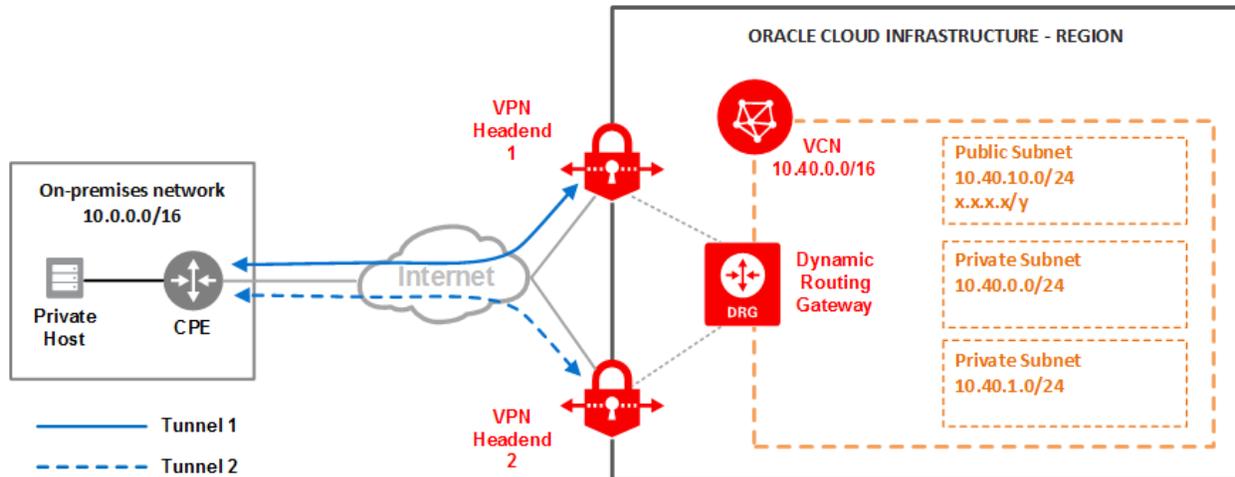
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

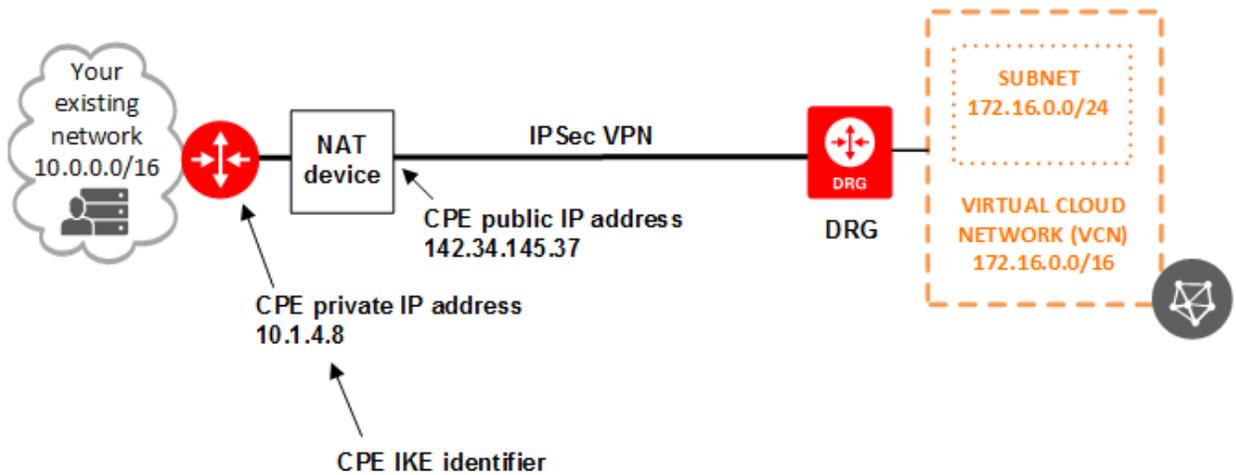
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

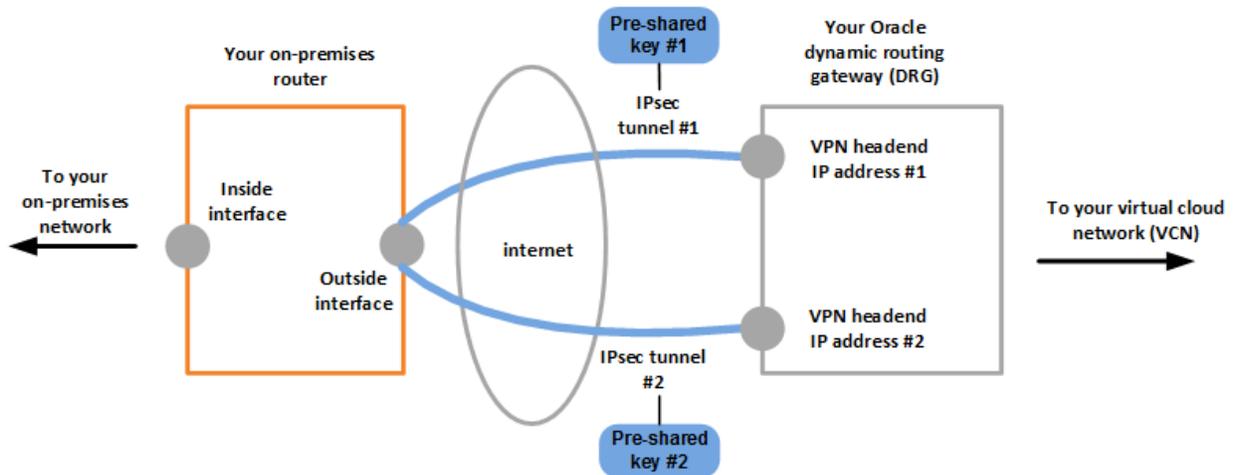
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

**Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPSec connection.



### DEFAULT LIBRESWAN CONFIGURATION FILES

The default Libreswan installation creates the following files:

- `etc/ipsec.conf`: The root of the Libreswan configuration.
- `/etc/ipsec.secrets`: The root of the location where Libreswan looks for secrets (the tunnel pre-shared keys).
- `/etc/ipsec.d/`: A directory for storing the `.conf` and `.secrets` files for your Oracle Cloud Infrastructure tunnels (for example: `oci-ipsec.conf` and `oci-ipsec.secrets`). Libreswan encourages you to create these files in this folder.

The default `etc/ipsec.conf` file includes this line:

```
include /etc/ipsec.d/*.conf
```

The default `etc/ipsec.secrets` file includes this line:

```
include /etc/ipsec.d/*.secrets
```

The preceding lines automatically merge all the `.conf` and `.secrets` files in the `/etc/ipsec.d` directory into the main configuration and secrets files that Libreswan uses.

### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

You specify the IKE version when setting up the IPSec configuration file in [task 3](#) in the next section. In that example file, there's a comment showing how to configure IKEv1 versus IKEv2.

### CONFIGURATION PROCESS

Libreswan supports both route-based and policy-based tunnels. The tunnel types can coexist without interfering with each other. The Oracle VPN headends use route-based tunnels. Oracle recommends that you configure Libreswan with the [Virtual Tunnel Interface \(VTI\) configuration syntax](#).

For details about the specific parameters used in this document, see [Supported IPSec Parameters](#).

### Task 1: Prepare the Libreswan instance

Depending on the Linux distribution you're using, you might need to enable IP forwarding on your interface to allow clients to send and receive traffic through Libreswan. In the `/etc/sysctl.conf` file, set the following values and apply the updates with `sudo sysctl -p`.

If you're using an interface other than `eth0`, change `eth0` in the following example to your interface (lines 5 and 7).

```
net.ipv4.ip_forward=1
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.send_redirects = 0
net.ipv4.conf.eth0.send_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
```

### Task 2: Determine the required configuration values

The Libreswan configuration uses the following variables. Determine the values before proceeding with the configuration.

- `${cpeLocalIP}`: The IP address of your Libreswan device.
- `${cpePublicIpAddress}`: The public IP address for Libreswan. This is the IP address of your outside interface. Depending on your network topology, the value might be different from `${cpeLocalIP}`.
- `${oracleHeadend1}`: For the first tunnel, the Oracle public IP endpoint obtained from the Oracle Console.
- `${oracleHeadend2}`: For the second tunnel, the Oracle public IP endpoint obtained from the Oracle Console.
- `${vti1}`: The name of the first VTI used. For example, `vti1`.
- `${vti2}`: The name of the second VTI used. For example, `vti2`.

- `${sharedSecret1}`: The pre-shared key for the first tunnel. You can use the default Oracle-provided pre-shared key, or provide your own when you set up the IPSec connection in the Oracle Console.
- `${sharedSecret2}`: The pre-shared key for the second tunnel. You can use the default Oracle-provided pre-shared key, or provide your own when you set up the IPSec connection in the Oracle Console.
- `${vcnCidrNetwork}`: The VCN IP range.

### Task 3: Set up the configuration file: `/etc/ipsec.d/oci-ipsec.conf`

Libreswan configuration uses the concept of *left* and *right* to define the configuration parameters for your local CPE device and the remote gateway. Either side of the connection (the *conn* in the Libreswan configuration) can be left or right, but the configuration for that connection must be consistent. In this example:

- **left:** Your local Libreswan CPE
- **right:** The Oracle VPN headend

Use the following template for your `/etc/ipsec.d/oci-ipsec.conf` file. The file defines the two tunnels that Oracle creates when you set up the IPSec connection.



#### Important

If your CPE is behind a 1-1 NAT device, uncomment the `leftid` parameter and set it equal to the `${cpePublicIpAddress}`.

```
conn oracle-tunnel-1
left=${cpeLocalIP}
leftid=${cpePublicIpAddress} # See preceding note about 1-1 NAT device
right=${oracleHeadend1}
authby=secret
leftsubnet=0.0.0.0/0
```

```
rightsubnet=0.0.0.0/0
auto=start
mark=5/0xffffffff # Needs to be unique across all tunnels
vti-interface=${vti1}
vti-routing=no
ikev2=no # To use IKEv2, change to ikev2=insist
ike=aes_cbc256-sha2_384;modp1536
phase2alg=aes_gcm256;modp1536
encapsulation=no
ikelifetime=28800s
salifetime=3600s
conn oracle-tunnel-2
left=${cpeLocalIP}
leftid=${cpePublicIpAddress} # See preceding note about 1-1 NAT device
right=${oracleHeadend2}
authby=secret
leftsubnet=0.0.0.0/0
rightsubnet=0.0.0.0/0
auto=start
mark=6/0xffffffff # Needs to be unique across all tunnels
vti-interface=${vti2}
vti-routing=no
ikev2=no # To use IKEv2, change to ikev2=insist
ike=aes_cbc256-sha2_384;modp1536
phase2alg=aes_gcm256;modp1536
encapsulation=no
ikelifetime=28800s
salifetime=3600s
```

### Task 4: Set up the secrets file: /etc/ipsec.d/oci-ipsec.secrets

Use the following template for your /etc/ipsec.d/oci-ipsec.secrets file. It contains two lines per IPSec connection (one line per tunnel).

```
${cpePublicIpAddress} ${oracleHeadend1}: PSK "${sharedSecret1}"
${cpePublicIpAddress} ${oracleHeadend2}: PSK "${sharedSecret2}"
```

### Task 5: Restart the Libreswan configuration

After setting up your configuration and secrets files, you must restart the Libreswan service.



#### Important

Restarting the Libreswan service may impact existing tunnels.

## CHAPTER 23 Networking

---

The following command rereads the config file and restarts the Libreswan service.

```
service ipsec restart
```

### Task 6: Configure IP routing

Use the following `ip` command to create static routes that send traffic to your VCN through the IPsec tunnels. If you're logged in with an unprivileged user account, you might need to use `sudo` before the command.



#### Important

Static routes created with the `ip route` command do not persist through a reboot. To determine how to make your routes persist, refer to the documentation of your Linux distribution of choice.

```
ip route add ${VcnCidrBlock} nexthop dev ${vti1} nexthop dev ${vti2}
ip route show
```

### Verification

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

#### CHECKING THE LIBRESWAN STATUS

Check the current state of your Libreswan tunnels by using the following command.

```
ipsec status
```

The tunnel is established if you see a line that includes the following:

```
STATE_MAIN_I4: ISAKMP SA established
```

If you're using IKEv2, you see the following:

## CHAPTER 23 Networking

---

```
STATE_V2_IPSEC_I (IPsec SA established)
```

In the future, if you need to open a support ticket with Oracle about your Libreswan tunnel, include the output of the preceding `ipsec status` command.

### CHECKING THE TUNNEL INTERFACE STATUS

Check if the virtual tunnel interfaces are up or down by using the `ifconfig` command or the `ip link show` command. You can also use applications such as `tcpdump` with the interfaces.

Here's an example of the `ifconfig` output with a working Libreswan implementation that shows the available VTIs.

```
ifconfig
<output trimmed>

vti01: flags=209<UP,POINTOPOINT,RUNNING,NOARP> mtu 8980
 inet6 fe80::5efe:a00:2 prefixlen 64 scopeid 0x20<link>
 tunnel txqueuelen 1000 (IPIP Tunnel)
 RX packets 0 bytes 0 (0.0 B)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 0 bytes 0 (0.0 B)
 TX errors 10 dropped 0 overruns 0 carrier 10 collisions 0

vti02: flags=209<UP,POINTOPOINT,RUNNING,NOARP> mtu 8980
 inet6 fe80::5efe:a00:2 prefixlen 64 scopeid 0x20<link>
 tunnel txqueuelen 1000 (IPIP Tunnel)
 RX packets 0 bytes 0 (0.0 B)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 0 bytes 0 (0.0 B)
 TX errors 40 dropped 0 overruns 0 carrier 40 collisions 0
```

Here's an example of the `ip link show` output:

```
ip link show
<output trimmed>

9: vti01@NONE: <POINTOPOINT,NOARP,UP,LOWER_UP> mtu 8980 qdisc noqueue
state UNKNOWN mode DEFAULT group default qlen 1000
 link/ipip 10.0.0.2 peer 129.213.240.52

10: vti02@NONE: <POINTOPOINT,NOARP,UP,LOWER_UP> mtu 8980 qdisc noqueue
```

## CHAPTER 23 Networking

---

```
state UNKNOWN mode DEFAULT group default qlen 1000
 link/ipip 10.0.0.2 peer 129.213.240.51
```

Also, in the Oracle Console, each IPsec tunnel should now be in the UP state.

### NEC IX Series

This configuration was validated using an IX3315 running Firmware Ver.10.1.16 and IX2106 running Firmware Ver.10.1.16.



#### **Important**

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

### **Before Starting**

Before configuring your CPE, make sure to:

- Configure your internet provider settings.
- Configure firewall rules to open UDP port 500, UDP port 4500, and ESP.

### **Supported Encryption Domain or Proxy ID**

The values for the encryption domain (also known as a proxy ID, security parameter index (SPI), or traffic selector) depend on whether your CPE supports route-based tunnels or policy-based tunnels. For more information about the correct encryption domain values to use, see [Supported Encryption Domain or Proxy ID](#).

### Parameters from API or Console

Get the following parameters from the Oracle Cloud Infrastructure Console or API.

#### **`${ipAddress#}` - one per tunnel**

- Oracle VPN headend IPsec tunnel endpoints.
- Example values: 129.146.12.52, 129.146.13.52

#### **`${sharedSecret#}` - one per tunnel**

- The IPsec IKE pre-shared-key.
- Example value: EXAMPLEDPfAMkD7nTH3SWr6OFabdT6exXn6enSlsKbE

#### **`${cpePublicIpAddress}`**

- The public IP address for the CPE (previously made available to Oracle via the Console).

#### **`${vcnCidrBlock}`**

- When creating the VCN, your company selected this CIDR to represent the IP aggregate network for all VCN hosts.
- Example Value: 10.0.0.0/20

### Parameters Based on Current CPE Configuration and State

The following parameters are based on your current CPE configuration.

#### **`${tunnelNumber#}` - one per tunnel**

- You will need one unused unit number per tunnel.
- Example values: 1, 2

#### **`${ikePolicy#}` - one per tunnel**

- You will need one unused IKE policy name per tunnel.
- Example values: ike-policy1, ike-policy2

#### **`${ipsecPolicy#}` - one per tunnel**

- You will need one unused autokey policy map name per tunnel.
- Example values: ipsec-policy1, ipsec-policy2

### **`${lanIpAddress}`**

- The local IP address of your CPE.
- Example value: 192.168.100.254

### **`${lanInterfaceNumber}`**

- The LAN interface of your CPE
- Example value: 1.0

### **Config Template Parameter Summary**

Each region has multiple Oracle IPSec headends. The template below allows you to set up multiple tunnels on your CPE, each to a corresponding headend. In the table below, "User" is you/your company.

<b>Parameter</b>	<b>Source</b>	<b>Example Value</b>
<code>\${ipAddress1}</code>	Console/API	129.146.12.52
<code>\${ipAddress2}</code>	Console/API	129.146.13.52
<code>\${sharedSecret1}</code>	Console/API	(long string)
<code>\${sharedSecret2}</code>	Console/API	(long string)
<code>\${cpePublicIpAddress}</code>	User	203.0.113.1
<code>\${vcnCidrBlock}</code>	User	10.0.0.0/20
<code>\${tunnelNumber1}</code>	User	1
<code>\${tunnelNumber2}</code>	User	2

Parameter	Source	Example Value
<code>\${ikePolicy1}</code>	User	ike-policy1
<code>\${ikePolicy2}</code>	User	ike-policy2
<code>\${ipsecPolicy1}</code>	User	ipsec-policy1
<code>\${ipsecPolicy2}</code>	User	ipsec-policy2
<code>\${lanInterfaceNumber}</code>	User	1.0
<code>\${lanIpAddress}</code>	User	192.168.100.254



**Important**

The following ISAKMP and IPsec policy parameter values are applicable to VPN Connect in the commercial cloud. For the [Government Cloud](#), you must use the values listed in [Required VPN Connect Parameters for Government Cloud](#).

**Commercial Cloud: ISAKMP Policy Options**

Parameter	Recommended Value
ISAKMP protocol version	Version 1
Exchange type	Main mode
Authentication method	Pre-shared keys
Encryption	AES-256-cbc

Parameter	Recommended Value
Authentication algorithm	SHA-256
Diffie-Hellman Group	Group 5
IKE session key lifetime	28800 seconds (8 hours)

### Commercial Cloud: IPsec Policy Options

Parameter	Recommended Value
IPsec protocol	ESP, tunnel-mode
Encryption	AES-256-cbc
Authentication algorithm	HMAC-SHA1-96
Diffie-Hellman Group	Group 5
Perfect Forward Secrecy	Enabled
IPsec session key lifetime	3600 seconds (1 hour)

### CPE Configuration

#### CONFIGURE ISAKMP AND IPSEC POLICIES

```
ip access-list sec-list permit ip src any dest any
ike nat-traversal
!
ike proposal ike-prop encryption aes-256 hash sha2-256 group 1536-bit
ike policy ${ikePolicy1} peer ${ipAddress1} key ${sharedSecret1} ike-prop
ike policy ${ikePolicy2} peer ${ipAddress2} key ${sharedSecret2} ike-prop
!
ipsec autokey-proposal ipsec-prop esp-aes-256 esp-sha lifetime time 3600
ipsec autokey-map ${ipsecPolicy1} sec-list peer ${ipAddress1} ipsec-prop pfs 1536-bit
ipsec autokey-map ${ipsecPolicy2} sec-list peer ${ipAddress2} ipsec-prop pfs 1536-bit
```

## CHAPTER 23 Networking

---

### CONFIGURE KEEPALIVE SETTING OF ICMP

```
watch-group watch_tunnel1 10
 event 20 ip unreachable-host ${lanIpAddress} Tunnel${tunnelNumber1} source
GigaEthernet${lanInterfaceNumber}
 action 10 ip shutdown-route ${vcnCidrBlock} Tunnel${tunnelNumber1}
 action 20 ipsec clear-sa Tunnel${tunnelNumber1}
!
network-monitor watch_tunnel1 enable
!
watch-group watch_tunnel2 10
 event 20 ip unreachable-host ${lanIpAddress} Tunnel${tunnelNumber2} source
GigaEthernet${lanInterfaceNumber}
 action 10 ip shutdown-route ${vcnCidrBlock} Tunnel${tunnelNumber2}
 action 20 ipsec clear-sa Tunnel${tunnelNumber2}
!
network-monitor watch_tunnel2 enable
```

### CONFIGURE VIRTUAL TUNNEL INTERFACES

```
interface Tunnel${tunnelNumber1}
 tunnel mode ipsec
 ip unnumbered GigaEthernet${lanInterfaceNumber}
 ip tcp adjust-mss auto
 ipsec policy tunnel ipsec-policy1 out
 no shutdown
!
interface Tunnel${tunnelNumber2}
 tunnel mode ipsec
 ip unnumbered GigaEthernet${lanInterfaceNumber}
 ip tcp adjust-mss auto
 ipsec policy tunnel ipsec-policy2 out
 no shutdown
```

### CONFIGURE STATIC ROUTES

```
ip ufs-cache enable
ip multipath per-flow
ip route ${vcnCidrBlock} Tunnel10.0
ip route ${vcnCidrBlock} Tunnel11.0
```

### Openswan

If you want to use Openswan to create an IPSec VPN to Oracle Cloud Infrastructure, see [Libreswan](#).

#### How Openswan and Libreswan Are Related

[Openswan](#) is a well-known IPSec implementation for Linux. It began as a fork of the now-defunct [FreeS/WAN project](#) in 2003. Unlike the FreeS/WAN project, it didn't exclusively target the GNU/Linux operation system, but expanded its use to other operating systems. In 2012, FreeS/WAN renamed itself to [The Libreswan Project](#) because of a lawsuit over a trademark of the name *openswan*.

As a result, when you try to install or query the Openswan package on Oracle Linux, by default the Libreswan package is installed or shown instead. The following yum search query command illustrates this behavior:

```
$ sudo yum search openswan
Loaded plugins: langpacks, ulninfo
Matched: openswan =====
NetworkManager-libreswan.x86_64 : NetworkManager VPN plugin for libreswan
NetworkManager-libreswan-gnome.x86_64 : NetworkManager VPN plugin for libreswan-GNOME files
libreswan.x86_64 : IPsec implementation with IKEv1 and IKEv2 keying protocols
```

Libreswan is maintained by The Libreswan Project and has been under active development for over 15 years, going back to the FreeS/WAN Project. For more information, see the [project's history](#).

### Palo Alto

This topic provides configuration for a Palo Alto device. The configuration was validated using PAN-OS version 8.0.0.

Palo Alto experience is required.



### Important

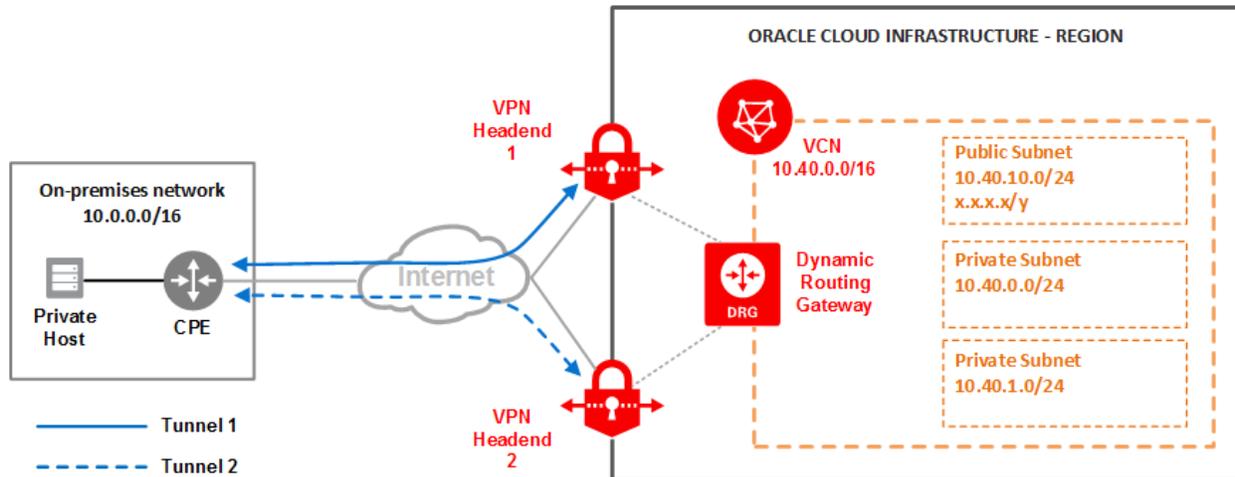
Oracle provides configuration instructions for a set of [vendors and devices](#). Make sure to use the configuration for the correct vendor.

If the device or software version that Oracle used to verify the configuration does not exactly match your device or software, the configuration might still work for you. Consult your vendor's documentation and make any necessary adjustments.

If your device is for a vendor not in the list of verified vendors and devices, or if you're already familiar with configuring your device for IPsec, see the list of [supported IPsec parameters](#) and consult your vendor's documentation for assistance.

VPN Connect is the IPsec VPN that Oracle Cloud Infrastructure offers for connecting your on-premises network to a virtual cloud network (VCN).

The following diagram shows a basic IPsec connection to Oracle Cloud Infrastructure with redundant tunnels. IP addresses used in this diagram are for example purposes only.



## Best Practices

This section covers general best practices and considerations for using VPN Connect.

### CONFIGURE ALL TUNNELS FOR EVERY IPSEC CONNECTION

Oracle deploys two IPsec headends for each of your connections to provide high availability for your mission-critical workloads. On the Oracle side, these two headends are on different routers for redundancy purposes. Oracle recommends configuring all available tunnels for maximum redundancy. This is a key part of the "Design for Failure" philosophy.

### HAVE REDUNDANT CPEs IN YOUR ON-PREMISES NETWORK LOCATIONS

Each of your sites that connects with IPsec to Oracle Cloud Infrastructure should have redundant edge devices (also known as customer-premises equipment (CPE)). You add each CPE to the Oracle Console and create a separate IPsec connection between your dynamic routing gateway (DRG) and each CPE. For each IPsec connection, Oracle provisions two tunnels on geographically redundant IPsec headends. For more information, see the [Connectivity Redundancy Guide \(PDF\)](#).

### ROUTING PROTOCOL CONSIDERATIONS

When you create an IPsec VPN, it has two redundant IPsec tunnels. Oracle encourages you to configure your CPE to use both tunnels (if your CPE supports it). Note that in the past, Oracle created IPsec VPNs that had up to four IPsec tunnels.

The following two routing types are available, and you choose the routing type separately for each tunnel in the IPsec VPN:

- **BGP dynamic routing:** The available routes are learned dynamically through BGP. The DRG dynamically learns the routes from your on-premises network. On the Oracle side, the DRG advertises the VCN's subnets.
- **Static routing:** When you set up the IPsec connection to the DRG, you specify the particular routes to your on-premises network that you want the VCN to know about. You also must configure your CPE device with static routes to the VCN's subnets. These routes are not learned dynamically.

For more information about routing with VPN Connect, including Oracle recommendations on how to manipulate the BGP best path selection algorithm, see [Routing for the Oracle IPsec VPN](#).

### OTHER IMPORTANT CPE CONFIGURATIONS

Ensure access lists on your CPE are configured correctly to not block necessary traffic from or to Oracle Cloud Infrastructure.

If you have multiple tunnels up simultaneously, ensure that your CPE is configured to handle traffic coming from your VCN on any of the tunnels. For example, you need to disable ICMP inspection, configure TCP state bypass, and so on. For more details about the appropriate configuration, contact your CPE vendor's support.

### Caveats and Limitations

This section covers general important characteristics and limitations of VPN Connect to be aware of.

### ASYMMETRIC ROUTING

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec connection. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

When you use multiple tunnels to Oracle Cloud Infrastructure, Oracle recommends that you configure your routing to deterministically route traffic through the preferred tunnel. If you want to use one IPsec tunnel as primary and another as backup, configure more-specific routes for the primary tunnel (BGP) and less-specific routes (summary or default route) for the backup tunnel (BGP/static). Otherwise, if you advertise the same route (for example, a default route) through all tunnels, return traffic from your VCN to your on-premises network will route to any of the available tunnels (because Oracle uses asymmetric routing).

For specific Oracle routing recommendations about how to force symmetric routing, see [Preferring a Specific Tunnel in the IPsec VPN](#).

### ROUTE-BASED OR POLICY-BASED IPSEC VPN

The IPsec protocol uses Security Associations (SAs) to determine how to encrypt packets. Within each SA, you define encryption domains to map a packet's source and destination IP address and protocol type to an entry in the SA database to define how to encrypt or decrypt a packet.



#### Note

Other vendors or industry documentation might use the term *proxy ID*, *security parameter index (SPI)*, or *traffic selector* when referring to SAs or encryption domains.

There are two general methods for implementing IPsec tunnels:

- **Route-based tunnels:** Also called *next-hop-based tunnels*. A route table lookup is performed on a packet's destination IP address. If that route's egress interface is an

IPSec tunnel, the packet is encrypted and sent to the other end of the tunnel.

- **Policy-based tunnels:** The packet's source and destination IP address and protocol are matched against a list of policy statements. If a match is found, the packet is encrypted based on the rules in that policy statement.

The Oracle VPN headends use route-based tunnels but can work with policy-based tunnels with some caveats listed in the following sections.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

### Encryption domain for route-based tunnels

If your CPE supports route-based tunnels, use that method to configure the tunnel. It's the simplest configuration with the most interoperability with the Oracle VPN headend.

Route-based IPSec uses an encryption domain with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** Any (0.0.0.0/0)
- **Protocol:** IPv4

If you need to be more specific, you can use a single summary route for your encryption domain values instead of a default route.

### Encryption domain for policy-based tunnels

If your CPE supports only policy-based tunnels, there are restrictions on the policy that you

can use on the CPE.

When you use policy-based tunnels, every policy entry that you define generates a pair of IPsec SAs. This pair is referred to as an *encryption domain*.



### Important

The Oracle VPN headend supports only a single encryption domain. If your policy includes multiple entries, the tunnel will flap or there will be connectivity problems in which only a single policy works at any one time.

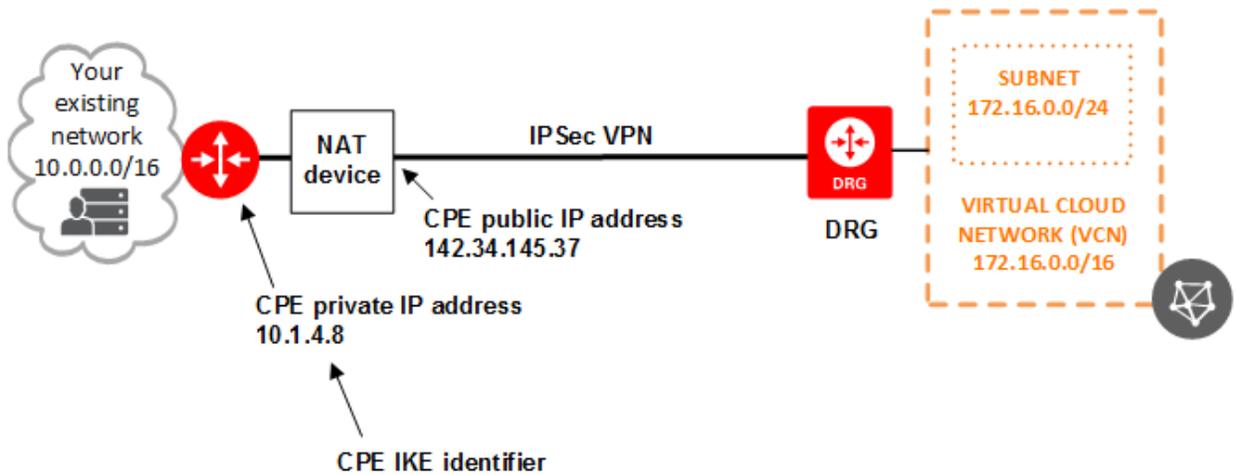
If you use policy-based IPsec, Oracle recommends using a *single encryption domain* with the following values:

- **Source IP address:** Any (0.0.0.0/0)
- **Destination IP address:** VCN CIDR (example: 10.120.0.0/20)
- **Protocol:** IPv4

Make sure the single encryption domain matches any traffic that needs to go from your on-premises network across the IPsec tunnel to the VCN. The VCN CIDR must not overlap with your on-premises network.

### IF YOUR CPE IS BEHIND A NAT DEVICE

In general, the CPE IKE identifier configured on your end of the connection must match the CPE IKE identifier that Oracle is using. By default, Oracle uses the CPE's *public* IP address, which you provide when you create the CPE object in the Oracle Console. However, if your CPE is behind a NAT device, the CPE IKE identifier configured on your end might be the CPE's *private* IP address, as show in the following diagram.



**Note**

Some CPE platforms do not allow you to change the local IKE identifier. If you cannot, you must change the remote IKE ID in the Oracle Console to match your CPE's local IKE ID. You can provide the value either when you set up the IPsec connection, or later, by editing the IPsec connection. Oracle expects the value to be either an IP address or a fully qualified domain name (FQDN) such as *cpe.example.com*. For instructions, see [Changing the CPE IKE Identifier That Oracle Uses](#).

**Supported IPsec Parameters**

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

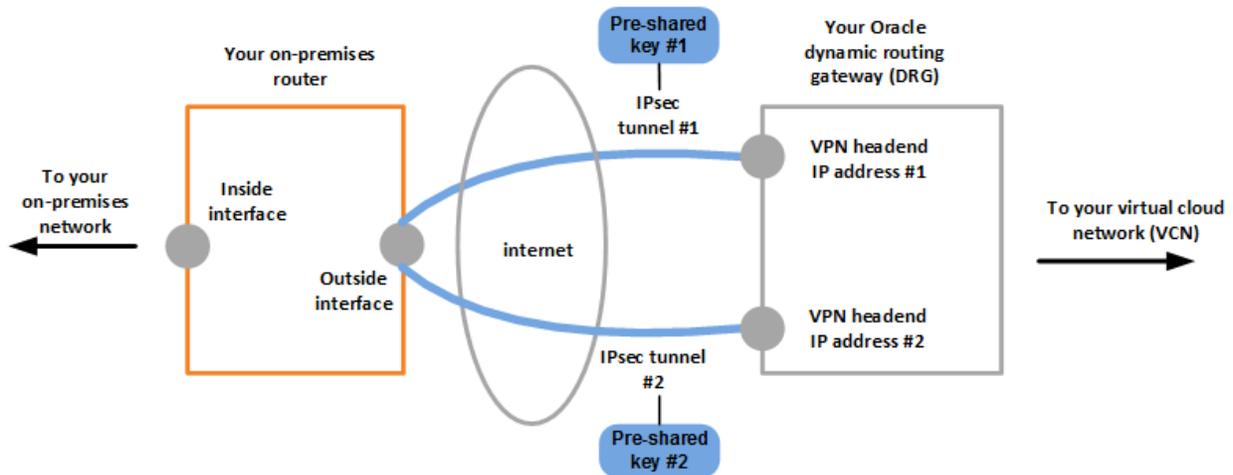
The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### CPE Configuration

**Important**

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for your CPE. If you need support or further assistance, contact your CPE vendor's support directly.

The following figure shows the basic layout of the IPsec connection.



### IMPORTANT DETAILS ABOUT THE CONFIGURATION INSTRUCTIONS

- **Commits:** For PAN to activate the configuration, you must perform the commit action after any configuration change.
- **Example IP addresses:** The example configuration uses IP addresses from class A 10.0.0.0/8 (RFC1918) and 198.51.100.0/24 (RFC5735). When you perform the configuration on the CPE, use the correct IP addressing plan for your networking topology.

The example configuration uses the following variables and values:

- **Inside tunnel1 interface - CPE:** 198.51.100.1/30
- **Inside tunnel2 interface - CPE:** 198.51.100.5/30
- **Inside tunnel1 interface - Oracle:** 198.51.100.2/30
- **Inside tunnel2 interface - Oracle:** 198.51.100.6/30
- **CPE ASN:** 64511
- **On-premises network :** 10.200.1.0/24
- **VCN CIDR block:** 10.200.0.0/24
- **CPE public IP address:** 10.100.0.100/24
- **Oracle VPN headend (DRG) IP address 1:** 10.150.128.1/32
- **Oracle VPN headend (DRG) IP address 2:** 10.150.127.1/32
- **Tunnel number 1:** tunnel.1
- **Tunnel number 2:** tunnel.2
- **Exit interface:** ethernet1/1

### ABOUT USING IKEV2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

If you want to use IKEv2, there are special variations of some steps presented in the next section. Here is a summary of the special steps:

- For [task 2 \(defining the ISAKMP peers\)](#), when you add the IKE gateway:
  - On the **General** tab, for the **Version**, select **IKEv2 only mode**.
  - On the **Advanced Options** tab, select the IKE crypto profile associated with the IKEv2 tunnel.
- For [task 5 \(configuring the IPSec sessions\)](#), configure the proxy ID.

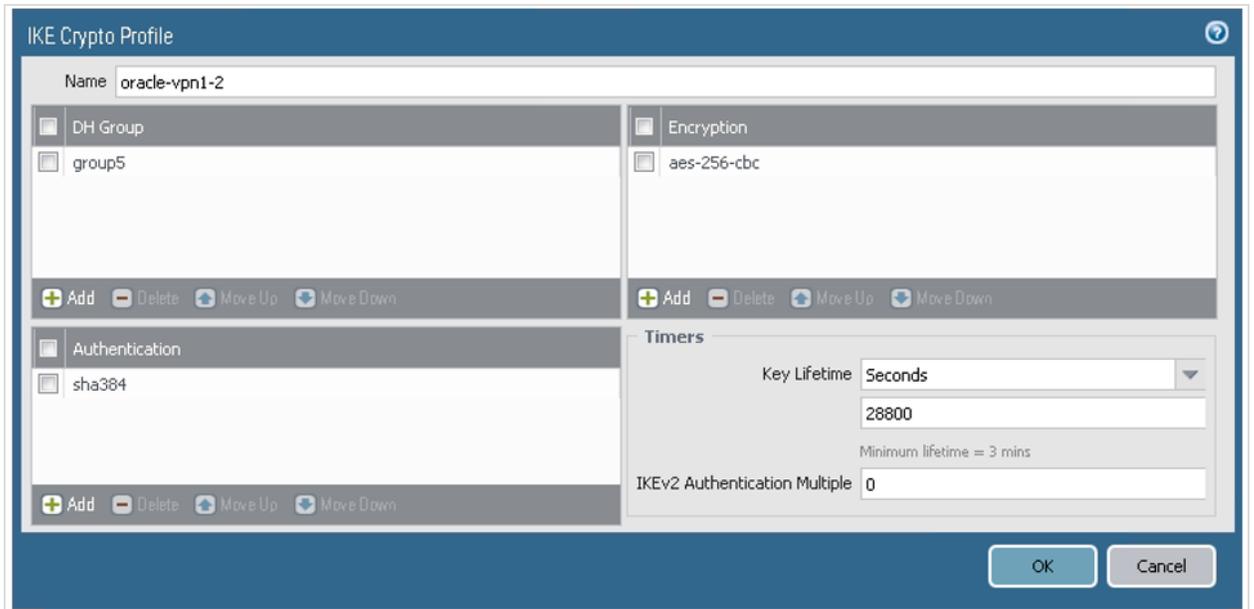
### CONFIGURATION PROCESS

The following process includes BGP configuration for the IPSec connection. If you instead want to use static routing, perform tasks 1-5, and then skip to [Configuring Static Routing](#).

### Task 1: Configure the ISAKMP Phase 1 policy

In this example, the same ISAKMP policy is used for both tunnels.

1. Go to **Network**, to **IKE Crypto**, and then click **Add**.
2. Configure the parameters as shown in the next screenshot. For a list of the values, see [Supported IPSec Parameters](#). If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#).



The next screenshot shows the final result for this task:

Name	Encryption	Authentication	DH Group	Key Lifetime
oracle-vpn1-2	aes-256-cbc	sha384	group5	28800 seconds

## Task 2: Define the ISAKMP peers

1. Go to **Network**, to **IKE Gateways**, and then click **Add**.
2. For peer 1, configure the parameters as shown in the next screenshots.
  - a. On the **General** tab:
    - **Version:** For IKEv1, select **IKEv1 only mode**. If you want to use IKEv2, select **IKEv2 only mode**. Notice that if you're using IKEv2, later in [task 5](#) you also add proxy IDs.

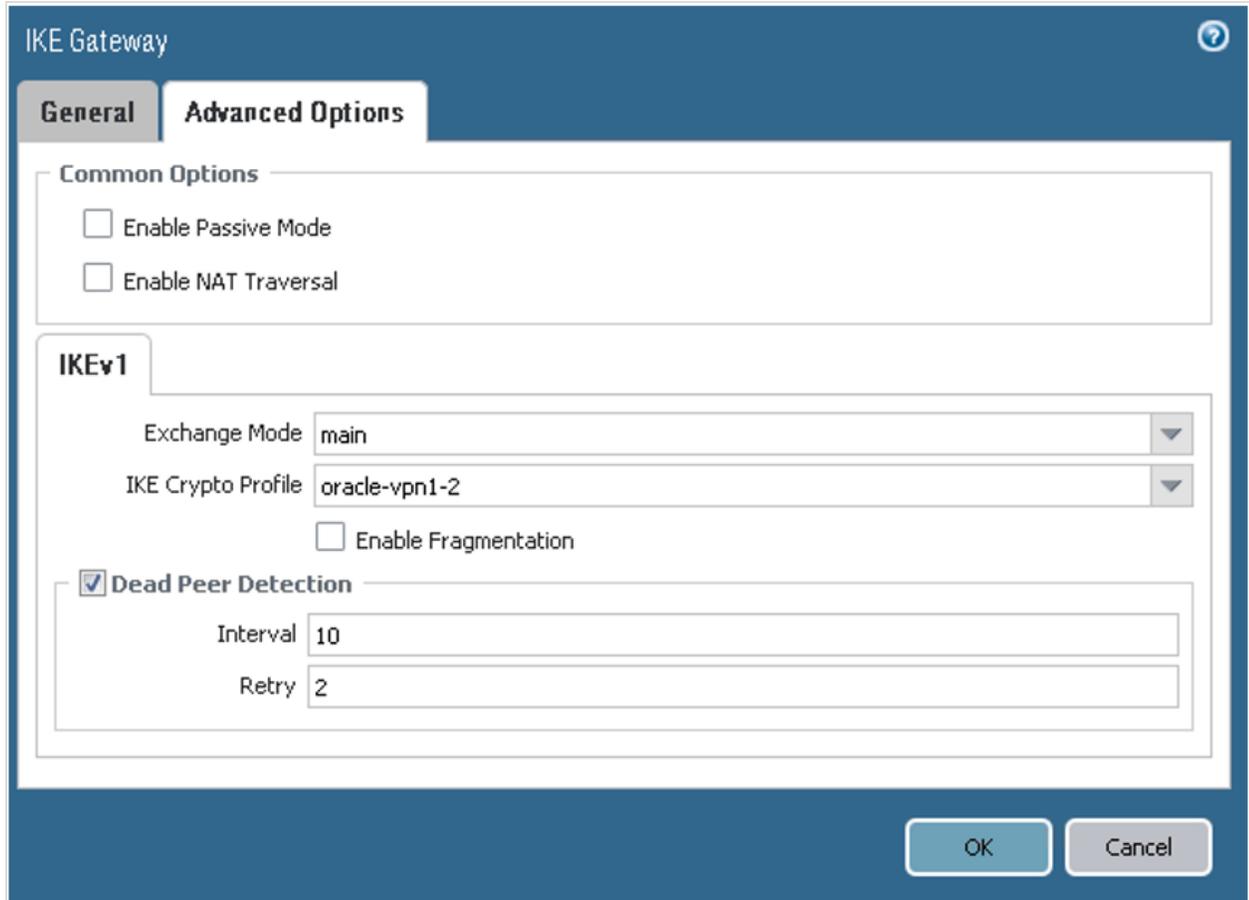
- **Interface:** The interface that owns the public IP address on the CPE. Change **ethernet1/1** to the particular value for your networking topology.
- **Peer IP addresses:** The public IP address that Oracle assigned to the Oracle headend of the tunnel. Change the value to the correct IP address for your first tunnel.
- **Pre-shared Key:** The shared secret that Oracle automatically assigned during IPsec tunnel creation. If you want, you can [specify a different value](#). Make sure to enter the same value here and in the Oracle Console.
- **Local Identification** and **Peer Identification:** The IKE IDs. The **Local Identification** is the CPE's public IP address. The **Remote Identification** is the Oracle VPN headend IP address for the first tunnel.

The screenshot shows the 'IKE Gateway' configuration window with the 'Advanced Options' tab selected. The configuration is as follows:

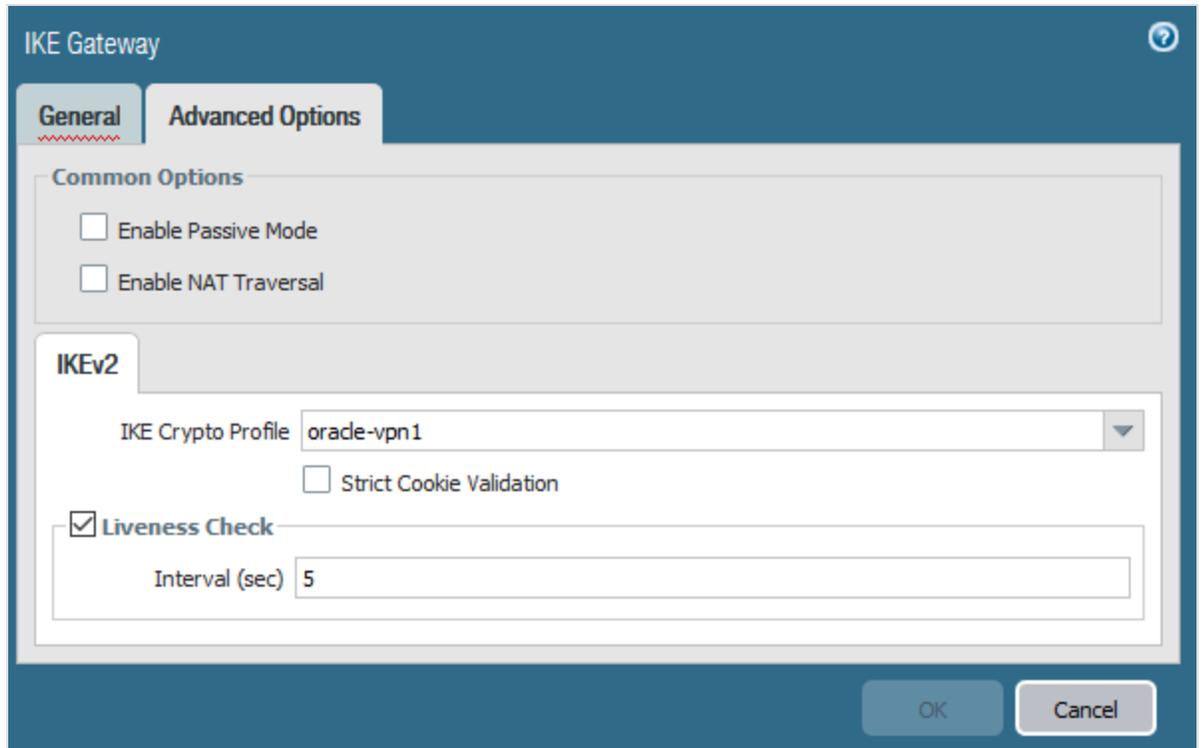
Field	Value
Name	oracle-vpn1
Version	IKEv1 only mode
Address Type	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6
Interface	ethernet1/1
Local IP Address	Internet_address
Peer IP Type	<input checked="" type="radio"/> Static <input type="radio"/> Dynamic
Peer IP Address	10.150.128.1
Authentication	<input checked="" type="radio"/> Pre-Shared Key <input type="radio"/> Certificate
Pre-shared Key	●●●●●●●●
Confirm Pre-shared Key	●●●●●●●●
Local Identification	IP address   10.100.0.100
Peer Identification	IP address   10.150.128.1

Buttons: OK, Cancel

- b. On the **Advanced Options** tab, ensure the values are set for the first peer according to the following screenshot.



If you are using IKEv2 instead, select the IKE crypto profile associated with the IKEv2 tunnel.



3. For peer 2, configure the parameters as shown in the next screenshots.
  - a. On the **General** tab:
    - **Version:** For IKEv1, select **IKEv1 only mode**. If you want to use IKEv2, select **IKEv2 only mode**. For IKEv2, notice that you also need to provide a proxy ID later in [task 5](#).
    - **Interface:** The interface that owns the public IP address on the CPE. Change **ethernet1/1** to the particular value for your networking topology.
    - **Peer IP addresses:** The public IP address that Oracle assigned to the Oracle headend of the tunnel. Change the value to the correct IP address for your second tunnel.

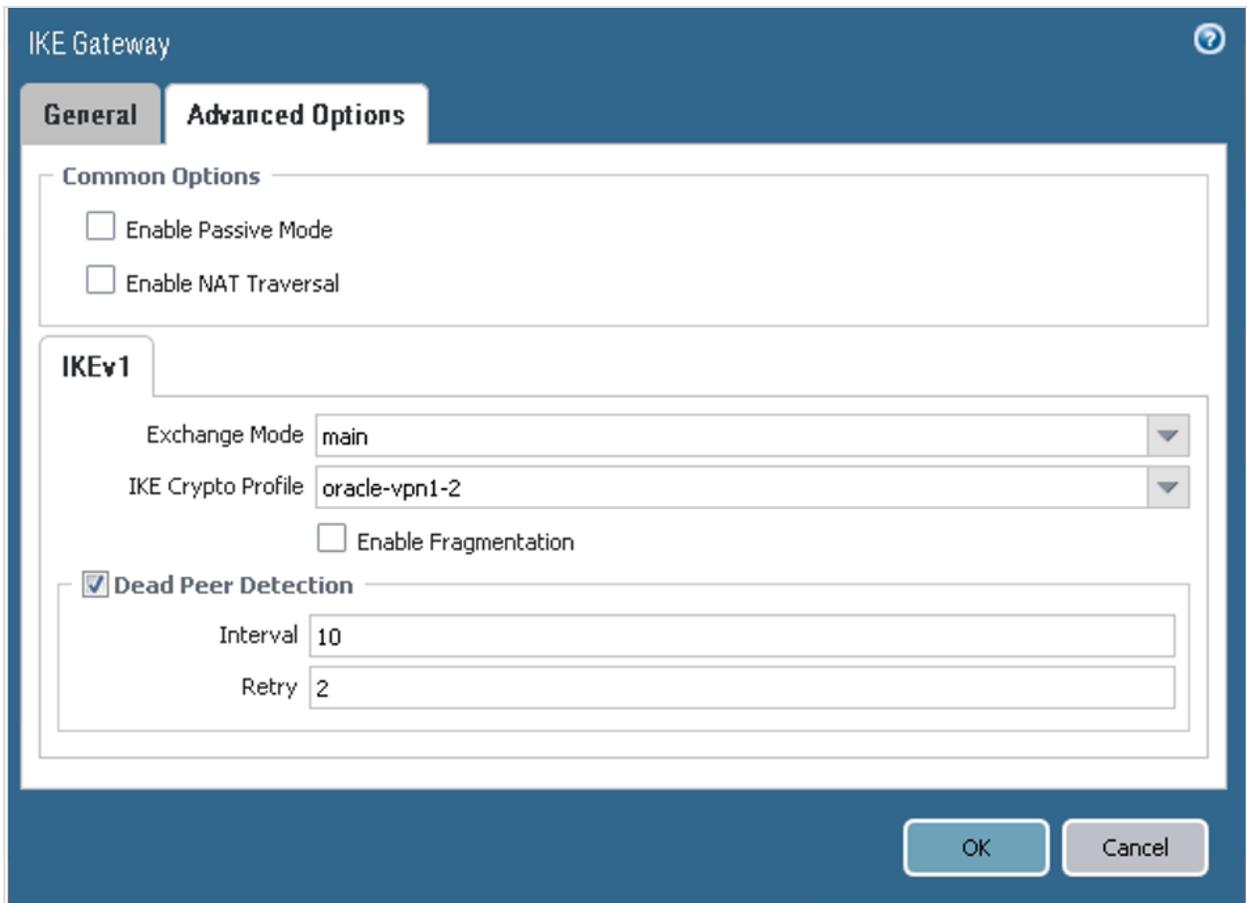
- **Pre-shared Key:** The shared secret that Oracle automatically assigned during IPSec tunnel creation. If you want, you can [specify a different value](#). Make sure to enter the same value here and in the Oracle Console.
- **Local Identification** and **Peer Identification:** The IKE IDs. The **Local Identification** is the CPE's public IP address. The **Remote Identification** is the Oracle VPN headend IP address for the second tunnel.

The screenshot shows the 'IKE Gateway' configuration window with the 'Advanced Options' tab selected. The configuration is as follows:

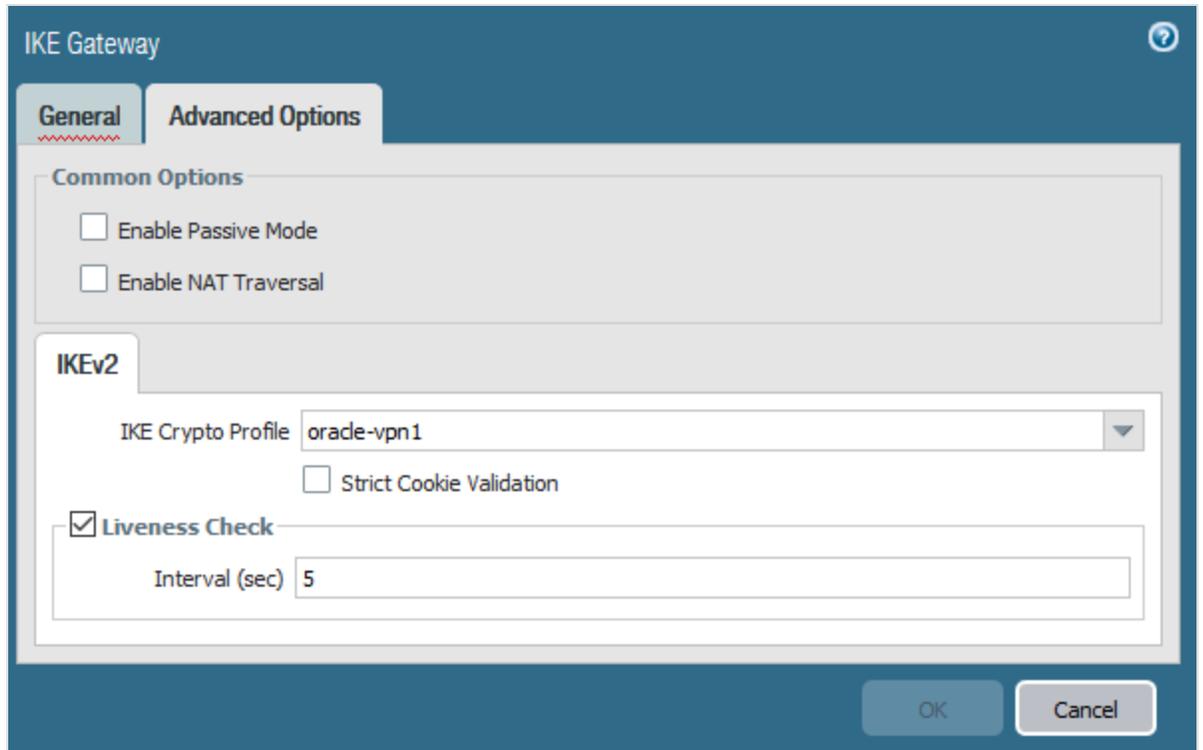
Field	Value
Name	oracle-vpn2
Version	IKEv1 only mode
Address Type	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6
Interface	ethernet1/1
Local IP Address	Internet_address
Peer IP Type	<input checked="" type="radio"/> Static <input type="radio"/> Dynamic
Peer IP Address	10.150.127.1
Authentication	<input checked="" type="radio"/> Pre-Shared Key <input type="radio"/> Certificate
Pre-shared Key	••••••••
Confirm Pre-shared Key	••••••••
Local Identification	IP address   10.100.0.100
Peer Identification	IP address   10.150.127.1

Buttons: OK, Cancel

- b. On the **Advanced Options** tab, ensure the values are set for the second peer according to this screenshot:



If you are using IKEv2 instead, select the IKE crypto profile associated with the IKEv2 tunnel.



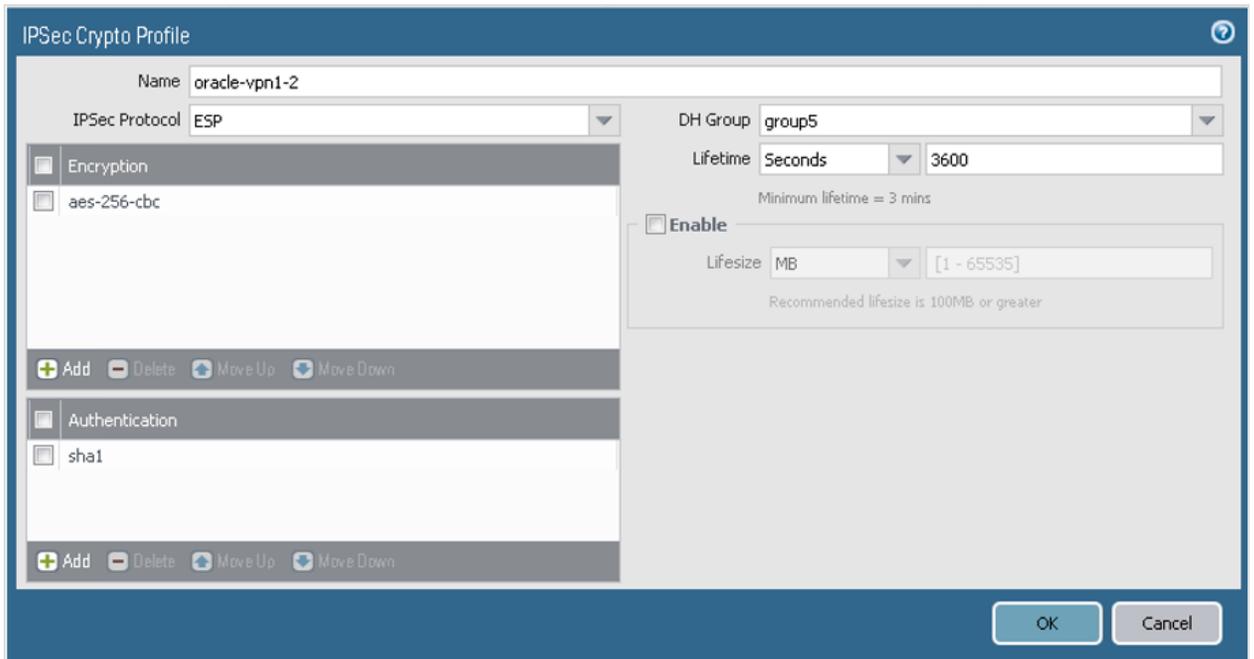
The next screenshot shows the final result for this task:

Name	Peer Address	Local Address		Peer ID		Local ID		Version	Mode	Passive Mode	NAT Traversal	IKE Advanced Options	
		Interface	IP	ID	Type	ID	Type					Crypto Profile	DPD
orade-vpn1	10.150.128.1	ethernet1/1	Internet_address	10.150.128.1	IP address	10.100.0.100	IP address	#ev1	main	<input type="checkbox"/>	<input type="checkbox"/>	orade-vpn1-2	enabled/102
orade-vpn2	10.150.127.1	ethernet1/1	Internet_address	10.150.127.1	IP address	10.100.0.100	IP address	#ev1	main	<input type="checkbox"/>	<input type="checkbox"/>	orade-vpn1-2	enabled/102

### Task 3: Define the IPSec Phase 2 policy

In this example, the same IPSec crypto profile is used for both tunnels.

1. Go to **Network**, to **IPSec Crypto**, and then click **Add**.
2. Configure the parameters as shown in the next screenshot.



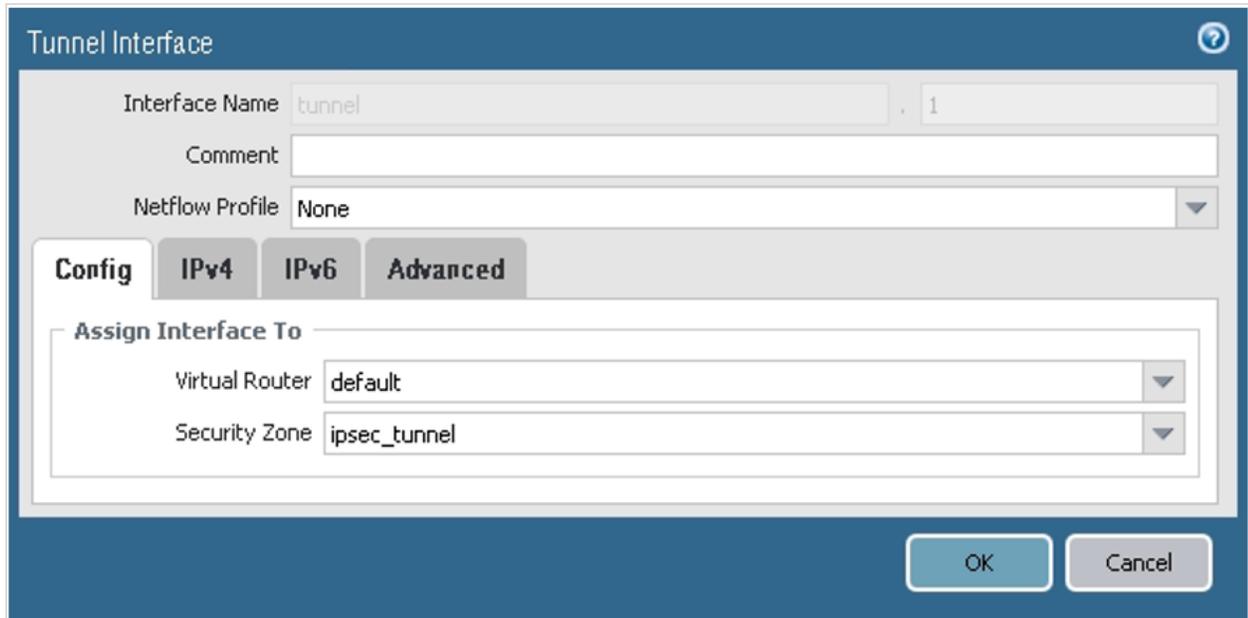
The next screenshot shows the final result for this task:

Name	ESP/AH	Encryption	Authentication	DH Group	Lifetime
oracle-vpn1-2	ESP	aes-256-cbc	sha1	group5	3600 seconds

#### Task 4: Configure the virtual tunnel interfaces

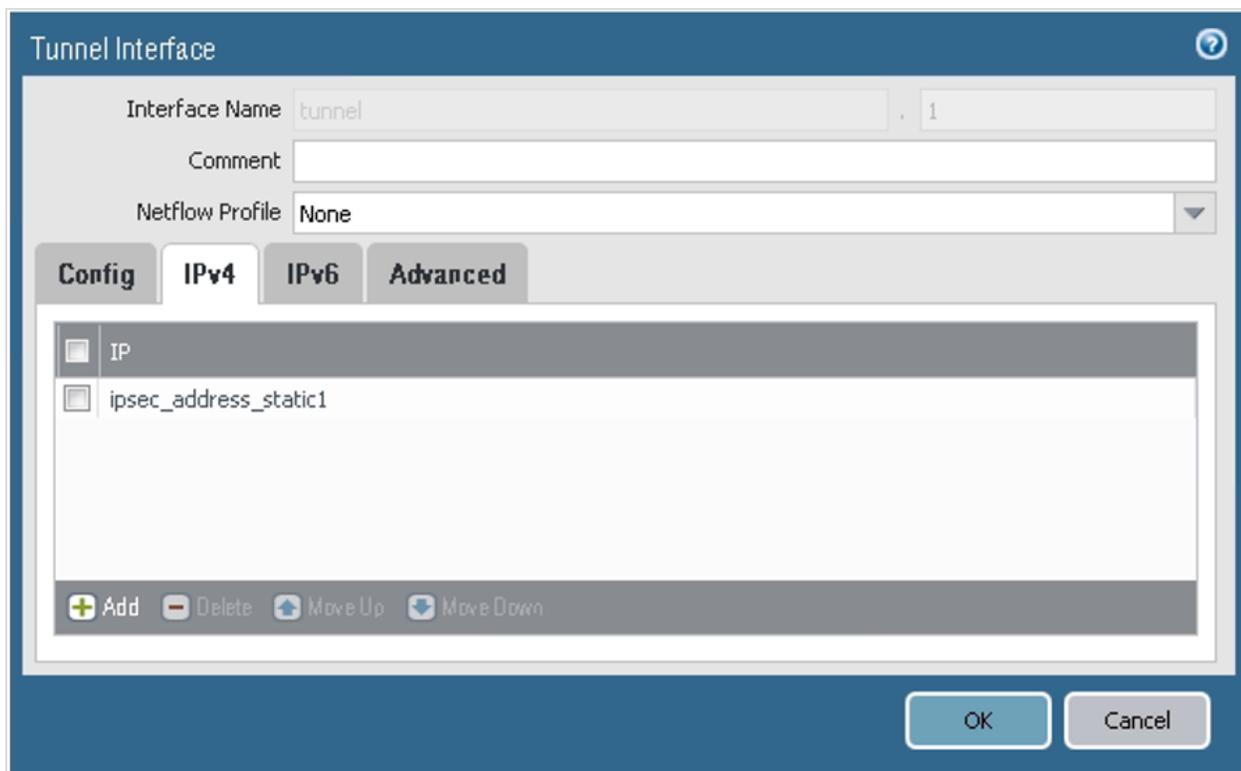
1. Go to **Network**, to **Interfaces**, to **Tunnel**, and then click **Add**.
2. For peer 1, configure the parameters as shown in the next screenshots.

- a. On the **Config** tab, assign the interface according to your virtual router and security zone configuration. In this example, the default virtual router and ipsec\_tunnel security zone are used.

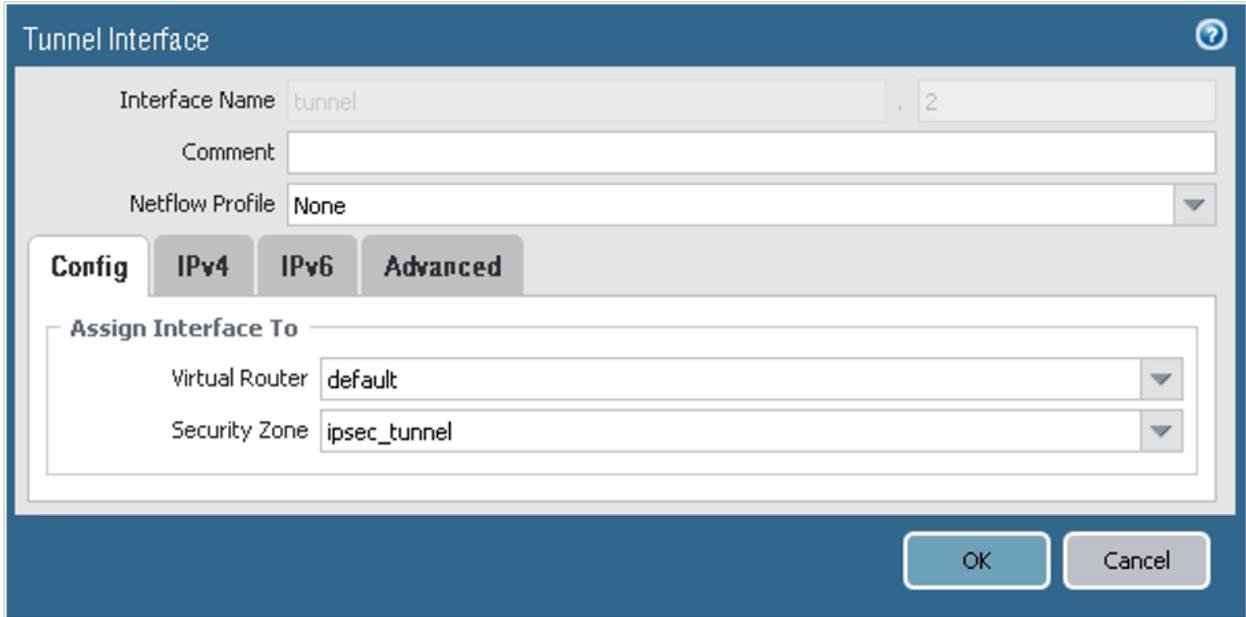


The screenshot shows the 'Tunnel Interface' configuration window. The 'Interface Name' is 'tunnel' and the ID is '1'. The 'Netflow Profile' is set to 'None'. The 'Config' tab is selected, showing the 'Assign Interface To' section with 'Virtual Router' set to 'default' and 'Security Zone' set to 'ipsec\_tunnel'. There are 'OK' and 'Cancel' buttons at the bottom right.

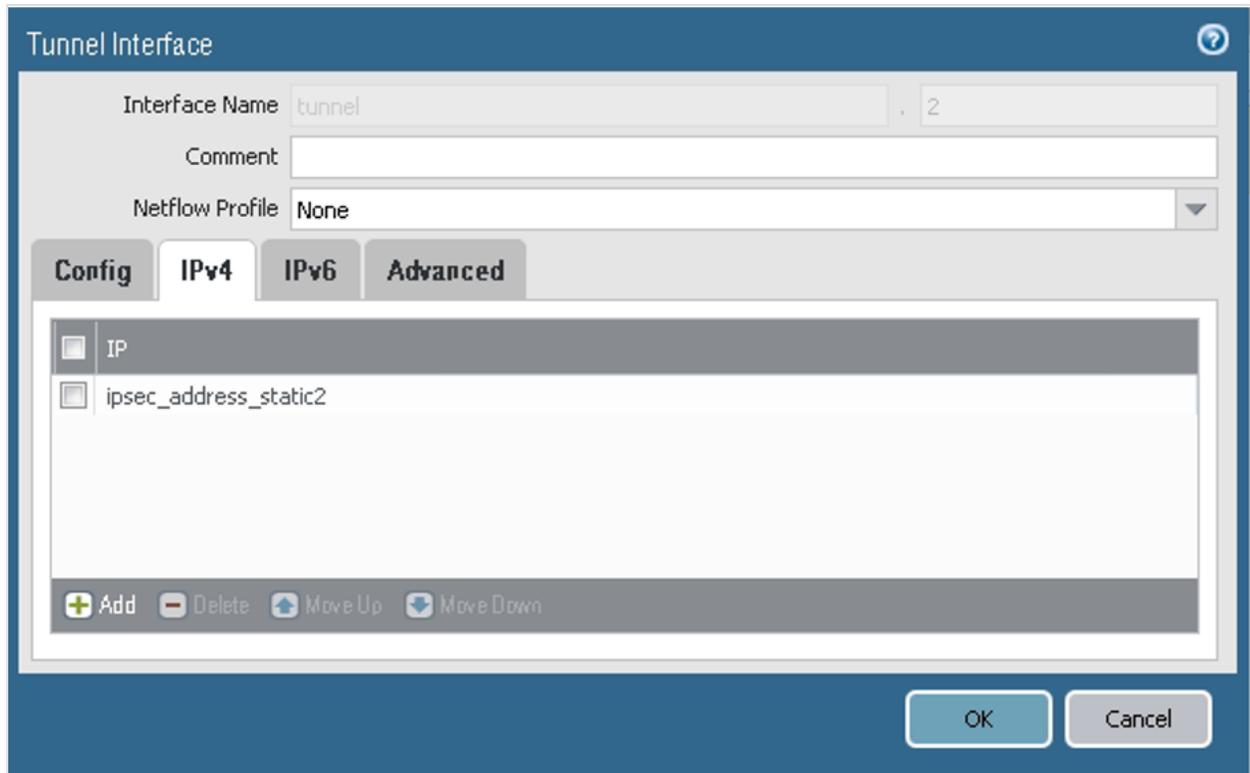
- b. On the **IPv4** tab, ensure the values are set for the first peer according to the following screenshot. In this example, the IP address for the tunnel interface is ipsec\_address\_static1 = 198.51.100.1/30. Configure your tunnel IP address according to your networking IP addressing plan.



3. For peer 2, configure the parameters as shown in the next screenshots.
  - a. On the **Config** tab, assign the interface according to your virtual router and security zone configuration. In this example, the default virtual router and ipsec\_tunnel security zone are used.



- b. On the **IPv4** tab, ensure the values are set for the second peer according to the following screenshot. In this example, the IP address for the tunnel interface is ipsec\_address\_static2 = 198.51.100.5/30. Configure your tunnel IP address according to your networking IP addressing plan.



The next screenshot shows the final result for this task:

Interface	Management Profile	IP Address	Virtual Router	Security Zone	Features	Comment
tunnel		none	none	none		
tunnel.1		ipsec_address_static1	default	ipsec_tunnel		
tunnel.2		ipsec_address_static2	default	ipsec_tunnel		

### Task 5: Configure the IPsec sessions

1. Go to **Network**, to **IPsec Tunnels**, and then click **Add**.
2. For peer 1, configure the parameters on the **General** tab as shown in the next screenshot.

Notice that if you're using IKEv1, you do not need to add specific proxy IDs to the **Proxy IDs** tab. They are not needed for an IKEv1 route-based VPN configuration.

However, for IKEv2, do add proxy IDs to the **Proxy IDs** tab for better interoperability. Ensure that you also configured the IKE gateway to use IKEv2 earlier in [task 2](#).



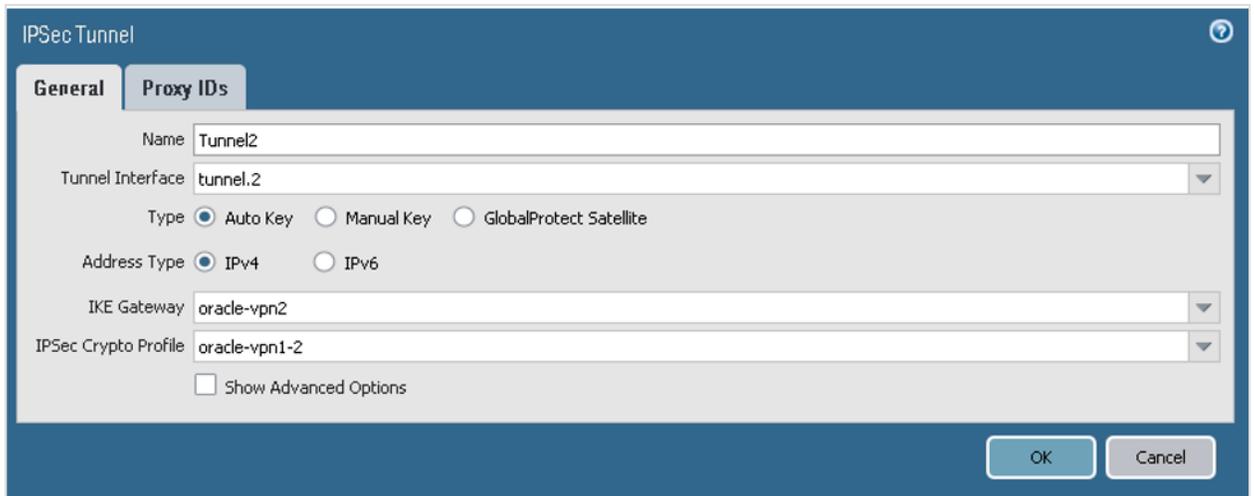
The screenshot shows the 'IPsec Tunnel' configuration dialog box with the 'General' tab selected. The fields are as follows:

- Name: Tunnel1
- Tunnel Interface: tunnel.1
- Type:  Auto Key,  Manual Key,  GlobalProtect Satellite
- Address Type:  IPv4,  IPv6
- IKE Gateway: oracle-vpn1
- IPsec Crypto Profile: oracle-vpn1-2
- Show Advanced Options

Buttons for 'OK' and 'Cancel' are visible at the bottom right.

3. For peer 2, configure the parameters on the **General** tab as shown in the next screenshot.

If you are using IKEv2, also add proxy IDs on the **Proxy IDs** tab.



The screenshot shows the 'IPsec Tunnel' configuration window. It has two tabs: 'General' and 'Proxy IDs'. The 'General' tab is active. The configuration fields are as follows:

- Name: Tunnel2
- Tunnel Interface: tunnel.2
- Type:  Auto Key,  Manual Key,  GlobalProtect Satellite
- Address Type:  IPv4,  IPv6
- IKE Gateway: oracle-vpn2
- IPsec Crypto Profile: oracle-vpn1-2
- Show Advanced Options

At the bottom right, there are 'OK' and 'Cancel' buttons.

### Task 6: Configure BGP over IPsec



#### Note

If you want to use static routing instead of BGP, skip task 6 and go to [Configuring Static Routing](#).

BGP over IPsec requires IP addresses on the tunnel interfaces on both ends. The screenshots in this example use these subnets for the tunnel interfaces:

- 198.51.100.0/30
  - CPE: 198.51.100.1/30
  - DRG: 198.51.100.2/30

- 198.51.100.4/30
  - CPE: 198.51.100.5/30
  - DRG: 198.51.100.6/30

Make sure to replace the example values with the BGP IP addresses you specified in the Oracle Console for the inside tunnel interfaces.

This task consists of three sub-tasks, each with multiple steps.

### Subtask 6-a: Configure the BGP parameters

1. Go to **Network**, to **Virtual Routers**, to **default**, and then to **BGP**. This example uses the default virtual router. Also, the example uses 10.200.1.10 for the router ID and 64511 for the ASN. Make sure to use the correct virtual router based on your networking configuration, and use the correct router ID and ASN for your environment.



The screenshot shows a configuration panel for BGP. It includes a checked 'Enable' checkbox, a 'Router ID' text box with the value '10.200.1.10', an 'AS Number' text box with the value '64511', and a 'BFD' dropdown menu currently set to 'None'.

2. On the **General** tab, configure the parameters as shown in the next screenshot.

## CHAPTER 23 Networking

**General** | **Advanced** | Peer Group | Import | Export | Conditional Adv | **Aggregate** | Redist Rules

**Options**

- Reject Default Route
- Install Route
- Aggregate MED
- Default Local Preference:
- AS Format:  2 Byte  4 Byte

**Path Selection**

- Always Compare MED
- Deterministic MED comparison

Auth Profiles

+ Add - Delete

3. On the **Advanced** tab, configure the parameters as shown in the next screenshot.

**General** | **Advanced** | Peer Group | Import | Export | Conditional Adv | **Aggregate** | Redist Rules

- ECMP Multiple AS Support
- Enforce First AS for EBGP
- Graceful Restart

Stale Route Time (sec):  Local Restart Time (sec):  Max Peer Restart Time (sec):

Reflector Cluster ID:  Confederation Member AS:

**Dampening Profiles**

<input type="checkbox"/>	Profile Name	Enable	Cutoff	Reuse	Max Hold Time (sec)	Decay Half Life Reachable (sec)	Decay Half Life Unreachable (sec)
<input checked="" type="checkbox"/>	default	<input type="checkbox"/>	1.25	0.5	900	300	900

+ Add - Delete

4. On the **Peer Group** tab:
  - a. Add the first Peer Group, and under the **Peer Group Name**, add the first session. Add the BGP session with the DRG.

Virtual Router - BGP - Peer Group/Peer

**Peer Group**

Name

Enable

Aggregated Confed AS Path

Soft Reset With Stored Info

Type

Import Next Hop  Original  Use Peer

Export Next Hop  Resolve  Use Self

Remove Private AS

- b. For the first tunnel, on the **Addressing** tab, configure the parameters as shown in the next screenshot. Oracle's BGP ASN in commercial regions is 31898. If you're configuring VPN Connect for the Government Cloud, see [Oracle's BGP ASN](#).

The screenshot shows a configuration window titled "Virtual Router - BGP - Peer Group - Peer". The window has a blue header bar with a help icon on the right. Below the header, there are several input fields and checkboxes:

- Name:** Session1
- Enable:**  Enable
- Peer AS:** 31898

The window has three tabs: "Addressing", "Connection Options", and "Advanced". The "Connection Options" tab is currently selected and highlighted in grey. It contains the following settings:

- Enable MP-BGP Extensions:**
- Address Family Type:**  IPv4  IPv6
- Subsequent Address Family:**  Unicast  Multicast

There are two main sections for address configuration:

- Local Address:** A section with a title bar. It contains two dropdown menus: "Interface" set to "tunnel.1" and "IP" set to "ipsec\_address\_static1".
- Peer Address:** A section with a title bar. It contains one text input field for "IP" set to "198.51.100.2".

At the bottom of the window, there are two buttons: "OK" and "Cancel".

- c. On the **Connection Options** tab, configure the parameters as shown in the next screenshot.

Virtual Router - BGP - Peer Group - Peer

Name

Enable

Peer AS

**Addressing** **Connection Options** **Advanced**

Auth Profile

Keep Alive Interval (sec)

Multi Hop

Open Delay Time (sec)

Hold Time (sec)

Idle Hold Time (sec)

**Incoming Connections**

Remote Port

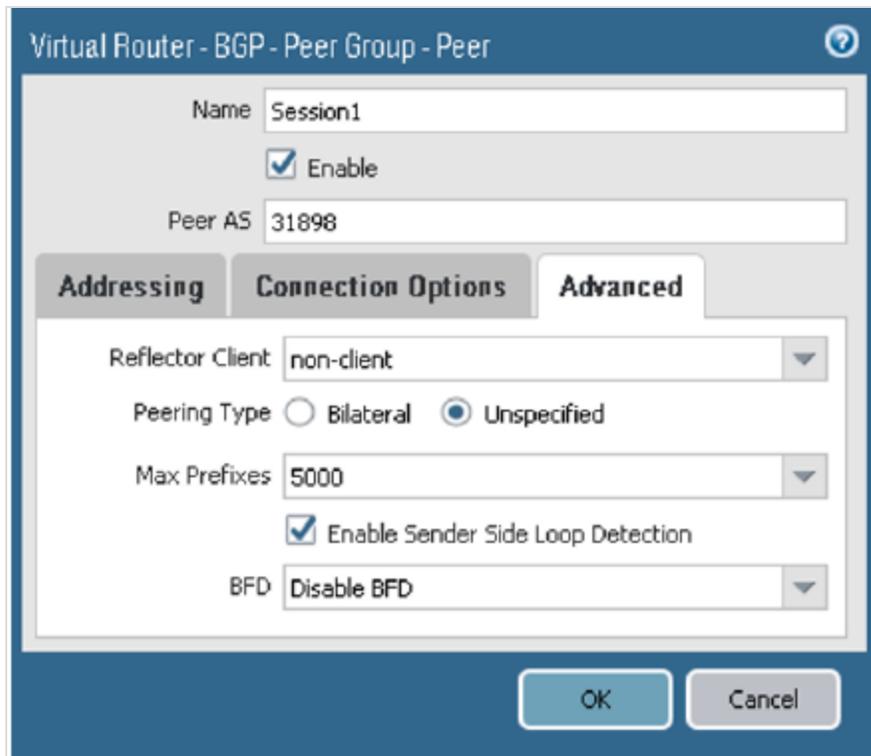
Allow

**Outgoing Connections**

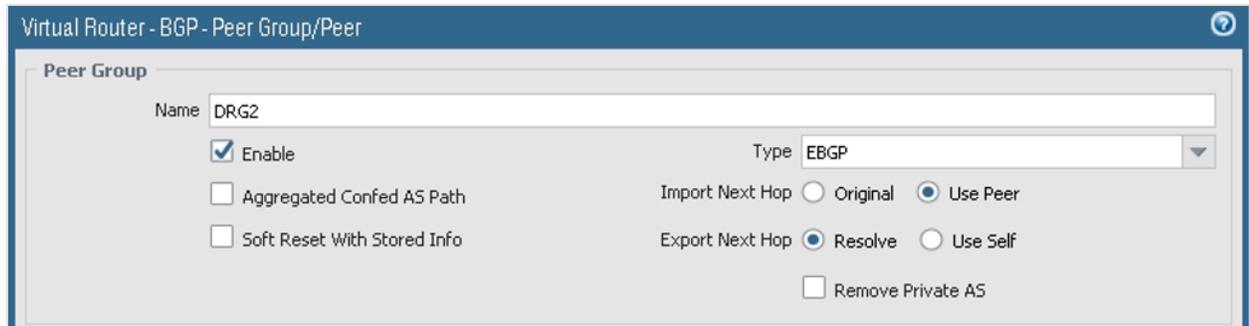
Local Port

Allow

- d. On the **Advanced** tab, configure the parameters as shown in the next screenshot.



- e. On the **Peer Group** tab, add the second Peer Group, and under the **Peer Group Name**, add the second session. Add the BGP session with the DRG.



- f. For the second tunnel, on the **Addressing** tab, configure the parameters as shown in the next screenshot.

Virtual Router - BGP - Peer Group - Peer

Name

Enable

Peer AS

**Addressing** **Connection Options** **Advanced**

Enable MP-BGP Extensions

Address Family Type  IPv4  IPv6

Subsequent Address Family  Unicast  Multicast

**Local Address**

Interface

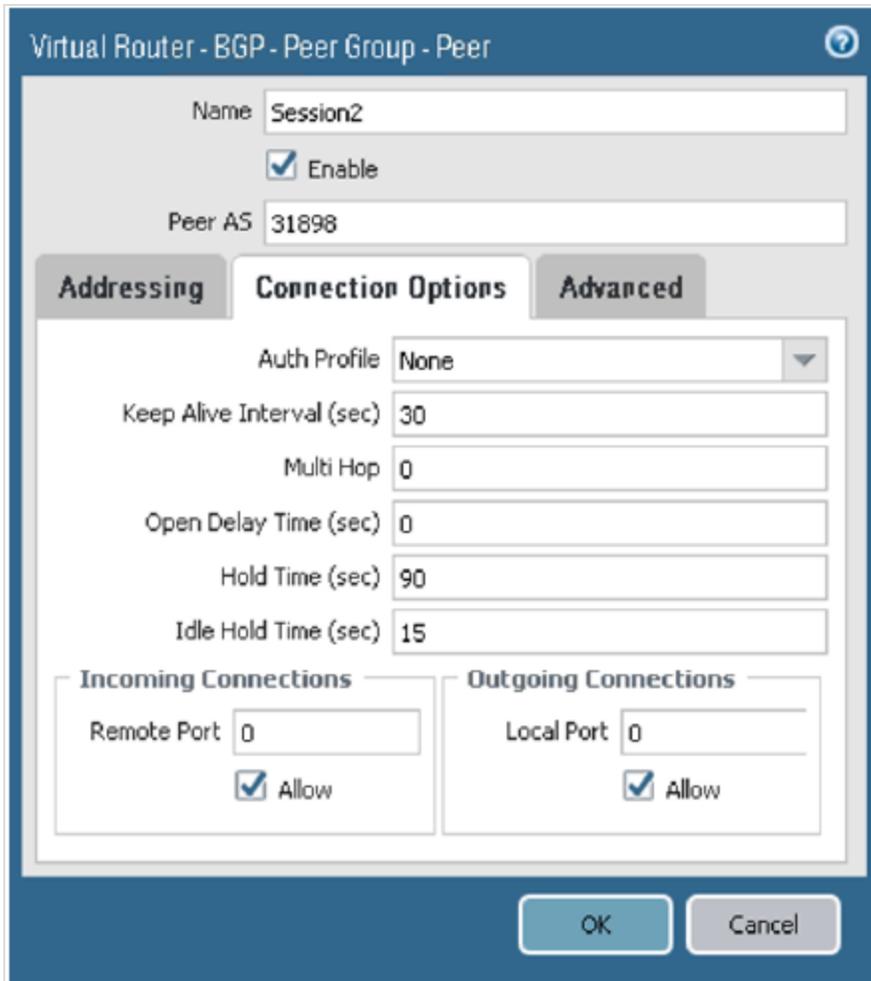
IP

**Peer Address**

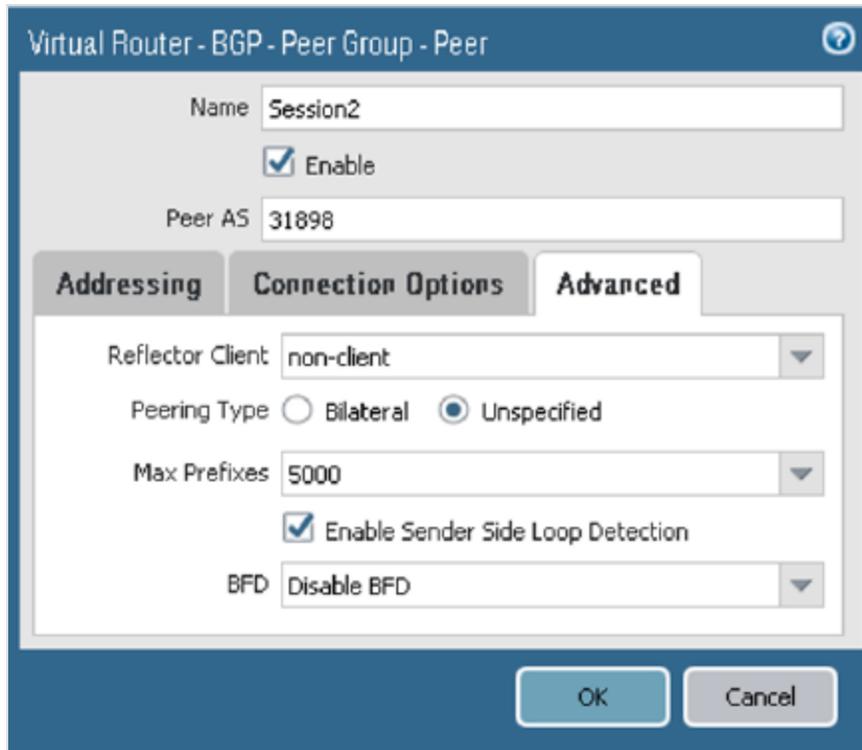
IP

OK Cancel

- g. On the **Connection Options** tab, configure the parameters as shown in the next screenshot.



- h. On the **Advanced** tab, configure the parameters as shown in the next screenshot.



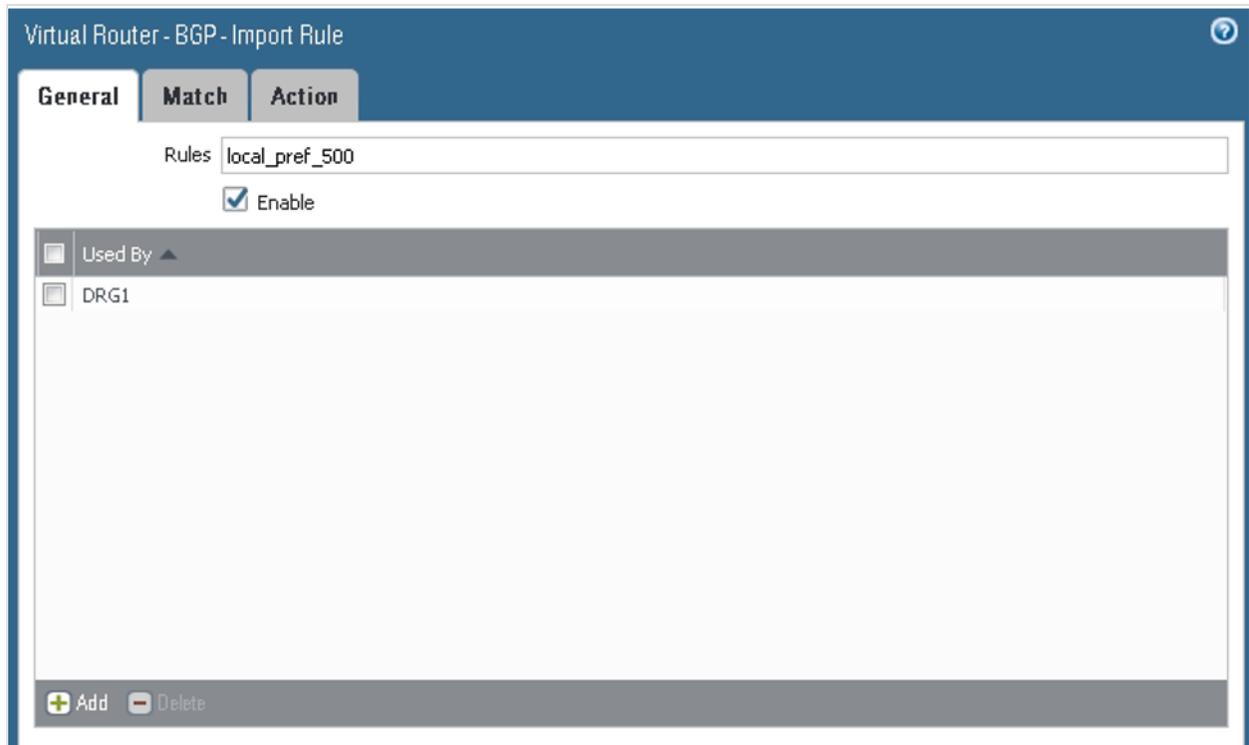
The next screenshot shows the final Peer Group configuration:

General   Advanced <b>Peer Group</b> Import   Export   Conditional Adv   Aggregate   Redist Rules						
Name	Enable	Type	Peers			
			Name	Peer Address	Local Address	
DRG1	<input checked="" type="checkbox"/>	ebgp	Session1	198.51.100.2	ipsec_address_static1	
DRG2	<input checked="" type="checkbox"/>	ebgp	Session2	198.51.100.6	ipsec_address_static2	

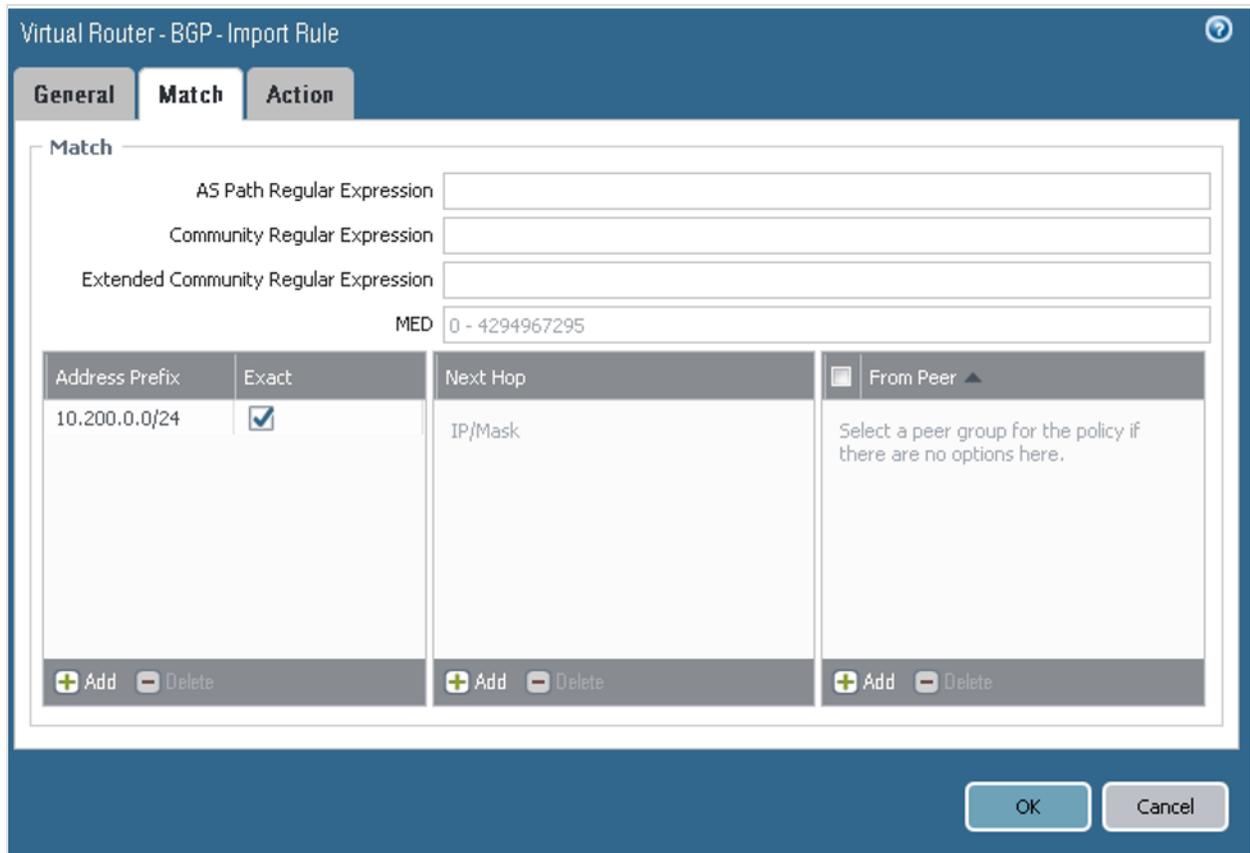
- On the **Import** tab, configure the parameters as shown in the next screenshots. Here you configure tunnel.1 as the primary and tunnel.2 as the backup for the VCN route received from the DRG by way of BGP (10.200.0.0/24). From the BGP perspective, both

tunnels will be in the Established state.

- a. For the first rule, on the **General** tab, configure the parameters as shown in the next screenshot.



- b. On the **Match** tab, configure the parameters as shown in the next screenshot.



- c. On the **Action** tab, configure the parameters as shown in the next screenshot.

The screenshot shows the 'Virtual Router - BGP - Import Rule' configuration window. The 'General' tab is active, displaying the following settings:

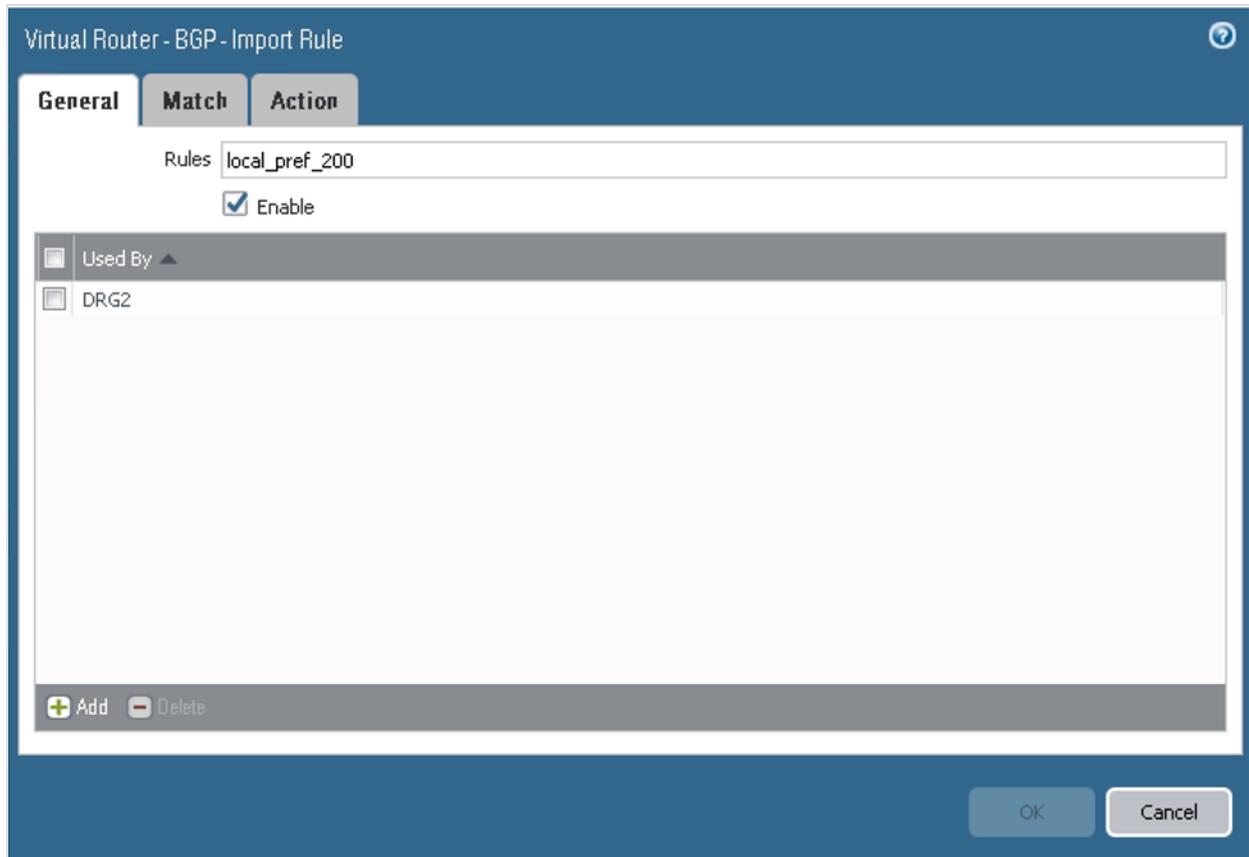
- Action: Allow
- Dampening: None
- Local Preference: 500
- MED: 0 - 4294967295
- Weight: 0 - 65535
- Next Hop: (empty)
- Origin: incomplete
- AS Path Limit: [1 - 255]

Below these settings are three sections:

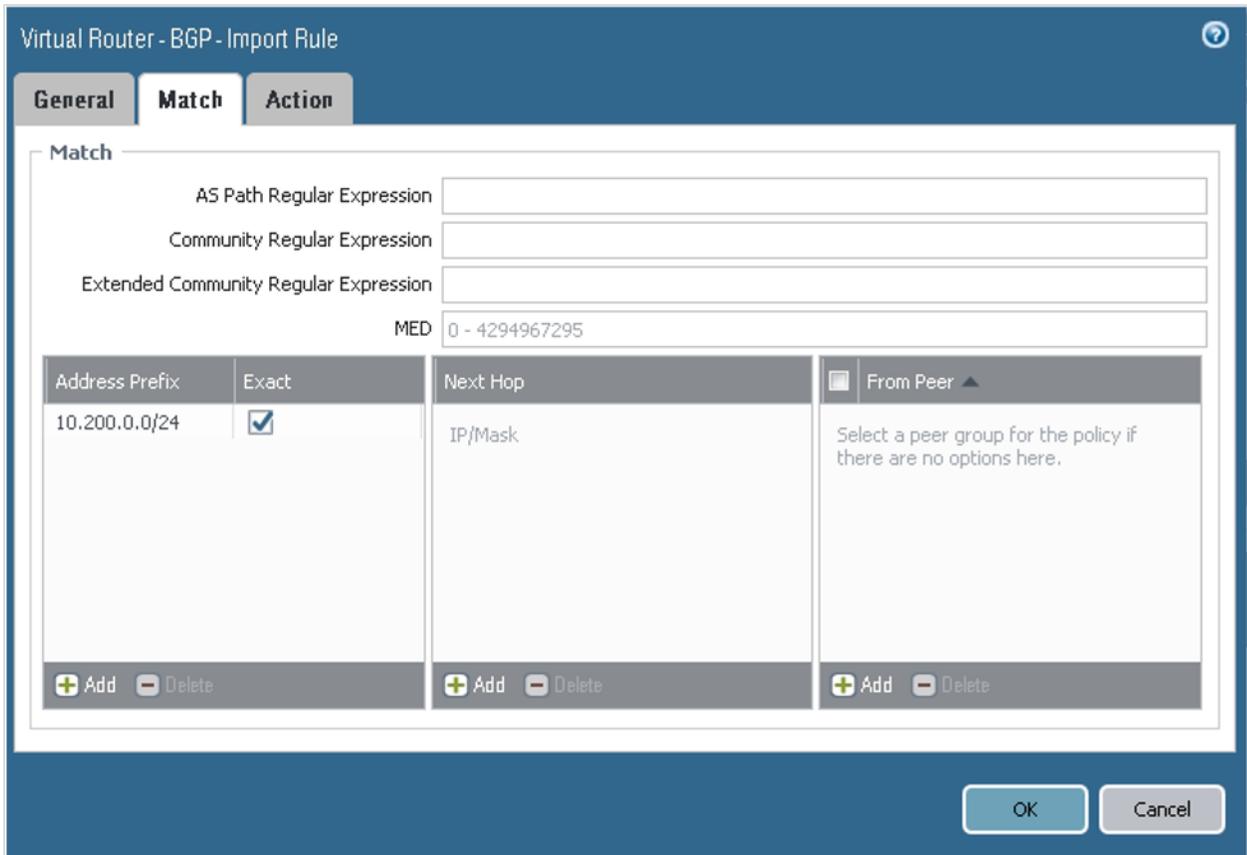
- AS Path:** Type  None  Remove
- Community:** Type None
- Extended Community:** Type None

At the bottom right, there are 'OK' and 'Cancel' buttons.

- d. For the second rule, on the **General** tab, configure the parameters as shown in the next screenshot.



- e. On the **Match** tab, configure the parameters as shown in the next screenshot.



- f. On the **Action** tab, configure the parameters as shown in the next screenshot.

Virtual Router - BGP - Import Rule

**General** **Match** **Action**

Action: Allow

Dampening: None

Local Preference: 200

MED: 0 - 4294967295

Weight: 0 - 65535

Next Hop:

Origin: incomplete

AS Path Limit: [1 - 255]

**AS Path**  
Type:  None  Remove

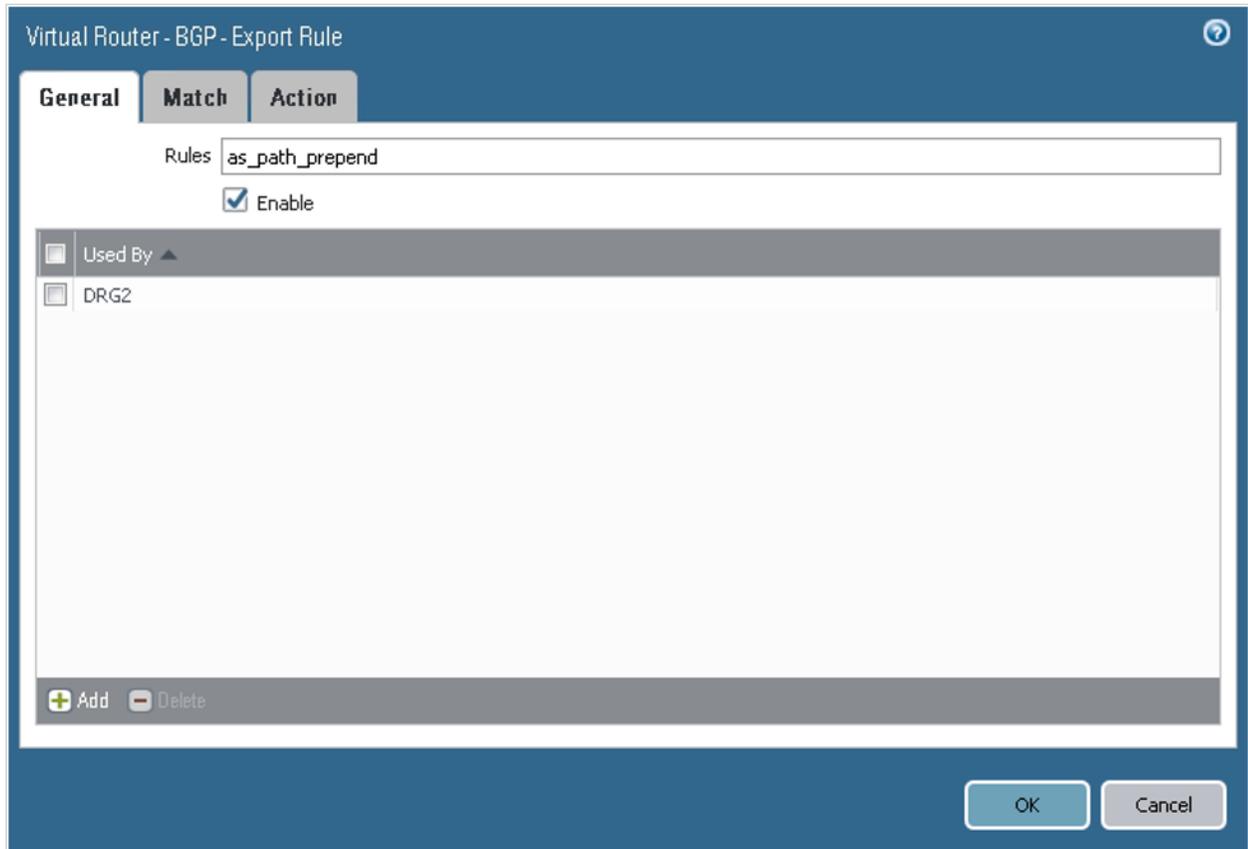
**Community**  
Type: None

**Extended Community**  
Type: None

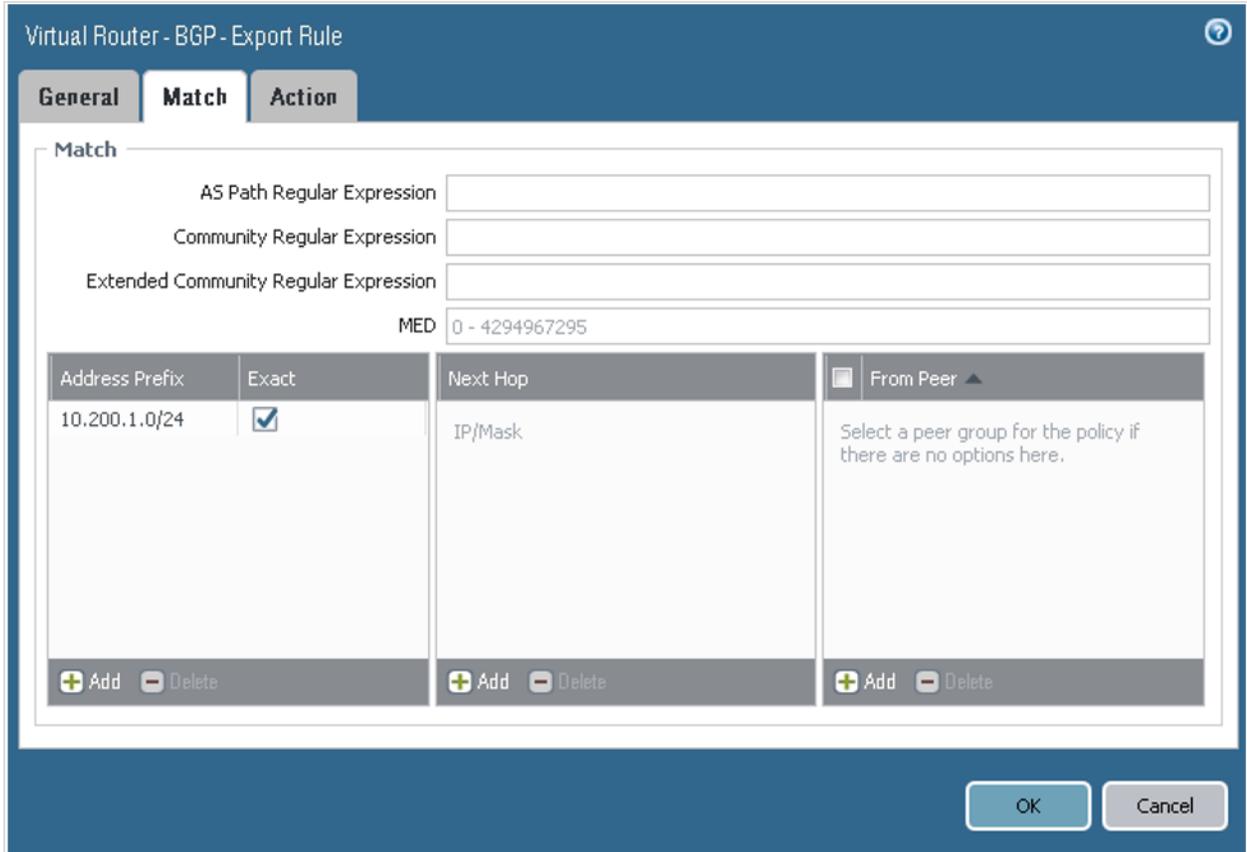
OK Cancel

- On the **Export** tab, configure the parameters as shown in the next screenshots. Here you configure a policy to force the DRG to prefer tunnel.1 for the returning path to the on-premises network CIDR (10.200.1.0/24).

- a. On the **General** tab, configure the parameters as shown in the next screenshot.



- b. On the **Match** tab, configure the parameters as shown in the next screenshot.



- c. On the **Action** tab, configure the parameters as shown in the next screenshot.

Virtual Router - BGP - Export Rule

**General** **Match** **Action**

Action: Allow

Local Preference: 0 - 4294967295

MED: 0 - 4294967295

Next Hop:

Origin: incomplete

AS Path Limit: [1 - 255]

**AS Path**  
 Type: Prepend  
 2

**Community**  
 Type: None

**Extended Community**  
 Type: None

OK Cancel

The next screenshot shows the final Export configuration:

General						
Advanced						
Peer Group						
Import						
Export						
Conditional Adv						
Aggregate						
Redist Rules						
			Peers			
Name	Enable	Type	Name	Peer Address	Local Address	
<input type="checkbox"/> DRG1	<input checked="" type="checkbox"/>	ebgp	Session1	198.51.100.2	ipsec_address_static1	
<input type="checkbox"/> DRG2	<input checked="" type="checkbox"/>	ebgp	Session2	198.51.100.6	ipsec_address_static2	

Notice that no configuration is required for the **Conditional Adv** or **Aggregate** tabs.

- On the **Redist Rules** tab, configure the parameters as shown in the next screenshot. Here you announce the on-premises network CIDR in BGP.

Virtual Router - BGP - Redistribute Rules - Rule

Address Family Type  IPv4  IPv6

Name    
Enter a IPv4 subnet or create a IPv4 redistribution profile first

Enable

Metric

Set Origin

Set MED

Set Local Preference

Set AS Path Limit

Set Community ▲

Select or enter 32-bit value in decimal or hex or in AS:VAL format - where AS and VAL are each in 0 - 65535 range. (Max 10 values)

Set Extended Community ▲

64-bit value in hex, or in TYPE:AS:VAL, TYPE:IP:VAL format. TYPE is 16-bit, the other two are 16-bit and 32-bit each. (Max 5 values)

Subtask 6-b: Wait for the BGP sessions to establish and then check the BGP status

1. Go to **Network**, to **IPSec Tunnels**, to the **Virtual Router** column, and then click **Show Routes**.

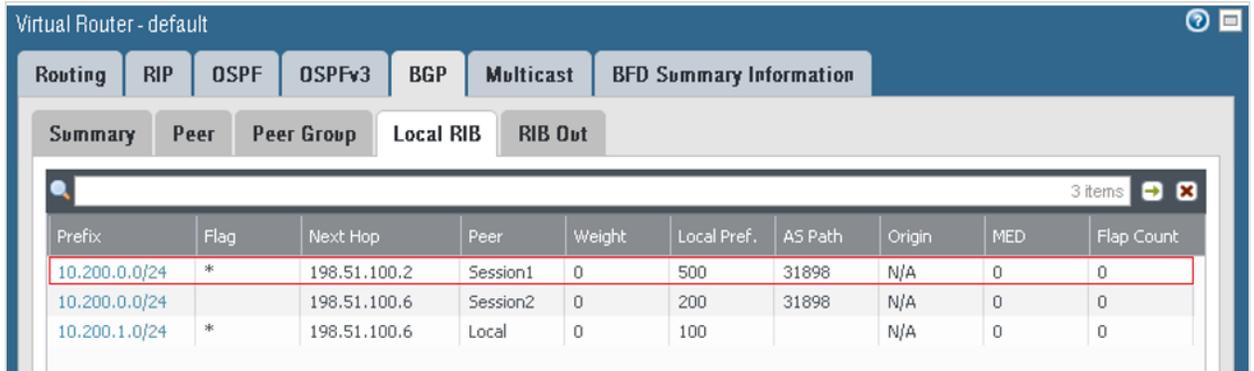
Name	Status	Type	IKE Gateway/Satellite			Tunnel Interface				
			Interface	Local IP	Peer IP	Interface	Virtual Router	Virtual System	Security Zone	Status
Tunnel1		Auto Key	ethernet1/1	Internet_address	10.150.128.1	tunnel.1	default <a href="#">(Show Routes)</a>	vsys1	ipsec_tunnel	
Tunnel2		Auto Key	ethernet1/1	Internet_address	10.150.127.1	tunnel.2	default <a href="#">(Show Routes)</a>	vsys1	ipsec_tunnel	

2. Go to **BGP**, and then to the **Peer** tab to verify that the BGP session is established. Any other value means that the BGP session has not been established successfully and route exchange will not occur.

Routing								
Routing	RIP	OSPF	OSPFv3	BGP	Multicast	BFD Summary Information		
Summary								
Peer								
Peer Group								
Local RIB								
RIB Out								
2 items								
Name	Group	Local IP	Peer IP	Peer AS	Password Set	Status	Status Duration (secs.)	
Session1	DRG1	198.51.100.1:...	198.51.100.2:57121	31898	no	Established	146313	
Session2	DRG2	198.51.100.5:...	198.51.100.6:43509	31898	no	Established	146313	

3. On the **Local RIB** tab: The prefixes are received from the DRG, with tunnel.1 being preferred.

## CHAPTER 23 Networking



Virtual Router - default

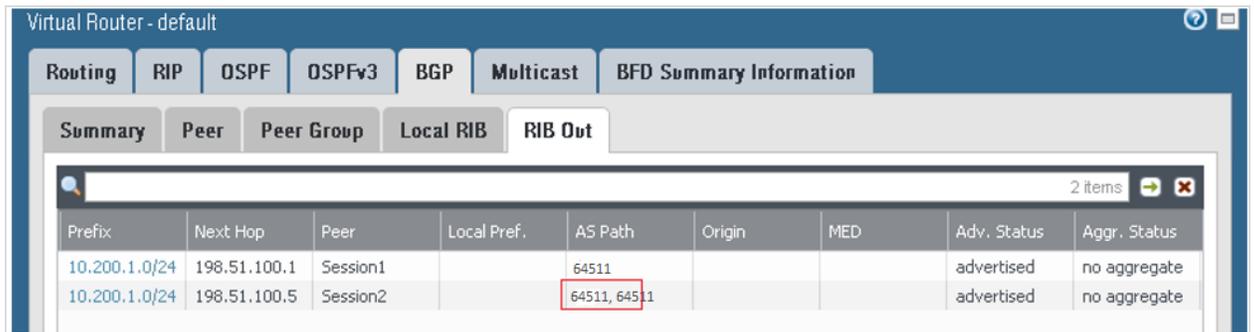
Routing RIP OSPF OSPFv3 BGP Multicast BFD Summary Information

Summary Peer Peer Group Local RIB RIB Out

3 items

Prefix	Flag	Next Hop	Peer	Weight	Local Pref.	AS Path	Origin	MED	Flap Count
10.200.0.0/24	*	198.51.100.2	Session1	0	500	31898	N/A	0	0
10.200.0.0/24		198.51.100.6	Session2	0	200	31898	N/A	0	0
10.200.1.0/24	*	198.51.100.6	Local	0	100		N/A	0	0

4. On the **RIB Out** tab: The on-premises network CIDR is sent by way of BGP to DRG1 with as\_path of 64511, and for DRG2, with an as\_path of 64511, 64511. In this way, based on the BGP Best Path Algorithm, the route preferred by the DRG to reach the on-premises network CIDR uses the connection over tunnel.1.



Virtual Router - default

Routing RIP OSPF OSPFv3 BGP Multicast BFD Summary Information

Summary Peer Peer Group Local RIB RIB Out

2 items

Prefix	Next Hop	Peer	Local Pref.	AS Path	Origin	MED	Adv. Status	Aggr. Status
10.200.1.0/24	198.51.100.1	Session1		64511			advertised	no aggregate
10.200.1.0/24	198.51.100.5	Session2		64511, 64511			advertised	no aggregate

Subtask 6-c: Confirm the BGP routes have been inserted in the routing table

Go to **Routing**, and then to the **Route Table** tab to view the routes.

Virtual Router - default

Routing | RIP | OSPF | OSPFv3 | BGP | Multicast | BFD Summary Information

Route Table | Forwarding Table | Static Route Monitoring

Route Table  Unicast  Multicast Display Address Family IPv4 and IPv6

11 items

Destination	Next Hop	Metric	Weight	Flags	Age	Interface
0.0.0.0/0	10.100.0.99	10		A S		ethernet1/1
10.100.0.0/24	10.100.0.100	0		A C		ethernet1/1
10.100.0.100/32	0.0.0.0	0		A H		
10.200.0.0/24	198.51.100.2			A?B	148970	
10.200.1.0/24	0.0.0.0	1		~		
10.200.1.0/24	10.200.1.10	0		A C		ethernet1/2
10.200.1.10/32	0.0.0.0	0		A H		
198.51.100.0/30	198.51.100.1	0		A C		tunnel.1
198.51.100.1/32	0.0.0.0	0		A H		
198.51.100.4/30	198.51.100.5	0		A C		tunnel.2
198.51.100.5/32	0.0.0.0	0		A H		

### CONFIGURING STATIC ROUTING

Use the instructions here if your CPE does not support BGP over IPsec, or you do not want to use BGP over IPsec.

In this task, you configure static routes to direct traffic through the tunnel interfaces to reach the DRG and finally the VCN hosts.

1. Follow tasks 1-5 in the preceding section.
2. Configure static routes:
  - a. Go to **Network**, to **Virtual Routers**, to **default**, to **Static Routes**, and then click **Add**.
  - b. For Route 1, configure the parameters as shown in the next image.

Virtual Router - Static Route - IPv4

Name: route1-VCN

Destination: 10.200.0.0/24

Interface: tunnel.1

Next Hop: None

Admin Distance: 10 - 240

Metric: 10

Route Table: Unicast

BFD Profile: Disable BFD

Path Monitoring

Failure Condition:  Any  All

Preemptive Hold Time (min): 2

<input type="checkbox"/>	Name	Enable	Source IP	Destination IP	Ping Interval(sec)	Ping Count

+ Add - Delete

OK Cancel

c. For Route 2, configure the parameters as shown in the next image.

Virtual Router - Static Route - IPv4

Name: route2-VCN

Destination: 10.200.0.0/24

Interface: tunnel.2

Next Hop: None

Admin Distance: 10 - 240

Metric: 10

Route Table: Unicast

BFD Profile: Disable BFD

Path Monitoring

Failure Condition:  Any  All

Preemptive Hold Time (min): 2

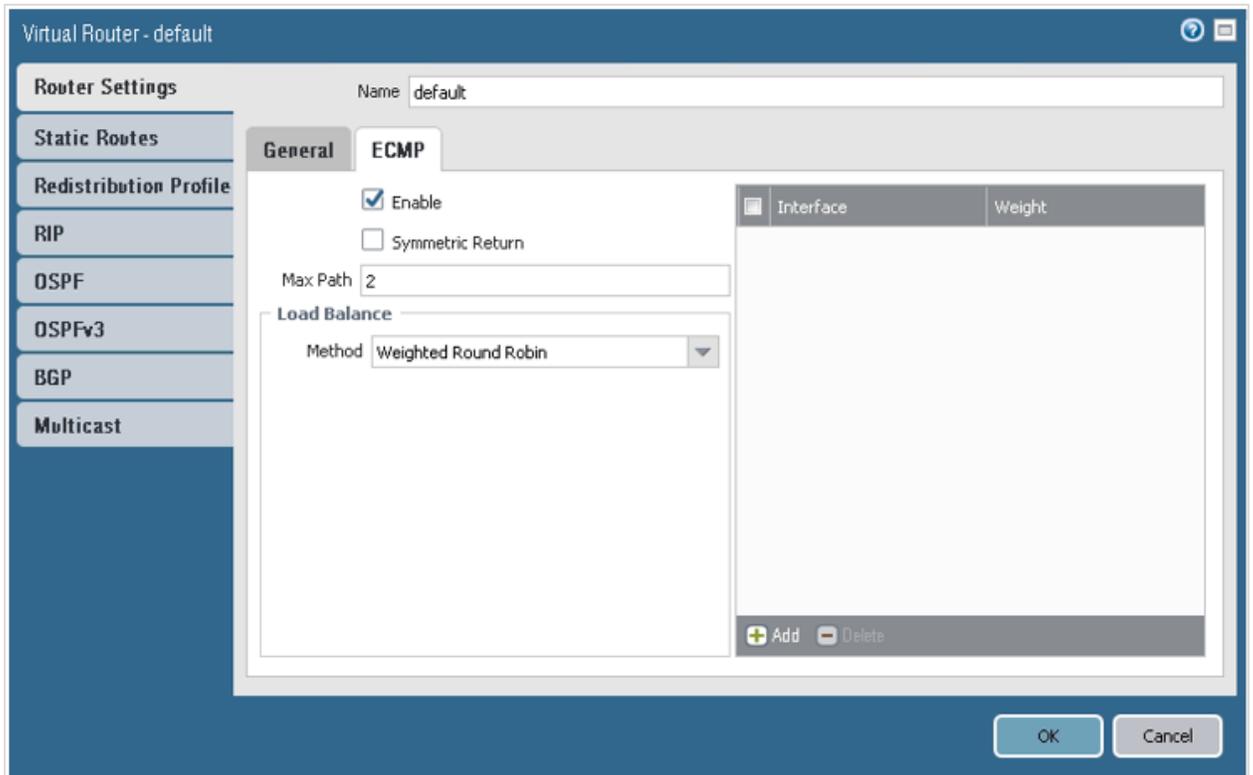
<input type="checkbox"/>	Name	Enable	Source IP	Destination IP	Ping Interval(sec)	Ping Count
+ Add - Delete						

OK Cancel

3. (Recommended) Enable ECMP for traffic sent through the two tunnels. The metric for both routes is set to 10. Here are some important notes about enabling ECMP:
  - First check to see if your networking design allows for ECMP.
  - Enabling or disabling ECMP on an existing virtual router causes the system to restart the virtual router. This might cause existing sessions to be terminated.

- This example uses the default virtual router. Use the correct virtual router for your network environment.

To enable ECMP, go to **Network**, to **Virtual Routers**, to **default**, to **Router Settings**, to **ECMP**, and then select the check box for **Enable**.



Here are screenshots that show the final configuration after this task is complete:

## CHAPTER 23 Networking

IPv4		IPv6							
3 items									
Name	Destination	Interface	Next Hop		Admin Distance	Metric	BFD	Route Table	
			Type	Value					
<input type="checkbox"/> default_route	0.0.0.0/0	ethernet1/1	ip-address	10.100.0.99	default	10	None	unicast	
<input checked="" type="checkbox"/> route1-VCN	10.200.0.0/24	tunnel.1			default	10	None	unicast	
<input checked="" type="checkbox"/> route2-VCN	10.200.0.0/24	tunnel.2			default	10	None	unicast	

IKE Gateway/Site-to-Site											
Name	Status	Type	Interface	Local IP	Peer IP	Status	Interface	Virtual Router	Virtual System	Security Zone	Status
Tunnel1	Tunnel Info	Auto Key	ethernet1/1	Internet_address	10.150.128.1	IKE Info	tunnel.1	default (Show Routes)	vsyst	pssec_tunnel	
Tunnel2	Tunnel Info	Auto Key	ethernet1/1	Internet_address	10.150.127.1	IKE Info	tunnel.2	default (Show Routes)	vsyst	pssec_tunnel	

### CHANGING THE IKE IDENTIFIER

If the CPE is behind a NAT device with a private IP address on the exit interface that the tunnel interfaces use as the source, you must specify the public IP address of the NAT device as the local IKE ID. You can do this by setting the **Local Identification** value in the **IKE Gateway** configuration:

The screenshot shows the 'IKE Gateway' configuration window with the 'Advanced Options' tab selected. The configuration is as follows:

Name	oracle-vpn1	
Version	IKEv1 only mode	
Address Type	<input checked="" type="radio"/> IPv4 <input type="radio"/> IPv6	
Interface	ethernet1/1	
Local IP Address	Internet_address	
Peer IP Type	<input checked="" type="radio"/> Static <input type="radio"/> Dynamic	
Peer IP Address	10.150.128.1	
Authentication	<input checked="" type="radio"/> Pre-Shared Key <input type="radio"/> Certificate	
Pre-shared Key	●●●●●●●●	
Confirm Pre-shared Key	●●●●●●●●	
Local Identification	IP address	10.100.0.100
Peer Identification	IP address	10.150.128.1

Buttons: OK, Cancel

**Verification**

To verify the IPsec tunnel status:

## CHAPTER 23 Networking



The screenshot shows the Palo Alto Networks GUI with the 'Network' tab selected. A table displays VPN Tunnel configurations. The table has columns for Name, Status, Type, Interface, Local IP, Peer IP, Status, Interface, Virtual Router, Virtual System, Security Zone, and Status. Two tunnels are listed: Tunnel1 and Tunnel2, both with a status of 'Tunnel Info' and 'IKE Info'.

Name	Status	Type	Interface	Local IP	Peer IP	Status	Interface	Virtual Router	Virtual System	Security Zone	Status
Tunnel1	Tunnel Info	Auto Key	ethernet1/1	Internet_address	10.150.128.1	IKE Info	tunnel.1	default (Show Routes)	vsys1	ipsec_tunnel	OK
Tunnel2	Tunnel Info	Auto Key	ethernet1/1	Internet_address	10.150.127.1	IKE Info	tunnel.2	default (Show Routes)	vsys1	ipsec_tunnel	OK

Use this command to verify the IKE SA:

```
show vpn ike-sa
```

Use this command to verify the IPsec tunnel configuration:

```
show vpn tunnel name <tunnel_name>
```

To verify the BGP status, look for **Established**:



The screenshot shows the Palo Alto Networks GUI for a Virtual Router. The 'BGP' tab is selected, and the 'Peer' sub-tab is active. A table displays BGP Peer configurations. The table has columns for Name, Group, Local IP, Peer IP, Peer AS, Password Set, Status, and Status Duration (secs.). Two peers are listed: Session1 and Session2, both with a status of 'Established'.

Name	Group	Local IP	Peer IP	Peer AS	Password Set	Status	Status Duration (secs.)
Session1	DRG1	198.51.100.1:51276	198.51.100.2:179	31898	no	Established	3810
Session2	DRG2	198.51.100.5:179	198.51.100.6:43509	31898	no	Established	155929

To verify the BGP status from the command line:

```
show routing protocol bgp peer peer-name <name>
```

To verify that the routes are installed in the routing table:

```
show routing route
```

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

### WatchGuard

You can configure a WatchGuard Firebox as your CPE device for an IPsec VPN.

Go to the [WatchGuard knowledge base article](#) to download the configuration instructions.

### Yamaha RTX Series

This configuration was validated using a RTX1210 running Firmware Rev.14.01.28 and RTX830 running Firmware Rev.15.02.03.



#### **Important**

Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

### **Before Starting**

Before configuring your CPE, make sure to:

- Configure your internet provider settings
- Configure firewall rules to open UDP port 500, UDP port 4500, and ESP.

### Supported Encryption Domain or Proxy ID

The values for the encryption domain (also known as a proxy ID, security parameter index (SPI), or traffic selector) depend on whether your CPE supports route-based tunnels or policy-based tunnels. For more information about the correct encryption domain values to use, see [Supported Encryption Domain or Proxy ID](#).

### Parameters from API or Console

Get the following parameters from the Oracle Cloud Infrastructure Console or API.

#### **`${ipAddress#}`**

- Oracle VPN headend IPsec tunnel endpoints. There is one value for each tunnel.
- Example value: 129.146.12.52

#### **`${sharedSecret#}`**

- The IPsec IKE pre-shared-key. There is one value for each tunnel.
- Example value: EXAMPLEDPfAMkD7nTH3SWr6OFabdT6exXn6enSlSkbE

#### **`${cpePublicIpAddress}`**

- The public IP address for the CPE (previously made available to Oracle via the Console).

#### **`${VcnCidrBlock}`**

- When creating the VCN, your company selected this CIDR to represent the IP aggregate network for all VCN hosts.
- Example Value: 10.0.0.0/20

### Parameters Based on Current CPE Configuration and State

The following parameters are based on your current CPE configuration.

#### **`${tunnelInterface#}`**

- An interface number to identify the specific tunnel.
- Example value: 1

### **`${ipsecPolicy#}`**

- The SA policy to be used for the selected inline interface.
- Example value: 1

### **`${localAddress}`**

- The public IP address of your CPE.
- Example value: 146.56.2.52

### **Config Template Parameter Summary**

Each region has multiple Oracle IPSec headends. The template below allows you to set up multiple tunnels on your CPE, each to a corresponding headend. In the table below, "User" is you/your company.

<b>Parameter</b>	<b>Source</b>	<b>Example Value</b>
<code>\${ipAddress1}</code>	Console/API	129.146.12.52
<code>\${sharedSecret1}</code>	Console/API	(long string)
<code>\${ipAddress2}</code>	Console/API	129.146.13.52
<code>\${sharedSecret2}</code>	Console/API	(long string)
<code>\${cpePublicIpAddress}</code>	User	1.2.3.4
<code>\${VcnCidrBlock}</code>	User	10.0.0.0/20

**Important**

The following ISAKMP and IPsec policy parameter values are applicable to VPN Connect in the commercial cloud. For the [Government Cloud](#), you must use the values listed in [Required VPN Connect Parameters for Government Cloud](#).

**ISAKMP Policy Options**

Parameter	Recommended Value
ISAKMP protocol version	Version 1
Exchange type	Main mode
Authentication method	Pre-shared keys
Encryption	AES-256-cbc
Authentication algorithm	SHA-256
Diffie-Hellman Group	Group 5
IKE session key lifetime	28800 seconds (8 hours)

**IPsec Policy Options**

Parameter	Recommended Value
IPsec protocol	ESP, tunnel-mode
Encryption	AES-256-cbc

Parameter	Recommended Value
Authentication algorithm	HMAC-SHA1-96
Diffie-Hellman Group	Group 5
Perfect Forward Secrecy	Enabled
IPSec session key lifetime	3600 seconds (1 hour)

### CPE Configuration

#### ISAKMP AND IPSEC CONFIGURATION

```
tunnel select 1
description tunnel OCI-VPN1
ipsec tunnel 1
 ipsec sa policy 1 1 esp aes256-cbc sha-hmac
 ipsec ike duration ipsec-sa 1 3600
 ipsec ike duration isakmp-sa 1 28800
 ipsec ike encryption 1 aes256-cbc
 ipsec ike group 1 modp1536
 ipsec ike hash 1 sha256
 ipsec ike keepalive log 1 off
 ipsec ike keepalive use 1 on dpd 5 4
 ipsec ike local address 1 ${cpePublicIpAddress}
 ipsec ike local id 1 0.0.0.0/0
 ipsec ike nat-traversal 1 on
 ipsec ike pfs 1 on
 ipsec ike pre-shared-key 1 text ${sharedSecret1}
 ipsec ike remote address 1 ${ipAddress1}
 ipsec ike remote id 1 0.0.0.0/0
ip tunnel tcp mss limit auto
tunnel enable 1

tunnel select 2
description tunnel OCI-VPN2
ipsec tunnel 2
 ipsec sa policy 2 2 esp aes256-cbc sha-hmac
 ipsec ike duration ipsec-sa 2 3600
 ipsec ike duration isakmp-sa 2 28800
```

## CHAPTER 23 Networking

---

```
ipsec ike encryption 2 aes256-cbc
ipsec ike group 2 modp1536
ipsec ike hash 2 sha256
ipsec ike keepalive log 2 off
ipsec ike keepalive use 2 on dpd 5 4
ipsec ike local address 2 ${cpePublicIpAddress}
ipsec ike local id 2 0.0.0.0/0
ipsec ike nat-traversal 2 on
ipsec ike pfs 2 on
ipsec ike pre-shared-key 2 text ${sharedSecret2}
ipsec ike remote address 2 ${ipAddress2}
ipsec ike remote id 2 0.0.0.0/0
ip tunnel tcp mss limit auto
tunnel enable 2

ipsec auto refresh on
```

### STATIC ROUTES CONFIGURATION

```
ip route ${VcnCidrBlock} gateway tunnel 1 hide gateway tunnel 2 hide
```

## Working with VPN Connect

This topic contains some details about working with VPN Connect and the related components. Also see these topics:

- [VPN Connect Overview](#)
- [VPN Connect Quickstart](#)
- [Setting Up VPN Connect](#)
- [CPE Configuration](#)
- [VPN Connect FAQ](#)
- [VPN Connect Metrics](#)
- [VPN Connect Troubleshooting](#)



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Viewing Tunnel Status and Configuration

When you successfully create the IPsec connection, Oracle produces important configuration information for each of the resulting IPsec tunnels. For example, see [task 2h](#) in the overall setup process. You can view that information and the status of the tunnels at any time. This includes the BGP status if the tunnel is configured to use BGP dynamic routing.

### To view the status and configuration information for the IPsec tunnels

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPsec Connections**.

A list of the IPsec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the IPsec connection you're interested in.  
Each tunnel's details are displayed, including the IPsec status, the BGP status (if the tunnel uses BGP dynamic routing), and the Oracle VPN IP address (the VPN headend).
3. To view a tunnel's shared secret:
  - a. Click the tunnel you're interested in.
  - b. Next to the **Shared Secret** field, click **Show**.

### Changing the Static Routes

You can change the static routes for an existing IPsec connection. You can provide up to 10 static routes.

Remember that an IPsec connection can use either static routing or BGP dynamic routing. You associate the static routes with the overall IPsec connection and not the individual tunnels. If an IPsec connection has static routes associated with it, Oracle uses them for routing a tunnel's traffic *only* if the tunnel itself is configured to use static routing. If it's configured to use BGP dynamic routing, the IPsec connection's static routes are ignored.



#### Important

The IPsec connection goes down while it is reprovisioned with your static route changes.

#### To edit the static routes

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPsec Connections**.

A list of the IPsec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. For the IPsec connection you're interested in, click the Actions icon (three dots), and then click **Edit**.

The current static routes are displayed.

3. Make your changes and click **Save Changes**.

### Changing the CPE IKE Identifier That Oracle Uses

If your [CPE is behind a NAT device](#), you might need to give Oracle your CPE IKE identifier. You can either specify it when you create the IPsec connection, or later edit the IPsec connection

and change the value. Oracle expects the value to be an IP address or fully qualified domain name (FQDN). When you specify the value, you also specify which type it is.



### Important

The IPSec connection goes down while it is reprovisioned to use your CPE IKE identifier.

### To change the CPE IKE identifier that Oracle uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

A list of the IPSec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. For the IPSec connection you're interested in, click the Actions icon (three dots), and then click **Edit**.

The current CPE IKE identifier that Oracle is using is displayed at the bottom of the dialog.

3. Enter your new values for **CPE IKE Identifier Type** and **CPE IKE Identifier**, and then click **Save Changes**.

### Using IKEv2

Oracle supports Internet Key Exchange (IKE) version 1 and [version 2](#) (IKEv2).

If you want to use IKEv2 and your CPE supports it, you must:

- Configure each IPSec tunnel to use IKEv2 in the Oracle Console. See the following procedures.

- Configure your CPE to use IKEv2 encryption parameters that the CPE supports. For a list of parameters that Oracle supports, see [Supported IPSec Parameters](#).

### New IPSec connection: using IKEv2



#### Note

If you [create a new IPSec connection manually](#), you can specify IKEv2 when you create the IPSec connection in the Oracle Console. See the procedure that immediately follows.

If you instead use the [VPN quickstart workflow](#), the IPSec connection is configured to use IKEv1 only. However, after the workflow is complete, you can edit the resulting IPSec tunnels in the Oracle Console and change them to use IKEv2.

To manually set up a new IPSec connection that uses IKEv2:

1. While [creating the IPSec connection](#) in the Oracle Console, in the **Advanced Options** section, click the **Tunnel 1** tab.
2. From the **IKE Version** menu, select **IKEv2**.
3. Repeat the preceding step for the **Tunnel 2** tab.
4. Later when configuring your CPE, configure it to use only IKEv2 and related IKEv2 encryption parameters that the CPE supports.

## Existing IPSec connection: upgrading to IKEv2



### Important

Oracle recommends performing the following process for one tunnel at a time to avoid disruption in your overall connection. If your connection is not redundant (for example, does not have multiple tunnels), expect downtime while you upgrade to IKEv2.

1. Change the tunnel's IKE version in the Oracle Console:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.
  - b. Click the IPSec connection you're interested in.
  - c. Click the tunnel to view its details.
  - d. Click **Edit**.
  - e. From the **IKE Version** menu, select **IKEv2**.
  - f. Click **Save Changes**.
2. Update your CPE configuration for the tunnel to use IKEv2 and the related encryption parameters that the CPE supports. For a list of parameters that Oracle supports, see [Supported IPSec Parameters](#).
3. If the security associations did not rekey immediately, force a rekey for that tunnel on your CPE. In other words, clear the phase 1 and phase 2 security associations and do not wait for them to expire. Some CPE devices wait for the SAs to expire before rekeying. Forcing the rekey lets you confirm immediately that the IKE version configuration is correct.
4. To verify, ensure that the security associations for the tunnel rekey correctly. If they don't, confirm that the correct IKE version is set in the Oracle Console and on your CPE, and that the CPE is using the desired parameters.

After you've confirmed the first tunnel is up and running again, repeat the preceding steps for the second tunnel.

### Changing the Shared Secret That an IPSec Tunnel Uses

When you set up an IPSec VPN, by default Oracle provides each tunnel's shared secret (also called the pre-shared key). You might have a particular shared secret that you want to use instead. You can specify each tunnel's shared secret when you create the IPSec connection, or you can edit the tunnels and provide each new shared secret then. For the shared secret, only numbers, letters, and spaces are allowed. Oracle recommends using a different shared secret for each tunnel.



#### Important

When you change a tunnel's shared secret, both the overall IPSec connection and the tunnel go into the Provisioning state while the tunnel is reprovisioned with the new shared secret. The other tunnel in the IPSec connection remains in the Available state. However, while the first tunnel is being reprovisioned, you cannot change the second tunnel's configuration.

### To change the shared secret that an IPSec tunnel uses

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

A list of the IPSec connections in the compartment that you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the IPSec connection you're interested in.
3. Click the tunnel you're interested in.

4. Next to the **Shared Secret** field, click **Edit**.
5. Enter your new value. Only numbers, letters, and spaces are allowed.
6. Click **Save Changes**.

### Changing from Static Routing to BGP Dynamic Routing

If you want to change an existing IPsec VPN from using static routing to using BGP dynamic routing, follow the process in this section.



#### Warning

When you change a tunnel's routing type, the tunnel's IPsec status does not change during reprovisioning. However, routing through the tunnel is affected. Traffic is temporarily disrupted until your network engineer configures your CPE device in accordance with the routing type change. **If your existing IPsec VPN is currently configured to use only a single tunnel, this process will disrupt your connection to Oracle.** If your IPsec VPN instead uses multiple tunnels, Oracle recommends reconfiguring one tunnel at a time to avoid disrupting your connection to Oracle.

### To change from static routing to BGP dynamic routing

Prerequisites:

- You've read this section: [Routing for the Oracle IPsec VPN](#)
- You've gathered the necessary BGP routing information:

- Your network's ASN. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- For each tunnel: The BGP IP address for each end of the tunnel (the two addresses for a given tunnel must be a pair from a /30 or /31 subnet, and they must be part of the IPsec VPN's encryption domain)

Repeat the following process for each tunnel in the IPsec connection:

1. Reconfigure the tunnel's routing type from static routing to BGP dynamic routing:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPsec Connections**.
  - b. Click the IPsec connection you're interested in.

The tunnels are listed, and the status for each tunnel is shown. The **BGP Status** for the tunnel you're interested in should show only a hyphen (no value), which means that the tunnel is currently configured to use static routing.
  - c. Click the tunnel to view all of its details.
  - d. Click **Edit**.
  - e. Do the following:
    - **Routing Type:** Select the radio button for **BGP Dynamic Routing**.
    - **BGP ASN:** Enter your network's BGP ASN.
    - **Inside Tunnel Interface - CPE:** Enter the BGP IP address with subnet mask (either /30 or /31) for the CPE end of the tunnel. For example: 10.0.0.16/31.
    - **Inside Tunnel Interface - Oracle:** Enter the BGP IP address with subnet mask (either /30 or /31) for the Oracle end of the tunnel. For example: 10.0.0.17/31.
  - f. Click **Save Changes**.

The tunnel's **BGP Status** changes to Down.

2. Have your network engineer update your CPE device's tunnel configuration to use BGP.

3. On your side of the connection, confirm that the tunnel's BGP session is in an established state. If it is not, make sure you've configured the correct IP addresses for the tunnel in the Oracle Console and also for your CPE device.
4. In the Oracle Console, confirm that the tunnel's **BGP Status** is now Up.
5. Confirm that your CPE device is learning routes from Oracle, and your CPE device is advertising routes to Oracle. If you want to re-advertise the Oracle routes from BGP back to your on-premises network, make sure your CPE device is configured accordingly. Your existing policy to advertise the static routes to your on-premises network may not work for the BGP learned routes.
6. Ping the Oracle BGP IP address from your side of the connection to confirm that traffic is flowing.

After you've confirmed the first tunnel is up and running with BGP, repeat the process for the second tunnel.



### Important

As noted in [Routing for the Oracle IPSec VPN](#), the static routes that are still configured for the overall IPSec connection do NOT override the BGP routing. Those static routes are ignored when Oracle routes traffic through a tunnel that is configured to use BGP.

Also, you can change a tunnel's routing type back to static routing if necessary. You might do this if the scheduled downtime window for the CPE device is ending soon and you're having trouble establishing the BGP session. When you switch back to static routing, make sure the overall IPSec connection still has your desired static routes configured.

### Monitoring Your IPsec VPN

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring your connection, see [VPN Connect Metrics](#).

### Disabling or Terminating the IPsec VPN

If you want to disable the IPsec VPN between your on-premises network and VCN, you can simply detach the DRG from the VCN instead of deleting the IPsec connection. If you're also using the DRG with [FastConnect](#), detaching the DRG would also interrupt the flow of traffic over FastConnect.

You can delete the IPsec connection. However, if you later want to re-establish it, your network engineer would have to configure your CPE device again with a new set of tunnel configuration information from Oracle.

If you want to permanently delete the entire IPsec VPN, you must first terminate the IPsec connection. Then you can delete the CPE object. If you're not using the DRG for another connection to your on-premises network, you can detach it from the VCN and then delete it.

### To delete an IPsec connection

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPsec Connections**.

A list of the IPsec connections in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the IPsec connection you're interested in.
3. Click **Terminate**.
4. Confirm the deletion when prompted.

The IPSec connection will be in the Terminating state for a short period while it's being deleted.

### To delete a CPE object

**Prerequisite:** There must not be an IPSec connection between the CPE object and a DRG.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.  
A list of the CPE objects in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. For the CPE object that you want to delete, click the Actions icon (three dots), and then click **Delete**.
3. Confirm the deletion when prompted.

The object will be in the Terminating state for a short period while it's being deleted.

### Managing Tags for an IPSec Connection or CPE Object

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

### To manage tags for an IPSec connection

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.  
A list of the IPSec connections in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the IPSec connection you're interested in.

3. Click the **Tags** tab to view or edit the existing tags. Or click **Add tags** to add new ones.

### To manage tags for a CPE object

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.

A list of the CPE objects in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the CPE object you're interested in.
3. Click the **Tags** tab to view or edit the existing tags. Or click **Apply tag(s)** to add new ones.

### Moving a CPE Object to a Different Compartment

You can move your resources from one compartment to another. After you move the resource to the new compartment, inherent policies apply immediately and affect access to the resource through the Console. Moving the CPE object to a different compartment does not affect the connection between your data center and Oracle Cloud Infrastructure. For more information, see [Moving Resources to a Different Compartment](#).

### To move a CPE object to a different compartment

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Customer-Premises Equipment**.
2. Find the CPE object in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.

### Managing Your DRG

For tasks related to DRGs, see [Dynamic Routing Gateways \(DRGs\)](#).

### VPN Connect FAQ

#### Can I customize configuration parameters such as IKE ID or IPsec VPN tunnel lifetime on the DRG?

No. Oracle has predetermined the configuration parameters that work with the IPsec VPN service. Your IPsec VPN can't be established if there is a mismatch.

If your CPE is behind a NAT device, you can provide Oracle with your CPE's IKE identifier so that Oracle can use the same value on the Oracle side. For more information, see [Changing the CPE IKE Identifier That Oracle Uses](#).

#### Can Oracle initiate the IPsec VPN connection?

No. You must initiate it from your end.

#### Is an IPsec VPN supported over FastConnect?

No, [FastConnect](#) and the IPsec VPN are two different services.

#### Can packets from the VCN be sourced with the public IP of the DRG? In other words, can the DRG source the NAT VCN traffic?

No. Packets from the VCN have the private IP address of the instance as a source IP address. Oracle cannot change the source IP address to the private or public IP address of the DRG.

Oracle has provided two VPN endpoints to build tunnels to. Does Oracle route the same network over both tunnels?

Yes. Traffic for all the subnets in the VCN attached to your DRG is routed over both tunnels.

How many VPN tunnels can I have from a single CPE device?

You can have a maximum of eight tunnels from a unique CPE IP address per region. If you want more than eight tunnels, either use a different IP address for the additional ones, or use a different CPE device (recommended).

I'm having trouble with my IPSec connection. What can I do?

See [VPN Connect Troubleshooting](#).

## Using the API for VPN Connect

This topic lists the Networking service API operations for managing VPN Connect components.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To [manage your VCN and subnets](#), use these operations:

- [ListVcns](#)
- [CreateVcn](#)
- [GetVcn](#)
- [UpdateVcn](#)
- [DeleteVcn](#)
- [ChangeVcnCompartment](#)
- [ListSubnets](#)
- [CreateSubnet](#)
- [GetSubnet](#)
- [UpdateSubnet](#)
- [DeleteSubnet](#)
- [ChangeSubnetCompartment](#)

To [manage your DRG](#), use these operations:

- [ListDrgs](#)
- [CreateDrg](#)
- [GetDrg](#)
- [UpdateDrg](#)
- [DeleteDrg](#)
- [ListDrgAttachments](#)
- [CreateDrgAttachment](#): This operation attaches a DRG to a VCN and results in a `DrgAttachment` object with its own OCID.
- [GetDrgAttachment](#)
- [UpdateDrgAttachment](#)
- [DeleteDrgAttachment](#): This operation detaches a DRG from a VCN by deleting the `DrgAttachment` object.

To [manage routing for your VCN](#), use these operations:

- [ListRouteTables](#)
- [GetRouteTable](#)
- [UpdateRouteTable](#)
- [CreateRouteTable](#)
- [DeleteRouteTable](#)

To [manage security lists for your VCN](#), use these operations:

- [ListSecurityLists](#)
- [GetSecurityList](#)
- [UpdateSecurityList](#)
- [CreateSecurityList](#)
- [DeleteSecurityList](#)

To manage your CPEs, use these operations:

- [ListCpes](#)
- [CreateCpe](#)
- [GetCpe](#)
- [UpdateCpe](#)
- [DeleteCpe](#)
- [ChangeCpeCompartment](#)

To manage your IPSec connections, use these operations:

- [ListIPSecConnections](#)
- [CreateIPSecConnection](#): Use this operation to get the configuration information for each tunnel, including the IP address of the DRG (the VPN headend) and the shared secret. See [CPE Configuration](#).
- [GetIPSecConnection](#)

- [UpdateIPSecConnection](#)
- [DeleteIPSecConnection](#)
- [ChangeIPSecConnectionCompartment](#)
- [GetIPSecConnectionDeviceStatus](#): Use this operation to determine the status of the IPSec tunnels (up or down).
- [GetIPSecConnectionDeviceConfig](#): Use this operation to get the configuration information for each tunnel.

## VPN Connect Metrics

You can monitor the health, capacity, and performance of your [VPN Connect](#) by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

This topic describes the metrics emitted by the metric namespace `oci_vpn`.

Resources: IPSec connections.

### Overview of Metrics: `oci_vpn`

The available metrics help you determine quickly if your [VPN Connect](#) is up, how much data is flowing over the connection, and if packets are being dropped for unexpected errors.

A VPN Connect includes these resources:

- An IPSec connection, which you can think of as the *parent resource* (identified by `parentResourceId` in the following discussion).
- One or more individual tunnels associated with that IPSec connection (each identified by the tunnel's `publicIp` in the following discussion).

### Required IAM Policy

To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources

being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_vpn

The metrics listed in the following table are automatically available for any VPN Connect that you create. You do not need to enable monitoring on the resource to get these metrics.

You also can use the Monitoring service to create [custom queries](#).

Each metric includes the following dimensions:

#### **PARENTRESOURCEID**

The OCID of the IPsec connection (the parent resource). The connection has multiple individual tunnels.

#### **PUBLICIP**

Although each tunnel has its own OCID, it can be easier to use the `publicIp` dimension to identify a specific IPsec tunnel in the connection. The value is the public IP address of the Oracle end of the tunnel (also known as the *Oracle VPN headend*).

Metric	Metric Display Name	Unit	Description	Dimensions
TunnelState	<b>IPSec Tunnel State</b>	Binary (1 or 0)	Whether the tunnel is up (1) or down (0).	parentResourceId publicIp
PacketsReceived	<b>Packets Received</b>	Packets	Number of packets received at the Oracle end of the connection.	
BytesReceived	<b>Bytes Received</b>	Bytes	Number of bytes received at the Oracle end of the connection.	
PacketsSent	<b>Packets Sent</b>	Packets	Number of packets sent from the Oracle end of the connection.	
BytesSent	<b>Bytes Sent</b>	Bytes	Number of bytes sent from the Oracle end of the connection.	
PacketsError	<b>Packets with Errors</b>	Packets	Number of packets dropped at the Oracle end of the connection. Dropped packets indicate a misconfiguration in some part of the overall system. Check if there's been a change to the configuration of your VCN, the IPSec VPN, or your CPE.	

### Using the Console

To view default metrics charts for an individual tunnel in an IPSec connection

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.
2. Click the IPSec connection to view its details.
3. Click the tunnel you're interested in to view its details and default metrics charts.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

To view default metric charts for all IPSec connections in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that contains the IPSec connection you're interested in.
3. For **Metric Namespace**, select **oci\_vpn**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.

Each IPSec tunnel is a single line in a given chart. The tunnel is identified in the chart by the public IP address of the Oracle end of the tunnel.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

## VPN Connect Troubleshooting

This topic covers troubleshooting techniques for an IPSec VPN that has issues.

Some of the troubleshooting techniques assume that you are a network engineer with access to your CPE device's configuration.

### General Issues

#### IPSec tunnel is DOWN

Check these items:

- **Basic configuration:** The IPSec tunnel consists of both [phase-1 \(ISAKMP\) and phase-2 \(IPSec\) configuration](#). Confirm that both are configured correctly on your CPE device. See the configuration appropriate for your CPE device:

#### List of configurations

- [Verified CPE Devices](#)
- Checkpoint:
  - [Check Point: Route-Based](#)
  - [Check Point: Policy-Based](#)
- Cisco ASA:
  - [Cisco ASA: Route-Based](#)
  - [Cisco ASA: Policy-Based](#)

- [Cisco IOS](#)
  - [FortiGate](#)
  - [Juniper MX](#)
  - [Juniper SRX](#)
  - [Libreswan](#)
  - [NEC IX Series](#)
  - [Openswan](#)
  - [Palo Alto](#)
  - [WatchGuard](#)
  - [Yamaha RTX Series](#)
- **Local and remote proxy IDs:** If you're using a policy-based configuration, check if your CPE is configured with more than one pair of local and remote proxy IDs (subnets). The Oracle VPN router supports only one pair. If your CPE has more than one pair, update the configuration to include only one pair, and choose one of the following two options:

Option	Local Proxy ID	Remote Proxy ID
1	ANY (or 0.0.0.0/0)	ANY (or 0.0.0.0/0)
2	On-premises CIDR (an aggregate that covers all the subnets of interest)	VCN's CIDR

- **NAT device:** If the CPE is behind a NAT device, the CPE IKE identifier configured on your CPE might not match the CPE IKE identifier Oracle is using (the public IP address of your CPE). If your CPE does not support setting the CPE IKE identifier on your end, you can provide Oracle with your CPE IKE identifier in the Oracle Console. For more information, see [If Your CPE Is Behind a NAT Device](#).

- **NAT-T:** Although NAT-T is supported in most regions, it is not yet fully supported in the US East (Ashburn) region (see this [known issue](#)). If your IPsec tunnels connect to that region, and your CPE is currently configured for NAT-T, disable NAT-T and then re-initiate the IPsec connection on your CPE.

### IPsec tunnel is UP, but no traffic is passing through

Check these items:

- **Phase 2 (IPsec) configuration:** Confirm that the [phase 2 \(IPsec\) parameters](#) are configured correctly on your CPE device. See the configuration appropriate for your CPE device:

#### List of configurations

- [Verified CPE Devices](#)
- Checkpoint:
  - [Check Point: Route-Based](#)
  - [Check Point: Policy-Based](#)
- Cisco ASA:
  - [Cisco ASA: Route-Based](#)
  - [Cisco ASA: Policy-Based](#)
- [Cisco IOS](#)
- [FortiGate](#)
- [Juniper MX](#)
- [Juniper SRX](#)
- [Libreswan](#)
- [NEC IX Series](#)
- [Openswan](#)

- [Palo Alto](#)
  - [WatchGuard](#)
  - [Yamaha RTX Series](#)
- **NAT-T:** Although NAT-T is supported in most regions, it is not yet fully supported in the US East (Ashburn) region (see this [known issue](#)). If your IPsec tunnels connect to that region, and your CPE is currently configured for NAT-T, disable NAT-T and then re-initiate the IPsec connection on your CPE.
  - **VCN security lists:** Ensure you've set up the [VCN security lists](#) to allow the desired traffic (both ingress and egress rules). Note that the VCN's [default security list](#) does not allow ping traffic (ICMP type 8 and ICMP type 0). You must add the appropriate ingress and egress rules to allow ping traffic.
  - **Firewall rules:** Ensure that your firewall rules allow both ingress and egress traffic with the Oracle VPN headend IPs and the VCN CIDR block.
  - **Asymmetric routing:** Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.
  - **Cisco ASA:** Do not use the `originate-only` option with an Oracle IPsec VPN tunnel. It causes the tunnel's traffic to be inconsistently blackholed. The command is only for tunnels between two Cisco devices. Here's an example of the command that you should NOT use for the Oracle IPsec VPN tunnels: `crypto map <map name> <sequence number> set connection-type originate-only`

### IPsec tunnel is UP, but traffic is passing in only one direction

Check these items:

- **Asymmetric routing:** Oracle uses asymmetric routing across the multiple tunnels that make up the IPsec VPN connection. Even if you configure one tunnel as primary and

another as backup, traffic from your VCN to your on-premises network can use any tunnel that is "up" on your device. Configure your firewalls accordingly. Otherwise, ping tests or application traffic across the connection will not reliably work.

- **Single tunnel preferred:** If you want to use only one of the tunnels, make sure that you have the proper policy or routing in place on the CPE to prefer that tunnel.
- **Multiple IPSec connections:** If you have multiple IPSec connections with Oracle, make sure to specify more specific static routes for the preferred IPSec connection.
- **VCN security lists:** Ensure that your [VCN security lists](#) allow traffic in both directions (ingress and egress).
- **Firewall rules:** Ensure that your firewall rules allow traffic in *both* directions with the Oracle VPN headend IPs and the VCN CIDR block.

### For an IPSec VPN with a Policy-Based Configuration

#### IPSec tunnel is DOWN

Check these items:

- **Basic configuration:** The IPSec tunnel consists of both [phase-1 \(ISAKMP\) and phase-2 \(IPSec\) configuration](#). Confirm that both are configured correctly on your CPE device. See the configuration appropriate for your CPE device:

#### List of configurations

- [Verified CPE Devices](#)
- Checkpoint:
  - [Check Point: Route-Based](#)
  - [Check Point: Policy-Based](#)

- Cisco ASA:
    - [Cisco ASA: Route-Based](#)
    - [Cisco ASA: Policy-Based](#)
  - [Cisco IOS](#)
  - [FortiGate](#)
  - [Juniper MX](#)
  - [Juniper SRX](#)
  - [Libreswan](#)
  - [NEC IX Series](#)
  - [Openswan](#)
  - [Palo Alto](#)
  - [WatchGuard](#)
  - [Yamaha RTX Series](#)
- **Local and remote proxy IDs:** If you're using a policy-based configuration, check if your CPE is configured with more than one pair of local and remote proxy IDs (subnets). The Oracle VPN router supports only one pair. If your CPE has more than one pair, update the configuration to include only one pair, and choose one of the following two options:

Option	Local Proxy ID	Remote Proxy ID
1	ANY (or 0.0.0.0/0)	ANY (or 0.0.0.0/0)
2	On-premises CIDR (an aggregate that covers all the subnets of interest)	VCN's CIDR

- **NAT device:** If the CPE is behind a NAT device, the CPE IKE identifier configured on your CPE might not match the CPE IKE identifier Oracle is using (the public IP address of your CPE). If your CPE does not support setting the CPE IKE identifier on your end, you can provide Oracle with your CPE IKE identifier in the Oracle Console. For more information, see [If Your CPE Is Behind a NAT Device](#).
- **Cisco ASA:** Do not use the `originate-only` option with an Oracle IPSec VPN tunnel. It causes the tunnel's traffic to be inconsistently blackholed. The command is only for tunnels between two Cisco devices. Here's an example of the command that you should NOT use for the Oracle IPSec VPN tunnels: `crypto map <map name> <sequence number> set connection-type originate-only`

### IPSec tunnel is UP but keeps flapping

Check these items:

- **Initiation of connection:** Ensure that your CPE device is initiating the connection.
- **Local and remote proxy IDs:** If you're using a policy-based configuration, check if your CPE is configured with more than one pair of local and remote proxy IDs (subnets). The Oracle VPN router supports only one pair. If your CPE has more than one pair, update the configuration to include only one pair, and choose one of the following two options:

Option	Local Proxy ID	Remote Proxy ID
1	ANY (or 0.0.0.0/0)	ANY (or 0.0.0.0/0)
2	On-premises CIDR (an aggregate that covers all the subnets of interest)	VCN's CIDR

- **Interesting traffic at all times:** In general, Oracle recommends having interesting traffic running through the IPSec tunnels at all times if your CPE supports it.

Certain Cisco ASA versions require the SLA monitor to be configured, which keeps interesting traffic running through the IPsec tunnels. For more information, see the section for "IP SLA Configuration" in the [Cisco ASA policy-based configuration template](#).

### IPsec tunnel is UP but traffic is unsteady

Check these items:

- **Local and remote proxy IDs:** If you're using a policy-based configuration, check if your CPE is configured with more than one pair of local and remote proxy IDs (subnets). The Oracle VPN router supports only one pair. If your CPE has more than one pair, update the configuration to include only one pair, and choose one of the following two options:

Option	Local Proxy ID	Remote Proxy ID
1	ANY (or 0.0.0.0/0)	ANY (or 0.0.0.0/0)
2	On-premises CIDR (an aggregate that covers all the subnets of interest)	VCN's CIDR

- **Interesting traffic at all times:** In general, Oracle recommends having interesting traffic running through the IPsec tunnels at all times if your CPE supports it. Certain Cisco ASA versions require the SLA monitor to be configured, which keeps interesting traffic running through the IPsec tunnels. For more information, see the section for "IP SLA Configuration" in the [Cisco ASA policy-based configuration template](#).

## BGP Session Troubleshooting

### BGP status is DOWN

Check these items:

- **IPSec status:** For the BGP session to be up, the IPSec tunnel itself must be up.
- **BGP address:** Verify that both ends of the tunnel are configured with the correct BGP peering IP address.
- **ASN:** Verify that both ends of the tunnel are configured with the correct BGP local ASN and Oracle BGP ASN. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- **MD5:** Verify that MD5 authentication is disabled or not configured on your CPE device. The Oracle IPSec VPN does not support MD5 authentication.
- **Firewalls:** Verify that your on-premises firewall or access control lists are not blocking the following ports:
  - TCP port 179 (BGP)
  - UDP port 500 (IKE)
  - IP protocol port 50 (ESP)

If your CPE device's firewall is blocking TCP port 179 (BGP), the BGP neighborhood will be down. Traffic cannot flow through the tunnel because the CPE device and Oracle router do not have any routes.

### BGP status is flapping

Check these items:

- **IPSec status:** For the BGP session to be up and not flapping, the IPSec tunnel itself must be up and not flapping.
- **Maximum prefixes:** Verify that you are advertising no more than 2000 prefixes. If you're advertising more, BGP won't be established.

### BGP status is UP, but no traffic is passing through

Check these items:

- **VCN security lists:** Ensure you've set up the [VCN security lists](#) to allow the desired traffic (both ingress and egress rules). Note that the VCN's [default security list](#) does not allow ping traffic (ICMP type 8 and ICMP type 0). You must add the appropriate ingress and egress rules to allow ping traffic.
- **Correct routes on both ends:** Verify that you have received the correct VCN routes from Oracle and the CPE device is using those routes. Likewise, verify that you are advertising the correct on-premises network routes over the IPSec VPN, and the VCN route tables use those routes.

### BGP status is UP, but traffic is passing in only one direction

Check these items:

- **VCN security lists:** Ensure that your [VCN security lists](#) allow traffic in both directions (ingress and egress).
- **Firewalls:** Verify that your on-premises firewall or access control lists are not blocking traffic to or from the Oracle end.
- **Asymmetric routing:** Oracle uses asymmetric routing. If you have multiple IPSec connections, make sure that your CPE device is configured for asymmetric route processing.
- **Redundant connections:** If you have redundant IPSec connections, make sure they're both advertising the same routes.

### Redundant Connections

Remember these important notes:

- FastConnect uses BGP dynamic routing. IPSec connections can use either static routing or BGP, or a combination.
- For important details about routing and preferred routes when using redundant connections, see [Route Advertisements and Path Preferences When You Have Multiple Connections](#).

- You can use two IPSec connections for redundancy. If both IPSec connections have only a default route (0.0.0.0/0) configured, traffic will route to either of those connections because Oracle uses asymmetric routing. If you want one IPSec connection as primary and another one as backup, configure more-specific routes for the primary connection and less-specific routes (or the default route of 0.0.0.0/0) on the backup connection.

### IPSec and FastConnect are both set up, but traffic is only passing through IPSec

Make sure to use more specific routes for the connection you want as primary. If you're using the same routes for both IPSec and FastConnect, see the discussion of routing preferences in [Route Advertisements and Path Preferences When You Have Multiple Connections](#).

### Two on-premises data centers each have an IPSec connection to Oracle, but only one is passing traffic

Verify that both IPSec connections are up and ensure that you have asymmetric route processing enabled on the CPE.

If both IPSec connections have only a default route (0.0.0.0/0) configured, traffic will route to either of those connections because Oracle uses asymmetric routing. If you want one IPSec connection as primary and another one as backup, configure more-specific routes for the primary connection and less-specific routes (or the default route of 0.0.0.0/0) on the backup connection.

For more information about this type of setup, see [Example Layout with Multiple Geographic Areas](#).

## FastConnect

The following topics have information about setting up Oracle Cloud Infrastructure FastConnect between your on-premises network and virtual cloud network (VCN):

- [FastConnect Overview](#)
- [FastConnect Requirements](#)
- [FastConnect Redundancy Best Practices](#)
- [Routing Details for Connections to Your On-Premises Network](#)
- [FastConnect: With an Oracle Provider](#)
- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)
- [FastConnect Metrics](#)
- [FastConnect Troubleshooting](#)

## FastConnect Overview

Oracle Cloud Infrastructure FastConnect provides an easy way to create a dedicated, private connection between your data center and Oracle Cloud Infrastructure. FastConnect provides higher-bandwidth options, and a more reliable and consistent networking experience compared to internet-based connections.

### Uses for FastConnect

With FastConnect, you can choose to use *private peering*, *public peering*, or both.

- **Private peering:** To extend your existing infrastructure into a virtual cloud network (VCN) in Oracle Cloud Infrastructure (for example, to implement a hybrid cloud, or a lift and shift scenario). Communication across the connection is with IPv4 private addresses (typically RFC 1918).
- **Public peering:** To access public services in Oracle Cloud Infrastructure without using the internet. For example, Object Storage, the Oracle Cloud Infrastructure Console and APIs, or public load balancers in your VCN. Communication across the connection is with IPv4 public IP addresses. Without FastConnect, the traffic destined for public IP addresses would be routed over the internet. With FastConnect, that traffic goes over your private physical connection. For a list of the services available with public peering,

see [FastConnect Supported Cloud Services](#). For a list of the public IP address ranges (routes) that Oracle advertises, see [FastConnect Public Peering Advertised Routes](#).

In general it's assumed you'll use private peering, and you *might* also use public peering. Most of this documentation is relevant to both, with specific details called out for private versus public.

If you choose to have multiple paths from your on-premises network to Oracle, see [Routing Details for Connections to Your On-Premises Network](#).

IPv6 addressing is currently supported only in the US [Government Cloud](#). For more information, see [IPv6 Addresses](#).

### How and Where to Connect

With FastConnect, there are different [connectivity models](#) to choose from.

#### Oracle Provider

- [List of Oracle Cloud Infrastructure FastConnect providers](#)
- Port speeds in 1-Gbps and 10-Gbps increments
- Instructions for integrating: [FastConnect: With an Oracle Provider](#)

#### Third-Party Provider

- Port speed of 1 Gbps or 10 Gbps per cross-connect
- Instructions for integrating: [FastConnect: With a Third-Party Provider](#)

#### Colocation with Oracle in an Oracle Cloud Infrastructure FastConnect Location

- [List of Oracle Cloud Infrastructure FastConnect locations](#) (see the [FAQ for specific addresses](#))
- Port speed of 1 Gbps or 10 Gbps per cross-connect
- Instructions for integrating: [FastConnect: Colocation with Oracle](#)

## CHAPTER 23 Networking

The following table summarizes several important requirements for each connectivity model.

<b>Requirement</b>	<b>With Oracle Provider</b>	<b>With Third-Party Provider</b>	<b>Colocation with Oracle</b>
Routing requirements	<a href="#">Yes</a>	<a href="#">Yes</a>	<a href="#">Yes</a>
BGP support	<a href="#">Yes</a>	<a href="#">Yes</a>	<a href="#">Yes</a>
Layer 3 support	<a href="#">Recommended</a>	<a href="#">Recommended</a>	<a href="#">Recommended</a>
Obtain a Letter of Authority (LOA) from Oracle	N/A	<a href="#">Yes</a>	Yes
Network connectivity	<a href="#">Yes</a>	<a href="#">Yes</a>	N/A
Cross-connect	<a href="#">Yes (from the provider)</a>	<a href="#">Yes</a>	<a href="#">Yes</a>
Redundant network connectivity	<a href="#">Recommended</a>	<a href="#">Recommended</a>	<a href="#">Recommended</a>
Cloud connectivity solution architecture support	Recommended	Recommended	Recommended
FastConnect SKU	Yes	Yes	Yes
Oracle Cloud Infrastructure Console user login (IAM policy unique setup)	<a href="#">Yes</a>	<a href="#">Yes</a>	<a href="#">Yes</a>
Tenancy established (see "Setting Up Your Tenancy" in the <i>Oracle Cloud Infrastructure Getting Started Guide</i> )	Yes	Yes	Yes

### Concepts

Here are some important concepts to understand (also see the following diagrams):

### **FASTCONNECT**

The general concept of a connection between your existing network and Oracle Cloud Infrastructure over a private physical network instead of the internet.

### **FASTCONNECT LOCATION**

A specific Oracle data center where you can connect to Oracle Cloud Infrastructure.

### **METRO AREA**

A geographical area (for example, Ashburn) with multiple FastConnect locations. All locations in a metro area connect to the same set of availability domains for resiliency in case of failure in a single location.

### **ORACLE PROVIDER**

A network service provider that has integrated with Oracle in a FastConnect location. See the list of the Oracle providers in [How and Where to Connect](#). If your provider is in the list, see [FastConnect: With an Oracle Provider](#).

### **THIRD-PARTY PROVIDER**

A network service provider that is NOT on the list of Oracle providers in [How and Where to Connect](#). If you have a third-party provider and want to use FastConnect, see [FastConnect: With a Third-Party Provider](#).

### **COLOCATION**

The situation where your equipment is deployed into a FastConnect location. If your network service provider is not on the list of Oracle providers in [How and Where to Connect](#), you must colocate.

### **CROSS-CONNECT**

In a colocation or third-party provider scenario, this is the physical cable connecting your existing network to Oracle in the FastConnect location.

### **CROSS-CONNECT GROUP**

In a colocation or third-party provider scenario, this is a link aggregation group (LAG) that contains at least one cross-connect. You can add additional cross-connects to a cross-connect group as your bandwidth needs increase. This is applicable only for colocation.

### **VIRTUAL CLOUD NETWORK (VCN)**

Your virtual network in Oracle Cloud Infrastructure. You can use a VCN to extend your infrastructure into the cloud. For more information, see [VCNs and Subnets](#).

### **DYNAMIC ROUTING GATEWAY (DRG)**

A virtual edge router attached to your VCN. Necessary for private peering. The DRG is a single point of entry for private traffic coming in to your VCN, whether it's over FastConnect or an [IPSec VPN](#). After creating the DRG, you must attach it to your VCN and add a route for the DRG in the VCN's route table to enable traffic flow. Instructions for everything are included in the sections that follow.

### **VIRTUAL CIRCUIT**

An isolated network path that runs over one or more physical network connections to provide a single, logical connection between the edge of your existing network and Oracle Cloud Infrastructure. *Private virtual circuits* support private peering, and *public virtual circuits* support public peering (see [Uses for FastConnect](#)). Each virtual circuit is made up of information shared between you and Oracle, as well as a provider (if you're connecting through an Oracle provider). You could have multiple private virtual circuits, for example, to isolate traffic from different parts of your organization (one virtual circuit for 10.0.1.0/24; another for 172.16.0.0/16), or to provide redundancy.

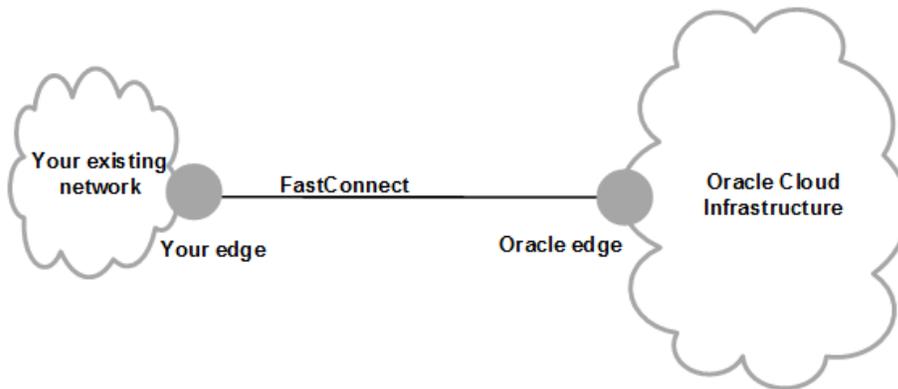
## Basic Network Diagrams

The diagrams in this section introduce the basic logical and physical connections involved in FastConnect. Details specific to private versus public peering are called out.

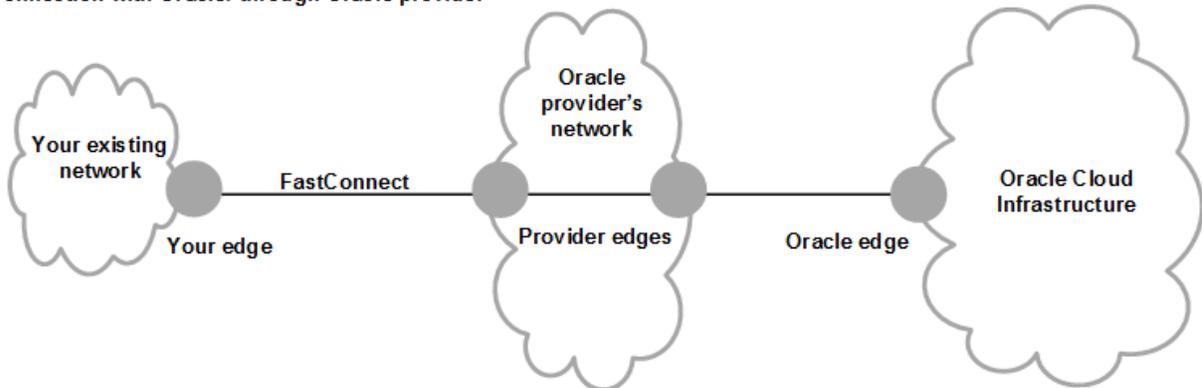
### General Concept of FastConnect

The following diagram illustrates the two ways to connect to Oracle with FastConnect: either through colocation with Oracle in the FastConnect location, or through an Oracle provider. In both cases, the connection goes between the edge of your existing network and Oracle.

Connection with Oracle: colocation in data center



Connection with Oracle: through Oracle provider



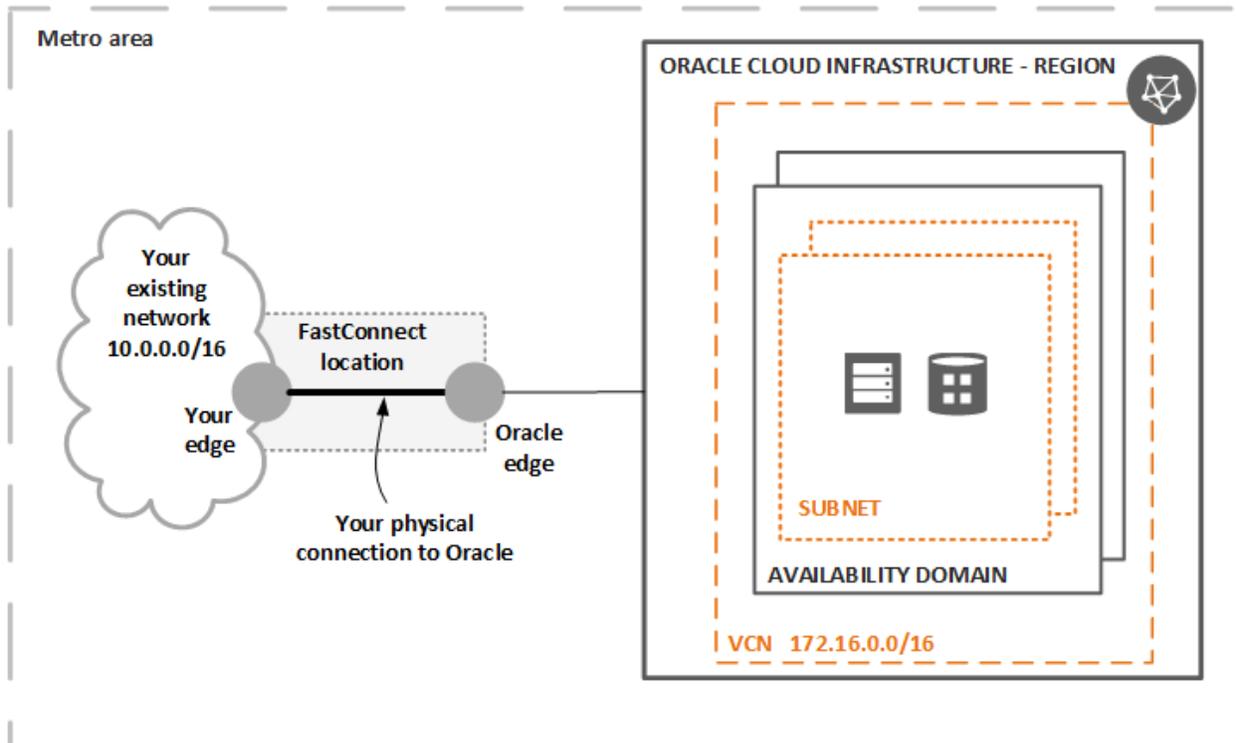
### Physical Connection

The next two diagrams give more detail about the physical connections. They also show the metro area that contains the FastConnect location, and a VCN within an Oracle Cloud

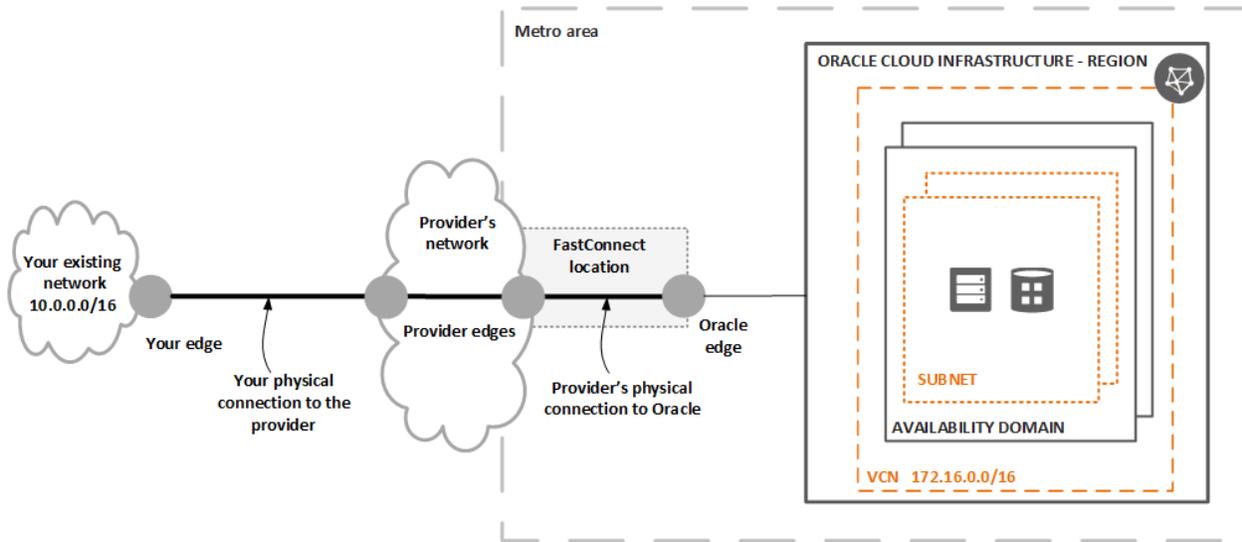
## CHAPTER 23 Networking

Infrastructure region.

The first diagram shows the colocation scenario, with your physical connection to Oracle within the FastConnect location.



The next diagram shows a scenario with either an Oracle provider or third-party provider. It shows your physical connection to the provider, and the provider's physical connection to Oracle within the FastConnect location.

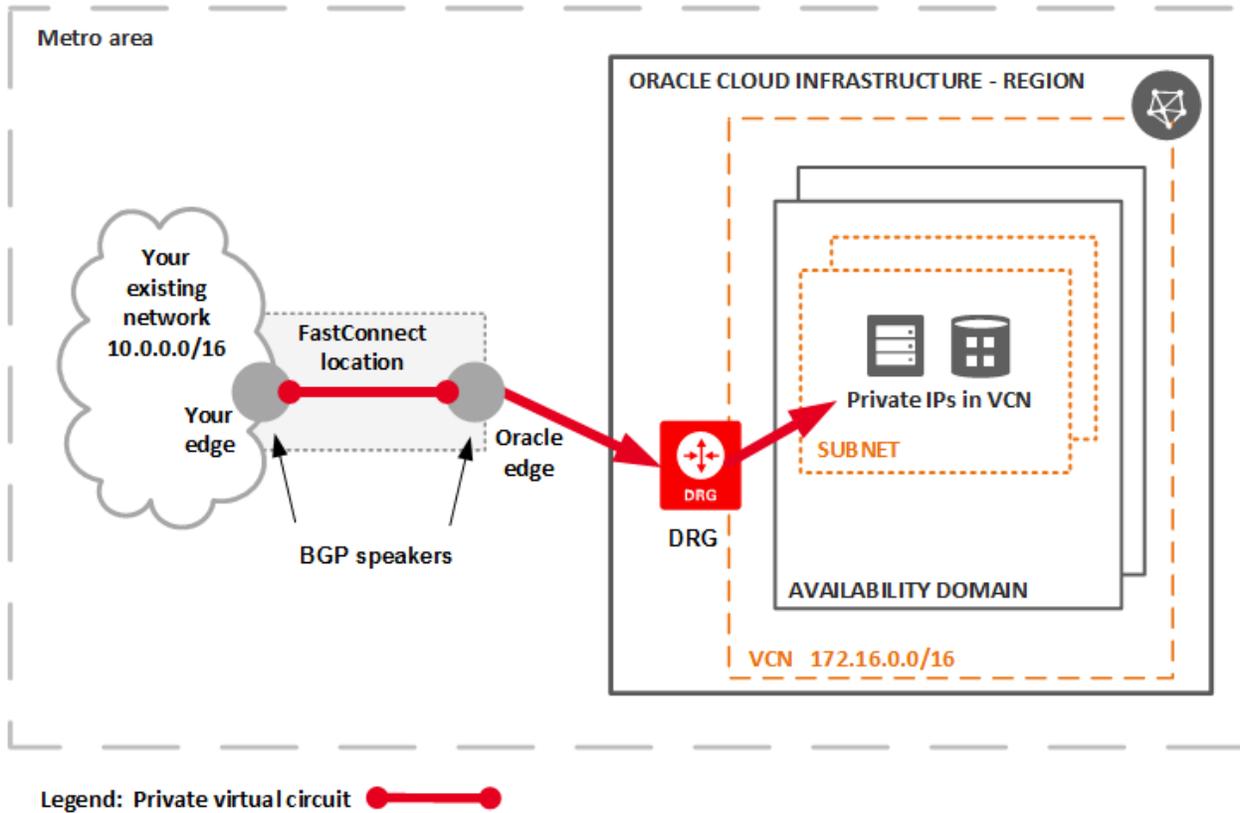


**Logical Connection: Private Virtual Circuit**

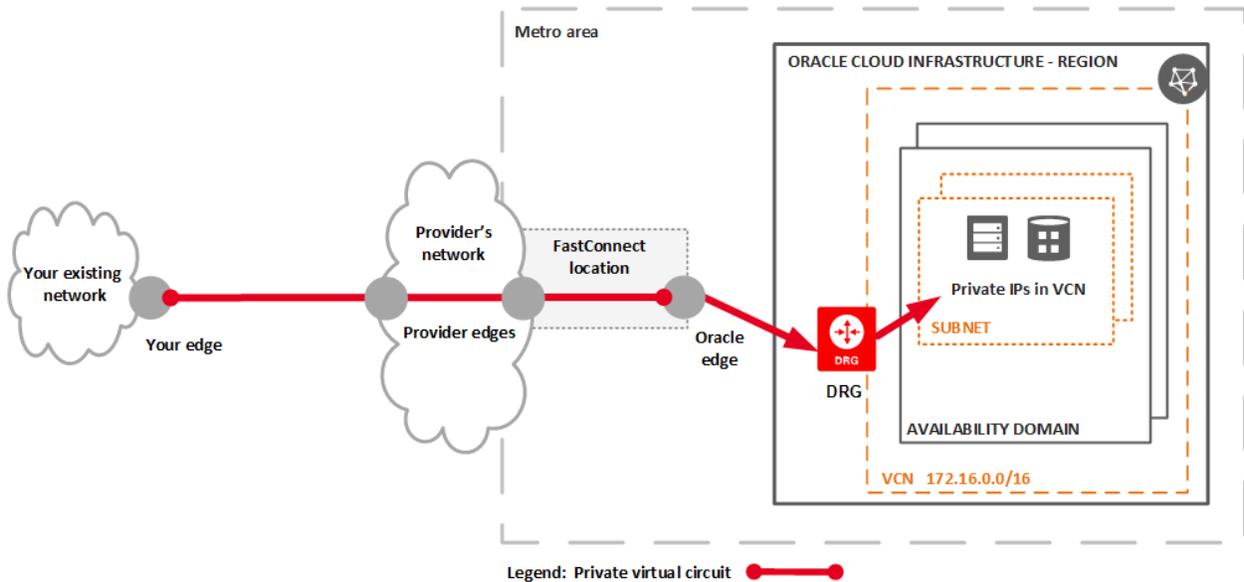
The next two diagrams show a private virtual circuit, which is a single, logical connection between your edge and Oracle Cloud Infrastructure by way of your DRG. Traffic is destined for private IP addresses in your VCN.

## CHAPTER 23 Networking

The first diagram shows the colocation scenario.



The next diagram shows the scenario with either an Oracle provider or third-party provider.

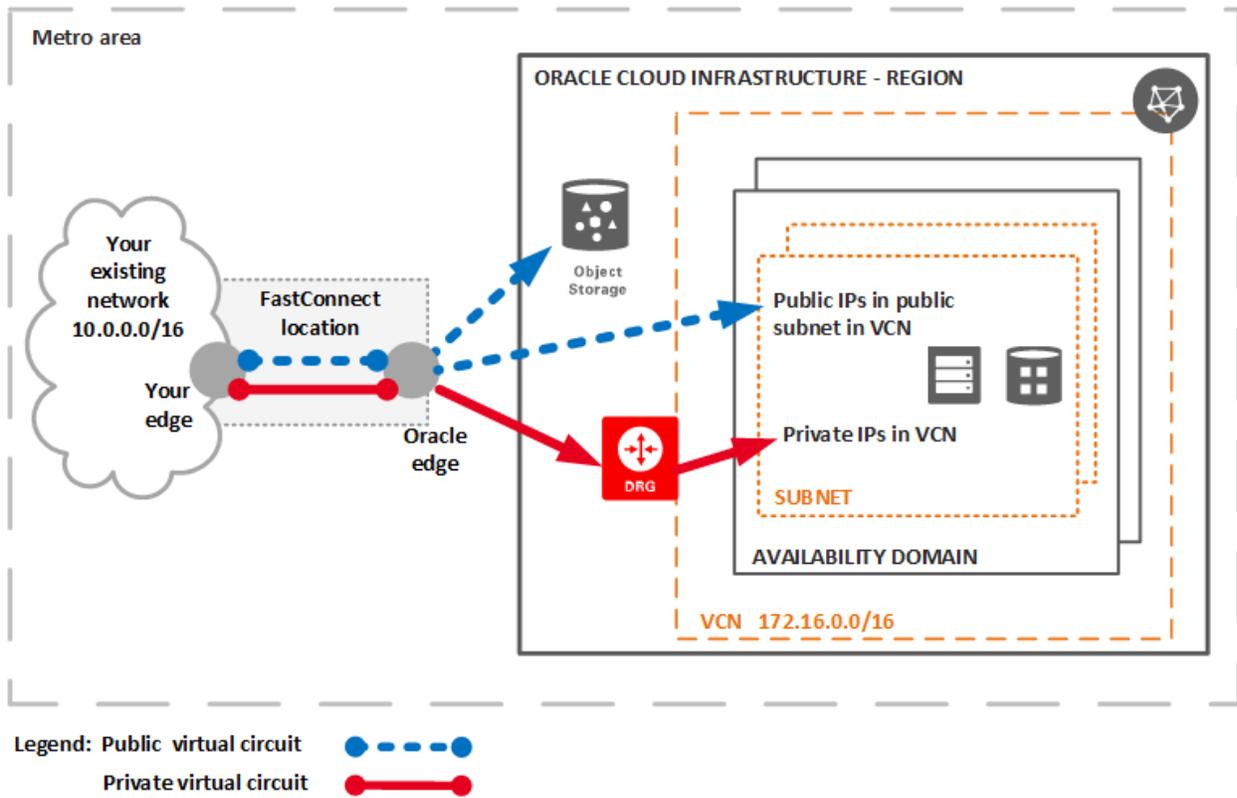


### Logical Connection: Public Virtual Circuit

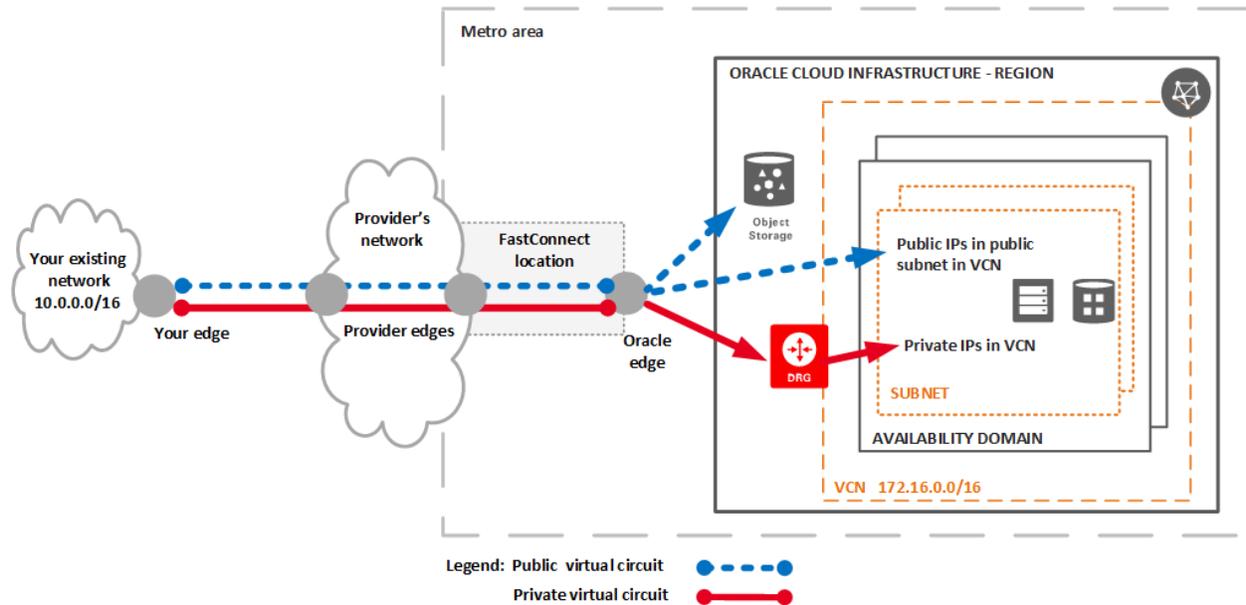
A public virtual circuit gives your existing network access to Oracle services in Oracle Cloud Infrastructure. For example, Object Storage, the Oracle Cloud Infrastructure Console and APIs, and public load balancers in your VCN. All communication across a public virtual circuit uses public IP addresses. For a list of services available with FastConnect public peering, see [FastConnect Supported Cloud Services](#). For a list of the public IP address ranges (routes) that Oracle advertises, see [FastConnect Public Peering Advertised Routes](#).

## CHAPTER 23 Networking

The first diagram shows the colocation scenario with both a private virtual circuit and a public virtual circuit. Notice that the DRG is not involved with the public virtual circuit, only the private virtual circuit.



The next diagram shows the scenario with either an Oracle provider or third-party provider.



Here are a few basics to know about public virtual circuits:

- You choose which of your organization's public IP prefixes you want to use with the virtual circuit. Each prefix must be /31 or less specific. Oracle verifies your organization's ownership of each prefix before sending any traffic for it across the connection. Oracle's verification for a given prefix can take up to three business days. You can get the status of the verification of each prefix in the Oracle Console or API. **Oracle begins advertising the Oracle Cloud Infrastructure public IP addresses across the connection only after successfully verifying at least one of your public prefixes.**
- You must configure your firewall rules to allow the desired traffic coming from the [Oracle public IP addresses](#).
- Your existing network can receive advertisements for Oracle's public IP addresses through multiple paths (for example: FastConnect and your internet service provider).

Make sure to give higher preference to FastConnect over your ISP. You must configure your edge appropriately so that traffic uses your desired path and you receive the benefits of FastConnect. This is particularly important if you decide to *also* set up your existing network with [private access to Oracle services](#). For important information about path preferences, see [Routing Details for Connections to Your On-Premises Network](#).

- You can add or remove public IP prefixes at any time by editing the virtual circuit. If you add a new prefix, Oracle first verifies your company's ownership before advertising it across the connection. If you remove a prefix, Oracle stops advertising the prefix within a few minutes of your editing the virtual circuit.

### **Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider**

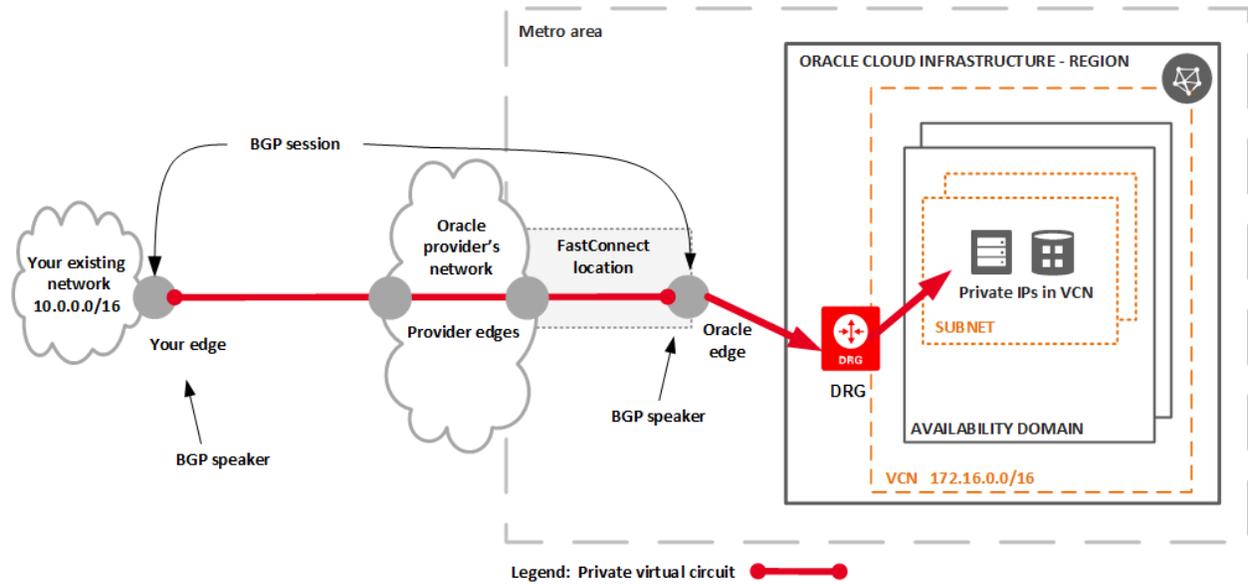
This section is applicable if you're using FastConnect through an Oracle provider. A Border Gateway Protocol (BGP) session is established from your edge, but where it goes to depends on which Oracle provider you use.



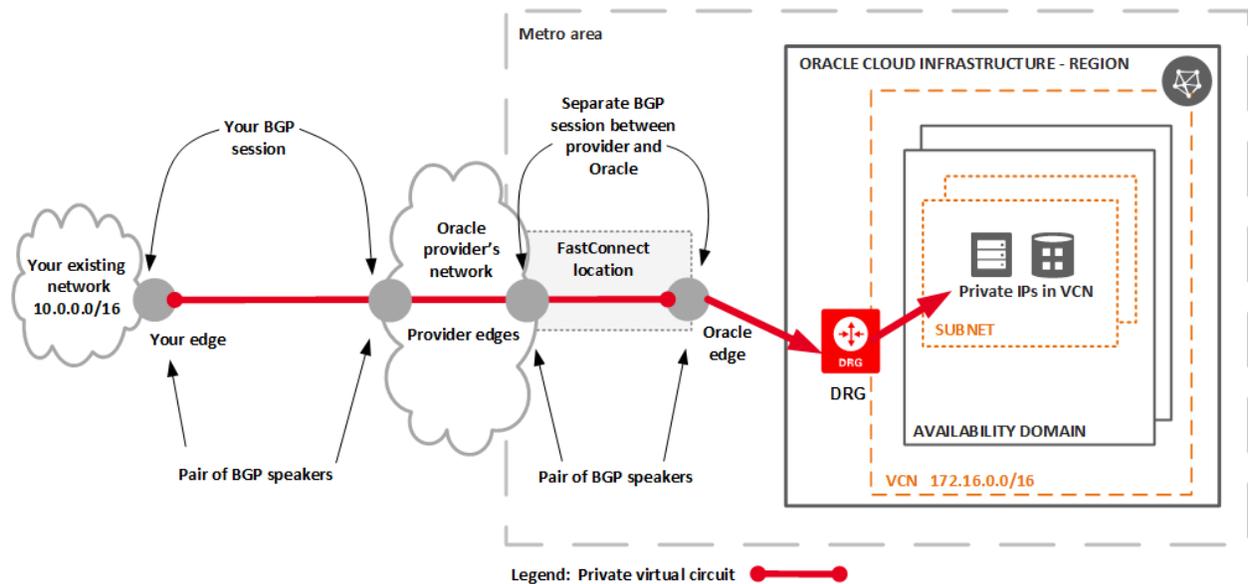
#### **Tip**

For simplicity, the following diagrams show only a private virtual circuit. However, the location of the BGP session is the same for a public virtual circuit.

**To Oracle:** With some of the Oracle providers, the BGP session goes from your edge to Oracle, as shown in the following diagram. When setting up the virtual circuit with Oracle, you are asked to provide basic BGP peering information (see [General Requirements](#)).



**To the Oracle provider:** With other Oracle providers, your BGP session goes from your edge to the provider's, as shown in the following diagram. When setting up the virtual circuit with Oracle, you are NOT asked for any BGP session information. Instead, you share BGP information with your Oracle provider. Notice that there's a separate BGP session that the provider establishes with Oracle.



### FastConnect with Access to Multiple VCNs

You can use a single FastConnect to access multiple VCNs. There are different network scenarios available depending on your needs and which FastConnect connectivity model you use. For more information, see these topics:

- [Transit Routing: Access to Multiple VCNs in the Same Region](#): This scenario can be used with either FastConnect or VPN Connect. It involves a single DRG, and multiple VCNs in a hub-and-spoke layout.

- [FastConnect with Multiple DRGs and VCNs](#): This scenario can be used only with FastConnect, and only if you're using a third-party provider or colocated with Oracle. It involves multiple DRGs and private virtual circuits.

### What's Next?

See these topics to get started:

- [FastConnect Requirements](#)
- [FastConnect Redundancy Best Practices](#)
- [Routing Details for Connections to Your On-Premises Network](#)
- [FastConnect: With an Oracle Provider](#)
- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)

## FastConnect Requirements

This topic covers the requirements for implementing FastConnect.

For general information about FastConnect, see [FastConnect](#).

### Before Getting Started: Learn and Plan

Here are basic things you need to do before getting started with FastConnect:

- **FastConnect concepts:** Make sure you're familiar with the basic concepts covered in [FastConnect](#).
- **Limits increase:** If you are colocated with Oracle, you must ask Oracle to increase your account limits for cross-connects. By default, these limits are initially set to 0, and a request to create a cross-connect will fail. For instructions, see [Requesting a Service Limit Increase](#). In your request, make sure to indicate the region where you need the resources. It can take a couple of business days for the limit increase to take effect.
- **Hardware or routing requirements:** Review the [requirements](#).

- **Tenancy setup and compartment design:** If you haven't yet, set up your tenancy. Think about who needs access to Oracle Cloud Infrastructure and how. For more information, see "Setting Up Your Tenancy" in the *Oracle Cloud Infrastructure Getting Started Guide*. Specifically for FastConnect, see [Required IAM Policy](#) to understand the policy required to work with FastConnect components.
- **Cloud network design:** Design your virtual cloud network (VCN), including how you want to allocate your VCN's [subnets](#), define [security list rules](#), define [route rules](#), set up [load balancers](#), and so on. For more information, see [Overview of Networking](#).
- **Redundancy:** Think through your overall redundancy model to ensure your network can handle planned maintenance by Oracle or your organization, and unexpected failures of the various components. For best practices, see [FastConnect Redundancy Best Practices](#).
- **Public IP prefixes:** If you plan to set up a public virtual circuit, get the list of the public IP prefixes that you want to use with the connection. Oracle must validate your organization's ownership of each of the prefixes before advertising each one over the connection.
- **Cloud network setup:** Set up your VCN, subnets, DRG, security lists, IAM policies, and so on, according to your design.

### General Requirements

Before getting started with FastConnect, make sure you meet the following requirements:

- Oracle Cloud Infrastructure account, with at least one user with appropriate Oracle Cloud Infrastructure Identity and Access Management (IAM) permissions (for example, a user in the Administrators group).
- Network equipment that supports Layer 3 routing using BGP.
- For colocation with Oracle: Ability to connect using single mode fiber in your selected FastConnect location. Also see [Hardware and Routing Requirements](#).
- For connection to an Oracle provider: At least one physical network connection with the provider. Also see [Hardware and Routing Requirements](#).

- For connection to a third-party provider: At least one physical connection with the provider. Also see [Hardware and Routing Requirements](#).
- For private peering only: At least one existing [DRG](#) set up for your VCN.
- For public peering only: The list of public IP address prefixes that you want to use with the connection. Oracle will validate your ownership of each prefix.



### Important

If you're colocating with Oracle, you must ask Oracle to increase your account limits for cross-connects. By default, these limits are initially set to 0, and a request to create one of these resources will fail. For instructions, see [Requesting a Service Limit Increase](#). In your request, make sure to **indicate the region where you need the resources**. It can take a couple of business days for the limit increase to take effect.

## Hardware and Routing Requirements

### If you're using an Oracle provider

Here are general routing requirements for FastConnect. These are particularly relevant if your BGP session is [between your edge and Oracle](#).

- **IP addressing supported:** IPv4. IPv6 addressing is currently supported only in the US [Government Cloud](#). For more information, see [IPv6 Addresses](#).
- **P2P IP addresses:**
  - For public virtual circuits, Oracle specifies the IP addresses.
  - For private virtual circuits where your BGP session is [between your edge and Oracle](#), you specify these addresses (/30 or /31, and one pair per virtual circuit).

If you set up multiple private virtual circuits that go to the same DRG, you must use a different address on your edge router for each virtual circuit.

- **Maximum IP MTU:** 9000
- **Routing protocol:** BGPv4
- **BGP prefix limit:** For public virtual circuits: 200 prefixes. For private virtual circuits: 2000 prefixes.
- **BGP ASN:** 2-byte or 4-byte ASNs are supported, except for those listed in [Special-Purpose Autonomous System \(AS\) Numbers](#). Public virtual circuits require a public ASN. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- **BGP MD5 authentication:** Optional to use with your virtual circuit. Oracle supports up to 128-bit MD5 authentication
- **BGP keep-alive interval:** 60s
- **BGP hold-time interval:** 180s



### Tip

By default, Oracle uses the default BGP timers of 60 seconds for keep-alive and 180 seconds for hold-time. If you need fast BGP convergence, you can use any value in these supported ranges: 6-60 seconds for keep-alive, and 18-180 seconds for hold-time.

If you're colocating in an FastConnect location or using a third-party provider

For the cross-connect group and cross-connects:

- **Bandwidth (two choices):**
  - 1 Gbps:
    - 1000Base-LX, 10 km range, 1310 nm optics
    - **You must configure your edge device so that auto-negotiation is OFF**
  - 10 Gbps:
    - 10 GbE, LR (10 km range), 1310 nm optics
- **General:**
  - Single Mode Fiber
  - Duplex LC connectors
  - Minimum Rx level > 12 dBm
- **Redundancy:**
  - Device redundancy highly recommended
  - In some regions, location redundancy is available and recommended
- **Capacity:**
  - Minimum: 1 x 1 GbE or 1 x 10 GbE
  - Maximum: 8 x 1 GbE or 8 x 10 GbE
- **LAG protocol:** LACP with short timers (3 @ 1s). If your router doesn't support LAG, you can set up a single non-LAG cross-connect.
- **VLAN tagging:** 802.1q (single tag)
- **VLAN range:** 100-4094 (you assign the VLANs)
- **Maximum interface MTU:** 9196 (include 4-byte FCS trailer)

For routing:

- **IP addressing supported:** IPv4. IPv6 addressing is currently supported only in the US [Government Cloud](#). For more information, see [IPv6 Addresses](#).
- **P2P IP addresses:**
  - For public virtual circuits, Oracle specifies the IP addresses.
  - For private virtual circuits, you specify the addresses (a /30 or /31). You need one pair of IP addresses per private virtual circuit. If you set up multiple private virtual circuits that go to the same DRG, you must use a different address on your edge router for each virtual circuit.
- **Maximum IP MTU:** 9000
- **Routing protocol:** BGPv4
- **BGP prefix limit:** For public virtual circuits: 200 prefixes. For private virtual circuits: 2000 prefixes.
- **BGP ASN:** 2-byte or 4-byte ASNs are supported, except for those listed in [Special-Purpose Autonomous System \(AS\) Numbers](#). Public virtual circuits require a public ASN. Oracle's BGP ASN is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- **BGP MD5 authentication:** Optional to use with your virtual circuit. Oracle supports up to 128-bit MD5 authentication
- **BGP keep-alive interval:** 60s
- **BGP hold-time interval:** 180s



### Tip

By default, Oracle uses the default BGP timers of 60 seconds for keep-alive and 180 seconds for hold-time. If you need fast BGP convergence, you can use any value in these supported ranges: 6-60 seconds for keep-alive, and 18-180 seconds for hold-time.

### Required IAM Policy

#### If you're using an Oracle provider

To work with Networking resources such as dynamic routing gateways (DRGs), VCNs, and virtual circuits, you need to have a user login to the Console, and your user needs sufficient authority (by way of an [IAM policy](#)) to perform all the instructions below. If your user is in the [Administrators group](#), you have the required authority.

If your user is not, then a policy like this would generally cover all the Networking resources:

```
Allow group NetworkAdmins to manage virtual-network-family in tenancy
```

To *only* create and manage a virtual circuit, you would need a policy like this:

```
Allow group VirtualCircuitAdmins to manage drgs in tenancy
```

```
Allow group VirtualCircuitAdmins to manage virtual-circuits in tenancy
```

The first statement (to manage DRGs) is necessary only for private virtual circuits.

For more information, see [Getting Started with Policies](#) and [Common Policies](#).

#### If you're colocating in a FastConnect location or using a third-party provider

To work with Networking resources such as dynamic routing gateways (DRGs), VCNs, virtual circuits, and cross-connects, you need to have a user login to the Console, and your user needs sufficient authority (by way of an [IAM policy](#)) to perform all the instructions that follow. If your user is in the [Administrators group](#), you have the required authority.

If your user is not, then a policy like this would generally cover all the Networking resources:

```
Allow group NetworkAdmins to manage virtual-network-family in tenancy
```

To *only* create and manage cross-connects, cross-connect groups, and virtual circuits, you would need a policy like this:

```
Allow group FastConnectAdmins to manage drgs in tenancy
Allow group FastConnectAdmins to manage cross-connects in tenancy
Allow group FastConnectAdmins to manage cross-connect-groups in tenancy
Allow group FastConnectAdmins to manage virtual-circuits in tenancy
```

The first statement (to manage DRGs) is necessary only for private virtual circuits.

For more information, see [Getting Started with Policies](#) and [Common Policies](#).

### Identifiers for Your FastConnect Resources

There are several identifiers for your resources:

- **Name for the overall connection:** When you create a new FastConnect connection, you can give it a descriptive name of your choice. If you don't specify one, Oracle automatically assigns a name to the connection.
- **Reference name for each cross-connect:** Each cross-connect has an optional reference name. If you set up a cross-connect, Oracle recommends that you fill in the reference name with the identifier for the cross-connect's physical fiber cable. That makes it easier for Oracle to help if future troubleshooting is required for the connection. After cabling is done and you have the identifier from the data center, you can add it to the cross-connect's information in the Oracle Console.
- **OCID for each resource:** Each cross-connect group, cross-connect, and virtual circuit has its own unique [Oracle-assigned identifier](#) called an OCID.

### What's Next?

Choose the topic that suits your situation:

- [FastConnect: With an Oracle Provider](#)
- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)

## FastConnect Redundancy Best Practices

This topic covers best practices for redundancy when implementing FastConnect.

For general information about FastConnect, see [FastConnect](#).

### Overview

In general, you should design your network to achieve high availability (HA). In addition, you should be prepared for these types of disruptions:

- Regularly scheduled maintenance by your organization, your provider (if you're using one), or Oracle.
- Unexpected failures on the part of your networking components, your provider, or Oracle. Failures are rare, but you should plan for them.

For redundancy, Oracle provides:

- Multiple providers for each region
- Two FastConnect locations for each of the following regions (all other regions have a single FastConnect location)
  - Germany Central (Frankfurt)
  - UK South (London)
  - US East (Ashburn)
  - US West (Phoenix)

- Two routers in each FastConnect location
- Multiple physical connections between each Oracle provider and Oracle (for a given region)

The redundancy best practices depend on which [connectivity model](#) you use. Also see [How and Where to Connect](#).

### If You Use an Oracle Provider

Connectivity model:

- [FastConnect: With an Oracle Provider](#)

Oracle handles redundancy of the physical connections between the provider and Oracle, and the redundancy of routers in the FastConnect locations. You should then handle redundancy of the physical connection between your existing network and the Oracle provider.

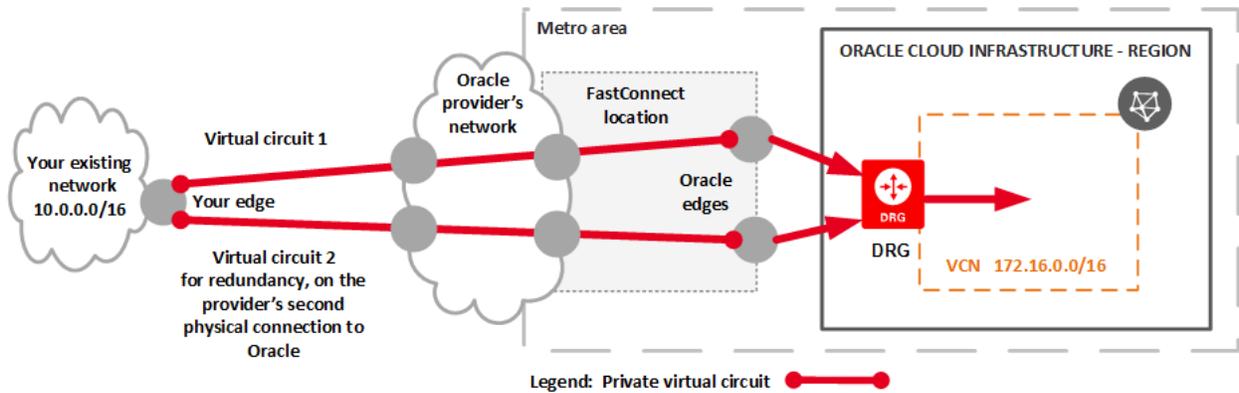
The remaining best practices depend on the provider you're using, and details of the BGP session from your edge:

- For some providers, the BGP session from your edge goes to Oracle. For redundancy best practices, see the next section.
- For other providers, the BGP session from your edge goes to the Oracle provider. For redundancy best practices, see [Oracle Provider Scenario: Your BGP Session Is to the Oracle Provider](#).

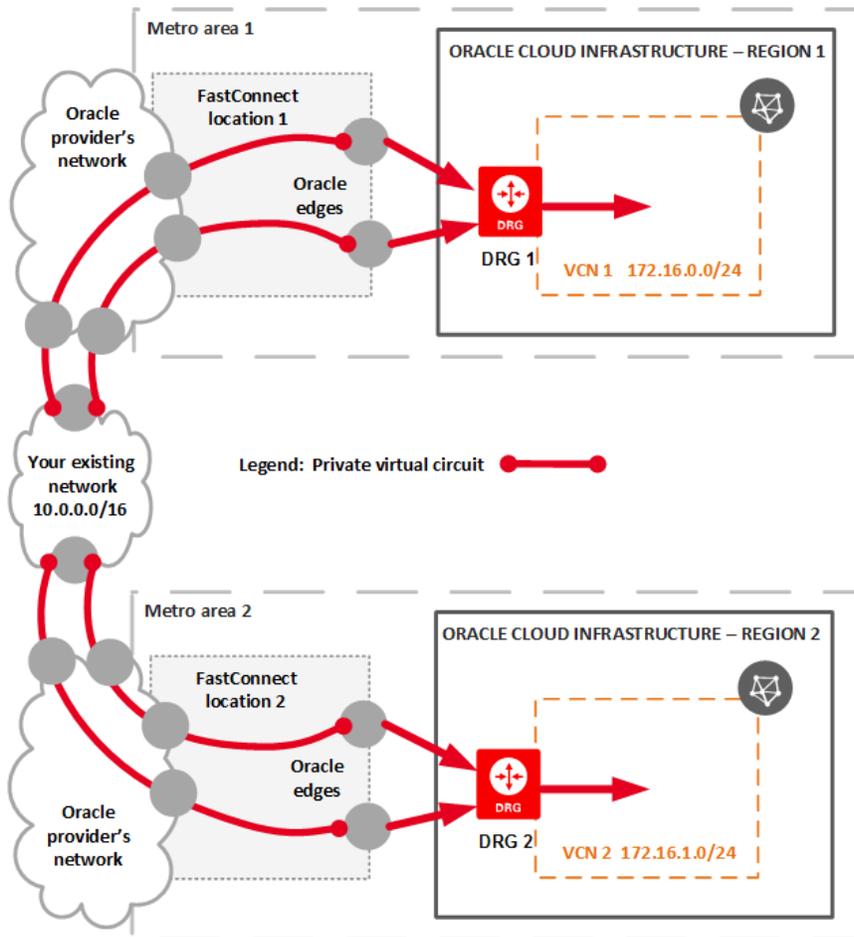
For information about the two scenarios, see [Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider](#).

#### **Oracle Provider Scenario: Your BGP Session Is to Oracle**

At a minimum, each Oracle provider has two separate physical connections to Oracle. Set up one virtual circuit on one physical connection (as primary), and the other virtual circuit on another physical connection (as secondary). The following diagram illustrates two virtual circuits, each going to a different router in a single FastConnect location. If the region has a second location, your provider's second physical connection might instead go to that location.



If you're working in a region that has only a single FastConnect location, you might also want location diversity. To achieve that, repeat the preceding setup of two virtual circuits with the same Oracle provider, but in a second FastConnect location in a nearby region. Notice that you must have a duplicate setup of your Oracle cloud resources in that second region, as shown in the following diagram.



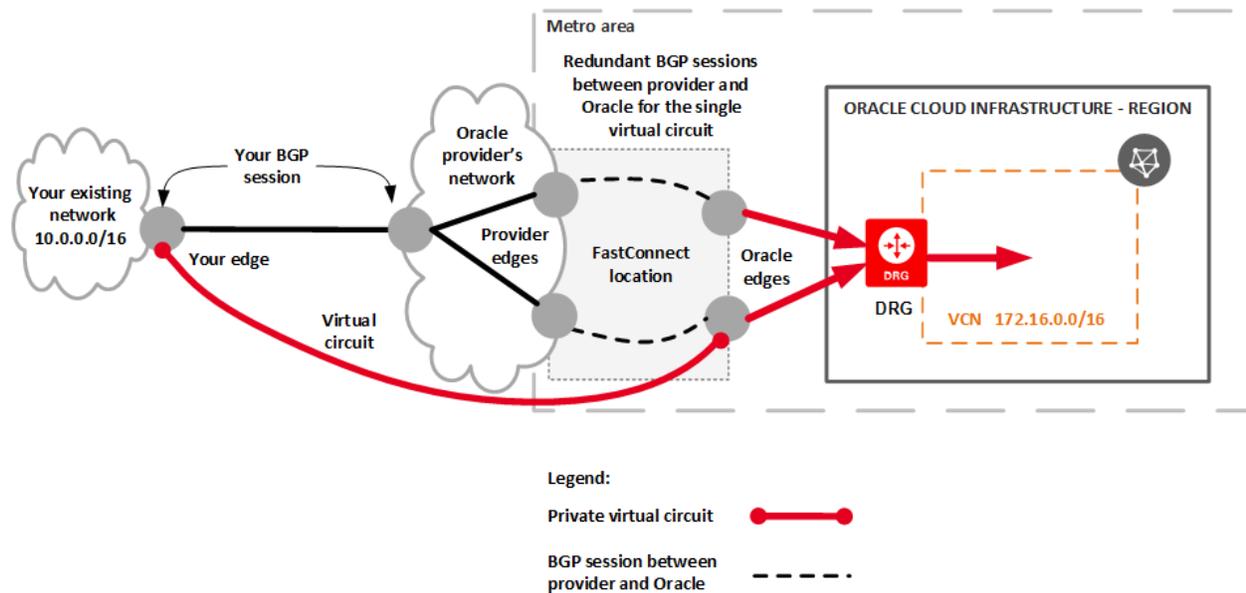
If you also want provider diversity, repeat your entire setup with another provider in each region that you're using.

### Oracle Provider Scenario: Your BGP Session Is to the Oracle Provider

In this scenario, the BGP session *from your edge* goes to the Oracle provider (as shown in the following diagram). Separate from your BGP session, the Oracle provider has *its own* BGP sessions with Oracle (between the provider's edge and Oracle's edge). The virtual circuit is a logical connection that goes from your edge to the Oracle edge.

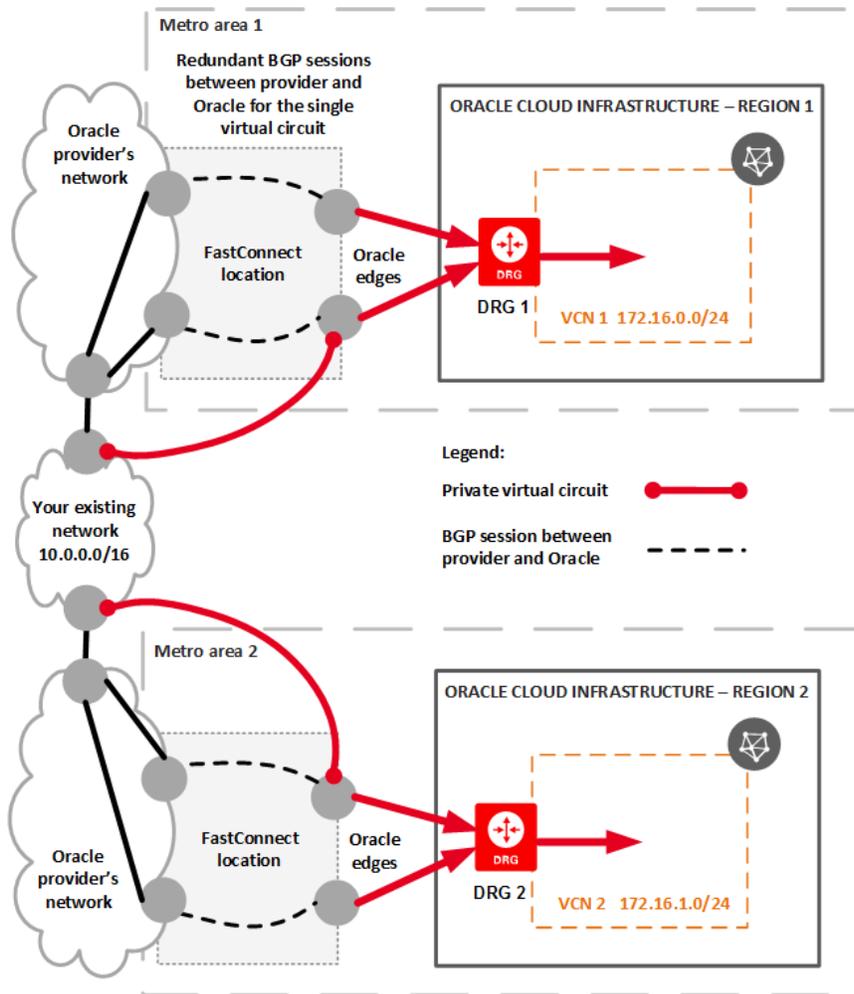
## CHAPTER 23 Networking

The provider has two separate physical connections to Oracle. You create one virtual circuit with the provider. In this scenario, the virtual circuit is automatically designed to be redundant and diverse. The virtual circuit has two separate BGP sessions between the provider and Oracle, each on a different physical connection. The following diagram shows the two separate BGP sessions for the single virtual circuit as dotted lines.



Separately, you must ensure that the connection between your edge and the provider is redundant and diverse.

If you're working in a region that has only a single FastConnect location, you might also want location diversity. To achieve that, repeat the preceding setup of a virtual circuit with the same Oracle provider, but in a second FastConnect location in a nearby region. Notice that you must also have a duplicate setup of your Oracle cloud resources in that second region, as shown in the following diagram.



## If You Use a Third-Party Provider or Colocate with Oracle

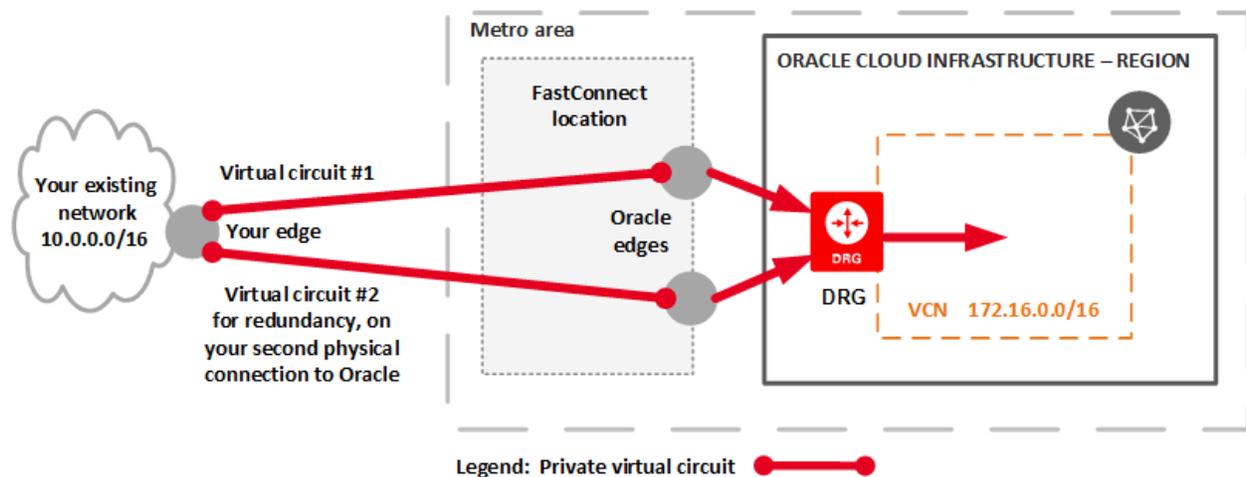
Connectivity models:

- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)

## CHAPTER 23 Networking

Oracle handles redundancy of the Oracle routers in the FastConnect locations. You should then handle redundancy of the physical connection between your existing network and Oracle.

To do this, create two physical connections to Oracle, one for each FastConnect location that serves the region. This means that in the Oracle Console, you set up two separate FastConnect connections. You then create two virtual circuits. Set up the first one on the first physical connection (the first FastConnect connection), and the second one on the second physical connection. The following diagram shows the general setup.



You might prefer to connect to only a single FastConnect location because of cost concerns, or if the region has only a single FastConnect location. In that case, create two physical connections and ensure each goes to a different Oracle router in that FastConnect location. You can do this in the Oracle Console when you set up the second physical connection. You can specify the proximity of that connection to other FastConnect connections in that location. For example, the following image shows how to request that your second physical connection (which is a cross-connect group) is created on a different router than your first connection in that FastConnect location (called MyConnection-1).

## CHAPTER 23 Networking

---

SPECIFY ROUTER PROXIMITY OPTIONAL  
For router redundancy, choose whether you want cross-connects in this group on the same or different router as another one.

CREATE CROSS-CONNECTS IN THIS GROUP ON:

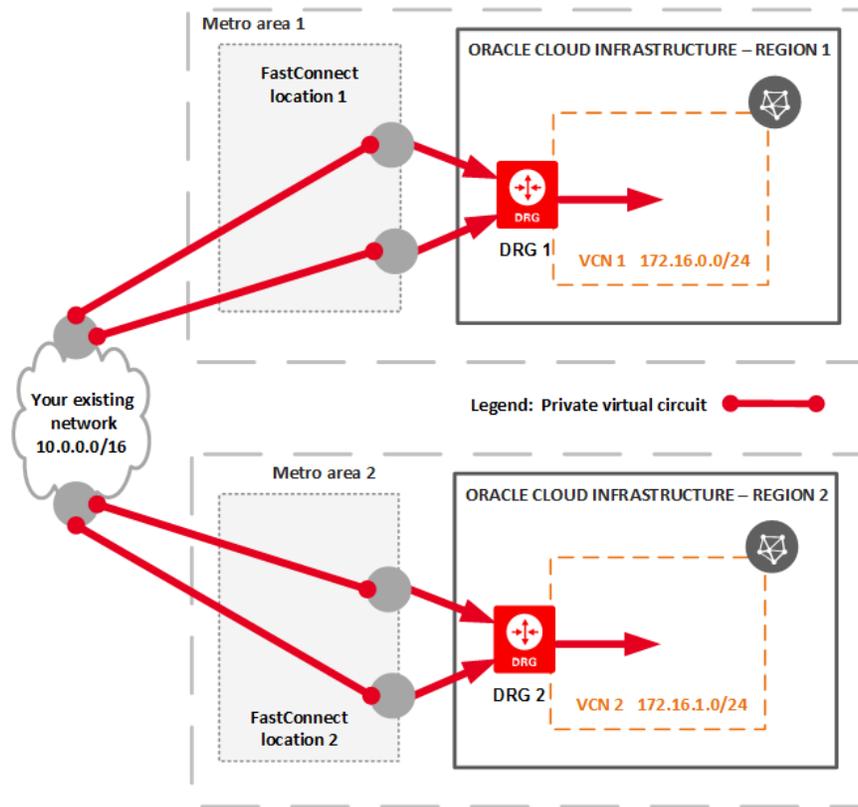
A different router than...

THE ROUTER USED FOR THIS EXISTING CONNECTION:

MyConnection-1

You must scale the bandwidth of both physical connections evenly, and by using a cross-connect group (LAG) for each connection. Imagine that you have two individual 10-Gbps cross-connects in a single FastConnect location (each to a different Oracle router for redundancy and diversity). If you need to have 20-Gbps bandwidth *at a given time*, you must ensure that each of your physical connections consists of a cross-connect group (LAG) to contain the cross-connect. Then you need to add another 10-Gbps cross-connect to each LAG, so that *each* redundant physical connection has two 10-Gbps cross-connects. FastConnect currently does not support equal-cost multi-path routing (ECMP).

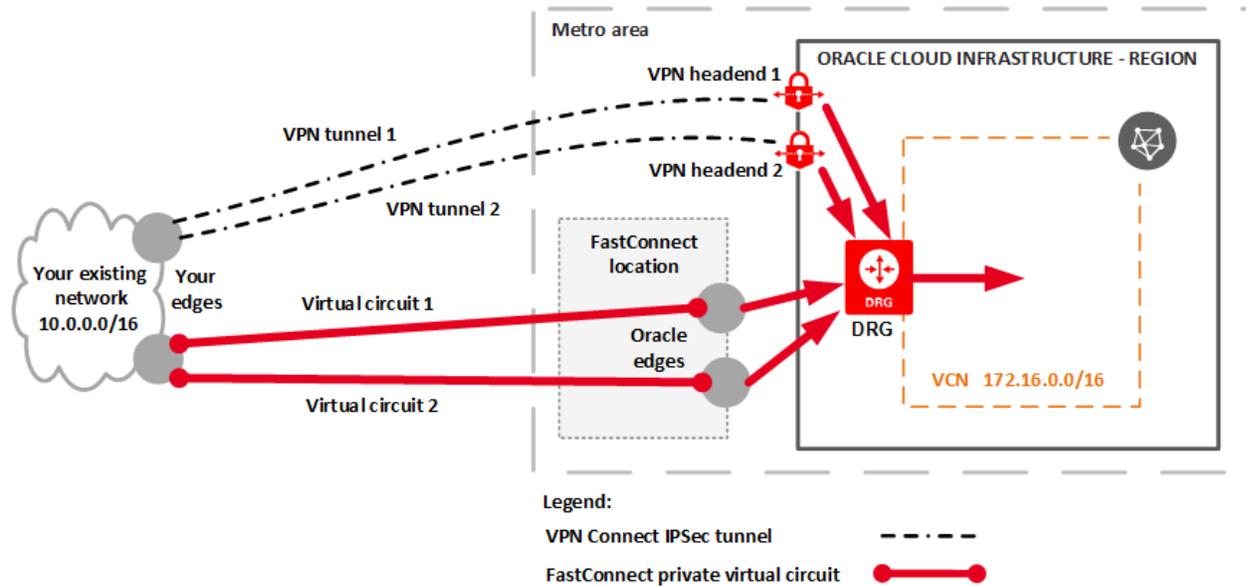
If you're working in a region that has only a single FastConnect location, you might also want location diversity. To achieve that, repeat your setup in a second FastConnect location in a nearby region. Notice that you must also have a duplicate setup of your Oracle cloud resources in that second region, as shown in the following diagram.



## VPN Connect as Backup for FastConnect

Oracle recommends using [VPN Connect](#) as a backup for your FastConnect connection. If you do, ensure that the VPN Connect tunnels are configured to use [BGP routing](#). Within your existing network, manipulate the routing to prefer routes learned through FastConnect over routes learned through VPN Connect. For example, use AS\_Path Prepend to influence egress traffic from Oracle, and use local preference to influence egress traffic from your network.

The following diagram shows a setup with redundant FastConnect virtual circuits and redundant VPN Connect tunnels.



## Related Resources

- [Routing Details for Connections to Your On-Premises Network](#)
- [Connectivity Redundancy Guide \(PDF\)](#)

## What's Next?

Choose the topic that suits your situation:

- [FastConnect: With an Oracle Provider](#)
- [FastConnect: With a Third-Party Provider](#)
- [FastConnect: Colocation with Oracle](#)

## FastConnect: With an Oracle Provider

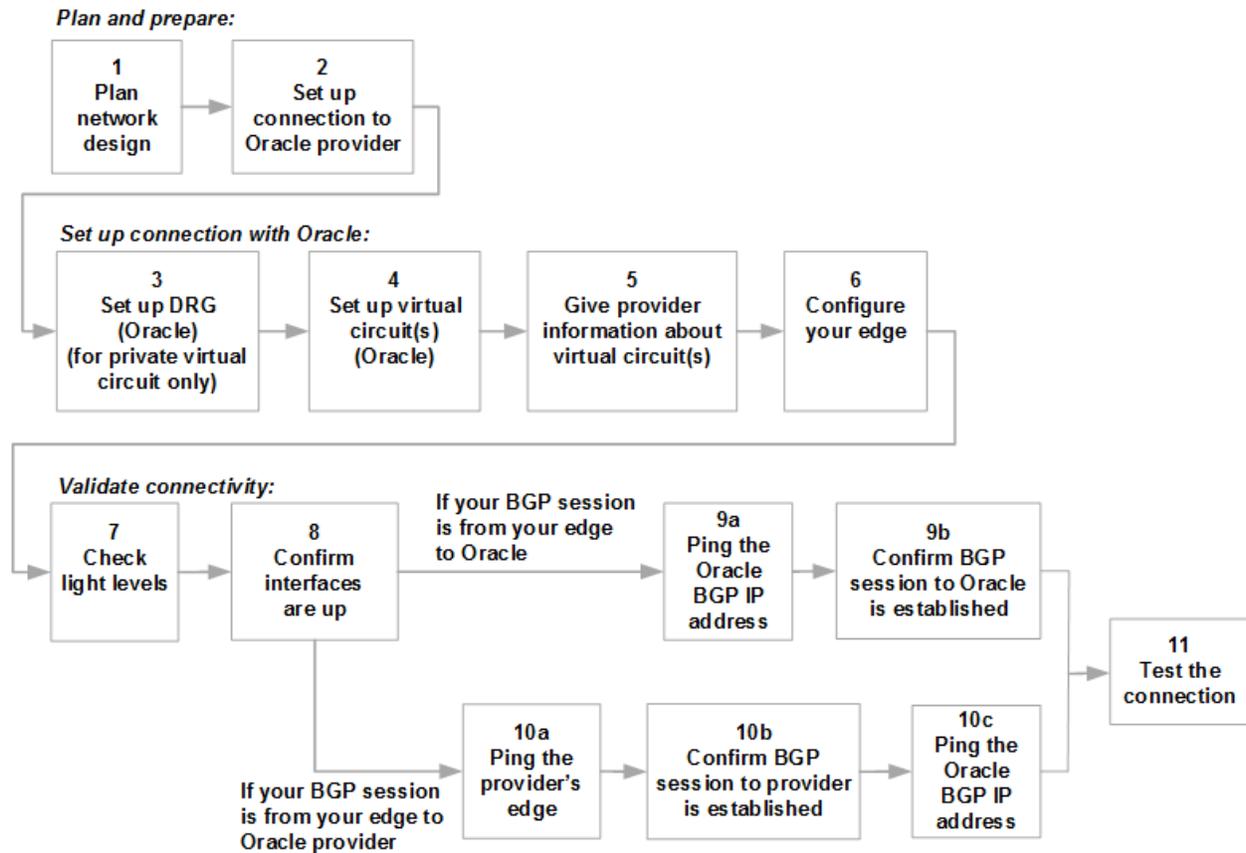
This topic is for customers who want to use Oracle Cloud Infrastructure FastConnect by connecting to an [Oracle provider](#). For a summary of the different ways to connect, see the [connectivity models](#).

If you instead want to use a network provider that is not on the list of Oracle providers, see [FastConnect: With a Third-Party Provider](#). Or if you want to use FastConnect by colocating with Oracle, see [FastConnect: Colocation with Oracle](#).

For general information about FastConnect, see [FastConnect](#).

## Getting Started with FastConnect

The following flow chart shows the overall process of setting up FastConnect.



Also see the sequence diagram in [To get the status of your virtual circuit.](#)



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Task 1: Learn and plan

If you haven't yet, walk through the planning in [Before Getting Started: Learn and Plan](#). Also see [FastConnect Redundancy Best Practices](#).

### Task 2: Set up connection to the Oracle provider

If you haven't already, start the process of ordering the connection from the Oracle provider, setting it up, and then testing it with the provider. It can take some time, depending on the provider.

### Task 3: Set up a DRG (private peering only)

**Summary:** If you plan to use a private virtual circuit (private peering), you need a DRG. If you haven't already, use the Oracle Cloud Infrastructure Console to set up a DRG, attach it to your VCN, and update routing in your VCN to include a route rule to send traffic to the DRG. It's easy to forget to update the route table. Without the route rule, no traffic will flow.

#### Instructions:

- [To create a DRG](#)
- [To attach a DRG to a VCN](#)
- [To update rules in an existing route table](#)

### Task 4: Set up your virtual circuit

**Summary:** Create one or more virtual circuits for your connection in the Oracle Console. If your network design includes more than one virtual circuit, complete the following steps for each one.

**Instructions:**

Repeat the following steps for each virtual circuit you want to create.

1. In the Console, confirm you're viewing the compartment that you want to work in. If you're not sure which one, use the compartment that contains the DRG that you'll connect to (for a private virtual circuit). This choice of compartment, in conjunction with a corresponding [IAM policy](#), controls who has access to the virtual circuit you're about to create.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.  
The resulting **FastConnect** page is where you'll create a new connection and later return to when you need to manage the connection.
3. Click **Create FastConnect**.
4. Select **Connect Through a Provider** and choose the provider from the list.
5. Click **Next**.
6. Enter the following for your virtual circuit:
  - **Name:** A friendly name that helps you keep track of your virtual circuits. The value does not need to be unique across your virtual circuits, and you can change it later. Avoid entering confidential information.
  - **Create in Compartment:** Leave as is (the compartment you're currently working in).
7. Choose the virtual circuit type (private or public). A private virtual circuit is for private peering (where your existing network receives routes for your VCN's private IP addresses). A public virtual circuit is for public peering (where your existing network

receives routes for the Oracle Cloud Infrastructure public IP addresses). Also see [Uses for FastConnect](#).

- For a private virtual circuit, enter the following:
  - **Dynamic Routing Gateway:** Select the DRG to route the FastConnect traffic to.
  - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).

If your BGP session goes to Oracle (see [Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider](#)), the dialog box includes additional fields for the BGP session:

- **Customer BGP IP Address:** The BGP peering IP address for your edge (your CPE), with either a /30 or /31 subnet mask.
  - **Oracle BGP IP Address:** The BGP peering IP address you want to use for the Oracle edge (the DRG), with either a /30 or /31 subnet mask.
  - **Enable IPv6 Address Assignment:** Available only in the Government Cloud. For more information, see [FastConnect and IPv6](#).
  - **Customer BGP ASN:** The public or private ASN for your network.
  - **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.
- For a public virtual circuit, enter the following:
    - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).
    - **Public IP Prefixes:** The public IP prefixes that you want Oracle to receive over the connection (each one must be /31 or less specific). You can enter a comma-separated list of prefixes, or one per line.

- **Customer BGP ASN:** The public ASN for your network. Present only if your BGP session goes to Oracle (see [Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider](#)). Note that Oracle specifies the BGP IP addresses for a public virtual circuit.
  - **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.
8. Click **Create**.

The virtual circuit is created. Its OCID and a link to the provider's portal are displayed in the resulting confirmation box at the top of the page. The OCID is also available with the other virtual circuit details.
  9. Copy and paste the OCID to another location. You give it to your provider in the next task.

The virtual circuit is listed on the FastConnect page.

Until you complete the next task and the provider does their provisioning work, the virtual circuit's Lifecycle State is PENDING PROVIDER and the BGP state is DOWN. After the provider does their work, the Lifecycle State switches to PROVISIONED. When the BGP session is established and working, the BGP state changes to UP.



### Tip

For a virtual circuit where your BGP session goes to the Oracle provider, the BGP state for the virtual circuit reflects the status of the *separate BGP session between the Oracle provider and Oracle*. For reference, see [Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider](#).

Also see the diagram in [To get the status of your virtual circuit](#).

### Task 5: Give the provider information about the virtual circuit

Contact the provider and give the OCID of each virtual circuit that you created, along with any other information the provider requests. Depending on the provider, you might do this in the provider's portal, or over the phone. The provider then configures each virtual circuit on their end to complete the connectivity.

If your provider is AT&T: After AT&T gives you the service key for your virtual circuit, create a ticket at [My Oracle Support](#) to request provisioning, and give Oracle the service key.

### Task 6: Configure your edge

**If your BGP session goes to Oracle:** (see [Oracle Provider Scenario: BGP Session to Either Oracle or the Oracle Provider](#)), configure your edge (your CPE) to use the BGP peering information (see [General Requirements](#)). Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#). By default, Oracle uses the default BGP timers of 60 seconds for keep-alive and 180 seconds for hold-time. If you need fast BGP convergence, you can use any value in these supported ranges: 6-60 seconds for keep-alive, and 18-180 seconds for hold-time. Also configure the router for redundancy according to the network design you decided on earlier (see [FastConnect Redundancy Best Practices](#)). After you successfully configure the BGP session, the virtual circuit's BGP session state changes to UP.

**If your BGP session instead goes to the Oracle provider:** You still need to configure your router if you haven't already. You may need to contact your provider to get the required BGP peering information. This BGP session must be up and running for FastConnect to work. Also configure your edge router for redundancy according to the network design you decided on earlier (see [FastConnect Redundancy Best Practices](#)).



### Important

For a public virtual circuit: Your existing network can receive advertisements for Oracle's public IP addresses through multiple paths (for example: FastConnect and your internet service provider). Make sure to give higher preference to FastConnect over your ISP. You must configure your edge appropriately so that traffic uses your desired path and you receive the benefits of FastConnect. This is particularly important if you decide to *also* set up your existing network with [private access to Oracle services](#). For important information about path preferences, see [Routing Details for Connections to Your On-Premises Network](#).

### Task 7: Check light levels

Confirm the light levels are good for each of your physical network connections to the provider. Don't proceed until they are.

### Task 8: Confirm your interfaces are up

Confirm your side of the interfaces for the connections to the provider are up. Don't proceed until they are.

### BGP Session Goes to Oracle

#### Task 9a: Ping the Oracle BGP IP address

For each virtual circuit, ping the Oracle BGP IP address assigned to the virtual circuit. Check

the error counters and look for any dropped packets. Don't proceed until you can successfully ping this IP address without errors.

### Task 9b: Confirm the BGP session is established

For each virtual circuit, confirm the BGP session is in an established state. When it is, the connection is ready to test (see [Task 11: Test the connection](#)).

### BGP Session Goes to the Provider

#### Task 10a: Ping the provider's edge

For each virtual circuit, ping the provider's edge. Check the error counters and look for any dropped packets. Don't proceed until you can successfully ping the provider's edge without errors.

#### Task 10b: Confirm the BGP session is established

Confirm the BGP session you have with the provider is in an established state. Don't proceed until it is.

#### Task 10c: Ping the Oracle BGP IP address

For each virtual circuit, ping the Oracle BGP IP address (which you can get from the provider). Check the error counters and look for any dropped packets. When you can successfully ping this IP address without errors, the connection is ready to test.

### Task 11: Test the connection

**For a private virtual circuit:** You should be able to launch an instance in your VCN and access it (for example, with SSH) from a host in your existing private network. See [Creating an Instance](#). If you can, your FastConnect private virtual circuit is ready to use.

### For a public virtual circuit:

1. Make sure that Oracle has successfully verified *at least one* of the public prefixes you've submitted. You can see the status of each prefix by viewing the virtual circuit's details in the Console. When one of the prefixes has been validated, Oracle starts advertising the regional Oracle Cloud Infrastructure public addresses over the connection.
2. Launch an instance with a public IP address.
3. Ping the public IP address from a host in your existing private network. You should see the packet on the FastConnect interface on the virtual circuit. If you do, your FastConnect public virtual circuit is ready to use. However, remember that *only the public prefixes that Oracle has successfully verified so far* are advertised on the connection.

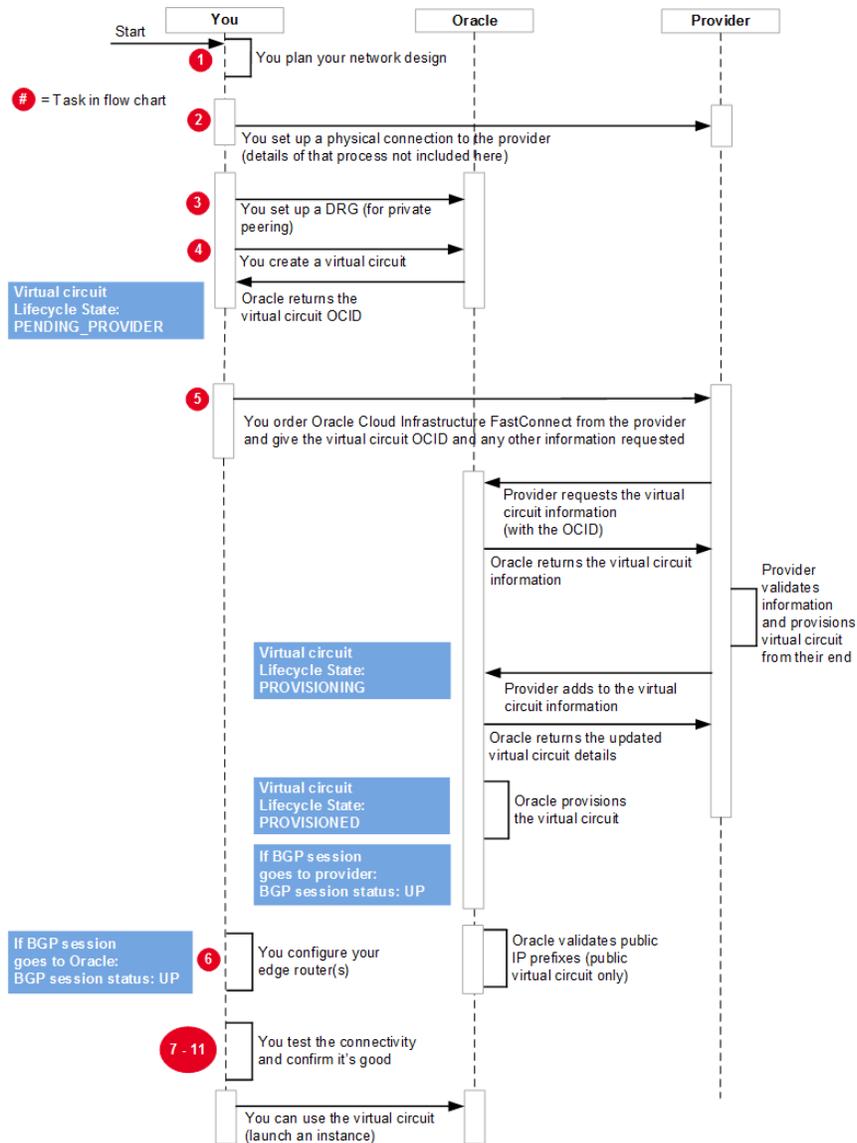
## Managing Your Virtual Circuit

### To get the status of your virtual circuit

1. In the Console, go to **Networking**, and then click **FastConnect** to view your list of connections.
2. Click the virtual circuit you're interested in to view its details.

The following diagram shows the different states of the virtual circuit when you're setting it up.

## CHAPTER 23 Networking



### To edit a virtual circuit

You can change these items for a virtual circuit:

- The name
- The bandwidth
- Which DRG it uses (for a private virtual circuit)
- The public IP prefixes (for a public virtual circuit)
- Depending on the situation, you might also have access to the BGP session information for the virtual circuit and thus can change it.



#### Important

If your virtual circuit is working and in the PROVISIONED state before you edit it, be aware that changing any of the properties besides the name, bandwidth, and public prefixes (for a public virtual circuit) causes the virtual circuit's state to switch to PROVISIONING and **may cause the related BGP session to go down**. After Oracle re-provisions the virtual circuit, its state returns to PROVISIONED. Make sure you confirm that the associated BGP session is back up.

If you change the public IP prefixes for a public virtual circuit, the BGP status is unaffected. Oracle begins advertising a new IP prefix over the connection only after verifying your ownership of it. The virtual circuit's state changes to PROVISIONING while Oracle implements any prefix changes.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Edit** and make your changes.
4. Click **Save Changes**.

To terminate a virtual circuit



### Important

Also terminate the connection with the provider, or else the provider may continue to bill you.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Delete**.
4. Confirm when prompted.

The virtual circuit's Lifecycle State changes to TERMINATING and then to TERMINATED.

To manage public IP prefixes for a public virtual circuit

For general information about the prefixes, see [Logical Connection: Public Virtual Circuit](#).

You can specify your public IP prefixes when you create the virtual circuit. See [Task 4: Set up your virtual circuit](#).

You can add or remove public IP prefixes later after creating the virtual circuit. See [To edit a virtual circuit](#). If you add a new prefix, Oracle first verifies your company's ownership before advertising it across the connection. If you remove a prefix, Oracle stops advertising the prefix within a few minutes of your editing the virtual circuit.

You can view the state of Oracle's verification of a given public prefix by viewing the virtual circuit's details in the Console. Here are the possible values:

- **In progress:** Oracle is in the process of verifying your organization's ownership of the prefix.
- **Failed:** Oracle could not verify your organization's ownership. Oracle will not advertise the prefix over the virtual circuit.
- **Completed:** Oracle successfully verified your organization's ownership. Oracle is advertising the prefix over the virtual circuit.

### To move a connection to a different compartment

You can move a connection from one compartment to another. After you move the connection to the new compartment, inherent policies apply immediately and affect access to the connection through the Console. Moving the connection to a different compartment does not affect the connection between your data center and Oracle Cloud Infrastructure. For more information, see [Moving Resources to a Different Compartment](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Find the connection in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.
5. If there are alarms monitoring the connection, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.

### Monitoring Your Connection

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring your connection, see [FastConnect Metrics](#).

### Troubleshooting

See [FastConnect Troubleshooting](#).

### FastConnect: With a Third-Party Provider

This topic is for customers who want to use Oracle Cloud Infrastructure FastConnect by connecting to a *third-party network provider* of their choice, and not an [Oracle provider](#). For a summary of the different ways to connect, see the [connectivity models](#).

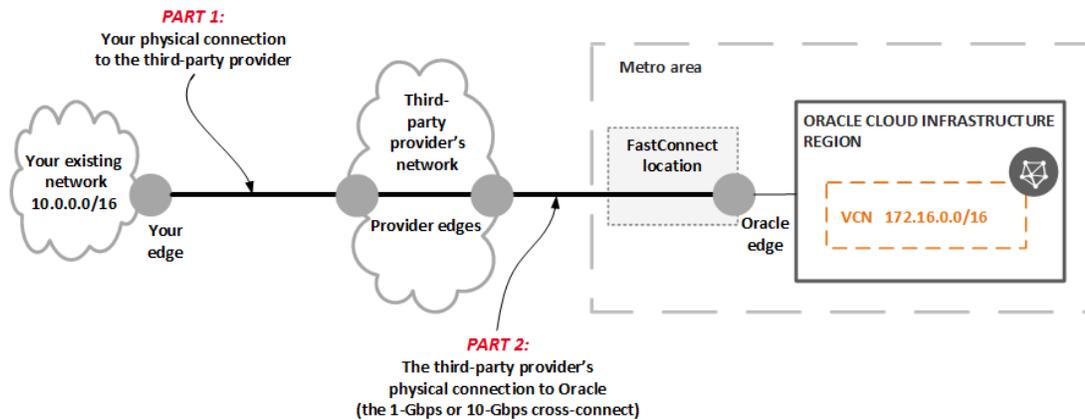
If you are using an Oracle provider, see [FastConnect: With an Oracle Provider](#). Or, if you want to use FastConnect by colocating with Oracle, see [FastConnect: Colocation with Oracle](#).

For general information about FastConnect, see [FastConnect](#).

### Important Points and Responsibilities

- You can use FastConnect by working with a third-party network service provider or carrier of your choice. The network provider must be capable of connecting to the Oracle routers in one of the [FastConnect data center locations](#) over single-mode fiber. For more detailed technical requirements, see [Hardware and Routing Requirements](#).
- Your overall connection with the third-party provider includes two parts, as illustrated in the following diagram:
  - **Part 1:** Your physical connection to the third-party provider. The rest of this topic assumes you've already set up this part of the overall connection.

- **Part 2:** The physical fiber connection (cross-connect) that the third-party provider sets up in the FastConnect location data center on your behalf.



- To obtain the Letter of Authorization (LOA) for the cross-connect, you must use the Oracle Console to [set up a cross-connect or cross-connect group](#). The resulting LOA from Oracle covers part 2 of the connection in the preceding diagram.
- You must forward the LOA to your third-party provider, who is responsible for working with the data center to set up the physical cross-connect on your behalf.
- The third-party provider issues a cross-connect order with the data center facility to run fiber optics to complete the connection from the third-party provider's cage to Oracle's patch panel as described in the LOA. Typically the data center colocation staff are the ones who run the fiber optics to complete the connection.
- Each LOA is valid for only a limited time. If the physical cross-connect is not set up before the LOA's expiration, the LOA is revoked.
- **The third-party provider is responsible for charging you for the entire connection (both parts 1 and 2).** Oracle does not set up this cross-connect in the data center, does not pay for it, and does not include it in your FastConnect charges.

- The LOA specifies an Oracle demarcation point. If your network provider is located at a different demarcation point in the data center cage, they must set up the cross-connect from their demarcation point to the Oracle demarcation point.

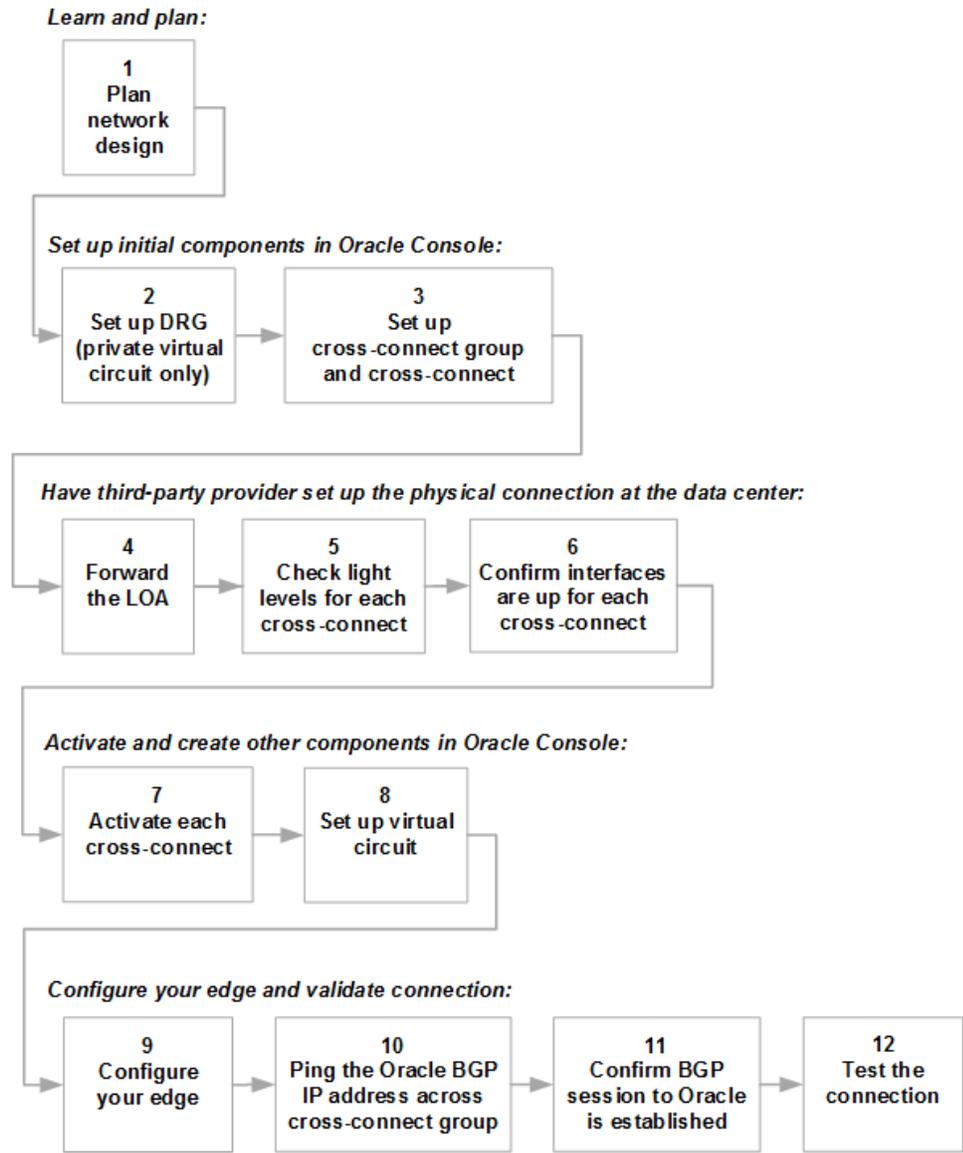
### Getting Started with FastConnect

The following flow chart shows the overall process of setting up FastConnect.



#### Note

In general, this topic assumes that your router supports link aggregation (LAG) and you will set up a cross-connect group (a LAG) with at least one cross-connect in it. The following procedures and screenshots reflect that. However, if your router doesn't support link aggregation, you can instead set up a single non-LAG cross-connect (with no cross-connect group). The procedures in this topic are still generally applicable. Instead you work only with a single cross-connect and not one or more in a cross-connect group.





### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Task 1: Learn and plan

If you haven't yet, walk through the planning in [Before Getting Started: Learn and Plan](#). Also see [FastConnect Redundancy Best Practices](#).

### Task 2: Set up a DRG (private peering only)

**Summary:** If you plan to use a private virtual circuit (private peering), you need a DRG. If you haven't already, use the Oracle Cloud Infrastructure Console to set up a DRG, attach it to your VCN, and update routing in your VCN to include a route rule to send traffic to the DRG. It's easy to forget to update the route table. Without the route rule, no traffic will flow.

#### Instructions:

- [To create a DRG](#)
- [To attach a DRG to a VCN](#)
- [To update rules in an existing route table](#)

### Task 3: Set up your cross-connect group and cross-connect

**Summary:** Create a connection in the Console, which consists of a cross-connect group (for link aggregation, or LAG) that contains at least one cross-connect. If you need more cross-connects in the group, you can [add them later](#). You can have a maximum of eight cross-connects in a group.

You have the option to set up a single non-LAG cross-connect (with no cross-connect group) if your router does not support link aggregation (LAG).

### Instructions:

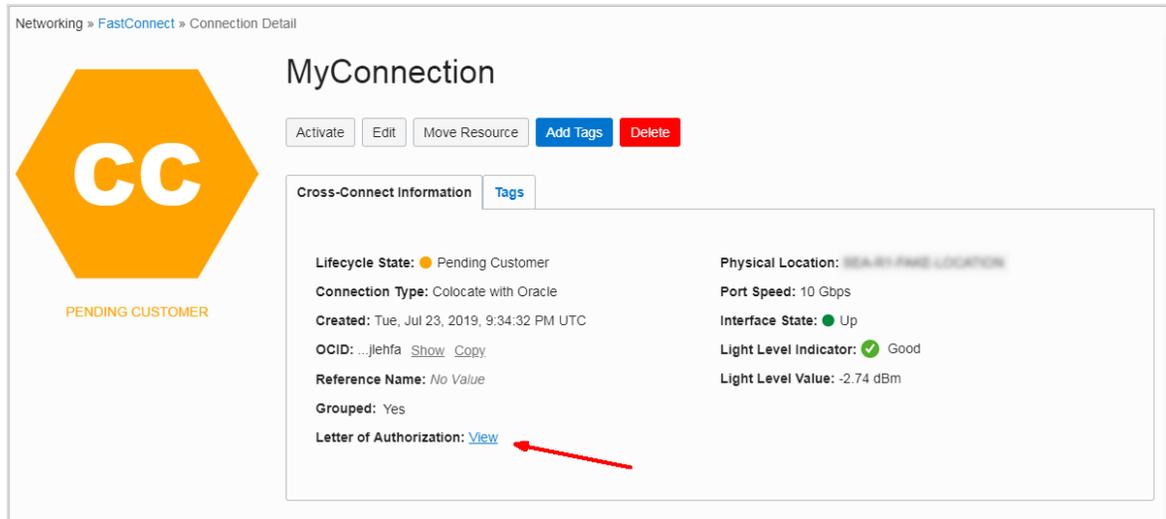
1. In the Console, confirm you're viewing the compartment that you want to work in. If you're not sure which one, use the compartment that contains the DRG that you'll connect to (for a private virtual circuit). This choice of compartment, in conjunction with a corresponding [IAM policy](#), controls who has access to the cross-connect group and each cross-connect you're about to create.

2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.

The resulting **FastConnect** page is where you'll create a new connection and later return to when you need to manage the connection and its components.

3. Click **Create FastConnect**.
4. Select **Colocate with Oracle** and click **Next**. Select this option even though a third-party provider will set up the physical connection to Oracle in the FastConnect location.
5. Enter the following items:
  - **Name:** A descriptive name that helps you keep track of this connection. You can't change the name later. Avoid entering confidential information. If you're creating a cross-connect group (LAG), the cross-connect group will use this name. Each cross-connect in this group will also use it, but with a hyphen and number appended (for example, MyName-1, MyName-2, and so on).
  - **Compartment:** Leave as is (the compartment you're currently working in).
  - **Cross-Connect Type:**
    - If your router supports LAG, select **Cross-Connect Group**. You will create a cross-connect group (a LAG) with at least one cross-connect.
    - If your router does not support link aggregation (LAG), select **Single Cross-Connect**. You will create a single non-LAG cross-connect with no cross-connect group.

- **Reference Name:** The ID for the physical LAG for the cross-connect group. This makes future connection troubleshooting easier. You might need to get this value from your third-party provider. If you don't have it, you can add it later. If you're creating a single non-LAG cross-connect, enter the ID for the physical fiber cable for the cross-connect.
  - **Number of cross-connects:** Available only if you're creating a cross-connect group. This is the number of individual cross-connects to create in the cross-connect group. In the Console, you can create three. If you need more, you can [add more cross-connects later](#) (total eight in a cross-connect group).
  - **Port speed:** All cross-connects must use a 10-Gbps port speed.
  - **Physical location:** The FastConnect location for this connection.
  - **Specify Router Proximity:** Optionally specify whether you want the new connection to be on the same or different router than one of your other connections.
6. Click **Create**.  
The new connection is created and listed on the FastConnect page.
  7. Click the new connection to see its details.
  8. **Print the LOA for each cross-connect:** Each cross-connect you just created has a Letter of Authorization (LOA). View each cross-connect's details, and then view and print the cross-connect's LOA. In the next task, you forward it to your third-party provider so they can request cabling at the FastConnect location. The cross-connect's status is PENDING CUSTOMER until you complete the next few tasks.



### Task 4: Forward the LOA to your third-party provider

Forward the LOA or LOAs from the preceding task to your third-party network provider so they can request cabling at the FastConnect location. Each LOA is valid for a limited time. The details are printed on the LOA.

### Task 5: Check light levels

After the third-party provider completes setup of the physical cross-connect in the FastConnect location, confirm from your side that the light levels for each physical connection (cross-connect) are good ( $> -15$  dBm). Don't proceed until they are.

In the Console, you can see the light levels that Oracle detects by viewing the details of the cross-connect, as shown in the following screenshot:

Networking » FastConnect » Connection Detail

### MyConnection

Activate Edit Move Resource Add Tags Delete

Cross-Connect Information Tags

**Lifecycle State:** Pending Customer  
**Connection Type:** Colocate with Oracle  
**Created:** Tue, Jul 23, 2019, 9:34:32 PM UTC  
**OCID:** ...jehfa [Show](#) [Copy](#)  
**Reference Name:** No Value  
**Grouped:** Yes  
**Letter of Authorization:** [View](#)

**Physical Location:** HEALTHY FIBER LOCATION  
**Port Speed:** 10 Gbps  
**Interface State:** Up  
**Light Level Indicator:** Good  
**Light Level Value:** -2.74 dBm

If they are not good, contact your third-party network provider.

### Task 6: Confirm your interfaces are up

For each cross-connect's physical fiber cable, confirm your side of the interfaces are up. Don't proceed until they are.

In the Console, you can see the status of Oracle's side of the interfaces (up or down) by viewing the details of the cross-connect (see the preceding screenshot).

If the interfaces are not up, contact your third-party network provider.

### Task 7: Activate each cross-connect

**Summary:** When your physical fiber cables in the FastConnect location are set up and ready to use, return to the Oracle Console and activate each cross-connect that you set up earlier. The process of activating a cross-connect informs Oracle that the corresponding physical fiber cable is ready. Oracle will then complete the router configuration for each cross-connect.

### Instructions:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click through to view the cross-connect's details, and then click **Activate**.
4. Confirm when prompted.
5. While still viewing the cross-connect's details, click **Edit** and enter the ID for the physical fiber cable for this cross-connect. Adding this value can help with any connection troubleshooting in the future. If you don't have the value available now, you can add it later.

If you have other cross-connects that are ready to use, wait for the first to be provisioned, and then activate the next one. Only one cross-connect in a group can be activated and then provisioned at a time.

After you complete this task, each cross-connect's status changes to PROVISIONING and then to PROVISIONED (typically within one minute).

### Task 8: Set up your virtual circuit

**Summary:** Create one or more virtual circuits for your connection in the Oracle Console. The cross-connect group (or your single non-LAG cross-connect) must first be in the PROVISIONED state.



### Important

If you want to use a *single* FastConnect to connect your existing network to *multiple* DRGs and VCNs, you must set up a different private virtual circuit for each VCN. Each virtual circuit must have a different VLAN and a different set of BGP IP addresses. For more information, see [FastConnect with Multiple DRGs and VCNs](#).

### Instructions:

1. In the Console, return to the connection you created earlier. Under **Resources**, click **Virtual Circuits**.
2. Click **Add Virtual Circuit**.
3. Enter the following for your virtual circuit:
  - **Name:** A descriptive name that helps you keep track of your virtual circuits. The value does not need to be unique across your virtual circuits, and you can change it later. Avoid entering confidential information.
  - **Compartment:** Select the compartment where you want to create the virtual circuit. If you're not sure, use the current compartment. This choice of compartment, in conjunction with a corresponding [IAM policy](#), controls who has access to the virtual circuit.
4. Choose the virtual circuit type (private or public). A private virtual circuit is for private peering (where your existing network receives routes for your VCN's private IP addresses). A public virtual circuit is for public peering (where your existing network receives routes for the Oracle Cloud Infrastructure public IP addresses). Also see [Uses for FastConnect](#).

- For a private virtual circuit, enter the following:
  - **Dynamic Routing Gateway:** Select the DRG to route the FastConnect traffic to.
  - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).
  - **VLAN:** The number of the VLAN to use for this virtual circuit. It must be a VLAN that is not already assigned to another virtual circuit.
  - **Customer BGP IP Address:** The BGP peering IP address for your edge (your CPE), with either a /30 or /31 subnet mask.
  - **Oracle BGP IP Address:** The BGP peering IP address you want to use for the Oracle edge (the DRG), with either a /30 or /31 subnet mask.
  - **Enable IPv6 Address Assignment:** Available only in the US Government Cloud. For more information, see [FastConnect and IPv6](#).
  - **Customer BGP ASN:** The public or private ASN for your network.
  - **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.
- For a public virtual circuit, enter the following:
  - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).
  - **Public IP Prefixes:** The public IP prefixes that you want Oracle to receive over the connection (each one must be /31 or less specific). You can enter a comma-separated list of prefixes, or one per line.
  - **VLAN:** The number of the VLAN to use for this virtual circuit. It must be a VLAN that is not already assigned to another virtual circuit.

- **Customer BGP ASN:** The public ASN for your network. Note that Oracle specifies the BGP IP addresses for a public virtual circuit.
- **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.

5. Click **Create**.

The virtual circuit is created.

The virtual circuit's status is PROVISIONING briefly while Oracle's system provisions the virtual circuit. The status then switches to DOWN if the BGP session between your edge and Oracle's edge is not yet correctly configured, if the VLAN isn't configured correctly, or if there any other problems. Otherwise the status switches to UP.

### Task 9: Configure your edge

Configure each of your edge routers to use the BGP information and VLAN for the virtual circuit. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#). By default, Oracle uses the default BGP timers of 60 seconds for keep-alive and 180 seconds for hold-time. If you need fast BGP convergence, you can use any value in these supported ranges: 6-60 seconds for keep-alive, and 18-180 seconds for hold-time.



### Important

For a public virtual circuit: Your existing network can receive advertisements for Oracle's public IP addresses through multiple paths (for example: FastConnect and your internet service provider). Make sure to give higher preference to FastConnect over your ISP. You must configure your edge appropriately so that traffic uses your desired path and you receive the benefits of FastConnect. This is particularly important if you decide to *also* set up your existing network with [private access to Oracle services](#). For important information about path preferences, see [Routing Details for Connections to Your On-Premises Network](#).

If you have a cross-connect group (a LAG) with one or more cross-connects in it, here are details to know about LACP:

- LACP is required on the network interface that is directly plugged in to Oracle's router.
- LACP is required even if you have only a single cross-connect in the cross-connect group.
- If the third-party provider is performing any media conversion, LACP must be configured on the provider's device instead of your device.

Also configure the router for redundancy according to the network design you decided on earlier. After you successfully configure BGP and the VLAN, the virtual circuit's status will switch to UP.

### Task 10: Ping the Oracle BGP IP address

Ping the Oracle BGP IP address assigned to the virtual circuit. Check the error counters and look for any dropped packets. Don't proceed until you can successfully ping this IP address

without errors.

If you've set up a cross-connect group: if the ping is not successful, and you're NOT learning MAC addresses, verify that you configured LACP as mentioned in Task 9.

### Task 11: Confirm the BGP session is established

For each virtual circuit you set up, confirm the BGP session is in an established state on your side of the connection.

### Task 12: Test the connection

**For a private virtual circuit:** You should be able to launch an instance in your VCN and access it (for example, with SSH) from a host in your existing private network. See [Creating an Instance](#). If you can, your FastConnect private virtual circuit is ready to use.

**For a public virtual circuit:**

1. Make sure that Oracle has successfully verified *at least one* of the public prefixes you've submitted. You can see the status of each prefix by viewing the virtual circuit's details in the Console. When one of the prefixes has been validated, Oracle starts advertising the regional Oracle Cloud Infrastructure public addresses over the connection.
2. Launch an instance with a public IP address.
3. Ping the public IP address from a host in your existing private network. You should see the packet on the FastConnect interface on the virtual circuit. If you do, your FastConnect public virtual circuit is ready to use. However, remember that *only the public prefixes that Oracle has successfully verified so far* are advertised on the connection.

### Managing Your Connection

#### To get the status of your connection

Look at the icon for the particular part of the connection that you're interested in (cross-connect group, cross-connect, or virtual circuit).

Here are reasons for particular status values:

#### Cross-Connect: PENDING CUSTOMER

- You need to forward the LOA to your third-party provider so they can request cabling at the FastConnect location. See [Task 4: Forward the LOA to your third-party provider](#).
- Or, you need to activate a cross-connect after confirming it's ready to use. See [Task 7: Activate each cross-connect](#), but make sure you've performed tasks 5 and 6 first.

#### Virtual circuit: DOWN

In general this means you've created a virtual circuit, but configuration is incomplete or incorrect:

- You need to configure your edge. See [Task 9: Configure your edge](#).
- Or, you've configured BGP or the VLAN incorrectly on your edge (make sure to configure the router to use the BGP and VLAN values assigned to the virtual circuit).

The following table summarizes the different states of each component involved in the connection at different points during setup:

Task in Setup Process	CCG Icon	CC Icon	VC Icon
<a href="#">Task 3: Set up your cross-connect group and cross-connect</a>	PENDING PROVISIONING	PENDING CUSTOMER	N/A
<a href="#">Task 7: Activate each cross-connect</a>	PROVISIONED	PROVISIONED	N/A
<a href="#">Task 8: Set up your virtual circuit</a>	PROVISIONED	PROVISIONED	PROVISIONING > DOWN
<a href="#">Task 9: Configure your edge</a>	PROVISIONED	PROVISIONED	DOWN > UP

### To add a new cross-connect to an existing cross-connect group

When you first create a cross-connect group in the Console, you're allowed to create three cross-connects in the group. You can later add more to increase the bandwidth and resiliency of the group. The total allowed number is eight.

1. Create the new cross-connect in the existing cross-connect group:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
  - b. Select the compartment where the connection resides, and then click the connection to view its details.
  - c. Click **Add Cross-Connect**.
  - d. Enter the following items:
    - **Name:** A descriptive name that helps you keep track of this cross-connect. The value does not need to be unique across your cross-connects. You can't change the name later. Avoid entering confidential information.

- **Reference Name:** Your ID for the physical fiber cable for the cross-connect. This makes future connection troubleshooting easier. If you don't have it, you can add it later.
  - e. Click **Add**.  
The cross-connect is created. The status of the cross-connect is PENDING CUSTOMER to indicate that you have more work to do.
  - f. Print the new cross-connect's LOA. You forward it to your third-party provider in the next step.
2. Perform tasks 4-7 in [Getting Started with FastConnect](#). In summary, you need to have the cabling set up for the new cross-connect, validate the light levels and interfaces are good, and then activate the cross-connect.

### To edit a virtual circuit

You can change these items for a virtual circuit:

- The name
- The bandwidth
- Which DRG it uses (for a private virtual circuit)
- Which VLAN it uses
- The BGP session information
- The public IP prefixes (for a public virtual circuit)



### Important

#### *Notes About Editing a Virtual Circuit*

If your virtual circuit is working and in the PROVISIONED state before you edit it, be aware that changing any of the properties besides the name, bandwidth, and public prefixes (for a public virtual circuit) causes the virtual circuit's state to switch to PROVISIONING and **may cause the related BGP session to go down**. After Oracle re-provisions the virtual circuit, its state returns to PROVISIONED. Make sure you confirm that the associated BGP session is back up.

If you change the public IP prefixes for a public virtual circuit, the BGP status is unaffected. Oracle begins advertising a new IP prefix over the connection only after verifying your ownership of it. The virtual circuit's state changes to PROVISIONING while Oracle implements any prefix changes.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Virtual Circuits**, and then click the virtual circuit to view its details.
4. Click **Edit** and make your changes.
5. Click **Save Changes**.

### To terminate a connection, or part of it

To stop being billed for a connection, you must terminate the virtual circuit, each cross-connect, and the cross-connect group associated with the connection (in that order).



#### **Important**

Also terminate the connection with the data center or third-party provider, or else they may continue to bill you.

### To terminate a virtual circuit

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Virtual Circuits**, and then click the virtual circuit to view its details.
4. Click **Delete**.
5. Confirm when prompted.

The virtual circuit's status changes to TERMINATING and then to TERMINATED.

### To terminate a cross-connect

If you have multiple cross-connects to delete in a cross-connect group, wait until the state of the first one changes to TERMINATED before deleting the next one. Also, you can't delete a cross-connect if it's the *last* provisioned cross-connect in a cross-connect group that's being used by a provisioned virtual circuit.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Cross-Connects**, and then click the cross-connect to view its details.
4. Click **Delete**.
5. Confirm when prompted.

The cross-connect's status changes to TERMINATING and then to TERMINATED.

### To terminate a cross-connect group

Prerequisite: The cross-connect group must have no virtual circuits running on it and contain no cross-connects.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Delete**.
4. Confirm when prompted.

The cross-connect group's status changes to TERMINATING and then to TERMINATED.

### To manage public IP prefixes for a public virtual circuit

For general information about the prefixes, see [Logical Connection: Public Virtual Circuit](#).

You can specify your public IP prefixes when you create the virtual circuit. See [Task 8: Set up your virtual circuit](#).

You can add or remove public IP prefixes later after creating the virtual circuit. See [To edit a virtual circuit](#). If you add a new prefix, Oracle first verifies your company's ownership before advertising it across the connection. If you remove a prefix, Oracle stops advertising the prefix within a few minutes of your editing the virtual circuit.

You can view the state of Oracle's verification of a given public prefix by viewing the virtual circuit's details in the Console. Here are the possible values:

- **In progress:** Oracle is in the process of verifying your organization's ownership of the prefix.
- **Failed:** Oracle could not verify your organization's ownership. Oracle will not advertise the prefix over the virtual circuit.
- **Completed:** Oracle successfully verified your organization's ownership. Oracle is advertising the prefix over the virtual circuit.

### To move a connection to a different compartment

You can move a connection from one compartment to another. After you move the connection to the new compartment, inherent policies apply immediately and affect access to the connection through the Console. Moving the connection to a different compartment does not affect the connection between your data center and Oracle Cloud Infrastructure. For more information, see [Moving Resources to a Different Compartment](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Find the connection in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.
5. If there are alarms monitoring the connection, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.

### Monitoring Your Connection

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring your connection, see [FastConnect Metrics](#).

### Troubleshooting

See [FastConnect Troubleshooting](#).

### FastConnect: Colocation with Oracle

This topic is for customers who are colocated with Oracle in a FastConnect location. For a summary of the different ways to connect, see the [connectivity models](#).

If you instead have a relationship with an [Oracle provider](#), see [FastConnect: With an Oracle Provider](#). Or if you have a relationship with a third-party provider, see [FastConnect: With a Third-Party Provider](#).

For general information about FastConnect, see [FastConnect](#).

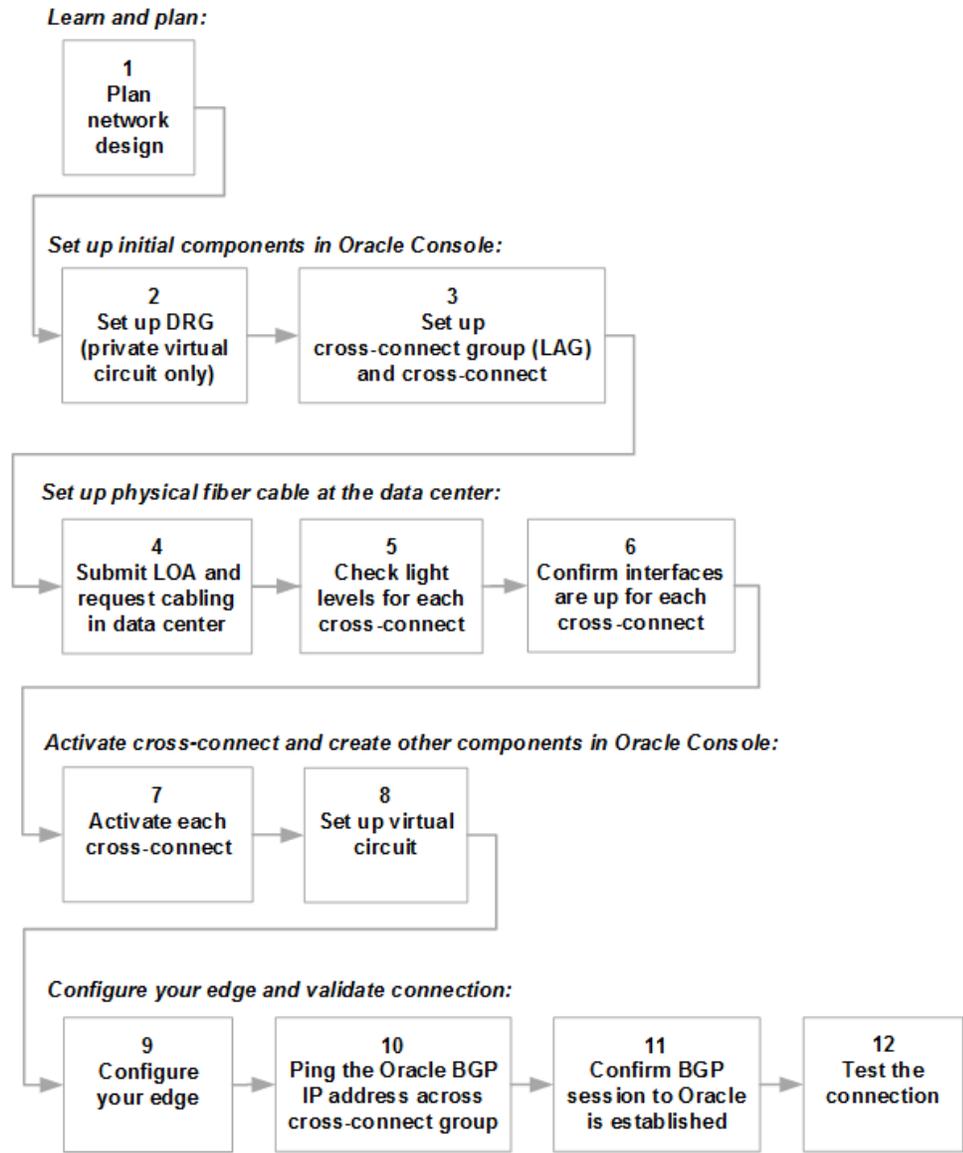
## Getting Started with FastConnect

The following flow chart shows the overall process of setting up FastConnect.



### Note

In general, this topic assumes that your router supports link aggregation (LAG) and you will set up a cross-connect group (a LAG) with at least one cross-connect in it. The following procedures and screenshots reflect that. However, if your router doesn't support link aggregation, you can instead set up a single non-LAG cross-connect (with no cross-connect group). The procedures in this topic are still generally applicable. Instead you work only with a single cross-connect and not one or more in a cross-connect group.





### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Task 1: Learn and plan

If you haven't yet, walk through the planning in [Before Getting Started: Learn and Plan](#). Also see [FastConnect Redundancy Best Practices](#).

### Task 2: Set up a DRG (private peering only)

**Summary:** If you plan to use a private virtual circuit (private peering), you need a DRG. If you haven't already, use the Oracle Cloud Infrastructure Console to set up a DRG, attach it to your VCN, and update routing in your VCN to include a route rule to send traffic to the DRG. It's easy to forget to update the route table. Without the route rule, no traffic will flow.

#### Instructions:

- [To create a DRG](#)
- [To attach a DRG to a VCN](#)
- [To update rules in an existing route table](#)

### Task 3: Set up your cross-connect group and cross-connect

**Summary:** Create a connection in the Console, which consists of a cross-connect group (for link aggregation, or LAG) that contains at least one cross-connect. If you need more cross-connects in the group, you can [add them later](#). You can have a maximum of eight cross-connects in a group.

You have the option to set up a single non-LAG cross-connect (with no cross-connect group) if your router does not support link aggregation (LAG).

### Instructions:

1. In the Console, confirm you're viewing the compartment that you want to work in. If you're not sure which one, use the compartment that contains the DRG that you'll connect to (for a private virtual circuit). This choice of compartment, in conjunction with a corresponding [IAM policy](#), controls who has access to the cross-connect group and each cross-connect you're about to create.

2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.

The resulting **FastConnect** page is where you'll create a new connection and later return to when you need to manage the connection and its components.

3. Click **Create FastConnect**.

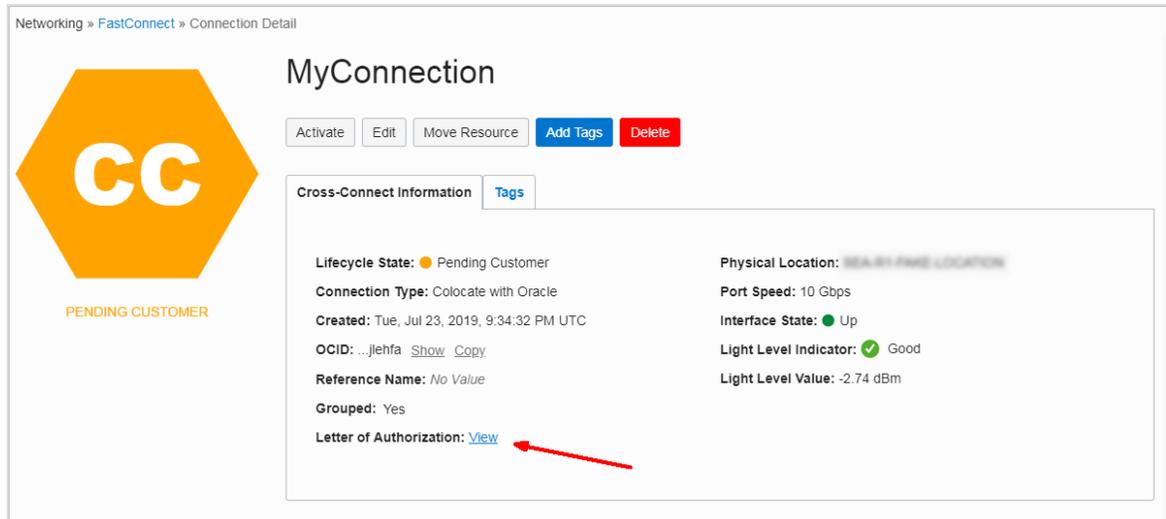
4. Select **Colocate with Oracle** and click **Next**.

5. Enter the following items:

- **Name:** A descriptive name that helps you keep track of this connection. You can't change the name later. Avoid entering confidential information. If you're creating a cross-connect group (LAG), the cross-connect group will use this name. Each cross-connect in this group will also use it, but with a hyphen and number appended (for example, MyName-1, MyName-2, and so on).
- **Compartment:** Leave as is (the compartment you're currently working in).
- **Cross-Connect Type:**
  - If your router supports LAG, select **Cross-Connect Group**. You will create a cross-connect group (a LAG) with at least one cross-connect.
  - If your router does not support link aggregation (LAG), select **Single Cross-Connect**. You will create a single non-LAG cross-connect with no cross-connect group.
- **Reference Name:** The ID for the physical LAG for the cross-connect group. This makes future connection troubleshooting easier. If you don't have it, you can add

it later. If you're creating a single non-LAG cross-connect, enter the ID for the physical fiber cable for the cross-connect.

- **Number of cross-connects:** Available only if you're creating a cross-connect group. This is the number of individual cross-connects to create in the cross-connect group. In the Console, you can create three. If you need more, you can [add more cross-connects later](#) (total eight in a cross-connect group).
  - **Port speed:** All cross-connects must use a 10-Gbps port speed.
  - **Physical location:** The FastConnect location for this connection.
  - **Specify Router Proximity:** Optionally specify whether you want the new connection to be on the same or different router than one of your other connections.
6. Click **Create**.  
The new connection is created and listed on the FastConnect page.
  7. Click the new connection to see its details.
  8. **Print the LOA for each cross-connect:** Each cross-connect you just created has a Letter of Authorization (LOA). View each cross-connect's details, and then view and print the cross-connect's LOA. In the next task, you submit it with your cabling request at the FastConnect location. The cross-connect's status is PENDING CUSTOMER until you complete the next few tasks.



### Task 4: Submit LOA and request cabling in the FastConnect location

At the FastConnect location, submit each LOA from the preceding task and request cabling for each cross-connect. Each LOA is valid for a limited time. The details are printed on the LOA.

### Task 5: Check light levels

For each cross-connect's physical fiber cable, confirm from your side that the light levels are good (> -15 dBm). Don't proceed until they are.

In the Console, you can see the light levels that Oracle detects by viewing the details of the cross-connect, as shown in the following screenshot:

Networking » FastConnect » Connection Detail

### MyConnection

Activate Edit Move Resource Add Tags Delete

Cross-Connect Information Tags

**Lifecycle State:** Pending Customer  
**Connection Type:** Colocate with Oracle  
**Created:** Tue, Jul 23, 2019, 9:34:32 PM UTC  
**OCID:** ...jehfa [Show](#) [Copy](#)  
**Reference Name:** No Value  
**Grouped:** Yes  
**Letter of Authorization:** [View](#)

**Physical Location:** REDACTED  
**Port Speed:** 10 Gbps  
**Interface State:** Up  
**Light Level Indicator:** Good  
**Light Level Value:** -2.74 dBm

### Task 6: Confirm your interfaces are up

For each cross-connect's physical fiber cable, confirm your side of the interfaces are up. Don't proceed until they are.

In the Console, you can see the status of Oracle's side of the interfaces (up or down) by viewing the details of the cross-connect (see the preceding screenshot).

### Task 7: Activate each cross-connect

**Summary:** When your physical fiber cables in the FastConnect location are set up and ready to use, return to the Oracle Console and activate each cross-connect that you set up earlier. The process of activating a cross-connect informs Oracle that the corresponding physical fiber cable is ready. Oracle will then complete the router configuration for each cross-connect.

### Instructions:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click through to view the cross-connect's details, and then click **Activate**.
4. Confirm when prompted.
5. While still viewing the cross-connect's details, click **Edit** and enter the ID for the physical fiber cable for this cross-connect. Adding this value can help with any connection troubleshooting in the future. If you don't have the value available now, you can add it later.

If you have other cross-connects that are ready to use, wait for the first to be provisioned, and then activate the next one. Only one cross-connect in a group can be activated and then provisioned at a time.

After you complete this task, each cross-connect's status changes to PROVISIONING and then to PROVISIONED (typically within one minute).

### Task 8: Set up your virtual circuit

**Summary:** Create one or more virtual circuits for your connection in the Oracle Console. The cross-connect group (or your single non-LAG cross-connect) must first be in the PROVISIONED state.



### Important

If you want to use a *single* FastConnect to connect your existing network to *multiple* DRGs and VCNs, you must set up a different private virtual circuit for each VCN. Each virtual circuit must have a different VLAN and a different set of BGP IP addresses. For more information, see [FastConnect with Multiple DRGs and VCNs](#).

### Instructions:

1. In the Console, return to the connection you created earlier. Under **Resources**, click **Virtual Circuits**.
2. Click **Add Virtual Circuit**.
3. Enter the following for your virtual circuit:
  - **Name:** A descriptive name that helps you keep track of your virtual circuits. The value does not need to be unique across your virtual circuits, and you can change it later. Avoid entering confidential information.
  - **Compartment:** Select the compartment where you want to create the virtual circuit. If you're not sure, use the current compartment. This choice of compartment, in conjunction with a corresponding [IAM policy](#), controls who has access to the virtual circuit.
4. Choose the virtual circuit type (private or public). A private virtual circuit is for private peering (where your existing network receives routes for your VCN's private IP addresses). A public virtual circuit is for public peering (where your existing network receives routes for the Oracle Cloud Infrastructure public IP addresses). Also see [Uses for FastConnect](#).

- For a private virtual circuit, enter the following:
  - **Dynamic Routing Gateway:** Select the DRG to route the FastConnect traffic to.
  - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).
  - **VLAN:** The number of the VLAN to use for this virtual circuit. It must be a VLAN that is not already assigned to another virtual circuit.
  - **Customer BGP IP Address:** The BGP peering IP address for your edge (your CPE), with either a /30 or /31 subnet mask.
  - **Oracle BGP IP Address:** The BGP peering IP address you want to use for the Oracle edge (the DRG), with either a /30 or /31 subnet mask.
  - **Enable IPv6 Address Assignment:** Available only in the US Government Cloud. For more information, see [FastConnect and IPv6](#).
  - **Customer BGP ASN:** The public or private ASN for your network.
  - **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.
- For a public virtual circuit, enter the following:
  - **Provisioned Bandwidth:** Choose your desired value. If your bandwidth needs increase later, you can update the virtual circuit to use a different value (see [To edit a virtual circuit](#)).
  - **Public IP Prefixes:** The public IP prefixes that you want Oracle to receive over the connection (each one must be /31 or less specific). You can enter a comma-separated list of prefixes, or one per line.
  - **VLAN:** The number of the VLAN to use for this virtual circuit. It must be a VLAN that is not already assigned to another virtual circuit.

- **Customer BGP ASN:** The public ASN for your network. Note that Oracle specifies the BGP IP addresses for a public virtual circuit.
- **Use a BGP MD5 Authentication Key (optional):** Select this check box and provide a key if your system requires MD5 authentication. Oracle supports up to 128-bit MD5 authentication.

5. Click **Create**.

The virtual circuit is created.

The virtual circuit's status is PROVISIONING briefly while Oracle's system provisions the virtual circuit. The status then switches to DOWN if the BGP session between your edge and Oracle's edge is not yet correctly configured, if the VLAN isn't configured correctly, or if there any other problems. Otherwise the status switches to UP.

### Task 9: Configure your edge

Configure each of your edge routers to use the BGP information and VLAN for the virtual circuit. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#). By default, Oracle uses the default BGP timers of 60 seconds for keep-alive and 180 seconds for hold-time. If you need fast BGP convergence, you can use any value in these supported ranges: 6-60 seconds for keep-alive, and 18-180 seconds for hold-time.



### Important

For a public virtual circuit: Your existing network can receive advertisements for Oracle's public IP addresses through multiple paths (for example: FastConnect and your internet service provider). Make sure to give higher preference to FastConnect over your ISP. You must configure your edge appropriately so that traffic uses your desired path and you receive the benefits of FastConnect. This is particularly important if you decide to *also* set up your existing network with [private access to Oracle services](#). For important information about path preferences, see [Routing Details for Connections to Your On-Premises Network](#).

If you have a cross-connect group (a LAG) with one or more cross-connects in it, here are details to know about LACP:

- LACP is required on the network interface that is directly plugged in to Oracle's router.
- LACP is required even if you have only a single cross-connect in the cross-connect group.

Also configure the router for redundancy according to the network design you decided on earlier. After you successfully configure BGP and the VLAN, the virtual circuit's status will switch to UP.

### Task 10: Ping the Oracle BGP IP address

Ping the Oracle BGP IP address assigned to the virtual circuit. Check the error counters and look for any dropped packets. Don't proceed until you can successfully ping this IP address without errors.

If you've set up a cross-connect group: if the ping is not successful, and you're NOT learning MAC addresses, verify that you configured LACP as mentioned in Task 9.

### Task 11: Confirm the BGP session is established

For each virtual circuit you set up, confirm the BGP session is in an established state on your side of the connection.

### Task 12: Test the connection

**For a private virtual circuit:** You should be able to launch an instance in your VCN and access it (for example, with SSH) from a host in your existing private network. See [Creating an Instance](#). If you can, your FastConnect private virtual circuit is ready to use.

**For a public virtual circuit:**

1. Make sure that Oracle has successfully verified *at least one* of the public prefixes you've submitted. You can see the status of each prefix by viewing the virtual circuit's details in the Console. When one of the prefixes has been validated, Oracle starts advertising the regional Oracle Cloud Infrastructure public addresses over the connection.
2. Launch an instance with a public IP address.
3. Ping the public IP address from a host in your existing private network. You should see the packet on the FastConnect interface on the virtual circuit. If you do, your FastConnect public virtual circuit is ready to use. However, remember that *only the public prefixes that Oracle has successfully verified so far* are advertised on the connection.

## Managing Your Connection

### To get the status of your connection

Look at the icon for the particular part of the connection that you're interested in (cross-

connect group, cross-connect, or virtual circuit).

Here are reasons for particular status values:

**Cross-Connect: PENDING CUSTOMER**

- You need to submit the LOA and request cabling at the FastConnect location. See [Task 4: Submit LOA and request cabling in the FastConnect location](#).
- Or, you need to activate a cross-connect after confirming it's ready to use. See [Task 7: Activate each cross-connect](#), but make sure you've performed tasks 5 and 6 first.

**Virtual circuit: DOWN**

In general this means you've created a virtual circuit, but configuration is incomplete or incorrect:

- You need to configure your edge. See [Task 9: Configure your edge](#).
- Or, you've configured BGP or the VLAN incorrectly on your edge (make sure to configure the router to use the BGP and VLAN values assigned to the virtual circuit).

The following table summarizes the different states of each component involved in the connection at different points during setup:

Task in Setup Process	CCG Icon	CC Icon	VC Icon
<a href="#">Task 3: Set up your cross-connect group and cross-connect</a>	PENDING PROVISIONING	PENDING CUSTOMER	N/A
<a href="#">Task 7: Activate each cross-connect</a>	PROVISIONED	PROVISIONED	N/A
<a href="#">Task 8: Set up your virtual circuit</a>	PROVISIONED	PROVISIONED	PROVISIONING > DOWN
<a href="#">Task 9: Configure your edge</a>	PROVISIONED	PROVISIONED	DOWN > UP

### To add a new cross-connect to an existing cross-connect group

When you first create a cross-connect group in the Console, you're allowed to create three cross-connects in the group. You can later add more to increase the bandwidth and resiliency of the group. The total allowed number is eight.

1. Create the new cross-connect in the existing cross-connect group:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
  - b. Select the compartment where the connection resides, and then click the connection to view its details.
  - c. Click **Add Cross-Connect**.
  - d. Enter the following items:
    - **Name:** A descriptive name that helps you keep track of this cross-connect. The value does not need to be unique across your cross-connects. You can't change the name later. Avoid entering confidential information.
    - **Reference Name:** Your ID for the physical fiber cable for the cross-connect. This makes future connection troubleshooting easier. If you don't have it, you can add it later.
  - e. Click **Add**.

The cross-connect is created. The status of the cross-connect is PENDING CUSTOMER to indicate that you have more work to do.
  - f. Print the new cross-connect's LOA. You submit it with your cabling order in the next step.
2. Perform tasks 4-7 in [Getting Started with FastConnect](#). In summary, you need to have the cabling set up for the new cross-connect, validate the light levels and interfaces are good, and then activate the cross-connect.

### To edit a virtual circuit

You can change these items for a virtual circuit:

- The name
- The bandwidth
- Which DRG it uses (for a private virtual circuit)
- Which VLAN it uses
- The BGP session information
- The public IP prefixes (for a public virtual circuit)



### Important

#### *Notes About Editing a Virtual Circuit*

If your virtual circuit is working and in the PROVISIONED state before you edit it, be aware that changing any of the properties besides the name, bandwidth, and public prefixes (for a public virtual circuit) causes the virtual circuit's state to switch to PROVISIONING and **may cause the related BGP session to go down**. After Oracle re-provisions the virtual circuit, its state returns to PROVISIONED. Make sure you confirm that the associated BGP session is back up.

If you change the public IP prefixes for a public virtual circuit, the BGP status is unaffected. Oracle begins advertising a new IP prefix over the connection only after verifying your ownership of it. The virtual circuit's state changes to PROVISIONING while Oracle implements any prefix changes.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Virtual Circuits**, and then click the virtual circuit to view its details.
4. Click **Edit** and make your changes.
5. Click **Save Changes**.

### To terminate a connection, or part of it

To stop being billed for a connection, you must terminate the virtual circuit, each cross-connect, and the cross-connect group associated with the connection (in that order).

### To terminate a virtual circuit

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Virtual Circuits**, and then click the virtual circuit to view its details.
4. Click **Delete**.
5. Confirm when prompted.

The virtual circuit's status changes to TERMINATING and then to TERMINATED.

### To terminate a cross-connect

If you have multiple cross-connects to delete in a cross-connect group, wait until the state of the first one changes to TERMINATED before deleting the next one. Also, you can't delete a cross-connect if it's the *last* provisioned cross-connect in a cross-connect group that's being used by a provisioned virtual circuit.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Cross-Connects**, and then click the cross-connect to view its details.

4. Click **Delete**.
5. Confirm when prompted.

The cross-connect's status changes to TERMINATING and then to TERMINATED.

### To terminate a cross-connect group

Prerequisite: The cross-connect group must have no virtual circuits running on it and contain no cross-connects.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection to view its details.
3. Click **Delete**.
4. Confirm when prompted.

The cross-connect group's status changes to TERMINATING and then to TERMINATED.

### To manage public IP prefixes for a public virtual circuit

For general information about the prefixes, see [Logical Connection: Public Virtual Circuit](#).

You can specify your public IP prefixes when you create the virtual circuit. See [Task 8: Set up your virtual circuit](#).

You can add or remove public IP prefixes later after creating the virtual circuit. See [To edit a virtual circuit](#). If you add a new prefix, Oracle first verifies your company's ownership before advertising it across the connection. If you remove a prefix, Oracle stops advertising the prefix within a few minutes of your editing the virtual circuit.

You can view the state of Oracle's verification of a given public prefix by viewing the virtual circuit's details in the Console. Here are the possible values:

- **In progress:** Oracle is in the process of verifying your organization's ownership of the prefix.
- **Failed:** Oracle could not verify your organization's ownership. Oracle will not advertise the prefix over the virtual circuit.
- **Completed:** Oracle successfully verified your organization's ownership. Oracle is advertising the prefix over the virtual circuit.

### To move a connection to a different compartment

You can move a connection from one compartment to another. After you move the connection to the new compartment, inherent policies apply immediately and affect access to the connection through the Console. Moving the connection to a different compartment does not affect the connection between your data center and Oracle Cloud Infrastructure. For more information, see [Moving Resources to a Different Compartment](#).

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Find the connection in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.
5. If there are alarms monitoring the connection, update the alarms to reference the new compartment. See [To update an alarm after moving a resource](#) for more information.

### Monitoring Your Connection

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about monitoring your connection, see [FastConnect Metrics](#).

## Troubleshooting

See [FastConnect Troubleshooting](#).

## FastConnect Public Peering Advertised Routes

This topic lists the public IP address ranges (routes) that are advertised to your on-premises network by way of FastConnect public peering (a public virtual circuit).

When you connect with FastConnect to Oracle Cloud Infrastructure in a particular region, the routes advertised over the public virtual circuit may include routes for *other* Oracle Cloud Infrastructure regions, and for specific Oracle Cloud Infrastructure Classic regions. The following sections list the regional routes that are advertised over the public virtual circuit.

### Americas

If you use FastConnect public peering to connect to this Oracle Cloud Infrastructure region...	These Oracle Cloud Infrastructure regional routes are advertised over the public virtual circuit	These Oracle Cloud Infrastructure Classic regional routes are advertised over the public virtual circuit
Brazil East (Sao Paulo)	<a href="#">Brazil East (Sao Paulo)</a>	<a href="#">Sao Paulo-Classic</a>
Canada Southeast (Toronto)	<a href="#">Canada Southeast (Toronto)</a> <a href="#">US East (Ashburn)</a> <a href="#">US West (Phoenix)</a>	<a href="#">Ashburn-Classic</a> <a href="#">Chicago-Classic</a>

<b>If you use FastConnect public peering to connect to this Oracle Cloud Infrastructure region...</b>	<b>These Oracle Cloud Infrastructure regional routes are advertised over the public virtual circuit</b>	<b>These Oracle Cloud Infrastructure Classic regional routes are advertised over the public virtual circuit</b>
US East (Ashburn)	<a href="#">Canada Southeast (Toronto)</a> <a href="#">US East (Ashburn)</a> <a href="#">US West (Phoenix)</a>	<a href="#">Ashburn-Classic</a> <a href="#">Chicago-Classic</a>
US West (Phoenix)	<a href="#">Canada Southeast (Toronto)</a> <a href="#">US East (Ashburn)</a> <a href="#">US West (Phoenix)</a>	<a href="#">Ashburn-Classic</a> <a href="#">Chicago-Classic</a>

## Asia-Pacific (APAC)

If you use FastConnect public peering to connect to this Oracle Cloud Infrastructure region...	These Oracle Cloud Infrastructure regional routes are advertised over the public virtual circuit	These Oracle Cloud Infrastructure Classic regional routes are advertised over the public virtual circuit
Australia East (Sydney)	<a href="#">Australia East (Sydney)</a> <a href="#">India West (Mumbai)</a> <a href="#">Japan East (Tokyo)</a> <a href="#">South Korea Central (Seoul)</a>	<a href="#">Sydney-Classic</a>
India West (Mumbai)	<a href="#">Australia East (Sydney)</a> <a href="#">India West (Mumbai)</a> <a href="#">Japan East (Tokyo)</a> <a href="#">South Korea Central (Seoul)</a>	<a href="#">Sydney-Classic</a>
Japan East (Tokyo)	<a href="#">Australia East (Sydney)</a> <a href="#">India West (Mumbai)</a> <a href="#">Japan East (Tokyo)</a> <a href="#">South Korea Central (Seoul)</a>	<a href="#">Sydney-Classic</a>
South Korea Central (Seoul)	<a href="#">Australia East (Sydney)</a> <a href="#">India West (Mumbai)</a> <a href="#">Japan East (Tokyo)</a> <a href="#">South Korea Central (Seoul)</a>	<a href="#">Sydney-Classic</a>

## Europe, Middle East, Africa (EMEA)

<b>If you use FastConnect public peering to connect to this Oracle Cloud Infrastructure region...</b>	<b>These Oracle Cloud Infrastructure regional routes are advertised over the public virtual circuit</b>	<b>These Oracle Cloud Infrastructure Classic regional routes are advertised over the public virtual circuit</b>
Germany Central (Frankfurt)	<a href="#">Germany Central (Frankfurt)</a> <a href="#">Switzerland North (Zurich)</a> <a href="#">UK South (London)</a>	<a href="#">Amsterdam-Classic</a> <a href="#">Slough-Classic</a>
Switzerland North (Zurich)	<a href="#">Germany Central (Frankfurt)</a> <a href="#">Switzerland North (Zurich)</a> <a href="#">UK South (London)</a>	<a href="#">Amsterdam-Classic</a> <a href="#">Slough-Classic</a>
UK South (London)	<a href="#">Germany Central (Frankfurt)</a> <a href="#">Switzerland North (Zurich)</a> <a href="#">UK South (London)</a>	<a href="#">Amsterdam-Classic</a> <a href="#">Slough-Classic</a>

## Oracle Cloud Infrastructure Regional Routes

## Australia East (Sydney)

- 134.70.92.0/23
- 134.70.94.0/23

- 140.91.38.0/23
- 140.91.212.0/23
- 140.204.20.0/23
- 140.204.22.0/23
- 140.238.192.0/20
- 140.238.192.0/21
- 152.67.96.0/20
- 192.29.144.0/21

### Brazil East (Sao Paulo)

- 134.70.84.0/23
- 134.70.86.0/23
- 140.91.34.0/23
- 140.91.208.0/23
- 140.204.12.0/23
- 140.204.14.0/23
- 140.238.176.0/20
- 140.238.176.0/21
- 192.29.128.0/21

### Canada Southeast (Toronto)

- 132.145.96.0/21
- 132.145.104.0/22
- 132.145.108.0/22
- 134.70.72.0/23

- 134.70.74.0/23
- 140.91.28.0/23
- 140.91.202.0/23
- 140.204.0.0/23
- 140.204.2.0/23
- 140.238.128.0/20
- 192.29.0.0/21
- 192.29.8.0/22
- 192.29.12.0/22
- 192.29.64.0/20

### Germany Central (Frankfurt)

- 130.61.0.0/23
- 130.61.2.0/23
- 130.61.4.0/23
- 130.61.6.0/24
- 130.61.7.0/24
- 130.61.8.0/21
- 130.61.16.0/20
- 130.61.32.0/20
- 130.61.48.0/20
- 130.61.64.0/21
- 130.61.72.0/21
- 130.61.80.0/21
- 130.61.88.0/21

- 130.61.96.0/23
- 130.61.98.0/23
- 130.61.100.0/22
- 130.61.104.0/21
- 130.61.112.0/21
- 130.61.120.0/21
- 130.61.128.0/17
- 132.145.224.0/21
- 132.145.232.0/21
- 132.145.240.0/21
- 132.145.248.0/21
- 134.70.40.0/23
- 134.70.42.0/23
- 134.70.44.0/23
- 134.70.46.0/23
- 134.70.48.0/23
- 134.70.50.0/23
- 138.1.0.0/22
- 138.1.4.0/22
- 138.1.8.0/22
- 138.1.12.0/22
- 138.1.40.0/21
- 138.1.64.0/22
- 138.1.68.0/22
- 138.1.72.0/22

## CHAPTER 23 Networking

---

- 138.1.76.0/22
- 138.1.96.0/22
- 138.1.100.0/22
- 138.1.104.0/22
- 138.1.160.0/20
- 138.1.176.0/20
- 138.1.192.0/20
- 140.91.16.0/23
- 140.91.18.0/23
- 140.91.20.0/23
- 140.91.198.0/23
- 144.25.48.0/22
- 144.25.52.0/22
- 144.25.56.0/22
- 144.25.60.0/22
- 147.154.128.0/20
- 147.154.144.0/20
- 147.154.160.0/20
- 147.154.176.0/20
- 147.154.192.0/21
- 147.154.200.0/21
- 147.154.208.0/21

### India West (Mumbai)

- 134.70.76.0/23

- 134.70.78.0/23
- 140.91.30.0/23
- 140.91.204.0/23
- 140.204.4.0/23
- 140.204.6.0/23
- 140.238.160.0/21
- 140.238.224.0/21
- 192.29.48.0/21

### Japan East (Tokyo)

- 132.145.112.0/22
- 132.145.116.0/22
- 132.145.120.0/21
- 134.70.80.0/23
- 134.70.82.0/23
- 140.91.32.0/23
- 140.91.206.0/23
- 140.204.8.0/23
- 140.204.10.0/23
- 140.238.32.0/20
- 140.238.48.0/20
- 158.101.128.0/19
- 158.101.128.0/20
- 192.29.32.0/20
- 192.29.32.0/22

## CHAPTER 23 Networking

---

- 192.29.36.0/22
- 140.238.192.0/20
- 140.238.160.0/21
- 140.238.224.0/21
- 140.238.240.0/20
- 132.145.80.0/20
- 140.238.0.0/20
- 134.70.92.0/22
- 140.91.38.0/23
- 140.204.20.0/23
- 192.29.144.0/23
- 140.204.4.0/23
- 192.29.48.0/22
- 192.29.160.0/21
- 134.70.96.0/22
- 140.91.40.0/23
- 140.204.24.0/23
- 192.29.20.0/22
- 134.70.76.0/23
- 134.70.78.0/23
- 140.204.4.0/23
- 140.204.6.0/23
- 132.145.84.0/22
- 132.145.88.0/21
- 134.70.98.0/23

- 140.204.24.0/23
- 140.204.26.0/23
- 140.238.0.0/20
- 192.29.16.0/22
- 134.70.94.0/23
- 140.204.20.0/23
- 140.204.22.0/23
- 152.67.0.0/20
- 140.238.192.0/21
- 152.67.96.0/20
- 129.91.16.0/21
- 129.91.176.0/20
- 129.154.0.0/16
- 129.154.0.0/24
- 129.154.2.0/24
- 160.34.48.0/20
- 160.34.74.0/23
- 160.34.83.0/24
- 160.34.112.0/24
- 160.34.113.0/24
- 205.223.86.0/23
- 205.223.86.0/24
- 205.223.87.0/24

### South Korea Central (Seoul)

- 132.145.80.0/22
- 132.145.84.0/22
- 132.145.88.0/21
- 134.70.96.0/23
- 134.70.98.0/23
- 140.91.40.0/23
- 140.91.214.0/23
- 140.204.24.0/23
- 140.204.26.0/23
- 140.238.0.0/20
- 192.29.16.0/22
- 192.29.20.0/22

### Switzerland North (Zurich)

- 134.70.88.0/23
- 134.70.90.0/23
- 140.91.36.0/23
- 140.91.210.0/23
- 140.204.16.0/23
- 140.204.18.0/23
- 140.238.168.0/21
- 140.238.208.0/21
- 192.29.56.0/21

### UK South (London)

- 130.35.112.0/22
- 132.145.0.0/23
- 132.145.2.0/23
- 132.145.4.0/23
- 132.145.6.0/24
- 132.145.7.0/24
- 132.145.8.0/21
- 132.145.16.0/20
- 132.145.32.0/20
- 132.145.48.0/20
- 132.145.64.0/23
- 132.145.66.0/23
- 132.145.68.0/22
- 132.145.72.0/21
- 134.70.56.0/23
- 134.70.58.0/23
- 134.70.60.0/23
- 134.70.62.0/23
- 134.70.64.0/23
- 134.70.66.0/23
- 138.1.16.0/22
- 138.1.20.0/22
- 138.1.24.0/22

- 138.1.28.0/22
- 138.1.80.0/22
- 138.1.84.0/22
- 138.1.88.0/22
- 138.1.92.0/22
- 138.1.208.0/20
- 138.1.224.0/20
- 138.1.240.0/20
- 140.91.22.0/23
- 140.91.24.0/23
- 140.91.26.0/23
- 140.91.200.0/23
- 140.238.64.0/19
- 144.25.64.0/22
- 144.25.68.0/22
- 144.25.72.0/22
- 144.25.76.0/22
- 147.154.224.0/20
- 147.154.240.0/20

### US East (Ashburn)

- 129.213.0.0/23
- 129.213.2.0/23
- 129.213.4.0/23
- 129.213.6.0/24

- 129.213.7.0/24
- 129.213.8.0/21
- 129.213.16.0/20
- 129.213.32.0/20
- 129.213.48.0/20
- 129.213.64.0/20
- 129.213.80.0/20
- 129.213.96.0/20
- 129.213.112.0/20
- 129.213.128.0/22
- 129.213.132.0/22
- 129.213.136.0/22
- 129.213.140.0/22
- 129.213.144.0/21
- 129.213.152.0/21
- 129.213.160.0/21
- 129.213.168.0/21
- 129.213.176.0/20
- 129.213.192.0/21
- 129.213.200.0/21
- 129.213.208.0/21
- 130.35.16.0/22
- 130.35.20.0/22
- 130.35.24.0/22
- 130.35.28.0/22

## CHAPTER 23 Networking

---

- 130.35.48.0/20
- 130.35.64.0/20
- 130.35.80.0/20
- 130.35.96.0/21
- 130.35.104.0/21
- 130.35.120.0/21
- 130.35.144.0/22
- 130.35.148.0/22
- 130.35.152.0/22
- 130.35.156.0/22
- 130.35.176.0/22
- 130.35.180.0/22
- 130.35.184.0/22
- 130.35.188.0/22
- 130.35.192.0/22
- 130.35.196.0/22
- 130.35.200.0/22
- 130.35.204.0/22
- 130.35.208.0/22
- 130.35.212.0/22
- 130.35.216.0/22
- 130.35.220.0/22
- 130.35.224.0/22
- 130.35.232.0/21
- 132.145.128.0/20

## CHAPTER 23 Networking

---

- 132.145.144.0/20
- 132.145.160.0/20
- 132.145.176.0/20
- 132.145.192.0/21
- 132.145.200.0/21
- 132.145.208.0/21
- 132.145.216.0/21
- 134.70.24.0/23
- 134.70.26.0/23
- 134.70.28.0/23
- 134.70.30.0/23
- 134.70.32.0/23
- 134.70.34.0/23
- 138.1.48.0/21
- 140.91.10.0/23
- 140.91.12.0/23
- 140.91.14.0/23
- 140.91.196.0/23
- 144.25.32.0/22
- 144.25.36.0/22
- 144.25.40.0/22
- 144.25.44.0/22
- 144.25.80.0/20
- 147.154.0.0/20
- 147.154.16.0/20

## CHAPTER 23 Networking

---

- 147.154.32.0/20
- 147.154.48.0/20
- 147.154.64.0/21
- 147.154.72.0/21
- 147.154.80.0/21
- 150.136.0.0/16

### US West (Phoenix)

- 129.146.0.0/22
- 129.146.4.0/22
- 129.146.8.0/23
- 129.146.10.0/23
- 129.146.12.0/24
- 129.146.13.0/24
- 129.146.14.0/24
- 129.146.16.0/23
- 129.146.18.0/23
- 129.146.20.0/22
- 129.146.24.0/22
- 129.146.28.0/22
- 129.146.32.0/22
- 129.146.36.0/22
- 129.146.40.0/22
- 129.146.44.0/22
- 129.146.48.0/21

## CHAPTER 23 Networking

---

- 129.146.56.0/21
- 129.146.64.0/21
- 129.146.72.0/21
- 129.146.80.0/21
- 129.146.88.0/21
- 129.146.96.0/20
- 129.146.112.0/20
- 129.146.128.0/20
- 129.146.144.0/20
- 129.146.160.0/22
- 129.146.164.0/22
- 129.146.168.0/22
- 129.146.172.0/22
- 129.146.176.0/20
- 129.146.192.0/20
- 129.146.208.0/21
- 129.146.216.0/21
- 129.146.224.0/21
- 129.146.232.0/21
- 129.146.240.0/20
- 130.35.0.0/22
- 130.35.4.0/22
- 130.35.8.0/22
- 130.35.12.0/22
- 130.35.128.0/22

- 130.35.132.0/22
- 130.35.136.0/22
- 130.35.140.0/22
- 130.35.240.0/20
- 134.70.8.0/23
- 134.70.10.0/23
- 134.70.12.0/23
- 134.70.14.0/23
- 134.70.16.0/23
- 134.70.18.0/23
- 138.1.32.0/21
- 138.1.128.0/20
- 138.1.144.0/20
- 140.91.4.0/23
- 140.91.6.0/23
- 140.91.8.0/23
- 140.91.194.0/23
- 144.25.16.0/22
- 144.25.20.0/22
- 144.25.24.0/22
- 144.25.28.0/22
- 147.154.96.0/20
- 147.154.112.0/20
- 192.29.96.0/20

## Oracle Cloud Infrastructure Classic Regional Routes

### Amsterdam-Classic

- 130.162.0.0/16
- 132.226.0.0/16
- 140.86.0.0/16
- 141.145.0.0/19
- 160.34.16.0/20
- 160.34.120.0/24
- 160.34.121.0/24
- 205.223.82.0/24
- 205.223.83.0/24

### Ashburn-Classic

- 68.233.64.0/21
- 68.233.72.0/21
- 74.117.200.0/23
- 74.117.203.0/24
- 74.117.206.0/24
- 129.144.0.0/16
- 129.145.16.0/21
- 129.145.24.0/23
- 129.145.28.0/23
- 129.145.39.0/24
- 129.145.40.0/22

- 129.149.0.0/17
- 129.149.128.0/17
- 129.150.0.0/15
- 129.152.32.0/20
- 129.152.60.0/22
- 129.152.80.0/20
- 129.152.128.0/17
- 129.156.64.0/18
- 129.157.0.0/22
- 129.157.4.0/22
- 129.157.8.0/21
- 129.157.112.0/20
- 129.157.128.0/17
- 129.158.0.0/15
- 129.191.0.0/16
- 142.0.160.0/21
- 142.0.170.0/24
- 144.25.128.0/17
- 160.34.0.0/20
- 160.34.72.0/23
- 160.34.82.0/24
- 160.34.86.0/24
- 160.34.88.0/23
- 160.34.100.0/22
- 160.34.104.0/24

## CHAPTER 23 Networking

---

- 160.34.105.0/24
- 160.34.107.0/24
- 160.34.108.0/23
- 160.34.110.0/23
- 160.34.124.0/23
- 192.18.192.0/23
- 199.167.172.0/24
- 208.72.89.0/24
- 208.72.91.0/24
- 208.72.92.0/23
- 208.72.94.0/24

### Chicago-Classic

- 68.233.72.0/21
- 74.117.200.0/23
- 74.117.203.0/24
- 74.117.206.0/24
- 129.145.24.0/23
- 129.145.28.0/23
- 129.145.39.0/24
- 129.145.40.0/22
- 129.149.0.0/17
- 129.149.128.0/17
- 129.150.0.0/15
- 129.152.80.0/20

- 129.152.128.0/17
- 129.191.0.0/16
- 160.34.0.0/20
- 160.34.72.0/23
- 160.34.82.0/24
- 160.34.86.0/24
- 160.34.88.0/23
- 160.34.104.0/24
- 160.34.108.0/23
- 160.34.110.0/23
- 199.167.172.0/24
- 208.72.89.0/24
- 208.72.91.0/24
- 208.72.92.0/23
- 208.72.94.0/24

### Sao Paulo-Classical

- 129.91.0.0/20
- 144.22.0.0/17

### Slough-Classical

- 74.117.207.0/24
- 129.152.64.0/22
- 129.156.0.0/18
- 141.144.0.0/16

- 141.144.32.0/19
- 141.145.32.0/20
- 141.145.48.0/20
- 141.145.82.0/23
- 141.145.85.0/24
- 141.145.96.0/20
- 141.145.112.0/20
- 144.21.0.0/16
- 144.24.0.0/16
- 160.34.64.0/23
- 160.34.66.0/23
- 160.34.78.0/24
- 160.34.79.0/24
- 160.34.87.0/24
- 160.34.122.0/24
- 160.34.126.0/23
- 199.167.173.0/24
- 199.167.174.0/24
- 199.167.175.0/24
- 208.72.90.0/24

### Sydney-Classic

- 140.238.160.0/21
- 140.238.224.0/21
- 140.238.240.0/20

## CHAPTER 23 Networking

---

- 132.145.112.0/20
- 140.238.32.0/20
- 140.238.48.0/20
- 132.145.80.0/20
- 140.238.0.0/20
- 140.204.4.0/23
- 192.29.48.0/22
- 192.29.160.0/21
- 134.70.80.0/22
- 140.91.32.0/23
- 140.204.8.0/23
- 192.29.36.0/22
- 134.70.96.0/22
- 140.91.40.0/23
- 140.204.24.0/23
- 192.29.20.0/22
- 134.70.76.0/23
- 134.70.78.0/23
- 140.204.4.0/23
- 140.204.6.0/23
- 132.145.116.0/22
- 132.145.120.0/21
- 134.70.82.0/23
- 140.204.8.0/23
- 140.204.10.0/23

## CHAPTER 23 Networking

---

- 158.101.128.0/19
- 158.101.128.0/20
- 192.29.32.0/20
- 192.29.32.0/22
- 132.145.84.0/22
- 132.145.88.0/21
- 134.70.98.0/23
- 140.204.24.0/23
- 140.204.26.0/23
- 140.238.0.0/20
- 192.29.16.0/22
- 140.204.20.0/23
- 129.91.16.0/21
- 129.91.176.0/20
- 129.154.0.0/16
- 129.154.0.0/24
- 129.154.2.0/24
- 160.34.48.0/20
- 160.34.74.0/23
- 160.34.83.0/24
- 160.34.112.0/24
- 160.34.113.0/24
- 205.223.86.0/23
- 205.223.86.0/24
- 205.223.87.0/24

# FastConnect Metrics

You can monitor the health, capacity, and performance of your [FastConnect connection](#) by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

This topic describes the metrics emitted by the metric namespace `oci_fastconnect`.

Resources: cross-connect groups, virtual circuits

## Overview of Metrics: `oci_fastconnect`

Metrics are available for multiple resources in the FastConnect connection. The metrics help you determine quickly whether your FastConnect connection is up, how much data is flowing over the connection, and whether packets are being dropped for unexpected errors.

FastConnect offers different [connectivity models](#):

- [Connect with an Oracle provider](#): Metrics are available for virtual circuits in the connection.
- [Connect with a third-party provider](#): Metrics are available for the cross-connect group (LAG) and virtual circuits in the connection. Metrics for cross-connects will be available in a future release.
- [Colocate with Oracle](#): Metrics are available for the cross-connect group (LAG) and virtual circuits in the connection. Metrics for cross-connects will be available in a future release.

A cross-connect group (LAG) contains one or more cross-connects. If there are multiple and one goes down, the cross-connect group stays up, but the group might experience a lower overall bandwidth.

### Required IAM Policy

To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources

being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: `oci_fastconnect`

The metrics listed in the following table are automatically available for each virtual circuit or cross-connect group that you create. You do not need to enable monitoring to get these metrics.

You also can use the Monitoring service to create [custom queries](#).

Each metric includes the following dimensions:

#### **COMPONENT**

Possible values are `crossconnectgroup` and `virtualcircuit`. If you [connect through an Oracle provider](#), only the `virtualcircuit` component is available.

#### **RESOURCEID**

The OCID of the resource (either a cross-connect group or virtual circuit).

Metric	Metric Display Name	Unit	Description	Dimensions
ConnectionState	<b>Connection State</b>	Binary (1 or 0)	<p>The values are up (1) or down (0).</p> <p>For a virtual circuit, this is the operational state of the virtual circuit's interface.</p> <p>For a cross-connect group, this reflects the overall operational state of the cross-connects that make up the cross-connect group (LAG). If at least one of the cross-connects is up, this value is up (1). If <i>all</i> the cross-connects in the group are down, this value is down (0).</p>	component resourceId
PacketsReceived	<b>Packets Received</b>	Packets	<p>Number of packets received on the FastConnect interface at the Oracle end of the connection.</p> <p>For a cross-connect group (LAG), the value is the sum across all cross-connects in the group.</p>	

Metric	Metric Display Name	Unit	Description	Dimensions
BytesReceived	<b>Bytes Received</b>	Bytes	<p>Number of bytes received on the FastConnect interface at the Oracle end of the connection.</p> <p>For a cross-connect group (LAG), the value is the sum across all cross-connects in the group.</p>	
PacketsSent	<b>Packets Sent</b>	Packets	<p>Number of packets sent from the FastConnect interface at the Oracle end of the connection.</p> <p>For a cross-connect group (LAG), the value is the sum across all cross-connects in the group.</p>	

Metric	Metric Display Name	Unit	Description	Dimensions
BytesSent	<b>Bytes Sent</b>	Bytes	<p>Number of bytes sent from the FastConnect interface at the Oracle end of the connection.</p> <p>For a cross-connect group (LAG), the value is the sum across all cross-connects in the group.</p>	
PacketsError	<b>Packets with Errors</b>	Packets	<p>Number of packets dropped at the Oracle end of the connection. Dropped packets indicate a misconfiguration in some part of the overall system. Check if there's been a change to the configuration of your VCN, the virtual circuit, or your CPE.</p> <p>For a cross-connect group (LAG), the value is the sum across all cross-connects in the group.</p>	

## Using the Console

The instructions depend on which FastConnect [connectivity model](#) you use.

### If You Use an Oracle Provider

To view default metric charts for a single FastConnect connection

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Click the connection you're interested in.
3. The default metrics charts for the connection's virtual circuit are displayed on the resulting page.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

To view default metric charts for all FastConnect connections in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that you're interested in.
3. For **Metric Namespace**, select **oci\_fastconnect**.

The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.

If there are multiple FastConnect connections in the compartment, by default the charts show a separate line for each one (each virtual circuit).

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### If You Use a Third-Party Provider or Colocate with Oracle

In this situation, you manage both the physical connection (cross-connects) and logical connection (virtual circuit).

For the physical connection, metrics are available for the cross-connect group (LAG), but not the individual cross-connects. If you are using only a single cross-connect with no cross-connect group, then no metrics are available for the physical connection.

For the logical connection, metrics are available for each virtual circuit.

### To view default metric charts for a single FastConnect connection

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Click the connection you're interested in.
3. View the metrics for the resource you're interested in:
  - For a cross-connect group: Under **Resources**, click **Metrics**. The default metrics charts are displayed on the resulting page.
  - For a virtual circuit:
    - a. Under **Resources**, click **Virtual Circuits**.
    - b. Click the virtual circuit you're interested in. If it's a private virtual circuit, the default metrics charts are displayed on the resulting page. If it's a public virtual circuit, click **Metrics** to view the charts.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts for all FastConnect connections in a compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Compartment**, select the compartment that you're interested in.
3. For **Metric Namespace**, select **oci\_fastconnect**.  
The **Service Metrics** page dynamically updates the page to show charts for each metric that is emitted by the selected metric namespace.

By default the charts show a separate line for each resource in the compartment (each cross-connect group and virtual circuit).

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

## FastConnect Troubleshooting

This topic covers troubleshooting techniques for a FastConnect connection that has issues.

Some of the troubleshooting techniques assume that you are a network engineer with access to your CPE's configuration.

### Microsoft Azure Connection Issues

#### Problems terminating the Azure connection

The connection components must be [terminated in a specific order](#). If you don't follow this order, the FastConnect virtual circuit switches to a "Failed" state and cannot be deleted.

To fix a virtual circuit in the "Failed" state, go to the Azure portal and confirm the following items:

- The ExpressRoute circuit is not in the "Failed" state. If it is, click the ExpressRoute circuit's **Refresh** button. The circuit should return to its normal state.
- The ExpressRoute circuit has no connections. You must [first delete all its connections](#).

After you've confirmed the preceding items, you can continue with the termination process in the following steps:

1. In the Oracle Console, delete your FastConnect virtual circuit. Make sure it is deleted before proceeding.
2. In the Azure portal, confirm that the private peering for the ExpressRoute circuit has been deleted. Also confirm that the ExpressRoute circuit's status has changed to "Not Provisioned".
3. In the Azure portal, [delete the ExpressRoute circuit](#).

### General Issues

#### FastConnect is DOWN



#### Important

If you're working with an [Oracle provider](#) or a [third-party provider](#), contact both the provider and Oracle for help troubleshooting the issue. If you're [colocated with Oracle](#), contact Oracle.

#### Cross-connect and physical connection (layer 1)

Check these items:

- **Port allocation:** Verify that your connection is using the correct port, and the port is UP and activated.
- **Optical signal:** Verify that your connection is using the correct optics and transceiver, and the port is sending and receiving an optimal signal. For more information, see [FastConnect Requirements](#).
- **Fiber strands:** Try rolling or flipping the Tx/Rx fiber strands.
- **End-to-end physical connectivity:** Verify the end-to-end physical connectivity. Also verify the Tx/Rx optic signal between your CPE, the provider's network device (if you're working with a provider), and the Oracle FastConnect router.

### Data-link (layer 2)

Check the following items on your CPE. If you're working with a provider, also have them check the items on their network device:

- **BGP address:** Verify that the router is configured with the correct BGP peering IP address under the correct VLAN on the interface.
- **MAC address:** Verify that the router is receiving the MAC address from the Oracle FastConnect router, and that the MAC address has an entry in the router's address resolution protocol (ARP) table.
- **LAG and LACP:** Verify that the router has LAG configured and LACP is enabled on the interface (the Oracle FastConnect router requires both). For more information, see [FastConnect Requirements](#).

### Network and transport (layers 3 and 4)

Check the following items on your CPE. If you're working with a provider, also have them check the items on their network device:

- **BGP address:** Verify that the router is configured with the correct BGP peering IP address.

- **ASN:** Verify that the router is configured with the correct BGP local ASN and Oracle BGP ASN. Oracle's BGP ASN for the commercial cloud is 31898. For the Government Cloud, see [Oracle's BGP ASN](#).
- **MD5:** If you're using MD5 authentication, verify that the authentication string (the password) is correct.
- **Maximum prefixes:** Verify that you are advertising no more than the maximum allowed number of prefixes for virtual circuits. If you're advertising more, BGP won't be established. Here are the limits:
  - Public virtual circuits: maximum 200 prefixes
  - Private virtual circuits: maximum 2000 prefixes
- **Firewalls:** Verify that your on-premises firewall or access control lists are not blocking TCP port 179 (BGP) or any high-numbered TCP ports.

### FastConnect virtual circuit is UP, but BGP session is DOWN

The Oracle Console displays an alert if the virtual circuit is in the PROVISIONED state, but the BGP session is DOWN.

Typically, the alert indicates one of the following issues:

- You have not yet configured your CPE with the required information for the FastConnect connection. After you configure the CPE, the alert should no longer appear.
- You have configured your CPE with incorrect information. Verify that your CPE is configured with the correct information.

The CPE configuration information includes these items:

- BGP address for each side of the connection
- ASN for your network and for Oracle's network
- MD5 authentication string (if you're using MD5 authentication)
- Maximum number of allowed prefixes

For more details, see the preceding information shown for network and transport (layers 3 and 4) in [FastConnect is DOWN](#).

**Exception:** The preceding information is not relevant if you're using an Oracle provider, and the BGP session from your CPE goes to that provider and not Oracle. In that case, contact your provider to confirm that the BGP session they have with Oracle is configured correctly.

### FastConnect virtual circuit is UP, but no traffic is passing through

Check these items:

- **VCN security lists:** Ensure you've set up the [VCN security lists](#) to allow the desired traffic (both ingress and egress rules). Note that the VCN's [default security list](#) does not allow ping traffic (ICMP type 8 and ICMP type 0). You must add the appropriate ingress and egress rules to allow ping traffic.
- **Correct routes on both ends:** Verify that you have received the correct VCN routes from FastConnect and the CPE is using those routes. Likewise, verify that you are advertising the correct on-premises network routes to FastConnect and the VCN route tables use those routes.

### FastConnect virtual circuit is UP, but traffic is passing in only one direction

Check these items:

- **VCN security lists:** Ensure that your [VCN security lists](#) allow traffic in both directions (ingress and egress).
- **Firewalls:** Verify that your on-premises firewall or access control lists are not blocking traffic to or from the Oracle end.
- **Asymmetric routing:** Oracle uses asymmetric routing. If you have multiple virtual circuits, make sure that your CPE is configured for asymmetric route processing.
- **Redundant connections:** If you have redundant FastConnect virtual circuits, make sure they're both advertising the same routes.

### Redundant Connections

Remember that FastConnect uses BGP dynamic routing, and IPSec connections can use either static routing or BGP, or a combination.

### IPSec and FastConnect are both set up, but traffic is only passing through IPSec

Make sure to use more specific routes for the connection you want as primary. If you're using the same routes for both IPSec and FastConnect, see the discussion of routing preferences in [Route Advertisements and Path Preferences When You Have Multiple Connections](#).

## Internet Gateway

This topic describes how to set up and manage an internet gateway to give your VCN internet access.



#### Tip

Oracle also offers a [NAT gateway](#), which is recommended for subnets in your VCN that do not require ingress connections from the internet.

### Highlights

- An internet gateway is an optional virtual router you can add to your VCN to enable direct connectivity to the internet.
- The gateway supports connections initiated from within the VCN (egress) and connections initiated from the internet (ingress).

- Resources that need to use the gateway for internet access must be in a [public subnet](#) and have [public IP addresses](#). Resources that have private IP addresses can instead use a [NAT gateway](#) to initiate connections to the internet.
- Each public subnet that needs to use the internet gateway must have a [route table rule](#) that specifies the gateway as the target.
- You use [security rules](#) to control the types of traffic allowed in and out of resources in that subnet. Make sure to allow only the desired types of internet traffic.
- The internet gateway can be used only by resources in the gateway's VCN. Hosts in the connected on-premises network or in a [peered VCN](#) cannot use that internet gateway.

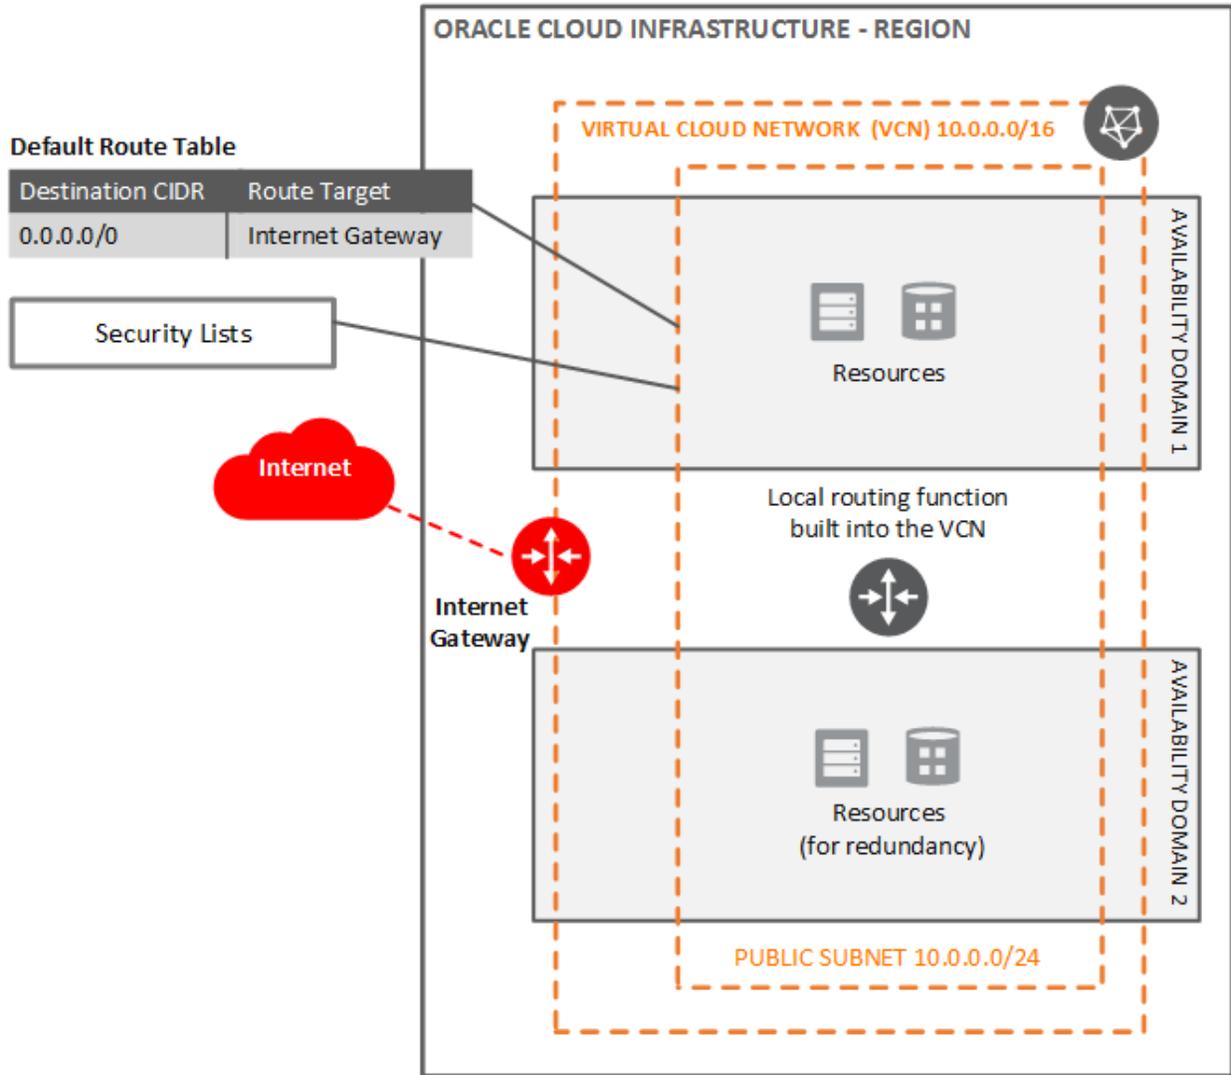
### Overview of Internet Gateways

Before continuing, make sure you've read [Access to the Internet](#) and also understand how to set up [security rules](#) for the resources in a subnet.

An internet gateway is an optional virtual router that connects the edge of the VCN with the internet. To use the gateway, the hosts on both ends of the connection must have public IP addresses for routing. Connections that originate in your VCN and are destined for a public IP address (either inside or outside the VCN) go through the internet gateway. Connections that originate outside the VCN and are destined for a public IP address inside the VCN go through the internet gateway.

A given VCN can have only one internet gateway. You control which public subnets in the VCN can use the gateway by configuring the subnet's associated route table. You use security rules to control the types of traffic allowed in and out of resources in those public subnets.

The following diagram illustrates a simple VCN setup with two public subnets. The VCN has an internet gateway, and the two public subnets are both configured to use the VCN's default route table. The table has a route rule that sends all egress traffic from the subnets to the internet gateway. The gateway allows any ingress connections from the internet with a destination IP address equal to the public IP address of a resource in the VCN. However, the public subnet's security list rules ultimately determine the specific *types* of traffic that are allowed in and out of the resources in the subnet. Those specific security rules are not shown in the diagram.





### Tip

When an internet gateway receives traffic from your VCN destined for a public IP address that is part of Oracle Cloud Infrastructure (such as Object Storage), the internet gateway routes the traffic to the destination without sending the traffic over the internet.

## Working with Internet Gateways

You create an internet gateway in the context of a specific VCN. In other words, the internet gateway is automatically attached to a VCN. However, you can disable and re-enable the internet gateway at any time. Compare this with a [dynamic routing gateway](#) (DRG), which you create as a standalone object that you then *attach* to a particular VCN. DRGs use a different model because they're intended to be modular building blocks for privately connecting VCNs to your on-premises network.

For traffic to flow between a subnet and an internet gateway, you must create a route rule accordingly in the subnet's route table (for example, destination CIDR = 0.0.0.0/0 and target = internet gateway). If the internet gateway is disabled, that means no traffic will flow to or from the internet even if there's a route rule that enables that traffic. For more information, see [Route Tables](#).

For the purposes of access control, you must specify the compartment where you want the internet gateway to reside. If you're not sure which compartment to use, put the internet gateway in the same compartment as the cloud network. For more information, see [Access Control](#).

You may optionally assign a friendly name to the internet gateway. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the internet gateway a unique identifier called an Oracle Cloud ID (OCID). For more information, see [Resource Identifiers](#).

To delete an internet gateway, it does not have to be disabled, but there must not be a route table that lists it as a target.

### Using the Console



#### **Warning**

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To set up an internet gateway

Prerequisites:

- You've determined which subnets in the VCN need access to the internet, and you've created those public subnets.
- You've determined the types of ingress and egress internet traffic that you want to enable for the resources in each public subnet (examples: ingress HTTPS connections, ingress ICMP ping connections).
- The required IAM policy is in place to allow you to work with Networking service resources. For administrators: see [IAM Policies for Networking](#).



### Important

If you've configured the public subnet to use the [default security list](#), remember that the list includes several helpful default rules that enable basic required access (examples: ingress SSH, egress access to all destinations). Oracle recommends that you become familiar with the basic access that these default rules provide. If you choose not to use the default security list, make sure to provide this basic access by implementing these security rules either in [network security groups](#) (NSGs) or custom [security lists](#).

The following procedure uses security lists, but you could instead implement the security rules in a network security group and then create all of the subnet's resources in that NSG.

1. For each public subnet that needs to use the internet gateway, set up the subnet's security list rules to allow the desired internet traffic.
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in (a security list associated with the public subnet).
  - c. Under **Resources**, click either **Ingress Rules** or **Egress Rules** depending on the type of rule you want to work with.
  - d. If you want to add a new rule, click **Add Ingress Rule** (or **Add Egress Rule**).

### Example

Imagine you have web servers in the public subnet. This example shows how to

add an ingress rule for HTTPS connections (TCP port 443) coming from the internet to the web server. Without this rule, inbound HTTPS connections are not allowed.

- i. Leave the **Stateless** check box unselected.
  - ii. **Source Type:** CIDR
  - iii. **Source CIDR:** 0.0.0.0/0
  - iv. **IP Protocol:** Leave as TCP.
  - v. **Source Port Range:** Leave as All.
  - vi. **Destination Port Range:** Enter 443.
- e. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Remove**.
  - f. If you wanted to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.
2. Create the VCN's internet gateway:
    - a. In the Console, while viewing the VCN you're interested in, click **Internet Gateways**
    - b. Click **Create Internet Gateway**.
    - c. Enter the following:
      - **Name:** A friendly name for the internet gateway. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
      - **Create in Compartment:** The compartment where you want to create the internet gateway, if different from the compartment you're currently working in.
      - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag

namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

- d. Click **Create Internet Gateway**.

Your internet gateway is created and displayed on the **Internet Gateways** page of the compartment you chose. It's already enabled, but you still need to add a route rule that allows traffic to flow to the gateway.

3. For each public subnet that needs to use the internet gateway, update the subnet's route table:
  - a. While viewing the VCN's details, click **Route Tables**.
  - b. Click the public subnet's route table to view its details.
  - c. Click **Add Route Rule**.
  - d. Enter the following:
    - **Target Type:** Internet Gateway
    - **Destination CIDR block:** 0.0.0.0/0 (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule)
    - **Compartment:** The compartment where the internet gateway is located.
    - **Target:** The internet gateway you just created.
  - e. Click **Save**.

An internet gateway is now enabled and working for your cloud network.

### To disable/enable an internet gateway

This is available only through the API. If you don't have access to the API and need to disable or enable an internet gateway, contact [Oracle Support](#). You can also easily delete and recreate the internet gateway if needed. Just make sure to update any route tables that refer to the internet gateway.

### To delete an internet gateway

Prerequisite: The internet gateway does not have to be disabled, but there must not be a route table that lists it as a target.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Internet Gateways**.
4. Click the Actions icon (three dots) for the internet gateway, and then click **Terminate**.
5. Confirm when prompted.

### To manage tags for an internet gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Internet Gateways**.
4. Click the Actions icon (three dots) for the internet gateway, and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move an internet gateway to a different compartment

You can move an internet gateway from one compartment to another. When you move an internet gateway to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.

3. Under **Resources**, click **Internet Gateways**.
4. Click the the Actions icon (three dots) for the internet gateway, and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your internet gateways, use these operations:

- [ListInternetGateways](#)
- [CreateInternetGateway](#)
- [GetInternetGateway](#)
- [UpdateInternetGateway](#)
- [DeleteInternetGateway](#)
- [ChangeInternetGatewayCompartment](#)

### NAT Gateway

This topic describes how to set up and manage a Network Address Translation (NAT) gateway. A NAT gateway gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Highlights

- You can add a NAT gateway to your VCN to give instances in a private subnet access to the internet.
- Instances in a private subnet don't have public IP addresses. With the NAT gateway, they can initiate connections to the internet and receive responses, but not receive inbound connections initiated from the internet.
- NAT gateways are highly available and support TCP, UDP, and ICMP ping traffic.

## Overview of NAT

NAT is a networking technique commonly used to give an entire private network access to the internet without assigning each host a public IPv4 address. The hosts can initiate connections to the internet and receive responses, but not receive inbound connections initiated from the internet.

When a host in the private network initiates an internet-bound connection, the NAT device's public IP address becomes the source IP address for the outbound traffic. The response traffic from the internet therefore uses that public IP address as the destination IP address. The NAT device then routes the response to the host in the private network that initiated the connection.

### Overview of NAT Gateways

The Networking service offers a reliable and highly available NAT solution for your VCN in the form of a NAT gateway.

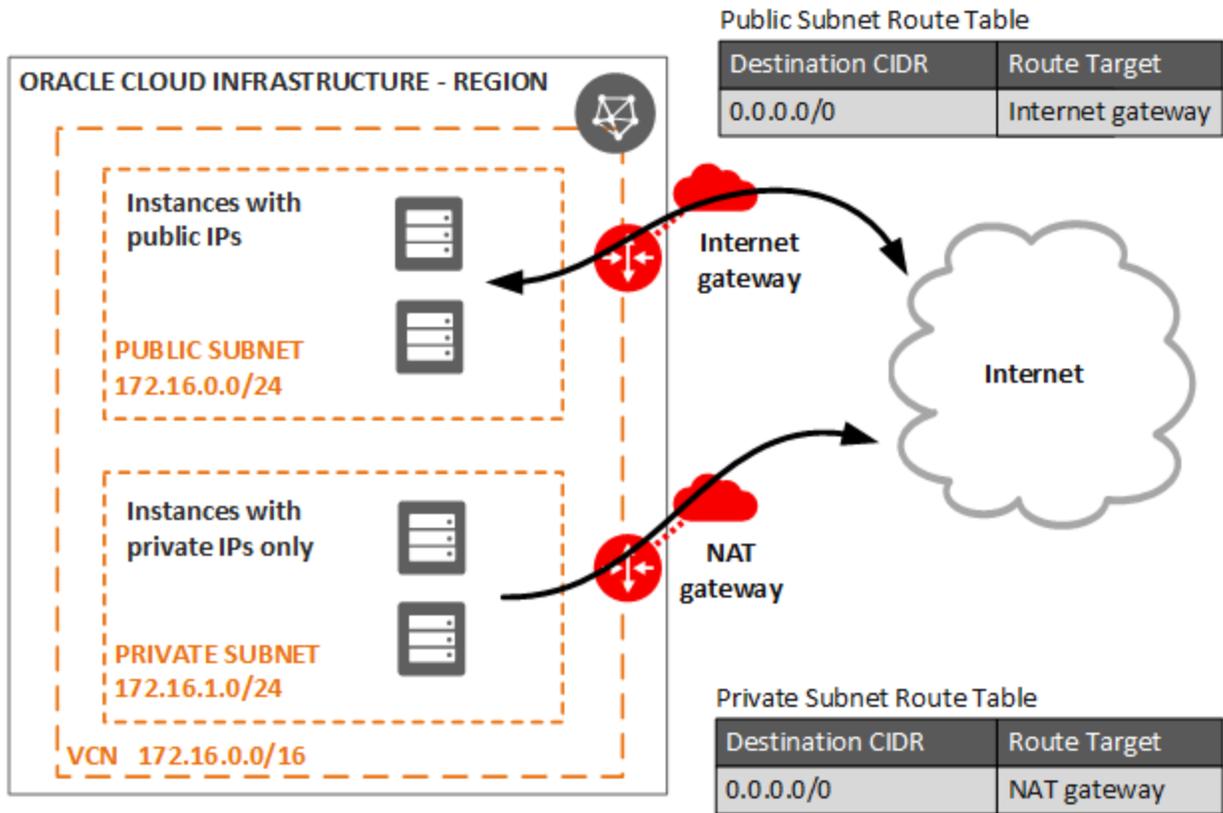
Example scenario: Imagine you have resources that need to receive inbound traffic from the internet (for example, web servers). You also have private resources that need to be protected from inbound traffic from the internet. All of these resources need to initiate connections to the internet to request software updates from sites on the internet.

You set up a VCN and add a public subnet to hold the web servers. When launching the instances, you assign public IP addresses to them so they can receive inbound internet traffic. You also add a private subnet to hold the private instances. They cannot have public IP addresses because they are in a private subnet.

You add an internet gateway to the VCN. You also add a route rule in the public subnet's route table that directs internet-bound traffic to the internet gateway. The public subnet's instances can now initiate connections to the internet and also receive inbound connections initiated from the internet. Remember that you can use [security rules](#) to control the types of traffic that are allowed in and out of the instances at the packet level.

You add a NAT gateway to the VCN. You also add a route rule in the private subnet's route table that directs internet-bound traffic to the NAT gateway. The private subnet's instances can now initiate connections to the internet. The NAT gateway allows responses, but it does not allow connections that are *initiated from the internet*. Without that NAT gateway, the private instances would instead need to be in the public subnet and have public IP addresses to get their software updates.

The following diagram illustrates the basic network layout for the example. The arrows indicate whether connections can be initiated in only one direction or both.



**Note**

A NAT gateway can be used only by resources in the gateway's own VCN. If the VCN is [peered with another](#), resources in the other VCN cannot access the NAT gateway.

Also, resources in an on-premises network connected to the NAT gateway's VCN with [FastConnect](#) or an [IPSec VPN](#) cannot use the NAT gateway.

Here are a few basics about NAT gateways:

- The NAT gateway supports TCP, UDP, and ICMP ping traffic.
- The gateway supports a maximum of approximately 20,000 concurrent connections to a single destination address and port.
- The Networking service automatically assigns a public IP address to the NAT gateway. You can't choose the public IP address or use one of your [reserved public IP addresses](#).
- There's a limit on the number of NAT gateways per VCN. You can request an increase to that limit. See [Service Limits](#).

### Routing for a NAT Gateway

You control routing in your VCN at the subnet level, so you can specify which subnets in your VCN use a NAT gateway. You can have more than one NAT gateway on a VCN (although you must [request an increase in your limits](#)). For example, if you want an external application to distinguish traffic from the VCN's different subnets, you could set up a different NAT gateway (and thus a different public IP address) for each subnet. A given subnet can route traffic to only a single NAT gateway.

### Blocking Traffic Through a NAT Gateway

You create a NAT gateway in the context of a specific VCN. In other words, the NAT gateway is automatically always attached to only one VCN of your choice. However, you can block or allow traffic through the NAT gateway at any time. By default, the gateway allows traffic upon creation. Blocking the NAT gateway prevents all traffic from flowing, regardless of any existing route rules or security rules in your VCN. For instructions on how to block traffic, see [To block/allow traffic for a NAT gateway](#).

### Transitioning to a NAT Gateway

If you're switching from using a [NAT instance in your VCN](#) to a NAT gateway, consider that the public IP address for your NAT device will change.

If you're switching from using an internet gateway to a NAT gateway, the instances with access to the NAT gateway no longer need public IP addresses to reach the internet. Also, the

instances no longer need to be in a public subnet. You can't switch a subnet from [public to private](#). However, you can [delete the ephemeral public IPs](#) from your instances if you like.

### Deleting a NAT Gateway

To delete a NAT gateway, its traffic does not have to be blocked, but there must not be a route table that lists it as a target. For instructions, see [To delete a NAT gateway](#).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

## Setting Up a NAT Gateway

### Task 1: Create the NAT gateway

1. In the Console, confirm you're viewing the compartment that contains the VCN that you want to add the NAT gateway to. For information about compartments and access control, see [Access Control](#).
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
3. Click the VCN you're interested in.
4. Under **Resources**, click **NAT Gateways**.
5. Click **Create NAT Gateway**.
6. Enter the following values:

- **Name:** A friendly name for the NAT gateway. It doesn't have to be unique. Avoid entering confidential information.
  - **Create in compartment:** The compartment where you want to create the NAT gateway, if different from the compartment you're currently working in.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create NAT Gateway**.

The NAT gateway is then created and displayed on the **NAT Gateways** page in the compartment you chose. The gateway allows traffic by default. At any time, you can [block or allow traffic](#) through it.

### Task 2: Update routing for the subnet

When you create a NAT gateway, you must also create a route rule that directs the desired traffic from the subnet to the NAT gateway. You do this for each subnet that needs to access the gateway.

1. Determine which subnets in your VCN need access to the NAT gateway.
2. For each of those subnets, update the subnet's route table to include a new rule:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.
  - c. Under **Resources**, click **Route Tables**.
  - d. Click the route table you're interested in.
  - e. Click **Add Route Rule** and enter the following values:

- **Target Type:** NAT Gateway.
  - **Destination CIDR Block:** 0.0.0.0/0
  - **Compartment:** The compartment where the NAT gateway is located.
  - **Target NAT Gateway:** The NAT gateway.
- f. Click **Add Route Rule**.

Any subnet traffic with a destination that matches the rule is routed to the NAT gateway. For more information about setting up route rules, see [Route Tables](#).

Later, if you no longer need the NAT gateway and want to delete it, you must first delete all the route rules in your VCN that specify the NAT gateway as the target.



### Tip

Without the required routing, traffic doesn't flow over the NAT gateway. If a situation occurs where you need to temporarily stop the traffic flow over the gateway, you can simply remove the route rule that enables traffic. Or you can [block traffic through the gateway](#) entirely. You do not need to delete it.

## Using the Console

### To create a NAT gateway

See the instructions in [Task 1: Create the NAT gateway](#).

### To block/allow traffic for a NAT gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click

### **Virtual Cloud Networks.**

2. Click the VCN you're interested in.
3. Under **Resources**, click **NAT Gateways**.
4. For the NAT gateway you're interested in, click the Actions icon (three dots) and then click **Block Traffic** (or **Allow Traffic** if you're enabling traffic for the NAT gateway).
5. Confirm when prompted.  
When the traffic is blocked, the NAT gateway's icon turns gray, and the label changes to **BLOCKED**. When the traffic is allowed, the NAT gateway's icon turns green, and the label changes to **AVAILABLE**.

### To update a NAT gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **NAT Gateways**.
4. For the NAT gateway you're interested in, click the Actions icon (three dots), and then click **Edit**.
5. Make your changes and click **Save Changes**.

### To delete a NAT gateway

Prerequisite: There must not be a route table that lists the NAT gateway as a target.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **NAT Gateways**.

4. For the NAT gateway you want to delete, click the Actions icon (three dots), and then click **Terminate**.
5. Confirm when prompted.

### To manage tags for a NAT gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **NAT Gateways**.
4. Click the Actions icon (three dots) for the NAT gateway, and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move a NAT gateway to a different compartment

You can move a NAT gateway from one compartment to another. When you move a NAT gateway to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. In **Resources**, click **NAT Gateways**.
4. Find the NAT gateway in the list, click the the Actions icon (three dots), and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

The NAT gateway moves to the new compartment immediately. Depending on your permissions, you can select the compartment in the left side menu to view the NAT gateway.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your NAT gateways, use these operations:

- [ListNatGateways](#)
- [CreateNatGateway](#)
- [GetNatGateway](#)
- [UpdateNatGateway](#)
- [DeleteNatGateway](#)
- [ChangeNatGatewayCompartment](#)

To manage route tables, see [Route Tables](#).

## Access to Oracle Services: Service Gateway

This topic describes how to set up and manage a service gateway. A service gateway enables cloud resources without public IP addresses to privately access Oracle services.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Access to Oracle Services

The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. These services have [public IP addresses](#) that you typically reach over the internet. However, you can access the Oracle Services Network *without the traffic going over the internet*. There are different ways, depending on which of your hosts need the access:

- **Hosts in your on-premises network:**
  - [Private access through a VCN with FastConnect private peering or VPN Connect](#): The on-premises hosts use private IP addresses and reach the Oracle Services Network by way of the VCN and the VCN's service gateway.
  - [Public access with FastConnect public peering](#): The on-premises hosts use public IP addresses.
- **Hosts in your VCN:**
  - Private access through a service gateway: This is the scenario covered in this topic. The VCN's hosts use private IP addresses.

### Highlights

- A service gateway lets your virtual cloud network (VCN) privately access specific Oracle services without exposing the data to the public internet. No internet gateway or NAT is required to reach those specific services. The resources in the VCN can be in a private subnet and use only private IP addresses. The traffic from the VCN to the Oracle service travels over the Oracle network fabric and never traverses the internet.
- The service gateway is regional and enables access only to supported Oracle services *in the same region* as the VCN.
- The service gateway allows access to supported Oracle services within the region to protect your data from the internet. Your workloads may require access to public endpoints or services not supported by the service gateway (for example, to download updates or patches). Ensure you have a [NAT gateway](#) or other access to the internet if necessary.

- The supported Oracle services are Oracle Cloud Infrastructure Object Storage and others in the Oracle Services Network. For a list, see [Service Gateway: Supported Cloud Services in Oracle Services Network](#).
- The service gateway uses the concept of a *service CIDR label*, which is a string that represents all the regional public IP address ranges for the service or group of services of interest (for example, *OCI PHX Object Storage* is the string for Object Storage in US West (Phoenix)). You use that service CIDR label when you configure the service gateway and related route rules to control traffic to the service. You can optionally use it when configuring [security rules](#). If the service's public IP addresses change in the future, you don't have to adjust those rules.
- You can set up a VCN so that your on-premises network has *private access* to Oracle services by way of the VCN and the VCN's service gateway. The hosts in your on-premises network communicate with their private IP addresses and the traffic does not go over the internet. For more information, see [Transit Routing: Private Access to Oracle Services](#)

### Overview of Service Gateways

A service gateway lets resources in your VCN privately access specific Oracle services, without exposing the data to an internet gateway or NAT. The resources in the VCN can be in a private subnet and use only private IP addresses. The traffic from the VCN to the service of interest travels over the Oracle network fabric and never traverses the internet.

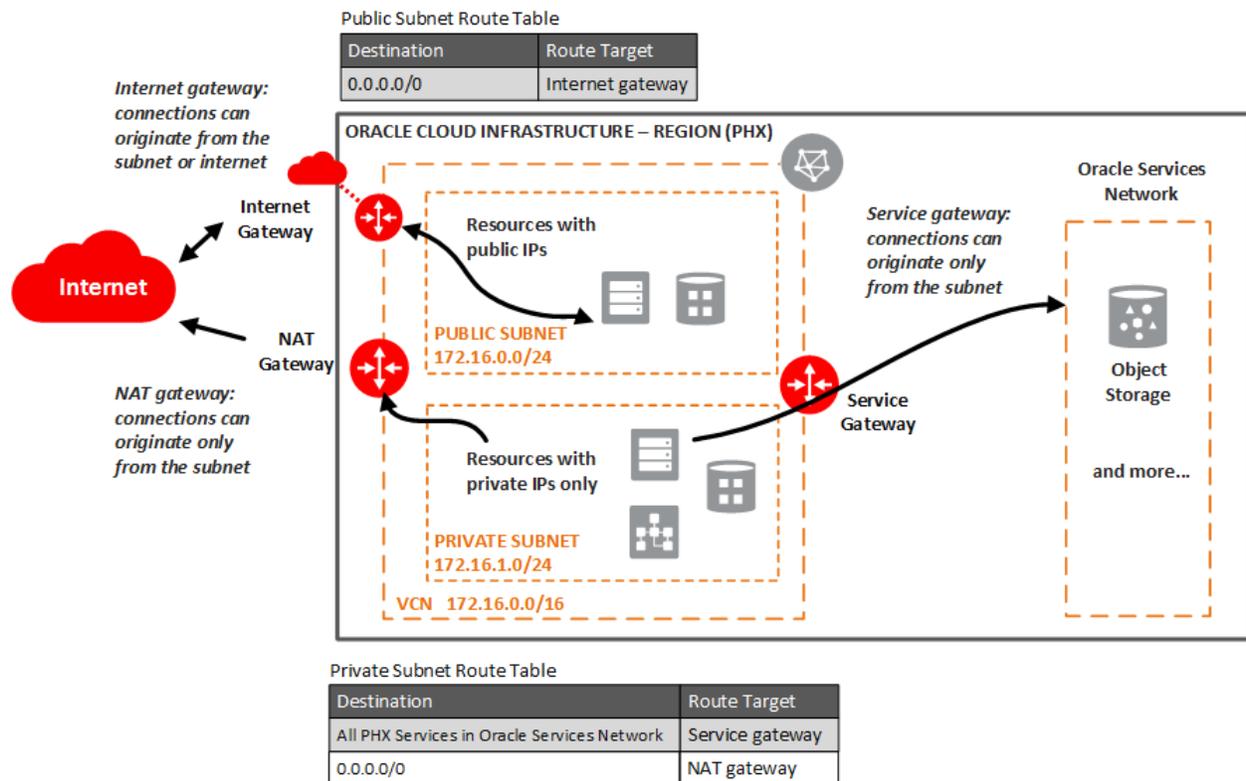
The following simple diagram illustrates a VCN that has both a public subnet and a private subnet. Resources in the private subnet have only private IP addresses.

The VCN has three gateways:

- **Internet gateway:** To provide the public subnet direct access to public endpoints on the internet. Connections can be initiated from the subnet or from the internet. The resources in the public subnet must have public IP addresses. For more information, see [Internet Gateway](#).
- **Service gateway:** To provide the private subnet with private access to supported Oracle services within the region. Connections can be initiated only from the subnet.

- **NAT gateway:** To provide the private subnet with private access to public endpoints on the internet. Connections can be initiated only from the subnet. For more information, see [NAT Gateway](#).

You control routing in your VCN at the subnet level, so you can specify which subnets in your VCN use each gateway. In the diagram, the route table for the public subnet sends non-local traffic through the internet gateway. The route table for the private subnet sends traffic destined for the Oracle services through the service gateway. It sends all remaining traffic to the NAT gateway.





### Important

See this [known issue](#) for information about configuring route rules with *service gateway* as the target on route tables associated with public subnets.

A service gateway can be used by resources in the gateway's own VCN. However, if the VCN is [peered with another](#), resources in the other VCN cannot access the service gateway.

Resources in your on-premises network that is connected to the service gateway's VCN with [FastConnect](#) or [VPN Connect](#) can also use the service gateway. For more information, see [Transit Routing: Private Access to Oracle Services](#).

Notice that your on-premises network can *also* use [FastConnect public peering](#) for private access to public Oracle services. That means that your on-premises network could have multiple paths to access Oracle services public IP address ranges. If that is the case, your edge device receives route advertisement of the Oracle services public IP address ranges over multiple paths. For important information about configuring your edge device correctly, see [Routing Details for Connections to Your On-Premises Network](#).

A VCN can have only one service gateway. For more information about limits, see [Service Limits](#).

For instructions on setting up a service gateway, see [Setting Up a Service Gateway in the Console](#).

### About Service CIDR Labels

Each Oracle service has a regional public endpoint that uses public IP addresses for access. When you set up a service gateway with access to an Oracle service, you also set up Networking service [route rules](#) and optionally [security rules](#) that control traffic with the service. That means you need to know the service's public IP addresses to set up those rules. To make it easier for you, the Networking service uses *service CIDR labels* to represent all the public CIDRs for a given Oracle service or a group of Oracle services. If a service's CIDRs change in the future, you don't have to adjust your route rules or security rules.

Examples:

- **OCI PHX Object Storage** is a service CIDR label that represents all the Object Storage CIDRs in the US West (Phoenix) region.
- **All PHX Services in Oracle Services Network** is a service CIDR label that represents all the CIDRs for the supported services in the Oracle Services Network in the US West (Phoenix) region. For a list of the services, see [Service Gateway: Supported Cloud Services in Oracle Services Network](#).

As you can see, a service CIDR label can be associated with either a single Oracle service (example: Object Storage), or multiple Oracle services.

The term *service* is often used in this topic in place of the more accurate term *service CIDR label*. The important thing to remember is that when you set up a service gateway (and related route rules), you specify the *service CIDR label* you're interested in. In the Console, you're presented with the available service CIDR labels. If you use the REST API, the [ListServices](#) operation returns the available `Service` objects. The `Service` object's `cidrBlock` attribute contains the service CIDR label (example: `all-phx-services-in-oracle-services-network`).

### Available Service CIDR Labels

Here are the available service CIDR labels:

- **OCI *<region>* Object Storage:** For information about the service, see [Overview of Object Storage](#)
- **All *<region>* Services in Oracle Services Network:** For a list of supported services, see [Service Gateway: Supported Cloud Services in Oracle Services Network](#).



#### Important

See this [known issue](#) for information about accessing Oracle YUM services through the service gateway.

### Enabling a Service CIDR Label for a Service Gateway

To give your VCN access to a given service CIDR label, you must *enable* that service CIDR label for the VCN's service gateway. You can do that when you create the service gateway, or later after it's created. You can also disable a service CIDR label for the service gateway at any time.



#### Important

Because Object Storage is covered by both **OCI <region> Object Storage** and **All <region> Services in Oracle Services Network**, a service gateway can use **only one of those service CIDR labels**. Likewise, a route table can have a single rule for one of the service CIDR labels. It cannot have two separate rules, one for each label.

If the service gateway is configured to use **All <region> Services in Oracle Services Network**, the route rule can use either CIDR label. However, if the service gateway is configured to use **OCI <region> Object Storage** and the route rule uses **All <region> Services in Oracle Services Network**, traffic to services in the Oracle Services Network except Object Storage will be blackholed. The Console prohibits you from configuring the service gateway and corresponding route table in that manner.

If you want to switch the service gateway to use a different service CIDR label, see [To switch to a different service CIDR label](#).

### Blocking Traffic Through a Service Gateway

You create a service gateway in the context of a specific VCN. In other words, the service gateway is always attached to that one VCN. However, you can block or allow traffic through the service gateway at any time. By default, the gateway allows traffic flow upon creation. Blocking the service gateway traffic prevents all traffic from flowing, regardless of what service CIDR labels are enabled, or any existing route rules or security rules in your VCN. For instructions on how to block traffic, see [To block or allow traffic for a service gateway](#).

### Route Rules and Security Rules for a Service Gateway

For traffic to be routed from a subnet in your VCN to a service gateway, you must add a rule accordingly to the subnet's route table. The rule must use the service gateway as the target. For the destination, you must use the [service CIDR label](#) that is enabled for the service gateway. This means that you don't have to know the specific public CIDRs, which could change over time.

Any traffic leaving the subnet and destined for the service's public CIDRs is then routed to the service gateway. If the service gateway traffic is blocked, no traffic flows through it even if there's a route rule that matches the traffic. For instructions on setting up route rules for a service gateway, see [Task 2: Update routing for the subnet](#).

The VCN's security rules must also allow the desired traffic. If you like, you can use a service CIDR label instead of a CIDR for the source or destination of the desired traffic. Again, this means that you don't have to know the specific public CIDRs for the service. For convenience, you can use a service CIDR label in security rules even if your VCN doesn't have a service gateway, and the traffic to the services uses an internet gateway.

You can use [stateful or stateless security rules](#) that use a service CIDR label:

- **For stateful rules:** Create an egress rule with the destination service = the service CIDR label of interest. As with any security rule, you can specify other items such as the IP protocol and source and destination ports.
- **For stateless rules:** You must have both egress and ingress rules. Create an egress rule with the destination service = the service CIDR label of interest. Also create an ingress rule with the source service = the service CIDR label of interest. As with any

security rule, you can specify other items such as the IP protocol and source and destination ports.

For instructions on setting up security rules that use a service CIDR label, see [Task 3: \(Optional\) Update security rules](#).

### **Object Storage: Allowing Bucket Access from Only a Particular VCN or CIDR Range**

If you use a service gateway to access Object Storage, you can write an IAM policy that allows access to a particular Object Storage bucket only if these requirements are met:

- The request goes through a service gateway.
- The request originates from the particular VCN or CIDR (for example, a subnet within a VCN) that is specified in the policy.

For examples of this particular type of IAM policy, and important caveats about its use, see [Task 4: \(Optional\) Update IAM Policies to Restrict Object Storage Bucket Access](#).

### **Deleting a Service Gateway**

To delete a service gateway, its traffic does not have to be blocked, but there must not be a route table that lists it as a target. See [To delete a service gateway](#).

### **Required IAM Policy**

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: see [IAM Policies for Networking](#).

### Setting Up a Service Gateway in the Console

Following is the process for setting up a service gateway. It assumes that you already have a VCN with a subnet (either private or public).



#### Important

The service gateway allows access to supported Oracle services within the region to protect your data from the internet. Your applications may require access to public endpoints or services not supported by the service gateway (for example, to download updates or patches). Ensure you have a [NAT gateway](#) or other access to the internet if necessary.

#### Task 1: Create the service gateway

1. In the Console, confirm you're viewing the compartment that contains the VCN that you want to add the service gateway to. For information about compartments and access control, see [Access Control](#).
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
3. Click the VCN you're interested in.
4. On the left side of the page, click **Service Gateways**.
5. Click **Create Service Gateway**.
6. Enter the following values:
  - **Name:** A descriptive name for the service gateway. It doesn't have to be unique. Avoid entering confidential information.
  - **Create in compartment:** The compartment where you want to create the service gateway, if different from the compartment you're currently working in.

- **Services:** Optionally select the [service CIDR label](#) you're interested in. If you don't select one now, you can later update the service gateway and add a service CIDR label then. Without at least one service CIDR label enabled for the gateway, no traffic flows through it.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Service Gateway**.

The service gateway is then created and displayed on the **Service Gateways** page in the compartment you chose. The gateway allows traffic through it by default. At any time, you can [block or allow the traffic through it](#).

### Task 2: Update routing for the subnet

When you configure a service gateway for a particular service CIDR label, you must also create a route rule that specifies that label as the destination and the target as the service gateway. You do this for each subnet that needs to access the gateway.

1. Determine which subnets in your VCN need access to the service gateway.
2. For each of those subnets, update the subnet's route table to include a new rule:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.
  - c. Under **Resources**, click **Route Tables**.
  - d. Click the route table you're interested in.
  - e. Click **Edit Route Rules**.
  - f. Click **Add Route Rule** and enter the following values:

- **Target Type:** Service Gateway.
  - **Destination Service:** The [service CIDR label](#) that is enabled for the gateway.
  - **Compartment:** The compartment where the service gateway is located.
  - **Target:** The service gateway.
- g. Click **Save**.

Any subnet traffic with a destination that matches the rule is routed to the service gateway. For more information about setting up route rules, see [Route Tables](#).

Later, if you no longer need the service gateway and want to delete it, you must first delete all the route rules in your VCN that specify the service gateway as the target.



### Tip

Without the required routing, traffic doesn't flow over the service gateway. If a situation occurs where you want to temporarily stop the traffic flow over the gateway to a particular service, you can simply remove the route rule that enables traffic. You can also disable that particular service CIDR label for the gateway. Or you can [block all traffic through the service gateway](#) entirely. You do not have to delete the gateway.

### Task 3: (Optional) Update security rules

When you configure a service gateway to access a service CIDR label, you must also ensure that the security rules are configured to allow the desired traffic. Your security rules might already allow this traffic, which is why this task is optional. The following procedure assumes you are using security lists to implement your security rules. The procedure describes how to set up a rule that uses the service CIDR label. You do this for each subnet that needs to access

the gateway.



### Tip

Security lists are one way to control traffic in and out of the VCN's resources. You can also use [network security groups](#), which let you apply a set of security rules to a set of resources that all have the same security posture.

1. Determine which subnets in your VCN need to connect to the services you're interested in.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic with the particular service:
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in.
  - c. Click **Edit All Rules** and create one or more rules, each for the specific type of traffic you want to allow. See the following example for more details.

### Example

Let's say you want to add a stateful rule that enables egress HTTPS (TCP port 443) traffic from the subnet to both Object Storage and Oracle YUM repos. Here are the basic steps you take when adding a rule:

- i. In the **Allow Rules for Egress** section, click **+Add Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Destination Type:** Service.
- iv. **Destination Service:** The [service CIDR label](#) that you're interested in. To access both Object Storage and Oracle YUM repos, it's **All <region> Services in Oracle Services Network**.

- v. **IP Protocol:** Leave as TCP.
  - vi. **Source Port Range:** Leave as All.
  - vii. **Destination Port Range:** Enter 443.
- d. Click **Save Security List Rules** at the bottom of the dialog box.

For more information about setting up security rules, see [Security Rules](#).

### Task 4: (Optional) Update IAM Policies to Restrict Object Storage Bucket Access

This task is applicable only if you're using a service gateway to access Object Storage. You can optionally write an IAM policy to allow only the resources in a specific VCN to write objects to a particular bucket.



#### Important

If you use one of the following IAM policies to restrict access to a bucket, the bucket is *not accessible from the Console*. It's accessible only from within the specific VCN or CIDR block.

Also, the IAM policies allow requests to Object Storage only if they go from the specified VCN or CIDR block *through the service gateway*. If they go through the internet gateway, the requests are denied.

The following example lets resources in the example ObjectBackup group write objects to an existing bucket called db-backup that resides in a compartment called ABC.

```
Allow group ObjectBackup to read buckets in compartment ABC
```

```
Allow group ObjectBackup to manage objects in compartment ABC where
 all {target.bucket.name='db-backup',
```

## CHAPTER 23 Networking

```
request.vcn.id='<VCN_OCID>',
any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}}
```

In addition to specifying a VCN's OCID, you can specify an IP address or ranges of addresses within the VCN. For example, the next version of the policy includes a `request.ipv4.ipaddress` variable with a value of `10.0.1.0/24`. If the resource making the request has an IP address that is not in this CIDR block, the request is denied.

Allow group ObjectBackup to read buckets in compartment ABC

```
Allow group ObjectBackup to manage objects in compartment ABC where
 all {target.bucket.name='db-backup',
 request.vcn.id='<VCN_OCID>',
 request.ipv4.ipaddress in ['10.0.1.0/24'],
 any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}}
```

You can specify multiple VCNs in the policy. The next example has OCIDs for two VCNs. You might do this if you've [set up your on-premises network with private access to Oracle services](#) through a VCN, and you've also set up one or more *other* VCNs with their own service gateways. For more information, see [Multiple VCNs with Private Access to Oracle Services](#).

Allow group ObjectBackup to read buckets in compartment ABC

```
Allow group ObjectBackup to manage objects in compartment ABC where
 all {target.bucket.name='db-backup',
 any {request.vcn.id='<VCN_OCID_1>', request.vcn.id='<VCN_OCID_2>'},
 any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}}
```

## Managing a Service Gateway in the Console

### To create a service gateway

See the instructions in [Task 1: Create the service gateway](#).

### To switch to a different service CIDR label

To avoid disrupting your Object Storage connections while switching between the **OCI**

**<region> Object Storage** service CIDR label and **All <region> Services in Oracle Services Network**, use the following process:

1. [Update the service gateway](#): Remove the existing service CIDR label, and then add the one you want to switch to. You can't enable both labels for the service gateway.
2. [Update relevant route rules](#): For each rule that uses the service gateway as the target, switch the rule's destination service from the existing service CIDR label to the one you want to switch to.
3. Update relevant security rules: Change any security rules that specify the existing service CIDR label to instead use the one you want to switch to. The rules can be in [network security groups](#) or [security lists](#).

If you instead delete your existing service gateway and create a new one, your Object Storage connections will be disrupted. Remember that before you can delete a service gateway, you must delete any route rules that specify that gateway as a target.

### To update a service gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Service Gateways**.
4. For the service gateway you're interested in, click the Actions icon (three dots), and then click **Edit**.
5. Make your changes and click **Save**.

### To block or allow traffic for a service gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.

3. Click **Service Gateways**.
4. For the service gateway you're interested in, click the Actions icon (three dots), and then click **Block Traffic** (or **Allow Traffic** if you're enabling traffic for the service gateway).  
When the traffic is blocked, the service gateway's icon turns gray, and the label changes to BLOCKED. When the traffic is allowed, the service gateway's icon turns green, and the label changes to AVAILABLE.

### To associate a route table with an existing service gateway

You perform this task only if you're implementing an [advanced transit routing scenario](#).

A service gateway can exist without a route table associated with it. However, after you associate a route table with a service gateway, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

**Prerequisite:** The route table must exist and belong to the VCN that the service gateway belongs to.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Service Gateways**.
4. For the service gateway you're interested in, click the Actions icon (•••), and then click either:
  - **Associate With Route Table:** If the service gateway has no route table associated with it yet.
  - **Associate Different Route Table:** If you're changing which route table is associated with the service gateway.
5. Select the compartment where the route table resides, and select the route table itself.
6. Click **Associate**.

### To delete a service gateway

Prerequisite: The service gateway does not have to block traffic, but there must not be a route table that lists it as a target.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Service Gateways**.
4. For the service gateway you're interested in, click the Actions icon (three dots), and then click **Delete**.
5. Confirm when prompted.

### To manage tags for a service gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. For the service gateway you're interested in, click the Actions icon (three dots), and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move a service gateway to a different compartment

You can move a service gateway from one compartment to another. When you move a service gateway to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.

3. In **Resources**, click **Service Gateways**.
4. Find the service gateway in the list, click the the Actions icon (three dots), and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

The service gateway moves to the new compartment immediately. Depending on your permissions, you can select the compartment in the left side menu to view the service gateway.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Managing a Service Gateway with the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).



#### Warning

If anyone in your organization implements a service gateway, be aware that you **may need to update any client code that works with Networking service route rules and security lists**. There are possible breaking API changes. For more information, see the [service gateway release notes](#).

To manage your service gateways, use these operations:

- [ListServiceGateways](#)
- [CreateServiceGateway](#)

- [GetServiceGateway](#)
- [UpdateServiceGateway](#)
- [DeleteServiceGateway](#)
- [ChangeServiceGatewayCompartment](#)
- [ListServices](#): Use this to determine the available [service CIDR labels](#).
- [GetService](#): Gets the details for a particular [service CIDR label](#).
- [AttachServiceId](#): Enables a [service CIDR label](#) for a service gateway.
- [DetachServiceId](#): Disables a [service CIDR label](#) for a service gateway.

To manage route tables, see [Route Tables](#). To manage security lists, see [Security Lists](#). To manage network security groups, see [Network Security Groups](#). To manage IAM policies, see [Managing Policies](#).

## Access to Other VCNs: Peering

VCN peering is the process of connecting multiple virtual cloud networks (VCNs). There are two types of VCN peering:

- [Local VCN peering \(within region\)](#)
- [Remote VCN peering \(across regions\)](#)

You can use VCN peering to divide your network into multiple VCNs (for example, based on departments or lines of business), with each VCN having direct, private access to the others. There's no need for traffic to flow over the internet or through your on-premises network by way of an [IPSec VPN](#) or [FastConnect](#). You can also place shared resources into a single VCN that all the other VCNs can access privately.

Because remote VCN peering crosses regions, you can use it (for example) to mirror or back up your databases in one region to another.

### Important Implications of Peering

This section summarizes some access control, security, and performance implications for peered VCNs. In general, you can control access and traffic between two peered VCNs by using IAM policies, route tables in each VCN, and security lists in each VCN.

#### Controlling the Establishment of Peerings

With IAM policies, you can control:

- Who can [subscribe your tenancy to another region](#) (required for remote VCN peering).
- Who in your organization has the authority to establish VCN peerings (for example, see the IAM policies in [Setting Up a Local Peering](#) and [Setting Up a Remote Peering](#)). Be aware that deletion of these IAM policies does not affect any existing peerings, only the ability for future peerings to be created.
- Who can [manage route tables](#) and [security lists](#).

#### Controlling Traffic Flow Over the Connection

Even if a peering connection has been established between your VCN and another, you can control the packet flow over the connection with route tables in your VCN. For example, you can restrict traffic to only specific subnets in the other VCN.

Without terminating the peering, you can stop traffic flow to the other VCN by simply removing route rules that direct traffic from your VCN to the other VCN. You can also effectively stop the traffic by removing any security list rules that enable ingress or egress traffic with the other VCN. This doesn't stop traffic flowing over the peering connection, but stops it at the [VNIC](#) level.

For more information about the routing and security lists, see the discussions in these sections:

Local VCN peering:

- [Important Local Peering Concepts](#)
- [Task E: Configure the route tables](#)

- [Task F: Configure the security rules](#)

Remote VCN peering:

- [Important Remote Peering Concepts](#)
- [Task E: Configure the route tables](#)
- [Task F: Configure the security rules](#)

### **Controlling the Specific Types of Traffic Allowed**

It's important that each VCN administrator ensure that all outbound and inbound traffic with the other VCN is intended/expected and well defined. In practice, this means implementing security list rules that explicitly state the types of traffic your VCN can send to the other and accept from the other.



### Important

Your instances running Oracle-provided Linux images or Windows images also have OS firewall rules that control access to the instance. When troubleshooting access to an instance, make sure that all of the following items are set correctly:

- The rules in the network security groups that the instance is in
- The rules in the security lists associated with the instance's subnet
- The instance's OS firewall rules

For more information, see [Oracle-Provided Images](#).

If your instance is running Oracle Linux 7, you need to use [firewalld](#) to interact with the iptables rules. For your reference, here are commands for opening a port (1521 in this example):

```
sudo firewall-cmd --zone=public --permanent --add-
port=1521/tcp

sudo firewall-cmd --reload
```

For instances with an iSCSI boot volume, the preceding `--reload` command can cause problems. For details and a workaround, see [Instances experience system hang after running firewall-cmd --reload](#).

In addition to security lists and firewalls, you should evaluate other OS-based configuration on the instances in your VCN. There could be default configurations that don't apply to your own VCN's CIDR, but inadvertently apply to the other VCN's CIDR.

### Using Default Security List Rules

If your VCN's subnets use the [default security list](#) with the default rules it comes with, be aware that there are two rules that allow ingress traffic from anywhere (that is, 0.0.0.0/0, and thus the other VCN):

- Stateful ingress rule that allows TCP port 22 (SSH) traffic from 0.0.0.0/0 and any source port
- Stateful ingress rule that allows ICMP type 3, code 4 traffic from 0.0.0.0/0 and any source port

Make sure to evaluate these rules and whether you want to keep or update them. As stated earlier, you should ensure that all inbound or outbound traffic that you permit is intended/expected and well defined.

### Preparing for Performance Impact and Security Risks

In general, you should prepare your VCN for the ways it could be affected by the other VCN. For example, the load on your VCN or its instances could increase. Or your VCN could experience a malicious attack directly from or by way of the other VCN.

Regarding performance: If your VCN is providing a service to another, be prepared to scale up your service to accommodate the demands of the other VCN. This might mean being prepared to launch additional instances as necessary. Or if you're concerned about high levels of network traffic coming to your VCN, consider using [stateless security list rules](#) to limit the level of connection tracking your VCN must perform. Stateless security list rules can also help slow the impact of a denial-of-service (DoS) attack.

Regarding security risks: You can't necessarily control whether the other VCN is connected to the internet. If it is, be aware that your VCN can be exposed to bounce attacks in which a malicious host on the internet can send traffic to your VCN but make it look like it's coming from the VCN you're peered with. To guard against this, as mentioned earlier, use your security lists to carefully limit the inbound traffic from the other VCN to expected and well-defined traffic.

## Local VCN Peering (Within Region)

This topic is about *local VCN peering*. In this case, *local* means that the VCNs reside *in the same region*. If the VCNs are in different regions, see [Remote VCN Peering \(Across Regions\)](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Overview of Local VCN Peering

*Local VCN peering* is the process of connecting two VCNs in the same region so that their resources can communicate using private IP addresses without routing the traffic over the internet or through your on-premises network. The VCNs can be in the same Oracle Cloud Infrastructure tenancy or different ones. Without peering, a given VCN would need an [internet gateway](#) and public IP addresses for the instances that need to communicate with another VCN.

For more information, see [Access to Other VCNs: Peering](#).

### Summary of Networking Components for Peering

At a high level, the Networking service components required for a local peering include:

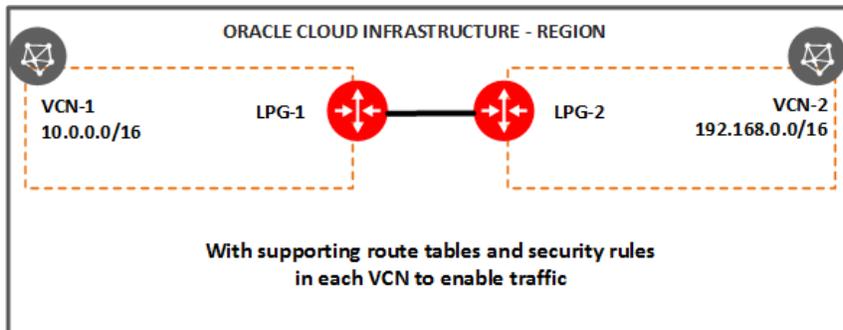
- Two VCNs with non-overlapping CIDRs, in the same region
- A *local peering gateway (LPG)* on each VCN in the peering relationship.
- A *connection* between those two LPGs.
- Supporting route rules to enable traffic to flow over the connection, and only to and from select subnets in the respective VCNs (if desired).

## CHAPTER 23 Networking

---

- Supporting security rules to control the types of traffic allowed to and from the instances in the subnets that need to communicate with the other VCN.

The following diagram illustrates the components.





### Note

A given VCN can use the peered LPGs to reach these resources:

- VNICs in the other VCN
- An on-premises network attached to the other VCN, if an advanced routing scenario called [transit routing](#) has been set up for the VCNs

A VCN can't use its peered VCN to reach other destinations outside of the VCNs (such as the internet). For example, if VCN-1 in the preceding diagram were to have an internet gateway, the instances in VCN-2 could not use it to send traffic to endpoints on the internet. However, be aware that VCN-2 could *receive* traffic from the internet by way of VCN-1. For more information, see [Important Implications of VCN Peering](#).

### Explicit Agreement Required from Both Sides

Peering involves two VCNs that might be owned by the same party or two different ones. The two parties might both be in your company but in different departments. Or the two parties might be in entirely different companies (for example, in a service-provider model).

Peering between two VCNs requires explicit agreement from both parties in the form of Oracle Cloud Infrastructure Identity and Access Management policies that each party implements for their own VCN's compartment or tenancy. If the VCNs are in different tenancies, each administrator must provide their tenancy [OCID](#) and put in place special policy statements to enable the peering.

### Advanced Scenario: Transit Routing

There's an advanced routing scenario called [transit routing](#) that enables communication between an on-premises network and multiple VCNs over a *single* [Oracle Cloud Infrastructure FastConnect](#) or [IPSec VPN](#). The VCNs must be in the same region and locally peered in a hub-and-spoke layout. As part of the scenario, the VCN that is acting as the hub has a route table *associated with each LPG* (typically route tables are associated with a VCN's subnets).

When you create an LPG, you can optionally associate a route table with it. Or if you already have an existing LPG without a route table, you [can associate a route table with it](#). The route table must belong to the LPG's VCN. A route table associated with an LPG can contain only rules that use the VCN's [attached DRG](#) as the target.

An LPG can exist without a route table associated with it. However, after you associate a route table with an LPG, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

### Important Local Peering Concepts

The following concepts help you understand the basics of VCN peering and how to establish a local peering.

#### PEERING

A *peering* is a single peering relationship between two VCNs. Example: If VCN-1 peers with three other VCNs, then there are three peerings. The *local* part of *local peering* indicates that the VCNs are in the same region. A given VCN can have a maximum of ten local peerings at a time.



### Warning

The two VCNs in the peering relationship must not have overlapping CIDRs. However, if VCN-1 is peered with three other VCNs, those three VCNs can have overlapping CIDRs with each other. You would set up the subnets in VCN-1 to have route rules that direct traffic to the targeted peered VCN.

### VCN ADMINISTRATORS

In general, VCN peering can occur only if both of the VCN administrators agree to it. In practice, this means that the two administrators must:

- Share some basic information with each other.
- Coordinate to set up the required Oracle Cloud Infrastructure Identity and Access Management policies to enable the peering.
- Configure their VCNs for the peering.

Depending on the situation, a single administrator might be responsible for both VCNs and the related policies.

For more information about the required policies and VCN configuration, see [Setting Up a Local Peering](#).

### ACCEPTOR AND REQUESTOR

To implement the IAM policies required for peering, the two VCN administrators must designate one administrator as the *requestor* and the other as the *acceptor*. The requestor must be the one to initiate the request to connect the two LPGs. In turn, the acceptor must create a particular IAM policy that gives the requestor permission to connect to LPGs in the acceptor's compartment. Without that policy, the requestor's request to connect fails.

### LOCAL PEERING GATEWAY (LPG)

A *local peering gateway (LPG)* is a component on a VCN for routing traffic to a locally peered VCN. As part of configuring the VCNs, each administrator must create an LPG for their VCN. A given VCN must have a separate LPG for each local peering it establishes (maximum ten LPGs per VCN). To continue with the previous example: VCN-1 would have three LPGs to peer with three other VCNs. In the API, a [LocalPeeringGateway](#) is an object that contains information about the peering. You can't reuse an LPG to later establish another peering.

### PEERING CONNECTION

When the requestor initiates the request to peer (in the Console or API), they're effectively asking to *connect the two LPGs*. The requestor must have information to identify each LPG (such as the LPG's compartment and name, or the LPG's OCID). Each administrator must put the required IAM policies in place for their compartment or tenancy.

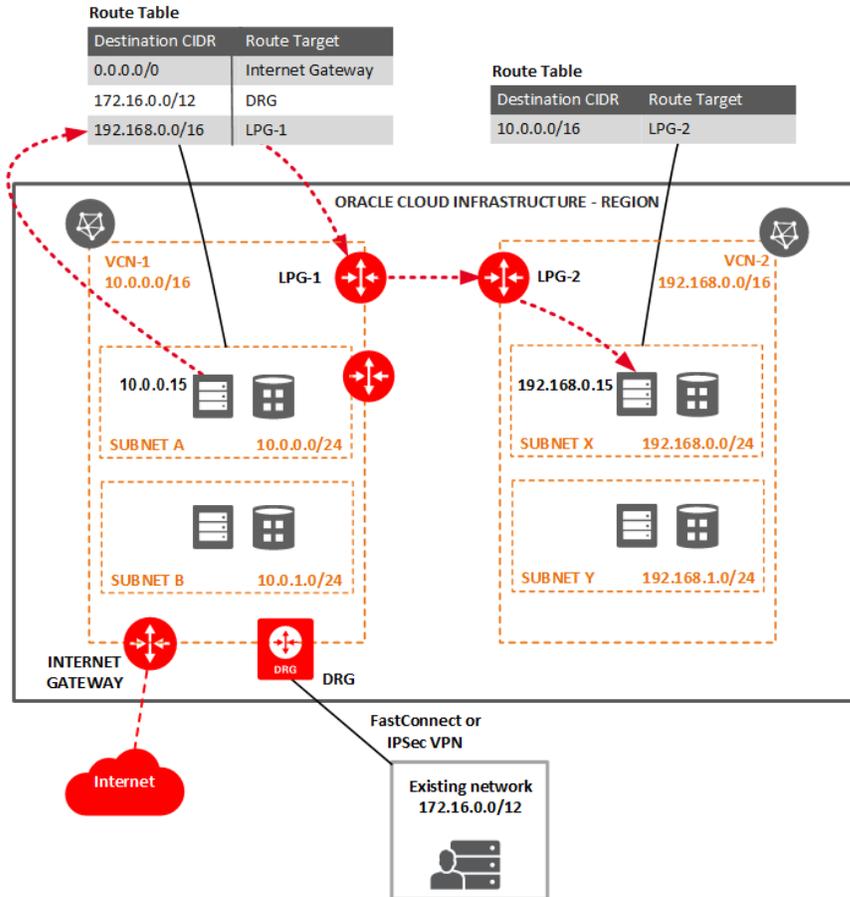
Either VCN administrator can terminate a peering by deleting their LPG. In that case, the other LPG's status switches to REVOKED. The administrator could instead render the connection non-functional by removing the route rules or security rules that enable traffic to flow across the connection (see the next sections).

### ROUTING TO THE LPG

As part of configuring the VCNs, each administrator must update the [VCN's routing](#) to enable traffic to flow between the VCNs. In practice, this looks just like routing you set up for any gateway (such as an [internet gateway](#) or [dynamic routing gateway](#)). For each subnet that needs to communicate with the other VCN, you update the subnet's route table. The route rule specifies the destination traffic's CIDR and your LPG as the target. Your LPG routes traffic that matches that rule to the other LPG, which in turn routes the traffic to the next hop in the other VCN.

In the following diagram, VCN-1 and VCN-2 are peered. Traffic from an instance in Subnet A (10.0.0.15) that is destined for an instance in VCN-2 (192.168.0.15) is routed to LPG-1 based on the rule in Subnet A's route table. From there the traffic is routed to LPG-2, and then from there, on to its destination in Subnet X.

## CHAPTER 23 Networking



### Note

As mentioned earlier, a given VCN can use the peered LPGs to reach VNICs in the other VCN, or the on-premises network if [transit routing](#) is set up for



the VCNs. But a VCN can't use the peered VCN to reach other destinations outside of the VCNs (such as the internet). For example, in the preceding diagram, VCN-2 cannot use the internet gateway attached to VCN-1.

### SECURITY RULES

Each subnet in a VCN has one or more [security lists](#) that control traffic in and out of the subnet's VNICs at the packet level. You can use security lists to control the type of traffic allowed with the other VCN. As part of configuring the VCNs, each administrator must determine which subnets in their own VCN need to communicate with VNICs in the other VCN and update their subnet's security lists accordingly.

If you use [network security groups](#) (NSGs) to implement security rules, notice that you have the option to write security rules for an NSG that specify *another* NSG as the source or destination of traffic. However, the two NSGs *must belong to the same VCN*.

## Important Implications of VCN Peering

If you haven't yet, read [Important Implications of Peering](#) to understand important access control, security, and performance implications for peered VCNs.

## Setting Up a Local Peering

Here's the general process for setting up a peering between two VCNs in the same region:

- A. **Create the LPGs:** Each VCN administrator creates an LPG for their own VCN.
- B. **Share information:** The administrators share the basic required information.
- C. **Set up the required IAM policies for the connection:** The administrators set up IAM policies to enable the connection to be established.
- D. **Establish the connection:** The requestor connects the two LPGs.

- E. **Update route tables:** Each administrator updates their VCN's route tables to enable traffic between the peered VCNs as desired.
- F. **Update security rules:** Each administrator updates their VCN's security rules to enable traffic between the peered VCNs as desired.

If desired, the administrators can perform tasks E and F *before* establishing the connection. In that case, each administrator must know the CIDR block or specific subnets from the other's VCN and share that in task B. After the connection is established, you can also get the CIDR block of the other VCN by viewing your own LPG's details in the Console. Look for **Peer Advertised CIDR**. Or if you're using the API, see the `peerAdvertisedCidr` parameter.

### Task A: Create the LPGs

Each administrator creates an LPG for their own VCN. "You" in the following procedure means an administrator (either the [acceptor or requestor](#)).



#### Note

##### **Required IAM Policy to Create LPGs**

If the administrators already have broad network administrator permissions (see [Let network admins manage a cloud network](#)), then they have permission to create, update, and delete LPGs. Otherwise, here's an example policy giving the necessary permissions to a group called LPGAdmins. The second statement is required because creating an LPG affects the VCN it belongs to, so the administrator must have permission to manage VCNs.

```
Allow group LPGAdmins to manage local-peering-gateways in
tenancy
Allow group LPGAdmins to manage vcns in tenancy
```

1. In the Console, confirm you're viewing the compartment that contains the VCN that you want to add the LPG to. For information about compartments and access control, see [Access Control](#).
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
3. Click the VCN you're interested in.
4. Under **Resources**, click **Local Peering Gateways**.
5. Click **Create Local Peering Gateway**.
6. Enter the following:
  - **Name:** A friendly name for the LPG. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
  - **Create in compartment:** The compartment where you want to create the LPG, if different from the compartment you're currently working in.
  - **Associate with Route Table (optional):** Only if you're setting up the advanced routing scenario called [transit routing](#). Select the compartment that contains the route table you want to associate with the LPG, and the route table itself. You can skip this part and associate the LPG with a route table later if you like.
  - **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click **Create Local Peering Gateway**.

The LPG is then created and displayed on the **Local Peering Gateways** page in the compartment you chose.

### Task B: Share information

If you're the **requestor**, give this information to the acceptor (for example, by email or other out-of-band method):

- If the VCNs are in the **same tenancy**: Name of the IAM group that should be granted permission to create a connection in the acceptor's compartment. In the example in the next task, the group is RequestorGrp.
- If the VCNs are in **different tenancies**: [OCID for your tenancy](#), and OCID for the IAM group that should be granted permission to create a connection in the acceptor's compartment. In the example in the next task, it's the OCID for the RequestorGrp.
- Optional: Your VCN's CIDR, or specific subnets for peering with the other VCN.

If you're the **acceptor**, give this information to the requestor:

- If the VCNs are in the **same tenancy**: OCID for your LPG. Optionally, also the names of your VCN, LPG, and the compartment each is in.
- If the VCNs are in **different tenancies**: OCID for your LPG, and OCID for your tenancy.
- Optional: Your VCN's CIDR, or specific subnets for peering with the other VCN.

### Task C: Set up the IAM policies (VCNs in same tenancy)

In this version of task C, both VCNs are in the same tenancy. If they're in different tenancies, instead see [Task C: Set up the IAM policies \(VCNs in different tenancies\)](#).

Both the requestor and acceptor must ensure that the right policies are in place:

- **Policy R (implemented by the requestor):**

```
Allow group RequestorGrp to manage local-peering-from in compartment RequestorComp
```

The requestor is in an IAM group called RequestorGrp. This policy lets anyone in the group initiate a connection from any LPG in the requestor's compartment (RequestorComp). Policy R can be attached to either the tenancy (root compartment) or

to RequestorComp. For information about why you would attach it to one versus the other, see [Policy Attachment](#).

- **Policy A (implemented by the acceptor):**

```
Allow group RequestorGrp to manage local-peering-to in compartment AcceptorComp

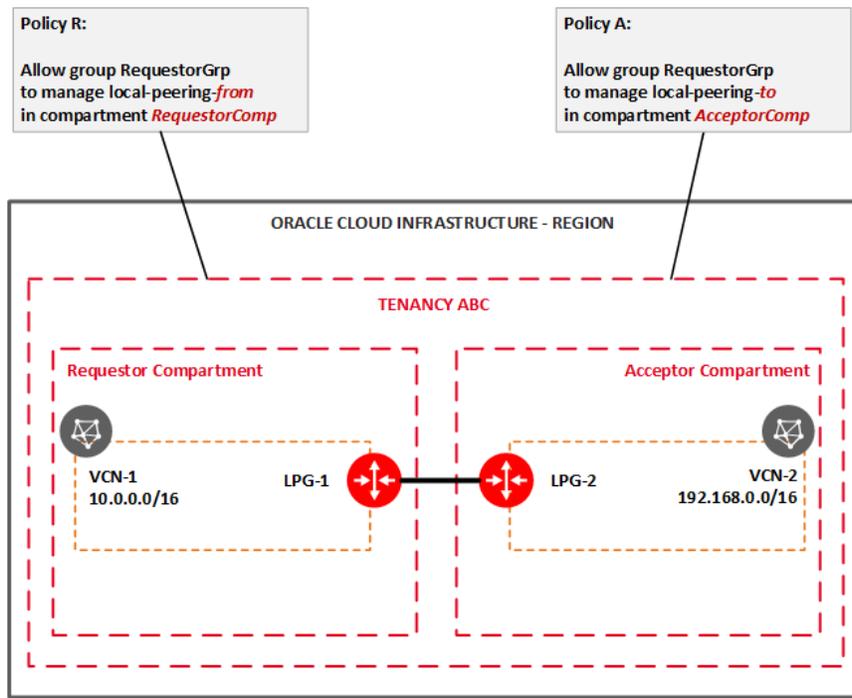
Allow group RequestorGrp to inspect vcns in compartment AcceptorComp
Allow group RequestorGrp to inspect local-peering-gateways in compartment AcceptorComp
```

The first statement in the policy lets the requestor connect to any LPG in the acceptor's compartment (AcceptorComp). This statement reflects the required agreement from the acceptor for the peering to be established. Policy A can be attached to either the tenancy (root compartment) or to AcceptorComp.



### Tip

The second and third statements in Policy A let the requestor list the VCNs and LPGs in AcceptorComp. The statements are required for the requestor to use the Console UI to select from a list of VCNs and LPGs in AcceptorComp and establish the connection. The following diagram focuses only on the first statement, which is the critical one that permits the connection.



Both Policy R and Policy A give RequestorGrp access. However, Policy R has a resource-type called *local-peering-from*, and Policy A has a resource-type called *local-peering-to*. Together, these policies let someone in RequestorGrp establish the connection *from* an LPG in the requestor's compartment *to* an LPG in the acceptor's compartment. The API call to create the connection specifies which two LPGs.

**Tip**

The permission granted by Policy R might already be in place if the requestor has permission in another policy to manage all of the Networking components in RequestorComp. For example, there might be a general Network Admin policy like this:

```
Allow group NetworkAdmin to manage virtual-network-family in
compartment RequestorComp
```

If the requestor is in the NetworkAdmin group, then they already have the required permissions covered in Policy R (the virtual-network-family includes LPGs). And further, if the policy is instead written to cover the *entire tenancy* instead of only compartment RequestorComp, then the requestor already has all the required permissions in both compartments to establish the connection. In that case, policy A is not required.

**Task C: Set up the IAM policies (VCNs in different tenancies)**

In this version of task C, the VCNs are in different tenancies (in other words, it's a *cross-tenancy* peering). If the VCNs are in the same tenancy, instead see [Task C: Set up the IAM policies \(VCNs in same tenancy\)](#).

Both the requestor and acceptor must ensure that the right policies are in place:

- **Policy R (implemented by the requestor):**

```
Define tenancy Acceptor as <acceptor_tenancy_OCID>

Allow group RequestorGrp to manage local-peering-from in compartment RequestorComp
```

## CHAPTER 23 Networking

---

```
Endorse group RequestorGrp to manage local-peering-to in tenancy Acceptor

Endorse group RequestorGrp to associate local-peering-gateways in compartment RequestorComp
with local-peering-gateways in tenancy Acceptor
```

The requestor is in an IAM group called RequestorGrp. This policy lets anyone in that group initiate a connection from any LPG in the requestor's compartment (RequestorComp).

The first statement is a "define" statement that assigns a friendly label to the acceptor's tenancy OCID. The statement happens to use "Acceptor" as the label, but it could be a value of the requestor's choice. **All "define" statements in a policy must be the first ones (at the top).**

The second statement lets the RequestorGrp establish a connection from an LPG in the requestor's compartment.

The third and fourth statements are special ones required because the LPGs are in different tenancies. They let the RequestorGrp connect an LPG in the requestor's tenancy to an LPG in the acceptor's tenancy.

If the desire is to give the RequestorGrp permission to connect to an LPG *in any tenancy*, the policy would instead look like this:

```
Allow group RequestorGrp to manage local-peering-from in compartment RequestorComp

Endorse group RequestorGrp to manage local-peering-to in any-tenancy

Endorse group RequestorGrp to associate local-peering-gateways in compartment RequestorComp
with local-peering-gateways in any-tenancy
```

Regardless, Policy R must be attached to the requestor's tenancy (root compartment), and not the requestor's compartment. Policies that enable cross-tenancy access must be attached to the tenancy. For more information about attachment of policies, see [Policy Attachment](#).

- **Policy A (implemented by the acceptor):**

```
Define tenancy Requestor as <requestor_tenancy_OCID>

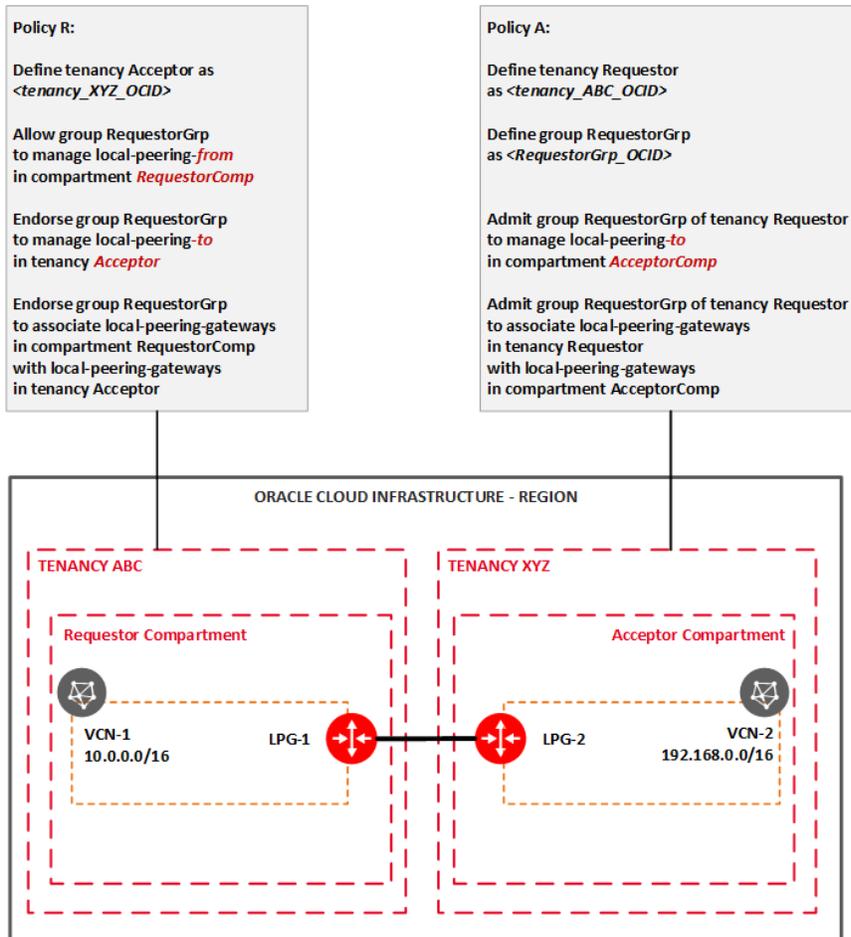
Define group RequestorGrp as <RequestorGrp_OCID>
```

```
Admit group RequestorGrp of tenancy Requestor to manage local-peering-to in compartment
AcceptorComp

Admit group RequestorGrp of tenancy Requestor to associate local-peering-gateways in tenancy
Requestor
 with local-peering-gateways in compartment AcceptorComp
```

Similar to the requestor's policy, this policy first uses "define" statements to assign friendly labels to the requestor's tenancy OCID and the RequestorGrp OCID. As mentioned earlier, the acceptor could use other values for those labels if desired. The third and fourth statements let the RequestorGrp connect to an LPG in the acceptor's compartment (AcceptorComp). **These statements reflect the critical agreement required for the peering to be established.** The word `Admit` indicates that the access applies to a group *outside the tenancy* where the policy resides. Policy A must be attached to the acceptor's tenancy (root compartment), and not the acceptor's compartment.

## CHAPTER 23 Networking



### Task D: Establish the connection

The requestor must perform this task.

Prerequisite: The requestor must have the OCID of the acceptor's LPG.



### Tip

If you're using the Console and the peering is between two VCNs in the same tenancy: Instead of specifying the acceptor's LPG OCID, you can instead pick the acceptor's VCN and LPG from lists of resources in the tenancy. However, you need to know both the name and compartment for the acceptor's VCN and LPG instead of the LPG's OCID. For reference, see [Task B: Share information](#).

1. In the Console, view the details for the requestor LPG that you want to connect to the acceptor LPG.
2. For the requestor LPG, click the Actions icon (three dots), and then click **Establish Peering Connection**.
3. Specify which LPG you want to peer with:
  - If the VCNs are in different tenancies: Select **Enter Local Peering Gateway OCID**, and enter the acceptor LPG's OCID.
  - If the VCNs are in the same tenancy: Do one of the following:
    - Select **Enter Local Peering Gateway OCID**, and enter the acceptor LPG's OCID.
    - Select **Browse Below**, and then select the acceptor's VCN and LPG from the lists provided. Remember that the VCN and LPG each might be in a different compartment than the one you're currently working in.
4. Click **Establish Peering Connection**.

The connection is established and the LPG's state changes to PEERED.

At this point, the details of each LPG update to show the **Peer VCN CIDR Block** for the other VCN. This is the CIDR of the other VCN across the peering from the LPG. Each administrator can use this information to update the route tables and security rules for their own VCN.

### Task E: Configure the route tables

As mentioned earlier, each administrator can do this task before or after the connection is established.

Prerequisite: Each administrator must have the CIDR block or specific subnets for the other VCN. If the connection is already established, look at the **Peer VCN CIDR Block** for your LPG in the Console. Otherwise, get the information from the other administrator by email or other method.

For your own VCN:

1. Determine which subnets in your VCN need to communicate with the other VCN.
2. Update the route table for each of those subnets to include a new rule that directs traffic destined for the other VCN's CIDR to your LPG:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.
  - c. Under **Resources**, click **Route Tables**.
  - d. Click the route table you're interested in.
  - e. Click **Add Route Rule** and enter the following:
    - **Target Type:** Local Peering Gateway.
    - **Destination CIDR Block:** The other VCN's CIDR block. If you want, you can specify a subnet or particular subset of the peered VCN's CIDR.
    - **Target Compartment:** The compartment where the LPG is located, if not the current compartment.
    - **Target:** The LPG.
  - f. Click **Add Route Rule**.

Any subnet traffic with a destination that matches the rule is routed to your LPG. For more information about setting up route rules, see [Route Tables](#).

Later, if you no longer need the peering and want to delete your LPG, you must first delete all the route rules in your VCN that specify the LPG as the target.



### Tip

Without the required routing, traffic doesn't flow between the peered LPGs. If a situation occurs where you need to temporarily stop the peering, you can simply remove the route rules that enable traffic. You don't need to delete the LPGs.

## Task F: Configure the security rules

As mentioned earlier, each administrator can do this task before or after the connection is established.

Prerequisite: Each administrator must have the CIDR block or specific subnets for the other VCN. In general, you should use the same CIDR block you used in the route table rule in [Task E: Configure the route tables](#).

What rules should you add?

- Ingress rules for the types of traffic you want to allow from the other VCN, specifically from the VCN's CIDR or specific subnets.
- Egress rule to allow outgoing traffic from your VCN to the other VCN. If the subnet already has a broad egress rule for all types of protocols to all destinations (0.0.0.0/0), then you don't need to add a special one for the other VCN.



### Note

The following procedure uses security lists, but you could instead implement the security rules in a [network security group](#) and then create all of the subnet's resources in that NSG.

For your own VCN:

1. Determine which subnets in your VCN need to communicate with the other VCN.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic specifically with the CIDR block or subnet of the other VCN:
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in.
  - c. Under **Resources**, click either **Ingress Rules** or **Egress Rules** depending on the type of rule you want to work with.
  - d. If you want to add a new rule, click **Add Ingress Rule** (or **Add Egress Rule**).

### Example

Let's say you want to add a stateful rule that enables ingress HTTPS (port 443) traffic from the other VCN's CIDR. Here are the basic steps you take when adding a rule:

- i. Leave the **Stateless** check box unselected.
- ii. **Source Type:** Leave as CIDR.
- iii. **Source CIDR:** Enter the same CIDR block that the route rules use (see [Task E: Configure the route tables](#)).
- iv. **IP Protocol:** Leave as TCP.

- v. **Source Port Range:** Leave as All.
  - vi. **Destination Port Range:** Enter 443.
- e. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Remove**.
  - f. If you wanted to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.

For more information about security rules, see [Security Rules](#).

### Using the Console

#### To create a local peering gateway

See the instructions in [Task A: Create the LPGs](#).

#### To associate a route table with an existing local peering gateway

This task is related to an advanced routing scenario called [transit routing](#).

An LPG can exist without a route table associated with it. However, after you associate a route table with an LPG, there must always be a route table associated with it. But, you can associate a *different* route table. You can also edit the table's rules, or delete some or all of the rules.

**Prerequisite:** The route table must exist and belong to the VCN that the LPG belongs to.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Local Peering Gateways**.
4. For the LPG you're interested in, click the Actions icon (⋮), and then click either:

- **Associate With Route Table:** If the LPG has no route table associated with it yet.
  - **Replace Route Table Association:** If you're changing which route table is associated with the LPG.
5. Select the compartment where the route table resides, and select the route table itself.
  6. Click **Associate**.

### To delete a local peering gateway

Prerequisite: First delete all the route rules in your VCN that specify the LPG as the target. Deleting those rules stops the routing in your VCN to the LPG. However, the LPG's state may still be PEERED if the other LPG still exists. Whenever an LPG is deleted, the LPG at the other side of the peering changes to the REVOKED state.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Click **Local Peering Gateways**.
4. For the LPG you want to delete, click the Actions icon (three dots), and then click **Terminate**.
5. Confirm when prompted.



#### Note

After deleting an LPG (and thus terminating a peering), it's recommended you review your security rules to remove any rules that enabled traffic with the other VCN.

### To manage tags for a local peering gateway

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Local Peering Gateways**.
4. Click the Actions icon (three dots) for the local peering gateway, and then click **View Tags**. From there you can view the existing tags, edit them, and apply new ones.

For more information, see [Resource Tags](#).

### To move a local peering gateway to a different compartment

You can move a local peering gateway from one compartment to another. When you move a local peering gateway to a new compartment, inherent policies apply immediately.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the VCN you're interested in.
3. Under **Resources**, click **Local Peering Gateways**.
4. Click the Actions icon (three dots) for the local peering gateway, and then click **Move Resource**.
5. Choose the destination compartment from the list.
6. Click **Move Resource**.

For more information about using compartments and policies to control access to your cloud network, see [Access Control](#). For general information about compartments, see [Managing Compartments](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your LPGs and create connections, use these operations:

- [ListLocalPeeringGateways](#)
- [CreateLocalPeeringGateway](#)
- [GetLocalPeeringGateway](#)
- [UpdateLocalPeeringGateway](#)
- [DeleteLocalPeeringGateway](#)
- [ConnectLocalPeeringGateways](#)
- [ChangeLocalPeeringGatewayCompartment](#)

### Remote VCN Peering (Across Regions)

This topic is about *remote VCN peering*. In this case, *remote* means that the VCNs reside *in different regions*. If the VCNs you want to connect are in the same region, see [Local VCN Peering \(Within Region\)](#).



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### Overview of Remote VCN Peering

*Remote VCN peering* is the process of connecting two VCNs in different regions (but the same tenancy). The peering allows the VCNs' resources to communicate using private IP addresses without routing the traffic over the internet or through your on-premises network. Without peering, a given VCN would need an [internet gateway](#) and public IP addresses for the instances that need to communicate with another VCN in a different region.

### Summary of Networking Components for Remote Peering

At a high level, the Networking service components required for a remote peering include:

- Two VCNs with non-overlapping CIDRs, in different regions that support remote peering. The VCNs must be in the same tenancy.



#### Note

*All VCN CIDRs Must Not Overlap*

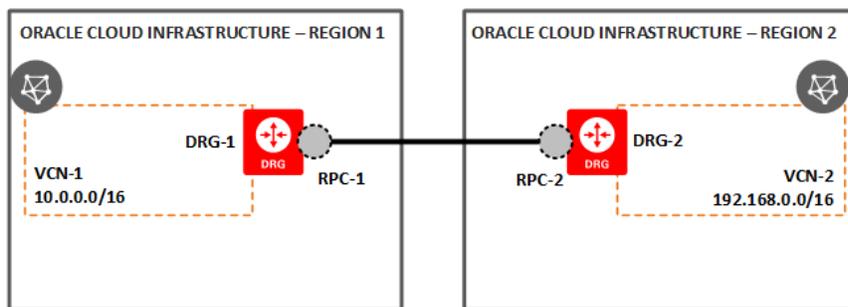
The two VCNs in the peering relationship must not have overlapping CIDRs. Also, if a particular VCN has multiple peering relationships, those other VCNs must not have overlapping CIDRs with each other. For example, if VCN-1 is peered with VCN-2 and also VCN-3, then VCN-2 and VCN-3 must not have overlapping CIDRs.

- A dynamic routing gateway (DRG) attached to each VCN in the peering relationship. Your VCN already has a DRG if you're using an [IPSec VPN](#) or an [Oracle Cloud Infrastructure FastConnect](#) private virtual circuit.
- A *remote peering connection (RPC)* on each DRG in the peering relationship.
- A *connection* between those two RPCs.

## CHAPTER 23 Networking

- Supporting [route rules](#) to enable traffic to flow over the connection, and only to and from select subnets in the respective VCNs (if desired).
- Supporting [security rules](#) to control the types of traffic allowed to and from the instances in the subnets that need to communicate with the other VCN.

The following diagram illustrates the components.



With supporting route tables and security rules  
in each VCN to enable traffic



### Note

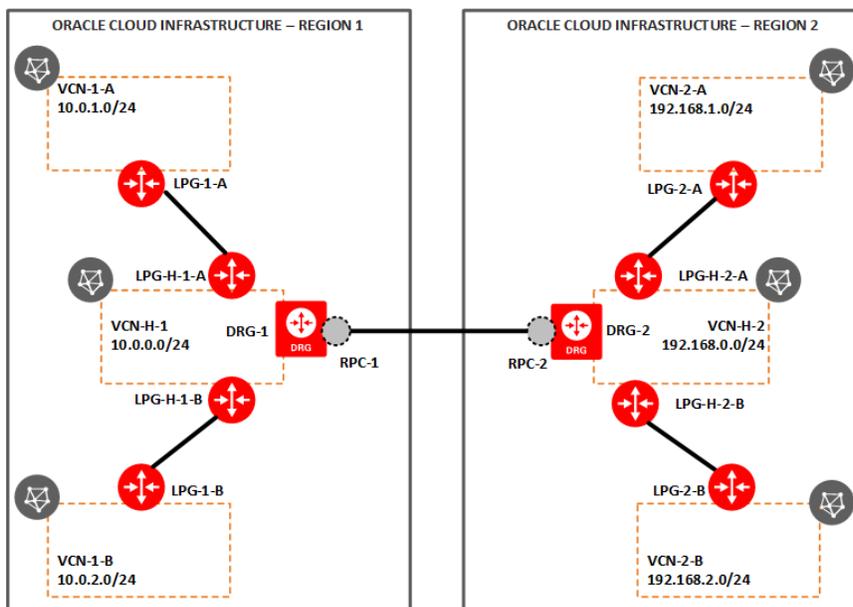
A given VCN can use the connected RPCs to reach only VNICs in the other VCN, and not destinations outside of the VCNs (such as the internet or your on-premises network). For example, if VCN-1 in the preceding diagram were to have an internet gateway, the instances in VCN-2 could NOT use it to send traffic to endpoints on the internet. However, be aware that VCN-2 could *receive* traffic from the internet via VCN-1. For more information, see [Important Implications of Peering](#).

### Spoke-to-Spoke: Remote Peering with Transit Routing

Imagine that in each region you have multiple VCNs in a hub-and-spoke layout, as shown in the following diagram. This type of layout within a region is discussed in detail in [Transit Routing: Access to Multiple VCNs in the Same Region](#). The spoke VCNs in a given region are [locally peered](#) with the hub VCN in the same region, using local peering gateways.

You can set up remote peering between the two hub VCNs. You can then also set up transit routing for the hub VCN's DRG and LPGs, as discussed in [Transit Routing: Access to Multiple VCNs in the Same Region](#). This setup allows a spoke VCN in one region to communicate with one or more spoke VCNs in the other region without needing a remote peering connection directly between those VCNs.

For example, you could configure routing so that resources in VCN-1-A could communicate with resources in VCN-2-A and VCN-2-B by way of the hub VCNs. That way, VCN 1-A is not required to have a *separate* remote peering with each of the spoke VCNs in the other region. You could also set up routing so that VCN-1-B could communicate with the spoke VCNs in region 2, without needing its own remote peerings to them.



### Explicit Agreement Required from Both Sides

Peering involves two VCNs in the same tenancy that might be administered by the same party or two different ones. The two parties might both be in your company but in different departments.

Peering between two VCNs requires explicit agreement from both parties in the form of Oracle Cloud Infrastructure Identity and Access Management policies that each party implements for their own VCN's compartment.

### Important Remote Peering Concepts

The following concepts help you understand the basics of VCN peering and how to establish a remote peering.

#### PEERING

A *peering* is a single peering relationship between two VCNs. Example: If VCN-1 peers with two other VCNs, then there are two peerings. The *remote* part of *remote peering* indicates that the VCNs are in different regions. For remote peering, the VCNs must be in the same tenancy.

#### VCN ADMINISTRATORS

In general, VCN peering can occur only if both of the VCN administrators agree to it. In practice, this means that the two administrators must:

- Share some basic information with each other.
- Coordinate to set up the required Oracle Cloud Infrastructure Identity and Access Management policies to enable the peering.
- Configure their VCNs for the peering.

Depending on the situation, a single administrator might be responsible for both VCNs and the related policies. The VCNs must be in the same tenancy.

For more information about the required policies and VCN configuration, see [Setting Up a Remote Peering](#).

### ACCEPTOR AND REQUESTOR

To implement the IAM policies required for peering, the two VCN administrators must designate one administrator as the *requestor* and the other as the *acceptor*. The requestor must be the one to initiate the request to connect the two RPCs. In turn, the acceptor must create a particular IAM policy that gives the requestor permission to connect to RPCs in the acceptor's compartment. Without that policy, the requestor's request to connect fails.

### REGION SUBSCRIPTION

To peer with a VCN in another region, your tenancy must first be subscribed to that region. For information about subscribing, see [Managing Regions](#).

### REMOTE PEERING CONNECTION (RPC)

A *remote peering connection (RPC)* is a component you create on the DRG attached to your VCN. The RPC's job is to act as a connection point for a remotely peered VCN. As part of configuring the VCNs, each administrator must create an RPC for the DRG on their VCN. A given DRG must have a separate RPC for each remote peering it establishes for the VCN (maximum 10 RPCs per tenancy). To continue with the previous example: the DRG on VCN-1 would have two RPCs to peer with two other VCNs. In the API, a [RemotePeeringConnection](#) is an object that contains information about the peering. You can't reuse an RPC to later establish another peering with it.

### CONNECTION BETWEEN TWO RPCS

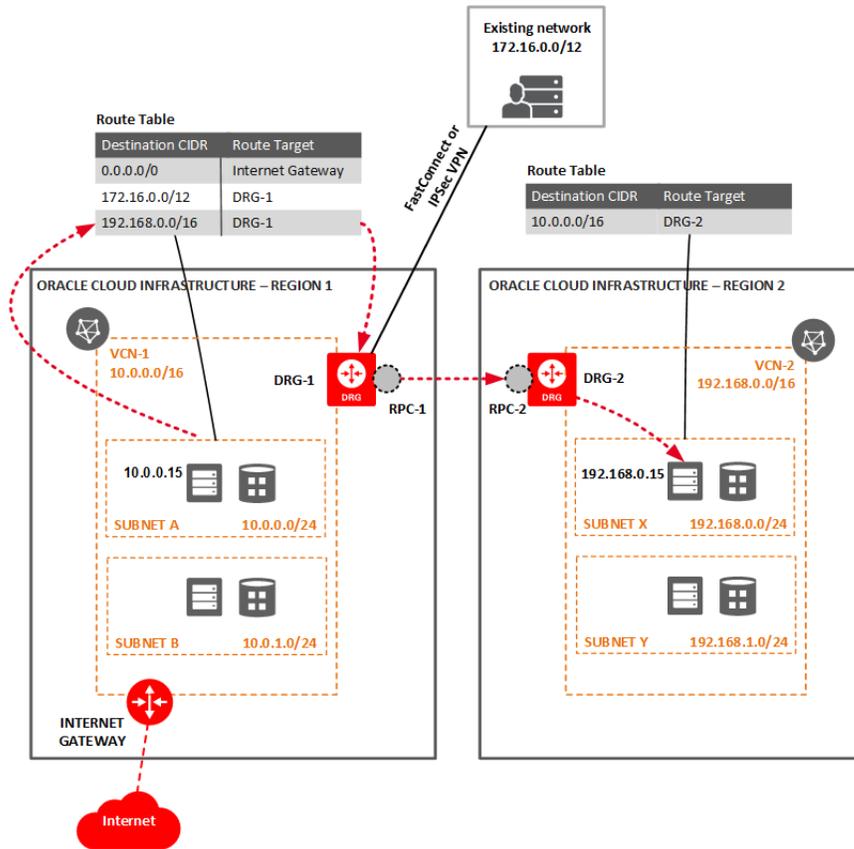
When the requestor initiates the request to peer (in the Console or API), they're effectively asking to *connect the two RPCs*. This means the requestor must have information to identify each RPC (such as the RPC's region and OCID).

Either VCN administrator can terminate a peering by deleting their RPC. In that case, the other RPC's status switches to REVOKED. The administrator could instead render the connection non-functional by removing the route rules that enable traffic to flow across the connection (see the next section).

### **ROUTING TO THE DRG**

As part of configuring the VCNs, each administrator must update the [VCN's routing](#) to enable traffic to flow between the VCNs. For each subnet that needs to communicate with the other VCN, you update the subnet's route table. The route rule specifies the destination traffic's CIDR and your DRG as the target. Your DRG routes traffic that matches that rule to the other DRG, which in turn routes the traffic to the next hop in the other VCN.

In the following diagram, VCN-1 and VCN-2 are peered. Traffic from an instance in Subnet A (10.0.0.15) that is destined for an instance in VCN-2 (192.168.0.15) is routed to DRG-1 based on the rule in Subnet A's route table. From there the traffic is routed through the RPCs to DRG-2, and then from there, on to the destination in Subnet X.



**Note**

As mentioned earlier, a given VCN can use the connected RPCs to reach only VNICs in the other VCN, and not destinations outside of the VCNs (such as the internet or your on-premises network). For example, in the preceding diagram, VCN-2 cannot



use the internet gateway attached to VCN-1.

### SECURITY RULES

Each subnet in a VCN has one or more [security lists](#) that control traffic in and out of the subnet's VNICs at the packet level. You can use security lists to control the type of traffic allowed with the other VCN. As part of configuring the VCNs, each administrator must determine which subnets in their own VCN need to communicate with VNICs in the other VCN and update their subnet's security lists accordingly.

If you use [network security groups](#) (NSGs) to implement security rules, notice that you have the option to write security rules for an NSG that specify *another* NSG as the source or destination of traffic. However, the two NSGs *must belong to the same VCN*.

## Important Implications of Peering

If you haven't yet, read [Important Implications of Peering](#) to understand important access control, security, and performance implications for peered VCNs.

## Setting Up a Remote Peering

This section covers the general process for setting up a peering between two VCNs in different regions.



### Important

The following procedure assumes that:

- Your tenancy is subscribed to the other VCN's region. If it's not, see [Managing Regions](#).
- You already have a DRG attached to your VCN. If you don't, see [Dynamic Routing Gateways \(DRGs\)](#).

- A. **Create the RPCs:** Each VCN administrator creates an RPC for their own VCN's DRG.
- B. **Share information:** The administrators share the basic required information.
- C. **Set up the required IAM policies for the connection:** The administrators set up IAM policies to enable the connection to be established.
- D. **Establish the connection:** The requestor connects the two RPCs.
- E. **Update route tables:** Each administrator updates their VCN's route tables to enable traffic between the peered VCNs as desired.
- F. **Update security rules:** Each administrator updates their VCN's security rules to enable traffic between the peered VCNs as desired.

If desired, the administrators can perform tasks E and F *before* establishing the connection. Each administrator needs to know the CIDR block or specific subnets from the other's VCN and share that in task B.

### Task A: Create the RPCs

Each administrator creates an RPC for their own VCN's DRG. "You" in the following procedure means an administrator (either the [acceptor or requestor](#)).



### Note

#### *Required IAM Policy to Create RPCs*

If the administrators already have broad network administrator permissions (see [Let network admins manage a cloud network](#)), then they have permission to create, update, and delete RPCs. Otherwise, here's an example policy giving the necessary permissions to a group called RPCAdmins. The second statement is required because creating an RPC affects the DRG it belongs to, so the administrator must have permission to manage DRGs.

```
Allow group RPCAdmins to manage remote-peering-connections in
tenancy
Allow group RPCAdmins to manage drgs in tenancy
```

1. In the Console, confirm you're viewing the compartment that contains the DRG that you want to add the RPC to. For information about compartments and access control, see [Access Control](#).
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
3. Click the DRG you're interested in.
4. Under **Resources**, click **Remote Peering Connections**.
5. Click **Create Remote Peering Connection**.
6. Enter the following:
  - **Name:** A friendly name for the RPC. It doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.

- **Create in compartment:** The compartment where you want to create the RPC, if different from the compartment you're currently working in.
7. Click **Create Remote Peering Connection**.  
The RPC is then created and displayed on the **Remote Peering Connections** page in the compartment you chose.
  8. If you're the acceptor, record the RPC's region and OCID to later give to the requestor.

### Task B: Share information

- If you're the acceptor, give this information to the requestor (for example, by email or other out-of-band method):
  - The region your VCN is in (the requestor's tenancy must be subscribed to this region).
  - Your RPC's OCID.
  - The CIDR blocks for subnets in your VCN that should be available to the other VCN. The requestor needs this information when setting up routing for the requestor VCN.
- If you're the requestor, give this information to the acceptor:
  - The region your VCN is in (the acceptor's tenancy must be subscribed to this region).
  - The name of the IAM group that should be granted permission to create a connection in the acceptor's compartment (in the example in the next task, the group is RequestorGrp).
  - The CIDR blocks for subnets in your VCN that should be available to the other VCN. The acceptor needs this information when setting up routing for the acceptor VCN.

### Task C: Set up the IAM policies

Both the requestor and acceptor must ensure the right policies are in place. These consist of:

- **Policy R (implemented by the requestor):**

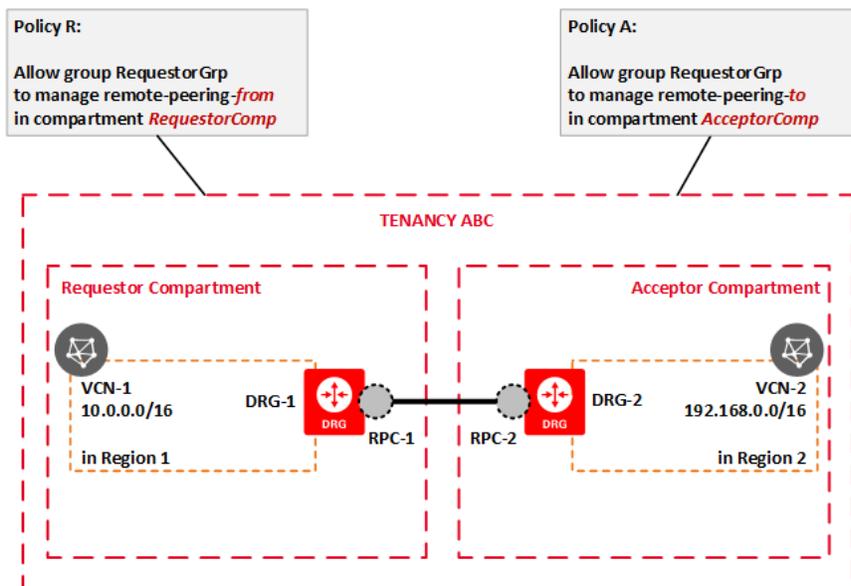
```
Allow group RequestorGrp to manage remote-peering-from in compartment RequestorComp
```

The requestor is in an IAM group called RequestorGrp. This policy lets anyone in the group initiate a connection from any RPC in the requestor's compartment (RequestorComp). Policy R can be attached to either the tenancy (root compartment) or to RequestorComp. For information about why you would attach it to one versus the other, see [Policy Attachment](#).

- **Policy A (implemented by the acceptor):**

```
Allow group RequestorGrp to manage remote-peering-to in compartment AcceptorComp
```

This policy lets the requestor connect to any RPC in the acceptor's compartment (AcceptorComp). This statement reflects the required agreement from the acceptor for the peering to be established. Policy A can be attached to either the tenancy (root compartment) or to AcceptorComp.



Both Policy R and Policy A give RequestorGrp access. However, Policy R has a resource-type called *remote-peering-from*, and Policy A has a resource-type called *remote-peering-to*. Together, these policies let someone in RequestorGrp establish the connection *from* an RPC in the requestor's compartment *to* an RPC in the acceptor's compartment. The API call to actually create the connection specifies which two RPCs.



### Tip

The permission granted by Policy R might already be in place if the requestor has permission in another policy to manage all of the Networking components in RequesterComp. For example, there might be a general Network Admin policy like this: `Allow group NetworkAdmin to manage virtual-network-family in compartment RequestorComp`. If the requestor is in the NetworkAdmin group, then they already have the required permissions covered in Policy R (the *virtual-network-family* includes RPCs). And further, if the policy is instead written to cover the *entire tenancy* (`Allow group NetworkAdmin to manage virtual-network-family in tenancy`), then the requestor already has all the required permissions in both compartments to establish the connection. In that case, policy A is not required.

### Task D: Establish the connection

The requestor must perform this task.

Prerequisite: The requestor must have:

- The region the acceptor's VCN is in (the requestor's tenancy must be subscribed to the region).
  - The OCID of the acceptor's RPC.
1. In the Console, view the details for the requestor RPC that you want to connect to the acceptor RPC.
  2. Click **Establish Connection**.
  3. Enter the following:
    - **Region:** The region that contains the acceptor's VCN. The drop-down list includes only those regions that both support remote VCN peering and your tenancy is subscribed to.
    - **Remote Peering Connection OCID:** The OCID of the acceptor's RPC.
  4. Click **Establish Connection**.

The connection is established and the RPC's state changes to PEERED.

### Task E: Configure the route tables

As mentioned earlier, each administrator can do this task before or after the connection is established.

Prerequisite: Each administrator must have the CIDR block or specific subnets for the other VCN.

For your own VCN:

1. Determine which subnets in your VCN need to communicate with the other VCN.
2. Update the route table for each of those subnets to include a new rule that directs traffic destined for the other VCN to your DRG:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.

- c. Under **Resources**, click **Route Tables**.
- d. Click the route table you're interested in.
- e. Click **Add Route Rule** and enter the following:
  - **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
  - **Destination CIDR Block:** The other VCN's CIDR block. If you want, you can specify a subnet or particular subset of the peered VCN's CIDR.
- f. Click **Add Route Rule**.

Any subnet traffic with a destination that matches the rule is routed to your DRG. For more information about setting up route rules, see [Route Tables](#).



### Tip

Without the required routing, traffic doesn't flow between the peered DRGs. If a situation occurs where you need to temporarily stop the peering, you can simply remove the route rules that enable traffic. You don't need to delete the RPCs.

## Task F: Configure the security rules

As mentioned earlier, each administrator can do this task before or after the connection is established.

Prerequisite: Each administrator must have the CIDR block or specific subnets for the other VCN. In general, you should use the same CIDR block you used in the route table rule in [Task E: Configure the route tables](#).

What rules should you add?

- Ingress rules for the types of traffic you want to allow from the other VCN, specifically from the VCN's CIDR or specific subnets.
- Egress rule to allow outgoing traffic from your VCN to the other VCN. If the subnet already has a broad egress rule for all types of protocols to all destinations (0.0.0.0/0), then you don't need to add a special one for the other VCN.



### Note

The following procedure uses security lists, but you could instead implement the security rules in a [network security group](#) and then create all of the subnet's resources in that NSG.

For your own VCN:

1. Determine which subnets in your VCN need to communicate with the other VCN.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic specifically with the CIDR block or subnet of the other VCN:
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in.
  - c. Under **Resources**, click either **Ingress Rules** or **Egress Rules** depending on the type of rule you want to work with.
  - d. If you want to add a new rule, click **Add Ingress Rule** (or **Add Egress Rule**).

### Example

Let's say you want to add a stateful rule that enables ingress HTTPS (port 443) traffic from the other VCN's CIDR. Here are the basic steps you take when adding a rule:

- i. In the **Allow Rules for Ingress** section, click **+Add Rule**.
  - ii. Leave the **Stateless** checkbox unchecked.
  - iii. **Source Type:** Leave as CIDR.
  - iv. **Source CIDR:** Enter the same CIDR block that the route rules use (see [Task E: Configure the route tables](#)).
  - v. **IP Protocol:** Leave as TCP.
  - vi. **Source Port Range:** Leave as All.
  - vii. **Destination Port Range:** Enter 443.
- e. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Remove**.
- f. If you wanted to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.

For more information about security rules, see [Security Rules](#).

## Using the Console

### To create a remote peering connection

See the instructions in [Task A: Create the RPCs](#).

### To delete a remote peering connection

Deleting an RPC terminates the peering. The RPC at the other side of the peering changes to the REVOKED state.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Dynamic Routing Gateways**.
2. Click the DRG you're interested in.

3. Under **Resources**, click **Remote Peering Connections**.
4. Click the RPC you're interested in.
5. Click **Terminate**.
6. Confirm when prompted.



### Note

After deleting an RPC (and thus terminating a peering), it's recommended you review your route tables and security rules to remove any rules that enabled traffic with the other VCN.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To manage your RPCs and create connections, use these operations:

- [ListAllowedPeerRegionsForRemotePeering](#)
- [ListRemotePeeringConnections](#)
- [CreateRemotePeeringConnection](#)
- [GetRemotePeeringConnection](#)
- [UpdateRemotePeeringConnection](#)
- [DeleteRemotePeeringConnection](#)
- [ConnectRemotePeeringConnections](#)

## Access to Oracle Cloud Infrastructure Classic

There are two ways to set up a connection between an Oracle Cloud Infrastructure Classic IP network and an Oracle Cloud Infrastructure virtual cloud network (VCN):

- [Option 1: Connection over the Oracle network](#)
  - You file a ticket with My Oracle Support and Oracle provisions a connection between the IP network's private gateway and the VCN's attached dynamic routing gateway (DRG). The connection runs over Oracle's network and not the internet.
  - The two environments must be in the same geographical area, and the connection is available only between the specific regions listed in [Overview](#).
  - The two environments must belong to the same company. Oracle validates this when setting up the connection.
- [Option 2: Connection over an IPSec VPN](#)
  - You set up an IPSec VPN between the IP network's VPN as a Service (VPNaaS) gateway and the VCN's attached DRG. The connection runs over the internet.
  - The two environments do not have to be in the same geographical area or regions.
  - The two environments do not have to belong to the same company.

## Connection Over Oracle Network

This topic describes one way to set up a connection between an Oracle Cloud Infrastructure Classic IP network and an Oracle Cloud Infrastructure virtual cloud network (VCN). The connection runs over Oracle's network.

Another option is to connect the two clouds with an IPSec VPN. For more information, see [Connection Over IPSec VPN](#).

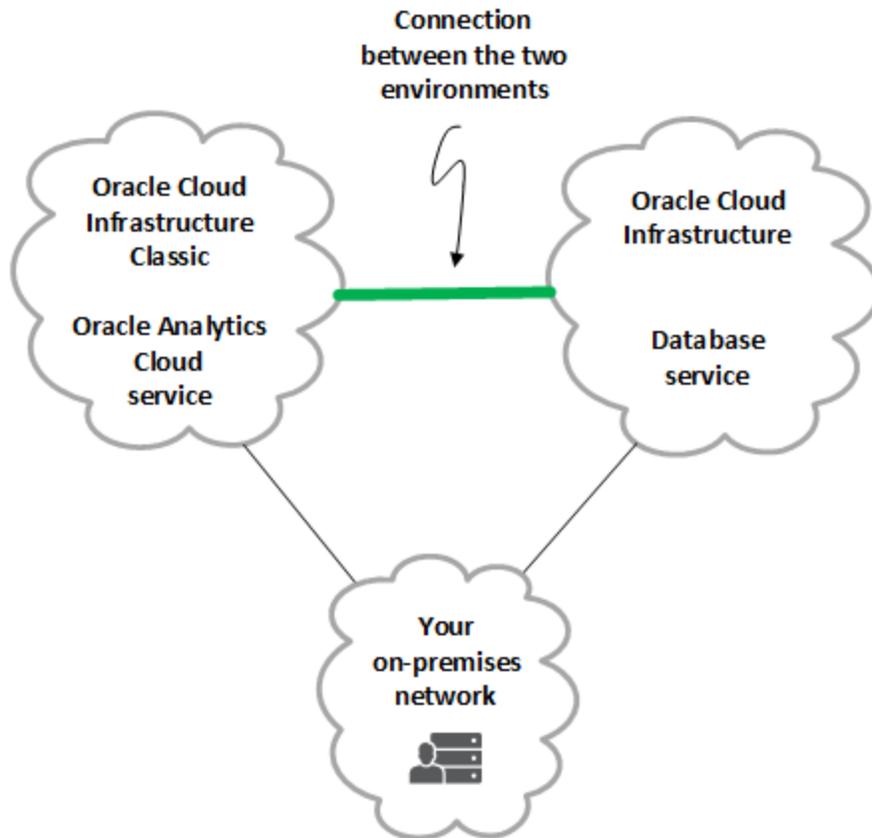
### Highlights

- You can run a hybrid workload between your Oracle Cloud Infrastructure Classic and Oracle Cloud Infrastructure environments.
- Oracle connects the IP network's private gateway to the VCN's attached dynamic routing gateway (DRG). The connection runs over the Oracle network. You configure routing and security rules in the environments to enable traffic.
- The two environments must belong to the same company and not have overlapping CIDRs. The cloud resources can communicate over the connection only with private IP addresses.
- The two environments must both be in the Ashburn area, the London area, or the Sydney area, and in specific regions listed in the next section. Connectivity to other regions is not supported.
- The connection is free of charge.

### Overview

You can request Oracle to provision a connection between your Oracle Cloud Infrastructure environment and your Oracle Cloud Infrastructure Classic environment. The connection facilitates a hybrid deployment with application components that are set up across the two environments. You can also use the connection to migrate workloads from Oracle Cloud Infrastructure Classic to Oracle Cloud Infrastructure. Compared to an IPsec VPN: the resources in the two environments have a more reliable and consistent network connection, with better throughput, because the traffic uses Oracle's internal links. Compared to FastConnect: you don't incur the additional cost and operational overhead of working with a FastConnect partner.

The following diagram shows an example of a hybrid deployment. Oracle Analytics Cloud is running in an Oracle Cloud Infrastructure Classic IP network and accessing the Database service in Oracle Cloud Infrastructure over the connection.



Here are other important details to know:

- The connection is supported only between these regions:
  - Oracle Cloud Infrastructure Australia East (Sydney) region and the Sydney Classic region
  - Oracle Cloud Infrastructure US East (Ashburn) region and the Ashburn Classic region
  - Oracle Cloud Infrastructure UK South (London) region and the Slough Classic region
- The connection enables communication that uses private IP addresses only.

## CHAPTER 23 Networking

---

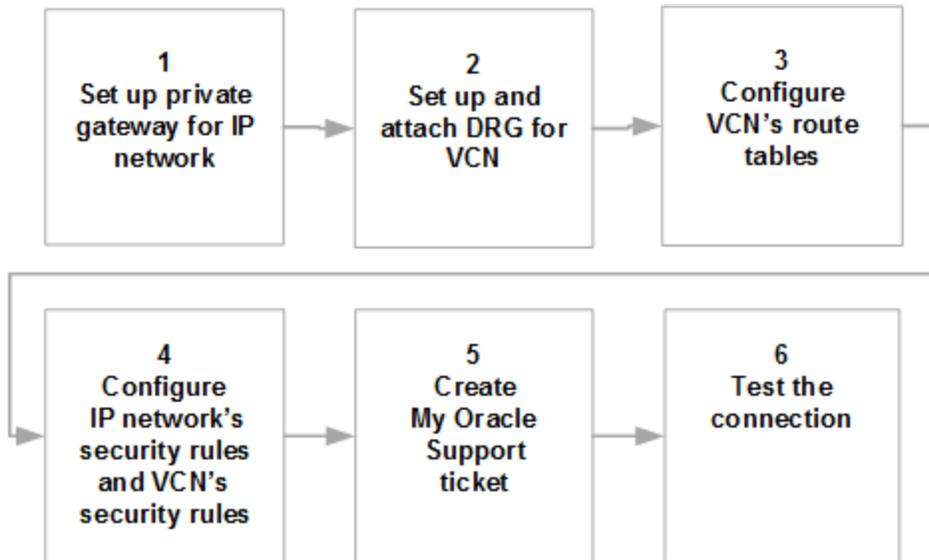
- The CIDR blocks of the IP network and VCN subnets that need to communicate must not overlap.
- The IP network and VCN must belong to the same company. Oracle validates this when setting up the connection.
- This connection enables communication only between resources in the Oracle Cloud Infrastructure Classic IP network and Oracle Cloud Infrastructure VCN. It does not enable traffic between your on-premises network through the IP network to the VCN, or from your on-premises network through the VCN to the IP network.
- The connection also does not enable traffic to flow from the IP network through the connected VCN to a peered VCN in the same Oracle Cloud Infrastructure region, or a different region.

The following table lists the comparable networking components required on each side of the connection.

<b>Component</b>	<b>Oracle Cloud Infrastructure Classic</b>	<b>Oracle Cloud Infrastructure</b>
Cloud network	IP network	VCN
Gateway	private gateway	dynamic routing gateway (DRG)
Routing	routes	route tables with route rules
Security rules	security rules	network security groups, security lists

## Connecting Your IP Network and VCN

The following flow chart shows the overall process of connecting your IP network and VCN.



### Prerequisites:

You must already have:

- An Oracle Cloud Infrastructure Classic [IP network](#).
- An Oracle Cloud Infrastructure [VCN with subnets](#).

### Task 1: Set up a private gateway for your IP network

If you do not already have a private gateway for your IP network, [create one](#).

### Task 2: Set up a dynamic routing gateway (DRG) for your VCN

If you do not already have a DRG attached to your VCN, create a DRG and attach it:

- [To create a DRG](#)
- [To attach a DRG to a VCN](#)

### Task 3: Configure route tables

#### For the IP network

When you create the private gateway and attach an IP network to it, traffic from cloud resources in the IP network uses the private gateway as the next hop. You do not need to update any routes for the IP network.

#### For the VCN

You must add a route rule that directs traffic from the VCN's subnets to the DRG:

1. Determine which subnets in your VCN need to communicate with the IP network.
2. Update the route table for each of those subnets to include a new rule that directs traffic destined for the IP network's CIDR to your DRG:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.
  - c. Under **Resources**, click **Route Tables**.
  - d. Click the route table you're interested in.
  - e. Click **Add Route Rule** and enter the following:
    - **Destination CIDR Block:** The IP network's CIDR block.
    - **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the

target yourself.

- f. Click **Add Route Rule**.

Any subnet traffic with a destination that matches the rule is routed to your DRG. For more information about setting up route rules, see [Route Tables](#).

Later, if you no longer need the connection and want to delete your DRG, you must first delete all the route rules in your VCN that specify the DRG as the target.

### Task 4: Configure the security rules

To ensure traffic flows between the IP network and VCN, make sure the IP network security rules and the VCN's security rules allow the desired traffic.

Here are the types of rules to add:

- Ingress rules for the types of traffic you want to allow into one cloud from the other, specifically from the other cloud's CIDR block.
- Egress rule to allow outgoing traffic from one cloud to the other. If the VCN's subnet already has a broad egress rule for all types of protocols to all destinations (0.0.0.0/0), then you don't need to add a special one for the IP network.

#### For the IP network

[Configure the network security rules](#) for the IP network to allow the desired traffic.

#### For the VCN



#### Note

The following procedure uses security lists, but you



could instead implement the security rules in one or more [network security groups](#) and then place the VCN's resources in NSGs.

1. Determine which subnets in your VCN need to communicate with the IP network.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic specifically with the CIDR block of the IP network:
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in.  
Under **Resources**, you can click **Ingress Rules** or **Egress Rules** to switch between the different types of rules.
  - c. Add one or more rules, each for the specific type of traffic you want to allow.

### Example:

Let's say you want to add a stateful rule that enables ingress HTTPS (port 443) traffic from the IP network's CIDR. Here are the basic steps you take when adding a rule:

- i. On the **Ingress Rules** page, click **Add Ingress Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Source CIDR:** Enter the same CIDR block that the route rules use (see [Task 3: Configure route tables](#)).
- iv. **IP Protocol:** Leave as TCP.
- v. **Source Port Range:** Leave as All.
- vi. **Destination Port Range:** Enter 443.
- vii. Click **Add Ingress Rule**.

For more information about setting up security rules, see [Security Rules](#).



### Important

The VCN's [default security list](#) does not allow ICMP echo reply and echo request (ping). You must add rules to enable that traffic. See [Rules to Enable Ping](#)

### Task 5: Create a My Oracle Support ticket

To have Oracle set up the connection, create a ticket at [My Oracle Support](#) and provide the following information:

- Ticket name: Create IP Network - VCN Connection - *<your\_company\_name>* - Ashburn
- OCI-C identity domain
- OCI-C private gateway name
- Region
- OCI tenancy OCID
- OCI DRG OCID

For example:

- Ticket name: Create IP Network - VCN Connection - ACME - Ashburn
- OCI-C identity domain: 123456789, uscom-east-1
- OCI-C private gateway name: Compute-acme/jack.jones@example.com/privategateway1
- Region: uscom-east-1 (OCI-C) / us-ashburn-1 (OCI)
- OCI tenancy OCID: ocid1.tenancy.oc1..examplefbpnk5cmdl7gkr6kcakfqmvhvbpcv
- OCI DRG OCID: ocid1.drg.oc1.iad.exampleutg6cmd3fqwqbea7ctadcatm

**It can take 3 to 4 business days before your My Oracle Support ticket is complete and the connection is ready to test.**

### Task 6: Test the connection

After you receive confirmation from your support person that the connection is ready, test the connection. Depending on how you've set up your IP network's security rules and VCN security rules, you should be able to [launch an instance](#) in your VCN and access it from an instance in the IP network. Or you should be able to connect from the VCN instance to an instance in the IP network. If you can, your connection is ready to use.

### Terminating the Connection

If you want to terminate the connection, file a ticket at [My Oracle Support](#).

## Connection Over IPSec VPN

This topic describes one way to set up a connection between an Oracle Cloud Infrastructure Classic IP network and an Oracle Cloud Infrastructure virtual cloud network (VCN). The connection runs over an IPSec VPN.

Another option is to have Oracle set up a connection over the Oracle network. For more information, see [Connection Over Oracle Network](#).

### Highlights

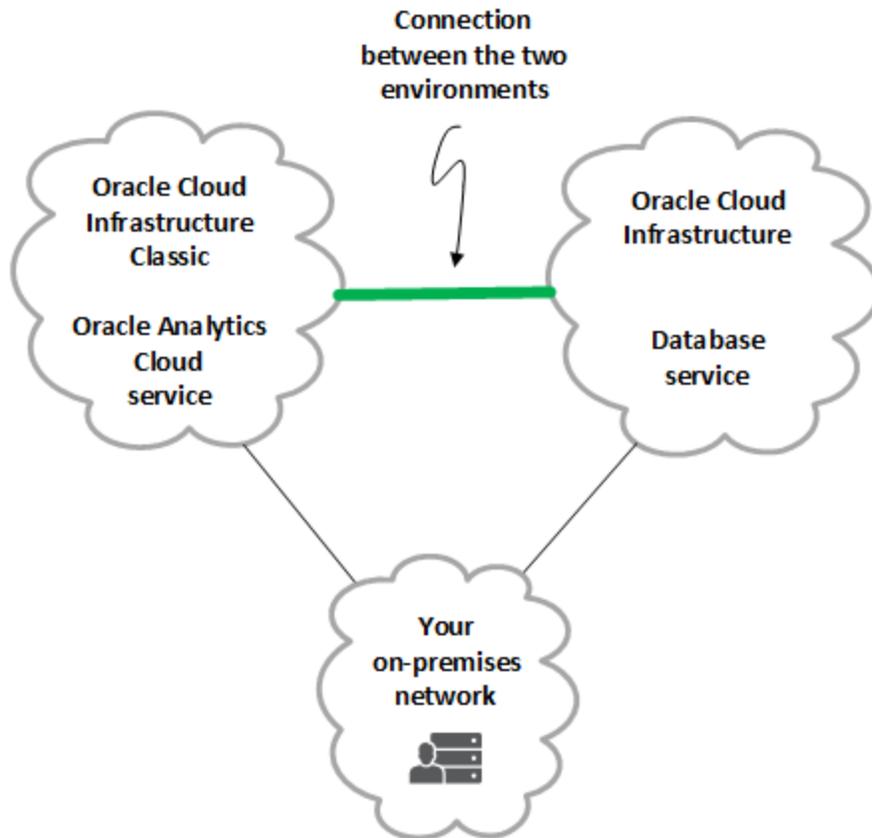
- You can run a hybrid workload between your Oracle Cloud Infrastructure Classic and Oracle Cloud Infrastructure environments.
- You set up an IPSec VPN between the IP network's VPN as a Service (VPNaaS) gateway and the VCN's attached dynamic routing gateway (DRG). The connection runs over the internet. You configure routing and security rules in the environments to enable traffic.

- The two environments must not have overlapping CIDRs. The cloud resources can communicate over the connection only with private IP addresses.
- The two environments do not have to be in the same geographical area or region.
- The connection is free of charge.

### Overview

You can connect your Oracle Cloud Infrastructure environment and your Oracle Cloud Infrastructure Classic environment with an IPsec VPN. The connection facilitates a hybrid deployment with application components that are set up across the two environments. You can also use the connection to migrate workloads from Oracle Cloud Infrastructure Classic to Oracle Cloud Infrastructure. Compared to using the Oracle network for the connection: you can set up the IPsec VPN yourself in a matter of minutes. Compared to FastConnect: you don't incur the additional cost and operational overhead of working with a FastConnect partner.

The following diagram shows an example of a hybrid deployment. Oracle Analytics Cloud is running in an Oracle Cloud Infrastructure Classic IP network and accessing the Database service in Oracle Cloud Infrastructure over the connection.



Here are other important details to know:

- The connection is supported in any of the Oracle Cloud Infrastructure and Oracle Cloud Infrastructure Classic regions. The two environments do not need to be in the same geographical area.
- The connection enables communication that uses private IP addresses only.
- The CIDR blocks of the IP network and VCN subnets that need to communicate must not overlap.
- This connection enables communication only between resources in the Oracle Cloud Infrastructure Classic IP network and Oracle Cloud Infrastructure VCN. It does not

## CHAPTER 23 Networking

enable traffic between your on-premises network through the IP network to the VCN, or from your on-premises network through the VCN to the IP network.

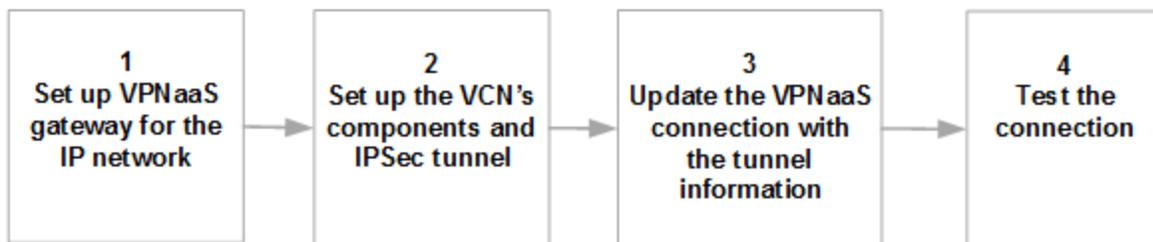
- The connection also does not enable traffic to flow from the IP network through the connected VCN to a peered VCN in the same Oracle Cloud Infrastructure region, or a different region.

The following table lists the comparable networking components required on each side of the connection.

Component	Oracle Cloud Infrastructure Classic	Oracle Cloud Infrastructure
Cloud network	IP network	VCN
Gateway	VPNaaS gateway	dynamic routing gateway (DRG)
Security rules	security rules	network security groups, security lists

### Setting Up the IPSec VPN Between Your IP Network and VCN

The following flow chart shows the overall process of connecting your IP network and VCN with an IPSec VPN.



#### Prerequisites:

You must already have:

- An Oracle Cloud Infrastructure Classic [IP network](#).
- An Oracle Cloud Infrastructure [VCN with subnets](#).

### Task 1: Set up a VPNaaS gateway for your IP network

1. Use these values when [setting up the VPNaaS gateway](#):
  - **IP Network:** The Oracle Cloud Infrastructure Classic IP network you want to connect to your VCN. Note that you can specify only a single IP network.
  - **Customer Gateway:** A placeholder value such as 129.213.240.51. Using this placeholder value lets you move forward in the process. You will change the value later with the Oracle Cloud Infrastructure VPN router's IP address.
  - **Customer Reachable Routes:** The CIDR block for the VCN. Note that you can specify only a single VCN.
  - **Specify Phase 2 ESP Proposal:** Check box selected.
  - **ESP Encryption:** AES 256
  - **ESP Hash:** SHA1
  - **IPSec Lifetime:** 1800
  - **Require Perfect Forward Secrecy:** Check box selected.
2. Record the resulting public IP address of the VPNaaS gateway.

### Task 2: Set up the VCN's components and IPSec tunnel

#### Task 2a: Set up a dynamic routing gateway (DRG) for your VCN

If you do not already have a DRG attached to your VCN, create a DRG and attach it:

- [To create a DRG](#)
- [To attach a DRG to a VCN](#)

### Task 2b: Configure routing to the DRG

You must add a route rule that directs traffic from the VCN's subnets to the DRG. Use the IP network's CIDR block as the destination for the rule.

1. Determine which subnets in your VCN need to communicate with the IP network.
2. Update the route table for each of those subnets to include a new rule that directs traffic destined for the IP network's CIDR to your DRG:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
  - b. Click the VCN you're interested in.
  - c. Under **Resources**, click **Route Tables**.
  - d. Click the route table you're interested in.
  - e. Click **Add Route Rule** and enter the following:
    - **Destination CIDR Block:** The IP network's CIDR block.
    - **Target Type:** Dynamic Routing Gateway. The VCN's attached DRG is automatically selected as the target, and you don't have to specify the target yourself.
  - f. Click **Add Route Rule**.

Any subnet traffic with a destination that matches the rule is routed to your DRG. For more information about setting up route rules, see [Route Tables](#).

Later, if you no longer need the connection and want to delete your DRG, you must first delete all the route rules in your VCN that specify the DRG as the target.

### Task 2c: Configure the security rules

To ensure traffic flows between the IP network and VCN, make sure the IP network security rules and the VCN's security rules allow the desired traffic.

Here are the types of rules to add:

- Ingress rules for the types of traffic you want to allow into one cloud from the other, specifically from the other cloud's CIDR block.
- Egress rule to allow outgoing traffic from one cloud to the other. If the VCN's subnet already has a broad egress rule for all types of protocols to all destinations (0.0.0.0/0), then you don't need to add a special one for the IP network.

### For the IP network

[Configure the network security rules](#) for the IP network to allow the desired traffic.

### For the VCN



#### Note

The following procedure uses security lists, but you could instead implement the security rules in one or more [network security groups](#) and then place the VCN's resources in NSGs.

1. Determine which subnets in your VCN need to communicate with the IP network.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic specifically with the CIDR block of the IP network:
  - a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
  - b. Click the security list you're interested in.  
Under **Resources**, you can click **Ingress Rules** or **Egress Rules** to switch between the different types of rules.
  - c. Add one or more rules, each for the specific type of traffic you want to allow.

### Example

Let's say you want to add a stateful rule that enables ingress HTTPS (port 443) traffic from the IP network's CIDR. Here are the basic steps you take when adding a rule:

- i. On the **Ingress Rules** page, click **Add Ingress Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Source CIDR:** Enter the same CIDR block that the route rules use (see [Task 2b: Configure routing to the DRG](#)).
- iv. **IP Protocol:** Leave as TCP.
- v. **Source Port Range:** Leave as All.
- vi. **Destination Port Range:** Enter 443.
- vii. Click **Add Ingress Rule**.

For more information about setting up security list rules, see [Security Lists](#).



#### Important

The VCN's [default security list](#) does not allow ICMP echo reply and echo request (ping). You must add rules to enable that traffic. See [Rules to Enable Ping](#)

### Task 2d: Create a CPE object

[Create a CPE object](#). You must provide an IP address. Use the VPNaaS gateway's public IP address.

### Task 2e: Create the IPSec connection

[From your DRG, create an IPSec connection to the CPE object.](#) You must provide one or more static routes. The values must match the IP network's subnets or aggregate.

The resulting IPSec connection consists of two tunnels. Record the IP address and shared secret for one of those tunnels. In the next task, you will provide those values.

### Task 3: Update the VPNaaS connection with the tunnel information

[Update the VPNaaS connection.](#) Use these values:

- **Customer Gateway:** The tunnel's IP address from the preceding task.
- **VPNaaS VPN Connection's Pre-shared Key:** The tunnel's shared secret from the preceding task.

After the VPNaaS connection is updated and provisioned, the state of your VCN's IPSec tunnel should change to Available. This might take a few minutes.

### Task 4: Test the connection

After the tunnel state changes to Available, test the connection. Depending on how you've set up your IP network's security rules and VCN security rules, you should be able to [launch an instance](#) in your VCN and access it from an instance in the IP network. Or you should be able to connect from the VCN instance to an instance in the IP network. If you can, your connection is ready to use.

## Terminating the Connection

If you want to terminate the connection, delete the IPSec connection:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **IPSec Connections**.

A list of the IPSec connections in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the IPSec connection you're interested in.
3. Click **Terminate**.
4. Confirm the deletion when prompted.

The IPSec connection will be in the Terminating state for a short period while it's being deleted.

## Access to Microsoft Azure

Oracle and Microsoft have created a cross-cloud connection between Oracle Cloud Infrastructure and Microsoft Azure in certain regions. This connection lets you set up cross-cloud workloads without the traffic between the clouds going over the internet. This topic describes how to set up virtual networking infrastructure resources to enable this kind of cross-cloud deployment.

### Highlights

- You can connect a Microsoft Azure virtual network (VNet) with an Oracle Cloud Infrastructure virtual cloud network (VCN) and run a cross-cloud workload. In the typical use case, you deploy your Oracle Database on Oracle Cloud Infrastructure, and deploy an Oracle, .NET, or custom application in Microsoft Azure.
- The two virtual networks must belong to the same company and not have overlapping CIDRs. The connection requires you to create an Azure ExpressRoute circuit and an Oracle Cloud Infrastructure FastConnect virtual circuit.
- The connection is currently available only in these areas:
  - Between the Oracle Cloud Infrastructure location in the US East (Ashburn) region and the [Azure Washington DC and Washington DC2 locations](#).

- Between the Oracle Cloud Infrastructure location in the UK South (London) region and the [Azure London location](#).

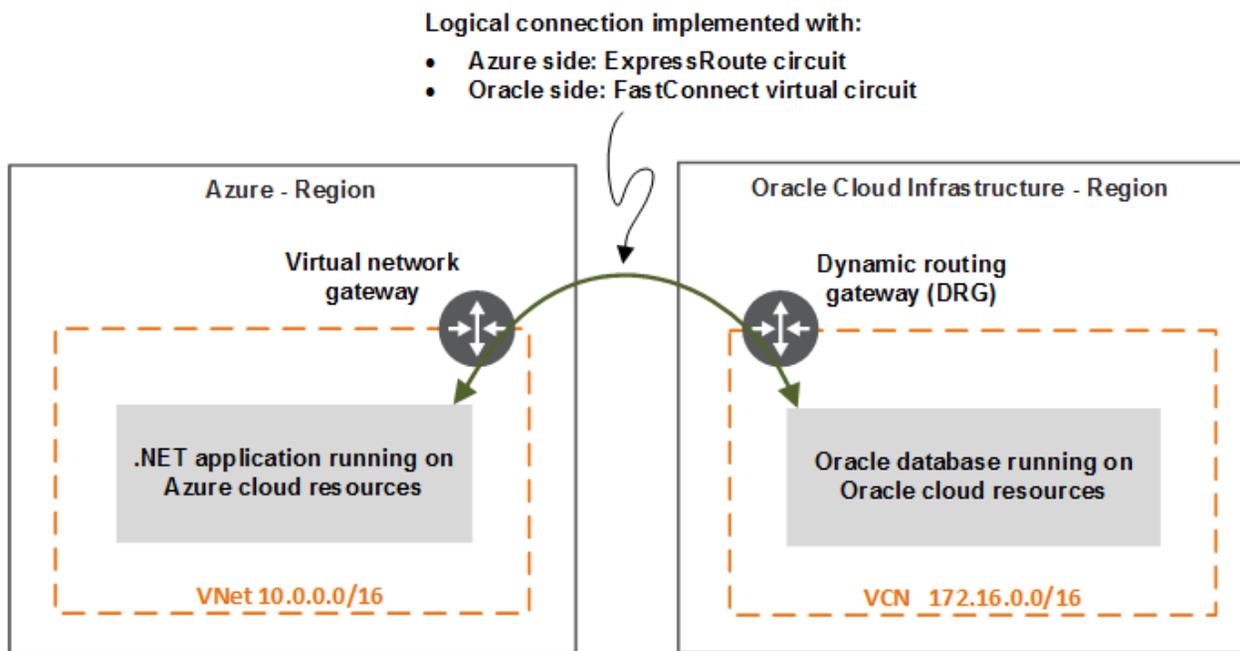
## Overview of Supported Traffic

Here are more details about the supported types of traffic.

### VNet-to-VCN Connection: Extension from One Cloud to Another

You can connect your VNet and VCN so that traffic that uses private IP addresses goes over the cross-cloud connection.

For example, the following diagram shows a VNet that is connected to a VCN. Resources in the VNet are running a .NET application that access an Oracle database that runs on Database service resources in the VCN. The traffic between the application and database uses a logical circuit that runs on the cross-cloud connection between Azure and Oracle Cloud Infrastructure.



To enable the connection between the VNet and VCN, you set up an Azure ExpressRoute circuit and an Oracle Cloud Infrastructure FastConnect virtual circuit. The connection has built-in redundancy, which means you need to set up only a single ExpressRoute circuit and single FastConnect virtual circuit. The bandwidth for the connection is the bandwidth value you choose for the ExpressRoute circuit.

For instructions, see [Setting Up a VNet-to-VCN Connection](#).

### Peered VCNs

The connection enables traffic to flow from the VNet through the connected VCN to a peered VCN in the same Oracle Cloud Infrastructure region, or a different region.

### Types of Traffic Not Supported by the Connection

This cross-cloud connection does not enable traffic between your on-premises network through the VCN to the VNet, or from your on-premises network through the VNet to the VCN.

## Important Implications of Connecting Clouds

This section summarizes some access control, security, and performance implications of connecting your VCN to a VNet. In general, you can control access and traffic by using IAM policies, route tables in the VCN, and security rules in the VCN.

The sections that follow discuss implications from the perspective of your VCN. There are similar implications for your VNet. As with your VCN, you can use Azure resources such as route tables and network security groups to secure your VNet.

### Controlling the Establishment of a Connection

With Oracle Cloud Infrastructure IAM policies, you can control:

- Who in your organization has the authority to create a FastConnect virtual circuit (see [Prerequisites: Required IAM Policy](#)). Be aware that deletion of the relevant IAM policy does not affect any existing connections to a VNet, only the ability for a future

connection to be created.

- Who can manage [route tables](#), [network security groups](#), and [security lists](#).

### **Controlling Traffic Flow Over the Connection**

Even if a connection has been established between your VCN and VNet, you can control the packet flow over the connection with route tables in your VCN. For example, you can restrict traffic to only specific subnets in the VNet.

Without terminating the connection, you can stop traffic flow to the VNet by simply removing route rules that direct traffic from your VCN to the VNet. You can also effectively stop the traffic by removing any security rules that enable ingress or egress traffic with the VNet. This doesn't stop traffic flowing over the connection, but stops it at the [VNIC](#) level.

### **Controlling the Specific Types of Traffic Allowed**

It's important that you ensure that all outbound and inbound traffic with the VNet is intended/expected and well defined. In practice, this means implementing Azure network security group and Oracle security rules that explicitly state the types of traffic one cloud can send to the other and accept from the other.



### Important

Your Oracle Cloud Infrastructure instances running Oracle-provided Linux images or Windows images also have firewall rules that control access to the instance. When troubleshooting access to an instance, make sure all of the following items are set correctly: the network security groups that the instance is in, the security lists associated with the instance's subnet, and the instance's firewall rules. For more information, see [Oracle-Provided Images](#).

If your instance is running Oracle Linux 7, you need to use [firewalld](#) to interact with the iptables rules. For your reference, here are commands for opening a port (1521 in this example):

```
sudo firewall-cmd --zone=public --permanent --add-
port=1521/tcp

sudo firewall-cmd --reload
```

For instances with an iSCSI boot volume, the preceding `--reload` command can cause problems. For details and a workaround, see [Instances experience system hang after running firewall-cmd --reload](#).

In addition to security rules and firewalls, you should evaluate other OS-based configuration on the instances in your VCN. There could be default configurations that don't apply to your own VCN's CIDR, but inadvertently apply to the VNet's CIDR.

### Using Default Security List Rules with Your VCN

If your VCN's subnets use the [default security list](#) with the default rules, be aware that there are two rules that allow ingress traffic from anywhere (that is, 0.0.0.0/0, and thus the VNet):

- Stateful ingress rule that allows TCP port 22 (SSH) traffic from 0.0.0.0/0 and any source port
- Stateful ingress rule that allows ICMP type 3, code 4 traffic from 0.0.0.0/0 and any source port

Make sure to evaluate these rules and whether you want to keep or update them. As stated earlier, you should ensure that all inbound or outbound traffic that you permit is intended/expected and well defined.

### Preparing for Performance Impact and Security Risks

In general, you should prepare your VCN for the ways it could be affected by the VNet. For example, the load on your VCN or its instances could increase. Or your VCN could experience a malicious attack directly from or by way of the VNet.

Regarding performance: If your VCN is providing a service to the VNet, be prepared to scale up your service to accommodate the demands of the VNet. This might mean being prepared to launch additional instances as necessary. Or if you're concerned about high levels of network traffic coming to your VCN, consider using [stateless security rules](#) to limit the level of connection tracking your VCN must perform. Stateless security rules can also help slow the impact of a denial-of-service (DoS) attack.

Regarding security risks: If the VNet is connected to the internet, be aware that your VCN can be exposed to bounce attacks in which a malicious host on the internet can send traffic to your VCN but make it look like it's coming from the VNet. To guard against this, as mentioned earlier, use your security rules to carefully limit the inbound traffic from the VNet to expected and well-defined traffic.

## Setting Up a VNet-to-VCN Connection

This section describes how to set up the logical connection between a VNet and VCN (for background, see [Overview of Supported Traffic](#)).

### Prerequisites: Resources You Need

You must already have:

- An [Azure VNet](#) with [subnets](#) and [virtual network gateway](#)
- An Oracle Cloud Infrastructure [VCN with subnets](#) and an [attached dynamic routing gateway \(DRG\)](#). It's easy to forget to [attach the DRG](#) to your VCN after you create it. If you already have an IPSec VPN or FastConnect between your on-premises network and VCN, then your VCN already has an attached DRG. You use that same DRG here when setting up the connection to Azure.

As a reminder, here is a table that lists the comparable networking components involved in each side of the connection.

Component	Azure	Oracle Cloud Infrastructure
Virtual network	VNet	VCN
Virtual circuit	ExpressRoute circuit	FastConnect private virtual circuit
Gateway	virtual network gateway	dynamic routing gateway (DRG)
Routing	route tables	route tables
Security rules	network security groups (NSGs)	network security groups (NSGs), security lists

### Prerequisites: BGP Information You Need

The connection between the VNet and VCN uses BGP dynamic routing. When you set up the Oracle virtual circuit, you provide the BGP IP addresses that will be used for the two

redundant BGP sessions between Oracle and Azure:

- A primary pair of BGP addresses (one IP address for the Oracle side, one IP address for the Azure side)
- A separate, secondary pair of BGP addresses (one IP address for the Oracle side, one IP address for the Azure side)

**For each pair, you must provide a separate /30 block of addresses** (each /30 has four IP addresses).

The second and third addresses in each /30 are used for the BGP IP address pair. Specifically:

- The second address in the block is for the Oracle side of the BGP session
- The third address in the block is for the Azure side of the BGP session

The first and last addresses in the block are used for other internal purposes.

For example, if the /30 is 10.0.0.20/30, then the addresses in the block are:

- 10.0.0.20/30
- 10.0.0.21/30: Use this for the Oracle side
- 10.0.0.22/30: Use this for the Azure side (also referred to as the "Customer" side in the Oracle Console)
- 10.0.0.23/30

Remember that you must also provide a /30 block for the secondary BGP addresses. For example: 10.0.0.24/30. In this case, 10.0.0.25/30 is for the Oracle side, and 10.0.0.26/30 is for the Azure side.

### **Prerequisites: Required IAM Policy**

It's assumed that you have the necessary Azure Active Directory access and Oracle Cloud Infrastructure IAM access to create and work with the relevant Azure and Oracle networking resources. Specifically for IAM: If your user is in the [Administrators group](#), you have the required authority.

## CHAPTER 23 Networking

If your user is not, then a policy like this one generally covers all the Networking resources:

```
Allow group NetworkAdmins to manage virtual-network-family in tenancy
```

To *only* create and manage a virtual circuit, you must have a policy like this:

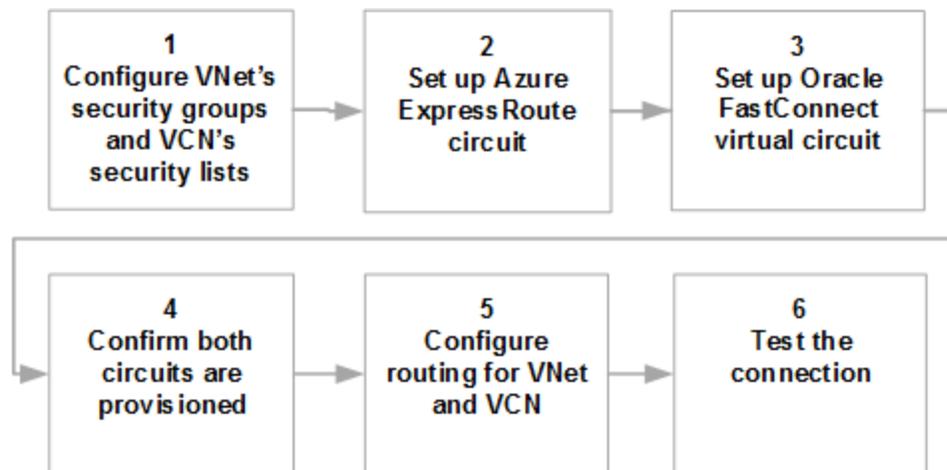
```
Allow group VirtualCircuitAdmins to manage drgs in tenancy
```

```
Allow group VirtualCircuitAdmins to manage virtual-circuits in tenancy
```

For more information, see [IAM Policies for Networking](#).

### Overall Process

The following flow chart shows the overall process of connecting your VNet and VCN.



### Task 1: Configure the network security groups and security rules

The first task is to determine what traffic needs to flow between the relevant subnets within the VNet and VCN, and then configure the VNet security groups and VCN security rules accordingly. Here are the general types of rules to add:

- Ingress rules for the types of traffic you want to allow into one cloud from the other, specifically from the other cloud's relevant subnets.
- Egress rule to allow outgoing traffic from one cloud to the other. If the VCN's subnet already has a broad egress rule for all types of protocols to all destinations (0.0.0.0/0), then you don't need to add a special one for the traffic to the VNet. The VCN's default security list includes a broad default egress rule like this.

More specifically, here are recommended types of traffic to allow between the VNet and VCN:

- [Ping traffic](#) in both directions for testing the connection from each side
- SSH (TCP port 22)
- Client connections to an Oracle database (SQL\*NET on TCP port 1521)

Make sure to only allow traffic to and from the specific address ranges of interest (for example, the other cloud's relevant subnets).

**For the VNet:** Determine which subnets in your VNet need to communicate with the VCN. Then [configure the network security groups](#) for those subnets to allow the desired traffic.

**For the VCN:**



### Note

The following procedure uses security lists, but you could instead implement the security rules in one or more [network security groups](#) and then place the VCN's relevant resources in NSGs.

1. Determine which subnets in your VCN need to communicate with the VNet.
2. Update the security list for each of those subnets to include rules to allow the desired egress or ingress traffic specifically with the VNet's CIDR block or a subnet of the VNet:

- a. In the Console, while viewing the VCN you're interested in, click **Security Lists**.
- b. Click the security list you're interested in.
- c. Click **Edit All Rules** and create one or more rules, each for the specific type of traffic you want to allow.

### Example: Outgoing ping from VCN to VNet

The following egress security rule lets an instance initiate a ping request to a host outside the VCN (echo request ICMP type 8). This is a stateful rule that automatically allows the response. No separate ingress rule for echo reply ICMP type 0 is required.

- i. In the **Allow Rules for Egress** section, click **+Add Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Destination CIDR:** The relevant subnet in the VNet (10.0.0.0/16 in the preceding diagram)
- iv. **IP Protocol:** ICMP
- v. **Type and Code:** 8

### Example: Incoming ping to VCN from VNet

The following ingress security rule lets an instance receive a ping request from a host in the VNet (echo request ICMP type 8). This is a stateful rule that automatically allows the response. No separate egress rule for echo reply ICMP type 0 is required.

- i. In the **Allow Rules for Ingress** section, click **+Add Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Source CIDR:** The relevant subnet in the VNet (10.0.0.0/16 in the preceding diagram)

- iv. **IP Protocol:** ICMP
- v. **Type and Code:** 8

### Example: Incoming SSH to VCN

The following ingress security rule lets an instance receive an SSH connection (TCP port 22) from a host in the VNet.

- i. In the **Allow Rules for Ingress** section, click **+Add Rule**.
- ii. Leave the **Stateless** check box unselected.
- iii. **Source CIDR:** The relevant subnet in the VNet (10.0.0.0/16 in the preceding diagram)
- iv. **IP Protocol:** TCP
- v. **Source Port Range:** All
- vi. **Destination Port Range:** 22

### Example: SQL\*Net connections to database

The following ingress security rule allows SQL\*Net connections (TCP port 1521) from hosts in the VNet.

- i. In the **Allow Rules for Ingress** section, click **+Add Rule**.
  - ii. Leave the **Stateless** check box unselected.
  - iii. **Source CIDR:** The relevant subnet in the VNet (10.0.0.0/16 in the preceding diagram)
  - iv. **IP Protocol:** TCP
  - v. **Source Port Range:** All
  - vi. **Destination Port Range:** 1521
- d. Click **Save Security List Rules** at the bottom of the dialog box.

For more information about setting up security rules, see [Security Rules](#).

### Task 2: Set up Azure ExpressRoute circuit

[Set up an ExpressRoute circuit](#) to *Oracle Cloud FastConnect*. During the circuit setup, you receive a service key from Microsoft. Record that service key, because you must provide it to Oracle in the next task.

Note that in the next task, you set up a FastConnect private virtual circuit to *Microsoft Azure: ExpressRoute*. When that virtual circuit finishes being provisioned, your ExpressRoute circuit will update to show that private peering is enabled.

### Task 3: Set up an Oracle Cloud Infrastructure FastConnect virtual circuit



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

1. In the Console, confirm you're viewing the compartment that you want to work in. If you're not sure which one, use the compartment that contains the DRG that you'll connect to. This choice of compartment, along with a corresponding [IAM policy](#), control who has access to the virtual circuit you're about to create.
2. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.  
The resulting **FastConnect** page is where you create a new virtual circuit and later return to when you need to manage the virtual circuit.
3. Click **Create Connection**.

4. Select **Connect through a provider** and choose **Microsoft Azure: ExpressRoute** from the list.
5. Enter the following for your virtual circuit:
  - **Name:** A friendly name of your choice. The value does not need to be unique across your virtual circuits, and you can change it later. Avoid entering confidential information.
  - **Create in Compartment:** Leave as is (the compartment you're currently working in).
  - **Virtual Circuit Type:** Select **Private Virtual Circuit**.
  - **Dynamic Routing Gateway Compartment:** Select the compartment where the DRG resides (it should already be selected).
  - **Dynamic Routing Gateway:** Select the DRG.
  - **Provisioned Bandwidth:** Choose the same bandwidth level you chose for the ExpressRoute circuit (or the closest value available).
  - **Provider Service Key:** Enter the key you received from Microsoft when you set up the ExpressRoute circuit.
  - **Customer Primary BGP IP Address:** This is the Azure primary BGP IP address. Enter the third address in the primary /30 block that you provide. For example: 10.0.0.22/30. For more information about this field and the next ones, see [Prerequisites: BGP Information You Need](#).
  - **Oracle Primary BGP IP address (optional):** You can leave this blank and Oracle will infer the address based on the /30 block you provided for the Azure BGP IP address. In this example, the correct value would be 10.0.0.21/30.
  - **Customer Secondary BGP IP Address:** This is the Azure secondary BGP IP address. Enter the third address in the secondary /30 block that you provide. For example: 10.0.0.26/30.

- **Oracle Primary BGP IP Address (optional):** You can leave this blank and Oracle will infer the address based on the /30 block you provided for the Azure BGP IP address. In this example, the correct value would be 10.0.0.25/30.
6. Click **Continue**.  
The virtual circuit is created.
  7. Click **Close**.

After you create the Oracle virtual circuit, you do not need to contact Azure to request provisioning of the circuit. It happens automatically.

### Task 4: Confirm that both circuits are provisioned

Within a few minutes, both circuits should be provisioned. To verify:

- For the ExpressRoute circuit, confirm that private peering is provisioned.
- For the FastConnect virtual circuit, confirm that its status is UP. See [To get the status of your FastConnect virtual circuit](#).

### Task 5: Configure the route tables

**For the VNet:** Determine which subnets in your VNet need to communicate with the VCN. Then [configure the route tables](#) for those subnets to route the desired traffic to the VNet gateway.

**For the VCN:**

1. Determine which subnets in your VCN need to communicate with the VNet.
2. Update the route table for each of those subnets to include a new rule that directs traffic destined for the VNet's CIDR to your DRG:
  - a. In the Console, while viewing the VCN you're interested in, click **Route Tables**.
  - b. Click the route table you're interested in.

- c. Click **Edit Route Rules**.
- d. Click **+ Another Route Rule** and enter the following:
  - **Target Type:** Dynamic Routing Gateway.
  - **Destination CIDR Block:** The relevant subnet in the VNet (10.0.0.0/16 in the preceding diagram).
  - **Compartment:** The compartment where the DRG is located, if not the current compartment.
  - **Target Dynamic Routing Gateway:** The DRG.
- e. Click **Save**.

Any subnet traffic with a destination that matches the rule is routed to your DRG. The DRG then knows to route the traffic to the VNet based on the virtual circuit's BGP session information.

Later, if you no longer need the connection and want to delete your DRG, you must first delete all the route rules in your VCN that specify the DRG as the target.

For more information about setting up route rules, see [Route Tables](#).

### Task 6: Test the connection

Depending on how you've set up your VNet security groups and VCN security rules, you should be able to [create an instance](#) in your VCN and access it from a host in the VNet. Or you should be able to connect from the instance to a host in the VNet. If you can, your connection is ready to use.



#### **Important**

If you decide to terminate the connection, you must follow a particular process. See [To terminate the connection to Azure](#).

## Managing a VNet-to-VCN Connection

### To get the status of your FastConnect virtual circuit

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection you're interested in. If the icon for the virtual circuit is green and says UP, the virtual circuit is provisioned and BGP has been correctly configured. The virtual circuit is ready to use.

### To edit a FastConnect virtual circuit

You can change these items for your virtual circuit:

- The name
- Which DRG it uses



#### Warning

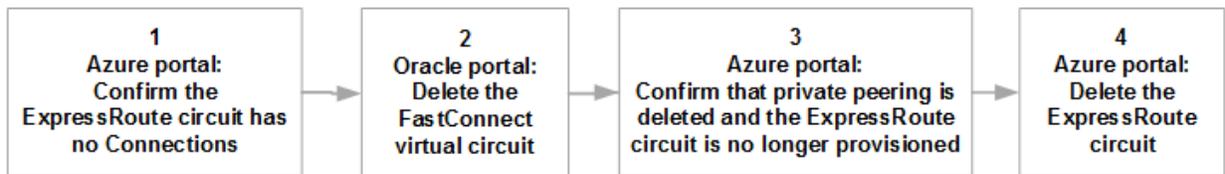
If your virtual circuit is in the PROVISIONED state, be aware that changing which DRG it uses switches the state to PROVISIONING and **may cause the connection to go down**. After Oracle reprovisions the virtual circuit, its state returns to PROVISIONED. Make sure to confirm that the connection is back up and working.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
2. Select the compartment where the connection resides, and then click the connection.

3. Click the virtual circuit.
4. Click **Edit** and make your changes.
5. Click **Save**.

### To terminate the connection to Azure

The following flow chart shows the overall process of terminating a VNet-to-VCN connection.



1. In the Azure portal, view the ExpressRoute circuit, and then view its **Connections**. Confirm that there are no **Connections** still in existence for the ExpressRoute circuit. All **Connections** [must first be deleted](#) before proceeding.
2. In the Oracle portal, delete your FastConnect virtual circuit:
  - a. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **FastConnect**.
  - b. Select the compartment where the connection resides, and then click the connection.
  - c. Click the virtual circuit.
  - d. Click **Delete**.
  - e. Confirm when prompted.  
The virtual circuit's Lifecycle State switches to TERMINATING.
3. In the Azure portal, confirm that the private peering for the ExpressRoute circuit has been deleted. Also confirm that the ExpressRoute circuit's status has changed to "Not

Provisioned".

4. In the Azure portal, [delete the ExpressRoute circuit](#).

The connection between Azure and Oracle Cloud Infrastructure is terminated.

## Access to Other Clouds with Libreswan

[Libreswan](#) is an open source IPsec implementation that is based on FreeS/WAN and Openswan. Most Linux distributions include Libreswan or make it easy to install. You can install it on hosts in either your on-premises network or a cloud provider network. This topic shows how to connect your Oracle Cloud Infrastructure virtual cloud network (VCN) with another cloud provider by using an IPsec VPN with a Libreswan VM as the customer-premises equipment (CPE).

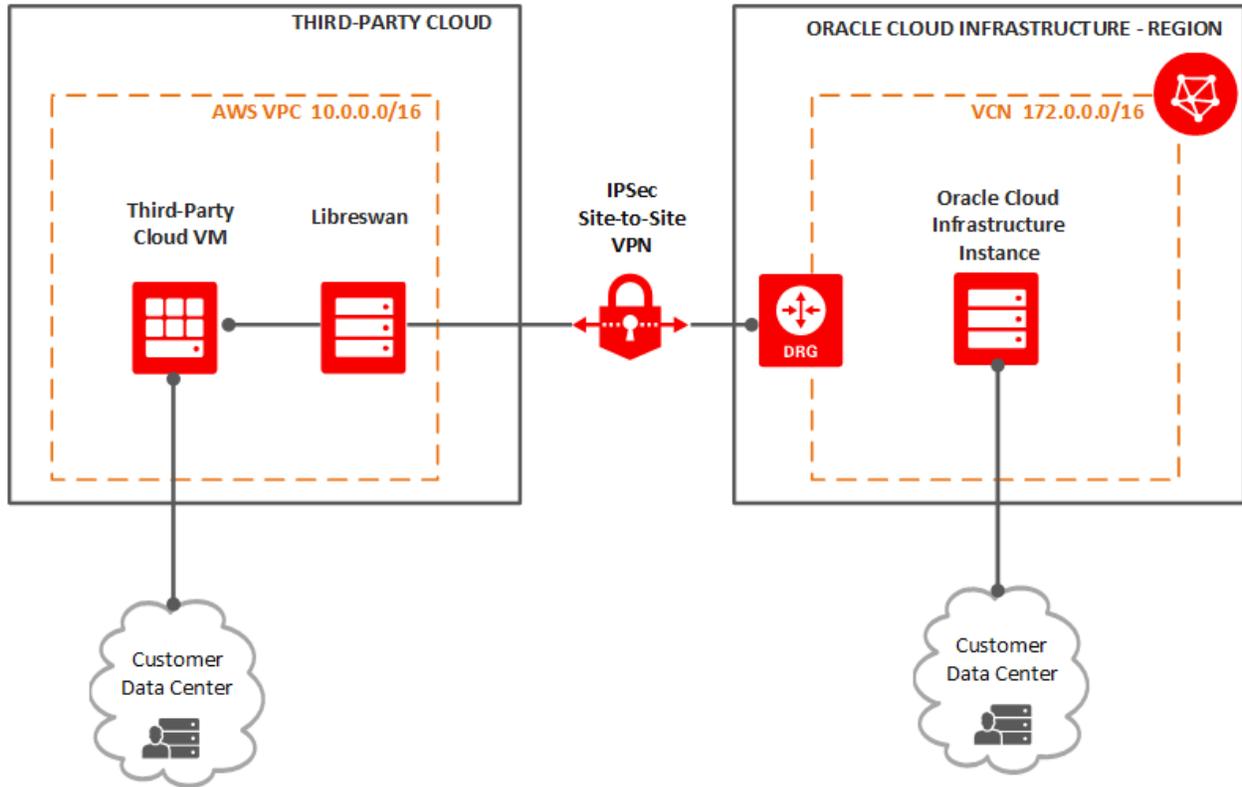
In the example shown here, the other cloud provider is Amazon Web Services (AWS). The connection is a secure and encrypted site-to-site IPsec VPN between the Oracle and Amazon environments. It enables resources in the two clouds to communicate with each other using their private IP addresses as if they are in the same network segment.

A [Libreswan CPE guide](#) is also available for all other use cases.

Virtual tunnel interface (VTI) support for this route-based configuration requires minimum Libreswan version 3.18 and a recent Linux 3.x or 4.x kernel. This configuration was validated using Libreswan version 3.29.

## Architecture

The following diagram shows the general layout of the connection.



## Supported IPsec Parameters

For a vendor-neutral list of supported IPsec parameters for all regions, see [Supported IPsec Parameters](#).

The Oracle BGP ASN for the commercial cloud is 31898. If you're configuring VPN Connect for the Government Cloud, see [Required VPN Connect Parameters for Government Cloud](#) and also [Oracle's BGP ASN](#).

### Configuration



#### Important

The configuration instructions in this section are provided by Oracle Cloud Infrastructure for Libreswan. If you need support or further assistance, consult the [Libreswan documentation](#).

Libreswan supports both route-based and policy-based tunnels. The tunnel types can coexist without interfering with each other. The Oracle VPN headends use route-based tunnels. Oracle recommends that you configure Libreswan with the [Virtual Tunnel Interface \(VTI\) configuration syntax](#).

Refer to [Supported IPSec Parameters](#) for more details about the specific parameters used in this document.

#### Default Libreswan Configuration Files

The default Libreswan installation creates the following files:

- `etc/ipsec.conf`: The root of the Libreswan configuration.
- `/etc/ipsec.secrets`: The root of the location where Libreswan looks for secrets (the tunnel pre-shared keys).
- `/etc/ipsec.d/`: A directory for storing the `.conf` and `.secrets` files for your Oracle Cloud Infrastructure tunnels (for example: `oci-ipsec.conf` and `oci-ipsec.secrets`). Libreswan encourages you to create these files in this folder.

The default `etc/ipsec.conf` file includes this line:

```
include /etc/ipsec.d/*.conf
```

The default `etc/ipsec.secrets` file includes this line:

```
include /etc/ipsec.d/*.secrets
```

The preceding lines automatically merge all the `.conf` and `.secrets` files in the `/etc/ipsec.d` directory into the main configuration and secrets files that Libreswan uses.

### About Using IKEv2

Oracle supports Internet Key Exchange version 1 (IKEv1) and version 2 (IKEv2). If you [configure the IPSec connection in the Console to use IKEv2](#), you must configure your CPE to use only IKEv2 and related IKEv2 encryption parameters that your CPE supports. For a list of parameters that Oracle supports for IKEv1 or IKEv2, see [Supported IPSec Parameters](#).

You specify the IKE version when setting up the IPSec configuration file in [task 4](#) in the next section. In that example file, there's a comment showing how to configure IKEv1 versus IKEv2.

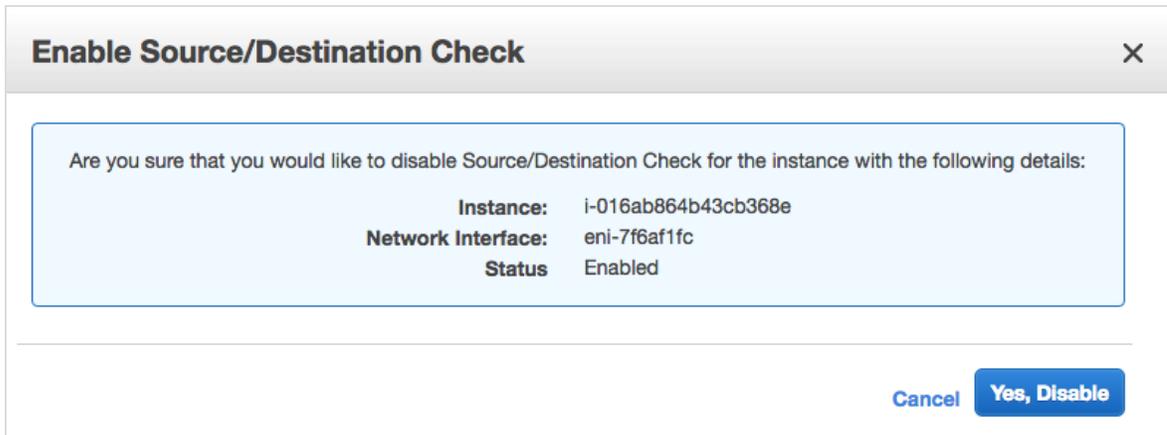
### Configuration Process

#### Task 1: Prepare the AWS Libreswan instance

1. Using the AWS Console or APIs, create a Libreswan VM by using its provisioning process. Use Oracle Linux, CentOS, or Red Hat as the main operating system.
2. After the new instance starts, connect to it with SSH and install the Libreswan package:

```
sudo yum -y install libreswan
```

3. In the AWS Console, disable the source and destination checks on the Libreswan VM instance by right-clicking the instance, clicking **Networking**, and then clicking **Change Source/Dest. Check**. When prompted, click **Yes, Disable**.



4. On the Libreswan VM, configure IP\_forward to allow AWS clients to send and receive traffic through the Libreswan VM. In the `/etc/sysctl.conf` file, set the following values and apply the updates with `sudo sysctl -p`.

```
net.ipv4.ip_forward=1
net.ipv4.conf.all.accept_redirects = 0
net.ipv4.conf.all.send_redirects = 0
net.ipv4.conf.default.send_redirects = 0
net.ipv4.conf.eth0.send_redirects = 0
net.ipv4.conf.default.accept_redirects = 0
net.ipv4.conf.eth0.accept_redirects = 0
```

- In the AWS Console, edit your AWS route table. Add a rule with the VCN CIDR (172.0.0.0/16) as the destination and the AWS Libreswan instance ID (i-016ab864b43cb368e in this example) as the target.

The screenshot shows the AWS console interface for a route table named 'rtb-2b5a7f57'. The 'Routes' tab is selected, displaying a table of route entries. The table has four columns: Destination, Target, Status, and Propagated. There are three rows of data, each representing a different route configuration.

Destination	Target	Status	Propagated
10.0.0.0/16	local	Active	No
0.0.0.0/0	<a href="#">igw-22029d5a</a>	Active	No
172.0.0.0/16	eni-7f6af1fc / i-016ab864b43cb368e	Active	No

- In the AWS Console, enable inbound TCP and UDP traffic on ports 4500 and 500 to allow Oracle Cloud Infrastructure IPSec VPN communication with the AWS Libreswan VM. This task includes editing both the AWS security groups and network ACLs. You can set the source value can be the Oracle public IP (the Oracle VPN headend IPSec tunnel endpoint) instead of 0.0.0.0/0.

## CHAPTER 23 Networking

For security groups:

The screenshot shows the AWS IAM console interface for a security group. At the top, there is a filter set to 'All security groups' and a search bar. Below this is a table listing security groups, with 'libreswan' selected. The details for 'libreswan' are shown below, including tabs for 'Summary', 'Inbound Rules', 'Outbound Rules', and 'Tags'. The 'Inbound Rules' tab is active, displaying a table of rules.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP (6)	4500	0.0.0.0/0	libreswan
ALL Traffic	ALL	ALL	sg-bcdb8df5	
SSH (22)	TCP (6)	22	0.0.0.0/0	
Custom TCP Rule	TCP (6)	500	0.0.0.0/0	libreswan
Custom UDP Rule	UDP (17)	4500	0.0.0.0/0	libreswan
Custom UDP Rule	UDP (17)	500	0.0.0.0/0	libreswan

## CHAPTER 23 Networking

For network ACLs:

The screenshot displays the Oracle Cloud Infrastructure console interface for a Network ACL. At the top, there is a search bar and a table with columns: Name, Network ACL ID, Associated With, Default, and VPC. The selected Network ACL is 'acl-5b68cb21', associated with '1 Subnet', and is the 'Default' ACL for VPC 'vpc-fd7a3486 | libreswanvpc'. Below this, the 'acl-5b68cb21' configuration page is shown, with tabs for Summary, Inbound Rules (selected), Outbound Rules, Subnet Associations, and Tags. A note states: 'Allows inbound traffic. Because network ACLs are stateless, you must create inbound and outbound rules.' An 'Edit' button is visible. A 'View:' dropdown is set to 'All rules'. Below is a table of rules:

Rule #	Type	Protocol	Port Range	Source	Allow / Deny
100	ALL Traffic	ALL	ALL	0.0.0.0/0	ALLOW
200	Custom TCP Rule	TCP (6)	4500	0.0.0.0/0	ALLOW
300	Custom TCP Rule	TCP (6)	500	0.0.0.0/0	ALLOW
400	Custom UDP Rule	UDP (17)	4500	0.0.0.0/0	ALLOW
500	Custom UDP Rule	UDP (17)	500	0.0.0.0/0	ALLOW
600	SSH (22)	TCP (6)	22	0.0.0.0/0	ALLOW
*	ALL Traffic	ALL	ALL	0.0.0.0/0	DENY

### Task 2: Configure the Oracle Cloud Infrastructure DRG and CPE object

1. In the Oracle Console, create a customer-premises equipment (CPE) object that points to the Libreswan AWS instance public IP address (34.200.255.174).

Customer-Premises Equipment *in* libreswan-demo *Compartment*

Create Customer-Premises Equipment

Name	IP Address	OCID	Created
<a href="#">AWS-CPE</a>	34.200.255.174	...mmw7ea <a href="#">Show</a> <a href="#">Copy</a>	Wed, Sep 11, 2019, 6:44:00 PM UTC

Showing 1 Item

- If you don't already have a DRG attached to your VCN: in the Oracle Console, create a DRG and then attach it to the VCN (172.0.0.0/16).

Networking » Dynamic Routing Gateways » Dynamic Routing Gateway Details » Attached Virtual Cloud Networks



## AWS-DRG

[Apply Tag\(s\)](#) [Terminate](#)

**Dynamic Routing Gateway Information** [Tags](#)

OCID: ...avco7a [Show](#) [Copy](#) Created: Wed, 11 Sep 2019 18:45:23 UTC

AVAILABLE

Resources

- IPSec Connections (1)
- Virtual Cloud Networks (1)**
- Virtual Circuits (0)
- Remote Peering Connections (0)

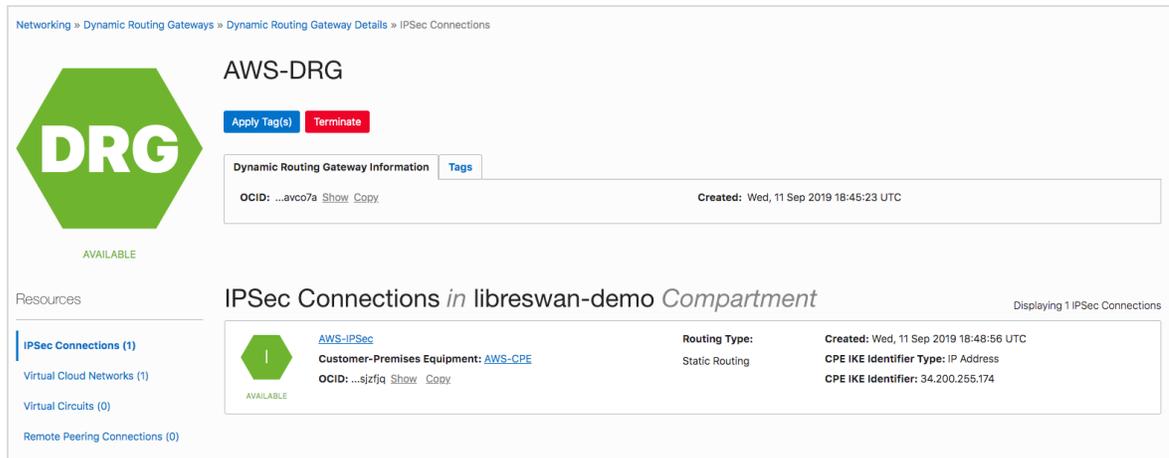
### Virtual Cloud Networks *in* libreswan-demo *Compartment*

Displaying 1 Attached Virtual Cloud Networks

[Attach to Virtual Cloud Network](#)

 VCN17Z ATTACHED	OCID: ...2pgxvq <a href="#">Show</a> <a href="#">Copy</a>	Compartment: libreswan-demo Attachment OCID: ...a4fcq <a href="#">Show</a> <a href="#">Copy</a>	CIDR Block: 172.0.0.0/16 Route Table: - <a href="#">ⓘ</a>
-----------------------------------------------------------------------------------------------------------	-----------------------------------------------------------	----------------------------------------------------------------------------------------------------	--------------------------------------------------------------

3. In the Oracle Cloud Console, create an IPsec connection and point it to the AWS VPC CIDR (10.0.0.0/16). In other words, when you create the IPsec connection, set its static route to the AWS VPC CIDR.



For each configured IPsec connection, Oracle creates two tunnels and assigns these items to each one:

- Oracle VPN headend IPsec tunnel endpoint
- Oracle VPN tunnel shared secret

You can view the IPsec tunnel status and Oracle VPN headend IP by clicking the Actions icon (three dots) for the IPsec connection, and then clicking **View Details**. Initially each tunnel is in the DOWN state (offline) because you still have some additional configuration to do later on the AWS Libreswan VM.

## CHAPTER 23 Networking

Networking » IPsec Connections » AWS-IPsec

### AWS-IPsec

[Edit](#) [Add Tags](#) [Terminate](#)

**IPsec Connection Information** [Tags](#)

**Static Route CIDR:** 10.0.0.0/16 ⓘ  
**Created:** Wed, Sep 11, 2019, 6:48:56 PM UTC  
**OCID:** ...sjzfq [Show](#) [Copy](#)

**DRG:** [AWS-DRG](#)  
**CPE:** [AWS-CPE](#)  
**CPE IKE Identifier Type:** IP Address  
**CPE IKE Identifier:** 34.200.255.174

Resources

[Tunnels \(2\)](#)

#### Tunnels in libreswan-demo Compartment

Name	Lifecycle State ⓘ	IPsec Status ⓘ	BGP Status ⓘ	Oracle VPN IP Address	Routing Type
<a href="#">AWS-IPsec-1</a>	● Available	● Down	—	129.146.12.51	Static Routing
<a href="#">AWS-IPsec-2</a>	● Available	● Down	—	129.146.13.49	Static Routing

Showing 2 Items

To view the shared secret, click the Actions icon (three dots) for an individual tunnel, and then click **View Details**. Next to **Shared Secret**, click **Show**.

Networking » IPsec Connections » AWS-IPsec » AWS-IPsec-1

### AWS-IPsec-1

[Edit](#)

**Tunnel Information**

**IPsec Status:** ● Down ⓘ  
**BGP Status:** —  
**Created:** Wed, Sep 11, 2019, 6:48:56 PM UTC  
**OCID:** ...76sfua [Show](#) [Copy](#)

**Routing Type:** Static Routing  
**BGP ASN:** —  
**Inside Tunnel Interface - CPE:** — ⓘ  
**Inside Tunnel Interface - Oracle:** — ⓘ  
**Shared Secret:** \*\*\*\*\* [Show](#) [Edit](#)

- In the Oracle Console, edit the VCN's security rules to enable ingress TCP and UDP traffic on ports 4500 and 500 like you did for the AWS security groups and network ACLs. You can use the AWS Libreswan VM public IP address instead of 0.0.0.0/0 if it's a persistent public IP. Also open all protocols and ports for ingress traffic from the AWS VPC CIDR (10.0.0.0/16). Remember: [Security lists](#) are associated with a subnet, so edit the security list associated with each subnet that needs to communicate with the AWS VPC. Or, if you're using VCN [network security groups](#), edit the rules in the relevant NSGs.

**Ingress Rules**

<input type="checkbox"/>	Stateless ▾	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
<input type="checkbox"/>	No	10.0.0.0/16	All Protocols				All traffic for all ports
<input type="checkbox"/>	No	0.0.0.0/0	TCP	All	4500		TCP traffic for ports: 4500
<input type="checkbox"/>	No	0.0.0.0/0	TCP	All	500		TCP traffic for ports: 500
<input type="checkbox"/>	No	0.0.0.0/0	UDP	All	4500		UDP traffic for ports: 4500
<input type="checkbox"/>	No	0.0.0.0/0	UDP	All	500		UDP traffic for ports: 500

0 Selected Showing 5 Items < Page 1 >

- In the Oracle Console, edit the VCN's route tables to add a rule that has the AWS VPC CIDR (10.0.0.0/16) as the destination CIDR block and the DRG you created earlier as the target. Remember: Route tables are associated with a subnet, so edit the route table associated with each subnet that needs to communicate with the AWS VPC. The following screenshot shows the route table for the VCN with an added route for the AWS VPC CIDR.

Route Rules		
Destination	Target Type	Target
<input type="checkbox"/> 10.0.0.0/16	Dynamic Routing Gateways	<a href="#">AWS-DRG</a>

0 Selected Showing 1 Item < Page 1 >

### Task 3: Determine the required configuration values

The Libreswan configuration uses the following variables. Determine the values before proceeding with the configuration.

- `${cpeLocalIP}`: The IP address of your Libreswan device.
- `${cpePublicIpAddress}`: The public IP address for Libreswan. This is the IP address of your outside interface. Depending on your network topology, the value might be different from `${cpeLocalIP}`.
- `${oracleHeadend1}`: For the first tunnel, the Oracle public IP endpoint obtained from the Oracle Console.
- `${oracleHeadend2}`: For the second tunnel, the Oracle public IP endpoint obtained from the Oracle Console.
- `${vti1}`: The name of the first VTI used. For example, vti1.
- `${vti2}`: The name of the second VTI used. For example, vti2.

- `${sharedSecret1}`: The pre-shared key for the first tunnel. You can use the default Oracle-provided pre-shared key, or provide your own when you set up the IPSec connection in the Oracle Console.
- `${sharedSecret2}`: The pre-shared key for the second tunnel. You can use the default Oracle-provided pre-shared key, or provide your own when you set up the IPSec connection in the Oracle Console.
- `${vcnCidrNetwork}`: The VCN IP range.

### Task 4: Set up the configuration file: `/etc/ipsec.d/oci-ipsec.conf`

Libreswan configuration uses the concept of *left* and *right* to define the configuration parameters for your local CPE device and the remote gateway. Either side of the connection (the *conn* in the Libreswan configuration) can be left or right, but the configuration for that connection must be consistent. In this example:

- **Left:** Your local Libreswan CPE
- **Right:** The Oracle VPN headend

Use the following template for your `/etc/ipsec.d/oci-ipsec.conf` file. The file defines the two tunnels that Oracle creates when you set up the IPSec connection.



#### Important

If your CPE is behind a 1-1 NAT device, uncomment the `leftid` parameter and set it equal to the `${cpePublicIpAddress}`.

```
conn oracle-tunnel-1
left=${cpeLocalIP}
leftid=${cpePublicIpAddress} # See preceding note about 1-1 NAT device
right=${oracleHeadend1}
authby=secret
leftsubnet=0.0.0.0/0
rightsubnet=0.0.0.0/0
auto=start
mark=5/0xffffffff # Needs to be unique across all tunnels
```

## CHAPTER 23 Networking

```
vti-interface=${vti1}
vti-routing=no
ikev2=no # To use IKEv2, change to ikev2=insist
ike=aes_cbc256-sha2_384;modp1536
phase2alg=aes_gcm256;modp1536
encapsulation=no
ikelifetime=28800s
salifetime=3600s
conn oracle-tunnel-2
left=${cpeLocalIP}
leftid=${cpePublicIpAddress} # See preceding note about 1-1 NAT device
right=${oracleHeadend2}
authby=secret
leftsubnet=0.0.0.0/0
rightsubnet=0.0.0.0/0
auto=start
mark=6/0xffffffff # Needs to be unique across all tunnels
vti-interface=${vti2}
vti-routing=no
ikev2=no # To use IKEv2, change to ikev2=insist
ike=aes_cbc256-sha2_384;modp1536
phase2alg=aes_gcm256;modp1536
encapsulation=no
ikelifetime=28800s
salifetime=3600s
```

### Task 5: Set up the secrets file: `/etc/ipsec.d/oci-ipsec.secrets`

Use the following template for your `/etc/ipsec.d/oci-ipsec.secrets` file. It contains two lines per IPSec connection (one line per tunnel).

```
${cpePublicIpAddress} ${ipAddress1}: PSK "${sharedSecret1}"
${cpePublicIpAddress} ${ipAddress2}: PSK "${sharedSecret2}"
```

### Task 6: Restart the Libreswan configuration

After setting up your configuration and secrets files, you must restart the Libreswan service with the following command.



#### Important

Restarting the Libreswan service may impact existing tunnels.

```
service ipsec restart
```

### Task 7: Configure IP routing

Use the following `ip` command to create static routes that send traffic to your VCN through the IPsec tunnels. If you're logged in with an unprivileged user account, you might need to use `sudo` before the command.



#### Important

Static routes created with the `ip route` command do not persist through a reboot. To determine how to make your routes persist, refer to the documentation of your Linux distribution of choice.

```
ip route add ${VcnCidrBlock} nexthop dev ${vti1} nexthop dev ${vti2}
ip route show
```

### Verification

A [Monitoring service](#) is also available from Oracle Cloud Infrastructure to actively and passively monitor your cloud resources. For information about monitoring your VPN Connect, see [VPN Connect Metrics](#).

If you have issues, see [VPN Connect Troubleshooting](#).

#### Checking the Libreswan Status

Check the current state of your Libreswan tunnels by using the following command:

```
ipsec status
```

The tunnel is established if you see a line that includes the following:

```
STATE_MAIN_I4: ISAKMP SA established
```

If you're using IKEv2, you see the following:

```
STATE_V2_IPSEC_I (IPsec SA established)
```

In the future, if you need to open a support ticket with Oracle about your Libreswan tunnel, include the output of the preceding `ipsec status` command.

### Checking the Tunnel Interface Status

Check if the virtual tunnel interfaces are up or down by using the `ifconfig` command or the `ip link show` command. You can also use applications such as `tcpdump` with the interfaces.

Here's an example of the `ifconfig` output with a working Libreswan implementation that shows the available VTIs.

```
ifconfig
<output trimmed>

vti01: flags=209<UP,POINTOPOINT,RUNNING,NOARP> mtu 8980
 inet6 fe80::5efe:a00:2 prefixlen 64 scopeid 0x20<link>
 tunnel txqueuelen 1000 (IPIP Tunnel)
 RX packets 0 bytes 0 (0.0 B)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 0 bytes 0 (0.0 B)
 TX errors 10 dropped 0 overruns 0 carrier 10 collisions 0

vti02: flags=209<UP,POINTOPOINT,RUNNING,NOARP> mtu 8980
 inet6 fe80::5efe:a00:2 prefixlen 64 scopeid 0x20<link>
 tunnel txqueuelen 1000 (IPIP Tunnel)
 RX packets 0 bytes 0 (0.0 B)
 RX errors 0 dropped 0 overruns 0 frame 0
 TX packets 0 bytes 0 (0.0 B)
 TX errors 40 dropped 0 overruns 0 carrier 40 collisions 0
```

Here's an example of the `ip link show` output:

```
ip link show
<output trimmed>

9: vti01@NONE: <POINTOPOINT,NOARP,UP,LOWER_UP> mtu 8980 qdisc noqueue
state UNKNOWN mode DEFAULT group default qlen 1000
 link/ipip 10.1.2.3 peer 129.146.12.51

10: vti02@NONE: <POINTOPOINT,NOARP,UP,LOWER_UP> mtu 8980 qdisc noqueue
state UNKNOWN mode DEFAULT group default qlen 1000
 link/ipip 10.1.2.3 peer 129.146.13.49
```

Also, in the Oracle Console, each IPSec tunnel should now be in the UP state.

## Network Performance

The content in the sections below apply to **Category 7** and **Section 3.c** of the Oracle PaaS and IaaS Public Cloud Services Pillar documentation, which you can download in PDF format from the [Oracle Cloud Infrastructure Service Level Agreement page](#).

Oracle Cloud Infrastructure provides a service-level agreement (SLA) for network throughput between instances in the same availability domain in a virtual cloud network (VCN). You might think of this as a measurement of LAN performance.



### Important

This SLA applies only to bare metal instances.

To meet the SLA, the network throughput for instances within the same availability domain and VCN must be at least 90% of the stated maximum for at least 99.9% of the billing month. Network throughput is measured in megabits per second (Mbps) or gigabits per second (Gbps).

For the stated maximum bandwidth by instance shape, see the "Network Bandwidth" column in [the "Shape" tables](#).

## Testing Methodology

**Summary:** Launch two bare metal instances in the same availability domain and VCN. Install and run the [iperf3](#) utility, with one instance as server and the other as client. Look at the `iperf3` bandwidth results to determine your VCN's network throughput.

### Instructions:

1. Launch two bare metal instances in the same availability domain in a single VCN. Designate one as the server and the other as the client. For launch instructions, see

### [Creating an Instance.](#)

2. Install `iperf3` on both instances. Example Linux command:

```
sudo yum install -y iperf3
```

3. Enable communication to the server instance on TCP port 5201 (for `iperf3`):
  - a. For the subnet that the server instance is in, add a rule to the subnet's security list to allow stateless ingress traffic on TCP port 5201 from any source IP address (0.0.0.0/0) and any source port. For instructions, see [To update rules in an existing security list](#). If you are instead using [network security groups \(NSGs\)](#) with the instance, add the rule to the instance's NSG.
  - b. On the instance itself, open the firewall to allow `iperf3` traffic. Example Linux commands:



#### Warning

For instances with an iSCSI boot volume, the following `--reload` command can cause problems. For details and a workaround, see [Instances experience system hang after running firewall-cmd --reload](#).

```
sudo firewall-cmd --zone=public --permanent --add-port 5201/tcp
sudo firewall-cmd --reload
```

4. Start the `iperf3` test:
  - a. On the server instance, run `iperf3` in server mode. Example Linux command:

```
iperf3 -s
```

- b. On the client instance, run `iperf3` in client mode and specify the private IP address of the server instance. Example Linux command:

```
iperf3 -c <server_instance_private_ip_address>
```

5. Look at the `iperf3` results on the client instance. The network throughput between the two instances is shown under "Bandwidth" in the last five lines of the client's `iperf3` test output. For example:

```

[ID] Interval Transfer Bandwidth Retr
[4] 0.00-10.00 sec XX.YY GBytes NN.NN Gbits/sec 752
[4] 0.00-10.00 sec XX.YY GBytes NN.NN Gbits/sec

iperf Done.
```

## Frequently Asked Questions

**Q:** My VCN isn't meeting the bandwidth SLA. What should I do?

**A:** Make sure that the CPU on the instance isn't loaded heavily with other services or applications. Confirm this with a utility such as `top` to look at the average CPU utilization. It should be less than one.

## Troubleshooting

These topics cover some common issues you might run into and how to address them:

- [Hanging Connection](#)
- [Subnet or VCN Deletion](#)
- [VPN Connect Troubleshooting](#)
- [FastConnect Troubleshooting](#)

### Hanging Connection

This topic covers one of the most common issues seen with communications between your cloud network and on-premises network: a hanging connection, even though you can ping hosts across the connection.

### Summary of Problem and Solutions

**Symptom:** Your virtual cloud network (VCN) is connected to your existing on-premises network via an IPsec VPN, or Oracle Cloud Infrastructure FastConnect. Hosts on one side of the connection can ping hosts on the other side, but the connection hangs. For example:

- You can SSH to a host across the connection, but after you log in to the host, the connection hangs.
- You can start a Virtual Networking Computing (VNC) connection, but the session hangs.
- You can start an SFTP download, but the download hangs.

**General problem:** *Path Maximum Transmission Unit Discovery (PMTUD)* is probably not working on one or both sides of the connection. It must be working on both sides of the connection so that both sides can know if they're trying to send packets that are too large for the connection and adjust accordingly. For a brief overview of Maximum Transmission Unit (MTU) and PMTUD, see [Overview of MTU](#) and [Overview of PMTUD](#).

### Solutions for fixing PMTUD:

1. **Make sure your hosts are configured to use PMTUD:** If the hosts in your on-premises network don't use PMTUD (that is, if they don't set the Don't Fragment flag in the packets), they have no way to discover if they're sending packets that are too large for the connection. Your instances on the Oracle side of the connection use PMTUD by default. Do not change that configuration on the instances.
2. **Make sure both the VCN security lists and the instance firewalls allow ICMP type 3 code 4 messages:** When PMTUD is in use, the sending hosts receive a special ICMP message if they send packets that are too large for the connection. Upon receipt of the message, the host can dynamically update the size of the packets to fit the connection. However, your instances can't receive these important ICMP messages if the security lists for the subnet in the VCN and/or the instance firewalls aren't configured to accept them.



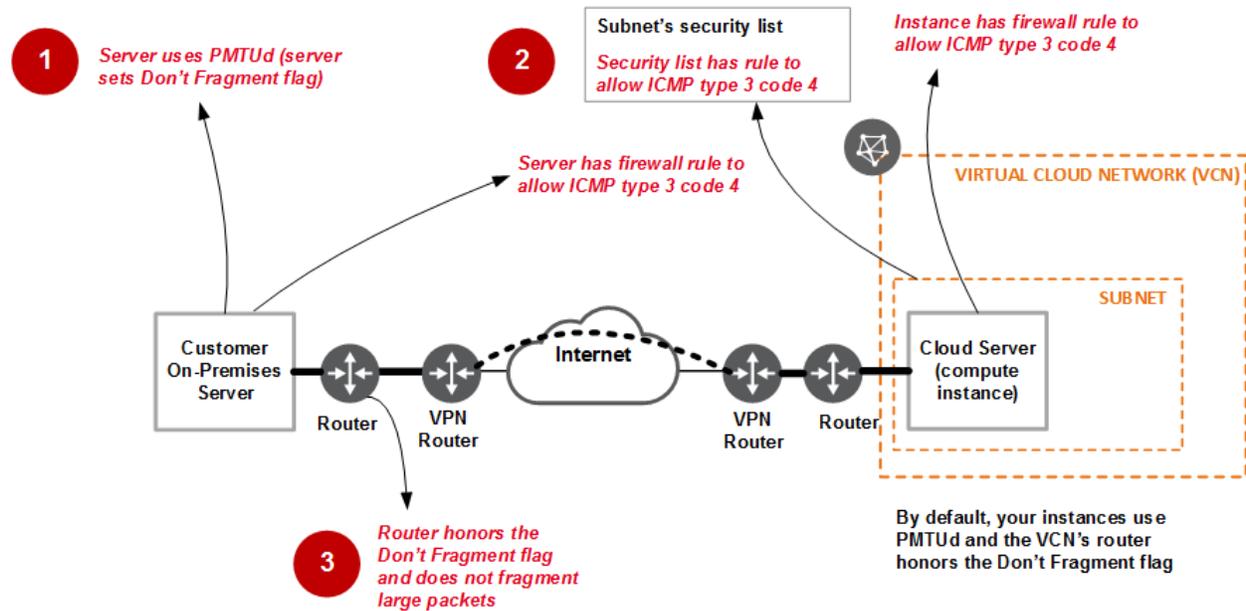
### Tip

If you're using [stateful security list rules](#) (for TCP, UDP, or ICMP traffic), you don't need to ensure that your security list has a rule to allow ICMP type 3 code 4 messages. With stateful rules, the Networking service tracks the connections and automatically allows corresponding ICMP type 3 code 4 messages without needing an explicit rule to allow them. Only if you're using stateless rules must you have an explicit ingress security list rule for ICMP type 3 code 4 messages. You must also confirm the instance firewalls are set up correctly.

To check to see if a host is receiving the messages, see [Finding Where PMTUD Is Broken](#).

3. **Make sure your router honors the *Don't Fragment* flag:** If the router doesn't honor the flag and thus ignores the use of PMTUD, it sends fragmented packets to the instances in the VCN, which is bad (see [Why Avoid Fragmentation?](#)). The VCN's security lists are most likely configured in such a way that they recognize only the initial fragment, and the remaining ones are dropped, causing the connection to hang. Instead, your router should use PMTUD and honor the Don't Fragment flag to determine the correct size of unfragmented packets to send through the connection.

The parts of the solution are numbered and called out in red italics in the following diagram. It shows an example scenario with your on-premises network connected to your VCN over an IPSec VPN.



Keep reading for a brief overview of MTU and PMTUD, and [how to check if PMTUD is working](#) on both sides of the network connection.

### Why Avoid Fragmentation?

You may be wondering why you want to avoid fragmentation. First, it adversely affects the performance of your application. Fragmentation requires reassembly of the fragments and retransmission if fragments are lost. Reassembly and retransmission require time and CPU resources.

Second, only the first fragment contains the source and destination port information. This means the other packets will probably be dropped by firewalls or your VCN's [security lists](#), which are typically configured to evaluate the port information. For fragmentation to work

with your firewalls and security lists, you would have to configure them to be more permissive than usual, which is not desirable.

### Overview of MTU

The communications between any two hosts across an Internet Protocol (IP) network use packets. Each packet has a source and destination IP address and a payload of data. Every network segment between the two hosts has a *Maximum Transmission Unit (MTU)* that represents the number of bytes that a single packet can carry.

Across the internet, the MTU is 1500 bytes. This is also true for most home networks and many corporate networks (and their Wi-Fi networks). Some data centers, including those for Oracle Cloud Infrastructure, have a larger MTU. The Compute instances use an MTU of 9000 by default. On a Linux host, you can use the `ifconfig` command to display the MTU of the host's network connection. For example, here's the `ifconfig` output from an Ubuntu instance (the MTU is highlighted in red italics):

```
ifconfig
ens3 Link encap:Ethernet HWaddr 00:00:17:01:17:83
inet addr:10.0.6.9 Bcast:10.0.6.31 Mask:255.255.255.224
inet6 addr: fe80::200:17ff:fe01:1783/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:9000 Metric:1
```

For comparison, here's the output from a machine connected to a corporate network:

```
ifconfig
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST>
mtu 1500
```

Notice that its MTU is the more typical 1500 bytes.

If the host is connected through a corporate VPN, the MTU is even smaller, because the VPN tunnel must encapsulate the traffic inside an IPsec packet and send it across the local network. For example:

```
ifconfig
utun0: flags=81d1<UP,POINTOPOINT,RUNNING,NOARP,PROMISC,MULTICAST>
mtu 1300
```

How do the two hosts figure out how large of a packet they can send to each other? For many types of network traffic, such as HTTP, SSH, and FTP, the hosts use the Transmission Control

## CHAPTER 23 Networking

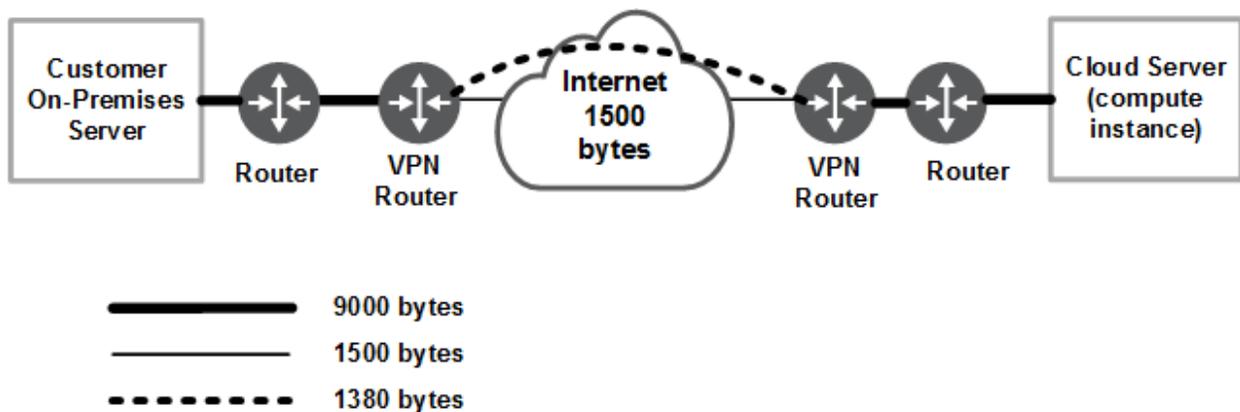
Protocol (TCP) to establish new connections. During the initial three-way handshake between two hosts, they each send the *Maximum Segment Size (MSS)* for how large their payload can be. This is smaller than the MTU. (TCP runs inside the Internet Protocol (IP), which is why it's referred to as TCP/IP. Segments are to TCP what packets are to IP.)

Using the [tcpdump](#) application, you can see the MSS value shared during the handshake. Here's an example from tcpdump (with the MSS highlighted in red italics):

```
12:11:58.846890 IP 129.146.27.25.22 > 10.197.176.19.58824: Flags [S.], seq
2799552952, ack 2580095593, win 26844, options [mss 1260, sackOK, TS val
44858491 ecr 1321638674, nop, wscale 7], length 0
```

The preceding packet is from an SSH connection to an instance from a laptop connected to a corporate VPN. The local network the laptop uses for its internet connection has an MTU of 1500 bytes. The VPN tunnel enforces an MTU of 1300 bytes. Then when the SSH connection is attempted, TCP (running inside the IP connection) tells the Oracle Cloud Infrastructure instance that it supports TCP segments that are less than or equal to 1260 bytes. With a corporate VPN connection, the laptop connected to the VPN typically has the smallest MTU and MSS compared to anything it's communicating with across the internet.

A more complex case is when the two hosts have a larger MTU than some network link between them *that is not directly connected to either of them*. The following diagram illustrates an example.



In this example, there are two servers, each directly connected to its own routed network that supports a 9000-byte MTU. The servers are in different data centers. Each data center is connected to the internet, which supports a 1500-byte MTU. There is an IPsec VPN tunnel between the two data centers. That tunnel crosses the internet, so the inside of the tunnel has a smaller MTU than the internet. In this diagram, the MTU is 1380 bytes.

If the two servers try to communicate (with SSH, for example), during the three-way handshake, they agree on an MSS around 8960. The initial SSH connection might succeed, because the maximum packet sizes during the initial SSH connection setup are usually less than 1380 bytes. When one side tries to send a packet larger than the smallest link between the two endpoints, Path MTU Discovery (PMTUD) becomes critical.

### Overview of PMTUD

Path MTU Discovery is defined in [RFC 1191](#). It works by requiring the two communicating hosts to set a *Don't Fragment* flag in the packets they each send. If a packet from one of these hosts reaches a router where the egress (or outbound) interface has an MTU smaller than the packet length, the router drops that packet. The router also returns an ICMP type 3 code 4 message to the host. This message specifically says "Destination Unreachable, Fragmentation Needed and Don't Fragment Was Set" (defined in [RFC 792](#)). Effectively the router tells the host: "You told me not to fragment packets that are too large, and this one's too large. I'm not sending it." The router also tells the host the maximum size packets allowed through that egress interface. The sending host then adjusts the size of its outbound packets so they're smaller than the value the router provided in the message.

Here's an example that shows the results when an instance tries to ping a host (8.8.8.8) over the internet with an 8000-byte packet and the Don't Fragment flag set (that is, with PMTUD in use). The returned ICMP message is highlighted in red italics:

```
ping 8.8.8.8 -M do -s 8000
PING 8.8.8.8 (8.8.8.8) 8000(8028) bytes of data.
From 4.16.139.250 icmp_seq=1 Frag needed and DF set (mtu = 1500)
```

The response is exactly what's expected. The destination host is across the internet, which has an MTU of 1500 bytes. Even though the sending host's local network connection has an MTU of 9000 bytes, the host can't reach the destination host with the 8000-byte packet and gets an ICMP message accordingly. PMTUD is working correctly.

For comparison, here's the same ping, but the destination host is across an IPsec VPN tunnel:

```
ping 192.168.6.130 -M do -s 8000
PING 192.168.6.130 (192.168.6.130) 8000(8028) bytes of data.
From 129.146.13.49 icmp_seq=1 Frag needed and DF set
```

Here the VPN router sees that to send this packet to its destination, the outbound interface is a VPN tunnel. That tunnel goes across the internet, so the tunnel must fit inside the internet's 1500-byte MTU link. The result is that the inside of the tunnel only allows packets up to 1360 bytes (which the router then lowered to 1358, which can make things more confusing).

### Finding Where PMTUD Is Broken

If PMTUD isn't working somewhere along the connection, you need to figure out why and where. Typically it's because the ICMP type 3 code 4 packet (from the router with the constrained link that can't fit the packet) never gets back to the sending host. This can happen if there's something blocking that kind of traffic between the host and the router. And it can happen on either side of the VPN tunnel (or other constrained MTU link).

#### TRY PINGING FROM EACH SIDE OF THE CONNECTION

To troubleshoot the broken PMTUD, you must determine if PMTUD is working on each side of the connection. In this scenario, let's assume the connection is an IPsec VPN.

**How to ping:** Like in [Overview of PMTUD](#), ping a host on the other side of the connection with a packet that you know is too large to fit through the VPN tunnel (for example, 1500 bytes or larger). Depending on which operating system the sending host uses, you might need to format the ping command slightly different to ensure the Don't Fragment flag is set. For both Ubuntu and Oracle Linux, you use the `-M` flag with the ping command.

Here's information about the `-M` flag:

```
-M pmtudisc_opt
Select Path MTU Discovery strategy. pmtudisc_option may be either do
(prohibit fragmentation, even local one), want (do PMTU discovery, fragment
locally when packet size is large), or dont (do not set DF flag).
```

Here's an example ping (with the -M flag and the resulting ICMP message highlighted in red italics)

```
ping -M do -s 1500 192.168.6.130
PING 192.168.6.130 (192.168.6.130) 1500(1528) bytes of data.
From 129.146.13.49 icmp_seq=1 Frag needed and DF set (mtu = 1358)
```

### Good: PMTUD is working

If the result includes the line "From x.x.x.x icmp\_seq=1 Frag needed and DF set (mtu = xxxx)", then PMTUD is working on that side of the tunnel. Note that the source address of the ICMP message is the public IP address of the tunnel the traffic is trying to go out (for example 129.146.13.49 in the preceding Ubuntu example).

Make sure to also ping from the other side of the connection to confirm PMTUD is working from that side. Both sides of the connection must recognize that there is a tunnel between them that can't fit the large packets.

### Bad: If you're testing your side of the connection and the ping succeeds

If you're sending the ping from a host in your on-premises network, and the ping succeeds, that probably means your edge router is not honoring the Don't Fragment flag. Instead the router is fragmenting the large packet. The first fragment reaches the destination host, so the ping succeeds, which is misleading. If you try to do more than just ping, the fragments after the first get dropped, and the connection will hang.

**Make sure to configure your router to honor the Don't Fragment flag.** The router's default configuration is to honor it, but someone might have changed the default.

### Bad: If you're testing the VCN side of the connection and you don't see the ICMP message

When testing from the VCN side of the connection, if you don't see the ICMP message in the response, there is probably something dropping the ICMP packet before it reaches your instance.

There could be two issues:

- **Security list:** The Networking [security list](#) could be missing an ingress rule that allows ICMP type 3 code 4 messages to reach the instance. This is an issue only if you're using [stateless security list rules](#). If you're using stateful rules, your connections are tracked and the ICMP message is automatically allowed without needing a specific security list rule to allow it. **If you're using stateless rules, make sure that the subnet the instance is in has a security list with an ingress rule that allows ICMP traffic type 3 code 4 from source 0.0.0.0/0 and any source port.** For more information, see [Security Lists](#), and specifically [To update rules in an existing security list](#).
- **Instance firewall:** The instance's firewall rules (set in the OS) could be missing a rule that allows ICMP type 3 code 4 messages to reach the instance. Specifically for a Linux instance, make sure that iptables or firewalld is configured to allow the ICMP type 3 code 4 messages.

### Avoiding the Need for PMTUD

Oracle recommends using PMTUD. However, in some situations it's possible to configure servers so they don't need to rely on it. Consider the case of the instances in your VCN communicating across an IPSec VPN to hosts in your on-premises network. You know the range of IP addresses for your on-premises network. You can add a special route to your instances that specifies the maximum MTU to use when communicating with hosts in that address range. The instance-to-instance communication within the VCN still uses an MTU of 9000 bytes.

The following information shows how to set that route on a Linux instance.

The default route table on the instance typically has two routes: the default route (for the default gateway), and a local route (for the local subnet). For example:

```
ip route show
default via 10.0.6.1 dev ens3
10.0.6.0/27 dev ens3 proto kernel scope link src 10.0.6.9
```

You can add another route that points to the same default gateway, but with the address range of the on-premises network and a smaller MTU. For example, in the following command, the

## CHAPTER 23 Networking

---

on-premises network is 1.0.0.0/8, the default gateway is 10.0.6.1, and the maximum MTU size is 1300 for packets being sent to the on-premises network.

```
ip route add 1.0.0.0/8 via 10.0.6.1 mtu 1300
```

The updated route table looks like this:

```
ip route show
default via 10.0.6.1 dev ens3
1.0.0.0/8 via 10.0.6.1 dev ens3 mtu 1300
10.0.6.0/27 dev ens3 proto kernel scope link src 10.0.6.9
```

Within the VCN, the instance-to-instance communication continues to use 9000 MTU. However, communication to the on-premises network uses a maximum of 1300. This example assumes there's no part of the connection between the on-premises network and VCN that uses an MTU smaller than 1300.



### Important

The preceding commands do not persist if you reboot the instance. You can make the route permanent by adding it to a configuration file in the OS. For example, for Oracle Linux, it's an interface-specific file called `/etc/sysconfig/network-scripts/route-<interface>`. For more information, see the documentation for your variant of Linux.

## Subnet or VCN Deletion

This topic covers reasons why deletion of a subnet or VCN might fail.

Remember:

- To delete a VCN, it must first be empty and have no related resources or attached gateways (for example: no [internet gateway](#), [dynamic routing gateway](#), and so on).
- To delete a VCN's subnets, they must first be empty.

### The Subnet Isn't Empty

The most common reason a subnet (and thus a VCN) can't be deleted is because the subnet contains one or more of these resources:

- [Load balancer](#)
- [Mount target](#)
- [DB system](#)



#### Note

When you create one of the preceding resources, you specify a VCN and subnet for it. The relevant service creates at least one [VNIC](#) in the subnet and attaches the VNIC to the resource. The service manages the VNICs on your behalf, so they are not readily apparent to you in the Console. The VNIC enables the resource to communicate with other resources over the network. Although this documentation commonly talks about the resource itself being in the subnet, it's actually the resource's attached VNIC. This documentation uses the term *parent resource* to refer to this type of resource.

If the subnet *is* empty when you try to delete it, its state changes to TERMINATING briefly and then to TERMINATED.

If the subnet is not empty, you instead get an error indicating that there are still resources that you must delete first. The error includes the OCID of a VNIC that is in the subnet (there could be more, but the error returns only a single VNIC's OCID).

You can use the [Oracle Cloud Infrastructure command line interface \(CLI\)](#) or another SDK or client to call the `GetVnic` operation with the VNIC OCID. The response includes the VNIC's *display name*. Depending on the type of parent resource, the display name can indicate which parent resource the VNIC belongs to. You can then delete that parent resource, or you can contact your administrator to determine who owns the resource. When the VNIC's parent

## CHAPTER 23 Networking

---

resource is deleted, the attached VNIC is also deleted from the subnet. If there are remaining VNICs in the subnet, repeat the process of determining and deleting each parent resource until the subnet is empty. Then you can delete the subnet.

For example, if you're using the CLI, use this command to get information about the VNIC.

```
oci network vnic get --vnic_id <VNIC_OCID>
```

### Load balancer example

Here is an example CLI response for a VNIC that belongs to a load balancer. The display name shows the load balancer's OCID:

```
{
 "data": {
 "availability-domain": "fooD:PHX-AD-1",
 "compartment-id": "ocidl.compartment.oc1..<unique_id_1>",
 "defined-tags": {},
 "display-name": "VNIC for LB ocidl.loadbalancer.oc1.phx.<unique_id_2>",
 "freeform-tags": {},
 "hostname-label": null,
 "id": "ocidl.vnic.oc1.phx.<unique_id_3>",
 "is-primary": false,
 "lifecycle-state": "AVAILABLE",
 "mac-address": "00:00:17:00:BB:CA",
 "private-ip": "10.0.0.6",
 "public-ip": null,
 "skip-source-dest-check": false,
 "subnet-id": "ocidl.subnet.oc1.phx.<unique_id_4>",
 "time-created": "2019-05-11T04:28:31.950000+00:00"
 },
 "etag": "5d8213fa"
}
```

### File Storage example

Here's an example for a VNIC that belongs to a File Storage mount target:

```
"display-name": "fss-<integer>",
```

Although the display name does not include an OCID, the `fss` characters indicate that the resource is for the File Storage service.

### Database example

Here's an example of the display name for a VNIC that belongs to a DB system:

```
"display-name": "ocidl.dbnode.oc1.phx.<unique_id>",
```

### A Network Security Group Isn't Empty

Another reason a VCN can't be deleted is because it contains a one or more [network security groups](#) (NSGs) that are not yet empty. To delete an NSG, it must not contain any VNICs (or parent resources with VNICs). You can determine what parent resources are in an NSG by using either the Console or REST API. For more information, see [Deleting NSGs](#).

### There Are Resources in Compartments You Don't Have Access To

You might not be able to see all the resources in a subnet or VCN. This is because subnets and VCNs can contain resources in multiple compartments, and you might not have access to all the compartments. For example, the subnet might contain instances that your team manages but also DB systems that another team manages. Another example: The VCN might have security lists or a gateway in a compartment that another team manages. You might need to contact your tenancy administrator to help you determine who owns the resources in the subnet or VCN.

# CHAPTER 24 Notifications

This chapter explains how to use the Notifications service.

## Notifications Overview

The Oracle Cloud Infrastructure Notifications service broadcasts messages to distributed components through a publish-subscribe pattern, delivering secure, highly reliable, low latency and durable messages for applications hosted on Oracle Cloud Infrastructure and externally. Use Notifications to get notified when event rules are triggered or alarms are breached, or to directly publish a message.



### Note

Notifications is not available in Oracle Cloud Infrastructure Government Cloud realms.

## How Notifications Works

The Notifications service enables you to set up communication channels for publishing messages using topics and subscriptions. When a message is published to a topic, the Notifications service sends the message to all of the topic's subscriptions.

When a subscriber's endpoint does not acknowledge receipt of the message, the Notifications service retries delivery. This situation can occur when the endpoint is offline. For example, the email server for an email address may be down.

### Delivery retry details

Notifications retries delivery following these steps until either (a) acknowledgement is received or (b) the subscription's retry duration is over. By default, the retry duration is two

hours.

1. Immediate retry.
2. Exponential backoff retry for the period of the subscription's retry duration, using the following timing:
  - a. 1 minute
  - b. 2 minutes
  - c. 4 minutes
  - d. 8 minutes
  - e. 16 minutes
  - f. 32 minutes
3. Discarding of the message at the end of the retry duration.

You can change the retry duration for a subscription. For instructions using the Console, see [To update the retry duration for a subscription](#). For the API, use the following operation: [UpdateSubscription](#).

### Notifications Concepts

The following concepts are essential to working with Notifications.

#### MESSAGE

The content that is published to a *topic*. Each message is delivered at least once per *subscription*. Every message sent out as email contains a link to *unsubscribe* from the related topic.

#### SUBSCRIPTION

An endpoint for a *topic*. Published *messages* are sent to each subscription for a *topic*. For supported subscription protocols, see [To create a subscription](#).

#### TOPIC

A communication channel for sending *messages* to the *subscriptions* in the topic. Each topic name is unique across the tenancy.



#### Note

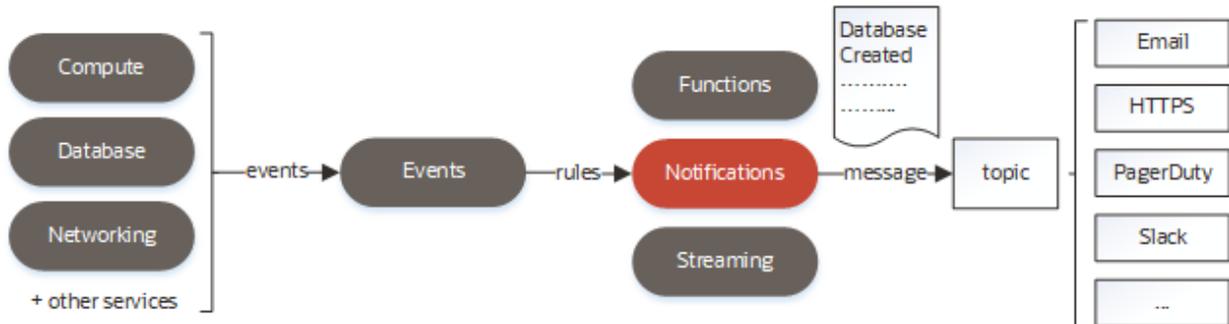
Messages sent out as email by the Oracle Cloud Infrastructure Notifications service are processed and delivered through Oracle resources in U.S.-based regions.

### Flow of Message Publication

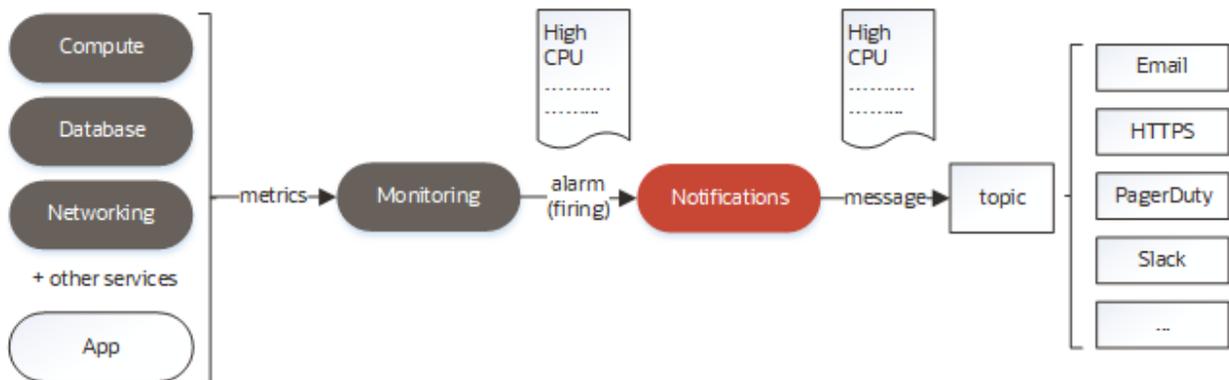
Notifications publishes messages when event rules are triggered, alarms are breached, or someone directly publishes a message.

Event rules: Notifications sends messages when [rules](#) are triggered. The message is sent to the topic specified in the rule. For example, a message might be configured for new databases. See [Managing Rules for Events](#).

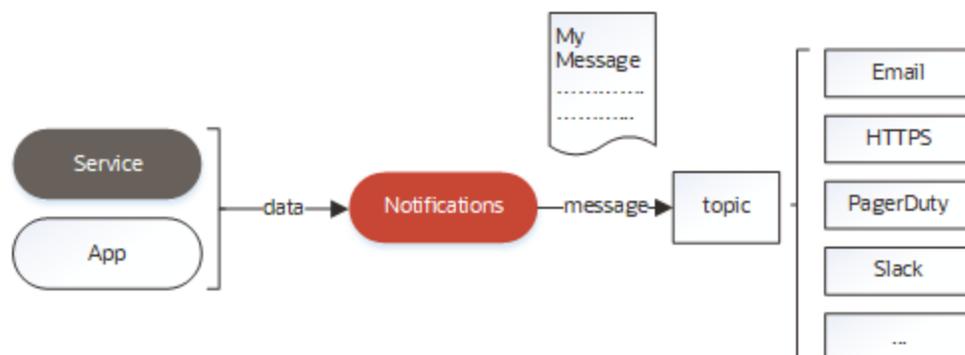
## CHAPTER 24 Notifications



Alarms: Notifications sends [alarm messages](#) when alarms are breached. The alarm message is sent to the topic specified in the alarm. For example, an alarm message might be configured for high CPU usage. See [Managing Alarms](#).



Direct publication: Notifications sends messages when you (or a service or app) publish the messages directly. The message is sent to the topic you specify. See [Publishing Messages](#).



## Availability

Notifications is currently available in the following regions:

Region Name	Region Location	Region Key
India West (Mumbai)	Asia-Pacific: Mumbai, India	BOM
South Korea Central (Seoul)	Asia-Pacific: Seoul, South Korea	ICN
Australia East (Sydney)	Asia-Pacific: Sidney, Australia	SYD
Japan East (Tokyo)	Asia-Pacific: Tokyo, Japan	NRT
Canada Southeast (Toronto)	Canada: Toronto	YYZ
Germany Central (Frankfurt)	Europe: Frankfurt, Germany	FRA
Switzerland North (Zurich)	Europe: Zurich, Switzerland	ZRH
Brazil East (Sao Paulo)	South America: Sao Paulo	GRU
UK South (London)	United Kingdom: London	LHR
US East (Ashburn)	United States: Ashburn, VA	IAD
US West (Phoenix)	United States: Phoenix, AZ	PHX

## Service Comparison for Sending Email Messages

Consider the following service features when deciding whether to use the Notifications service or the Email Delivery service to send your email messages. For more information about Email Delivery, see [Overview of the Email Delivery Service](#).

<b>Service Feature</b>	<b>Notifications service</b>	<b>Email Delivery service</b>
Requires confirmation before sending email.	Yes	No
Allows email decorations, such as signatures.	Yes	No
Allows raw email messages.	No	Yes
Supports MIME attachments.	No	Yes
Supports special handling for failed email delivery.	No	Yes
Priced for small messages (less than 32 KB, with a 64-KB limit).	Yes	No
Priced for large messages (greater than 32 KB, with a 2-MB limit).	No	Yes

## Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Moving Topics and Subscriptions to a Different Compartment

You can [move topics](#) and [subscriptions](#) from one compartment to another. When you move a topic to a new compartment, its associated subscriptions remain in their existing compartment. The same consideration applies when moving a subscription: its associated topic remains in its existing compartment.

After you move the topic or subscription to the new compartment, inherent policies apply immediately and affect access to the moved topic or subscription through the Console. For more information, see [Moving Resources to a Different Compartment](#).



#### Important

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Notifications resources, see [Details for the Notifications Service](#).

### Ways to Access Notifications

You can access the Notifications service using the Console (a browser-based interface) or the REST API. Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

Console: To access Notifications using the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.

API: To access Notifications through API, use [Notifications API](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Administrators: For common policies that give groups access to Notifications, see [Allow a group to manage topics](#), [Allow a group to manage topic subscriptions](#), and [Allow a group to publish messages to topics](#).

### Limits on Notifications

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

### Managing Topics and Subscriptions

This section describes how to manage topics and their subscriptions.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

A topic is a communication channel for sending messages to its subscriptions. A topic can have zero, one, or multiple subscriptions that are notified whenever a message is published to a topic.

## Prerequisites

**IAM policies:** To use Notifications, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a response that you don't have permission or are unauthorized, check with your administrator. You may not have the required type of access in the current compartment. For more information on user authorizations, see [Notifications Overview](#).

## Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

## Using the Console

### To create a topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.

2. Click **Create Topic** at the top of the topic list.
3. In the **Create Topic** dialog box, configure your topic.
  - **Name:** Required. Specify a friendly name for the topic. It must be unique across the tenancy; validation is case-sensitive. Avoid entering confidential information.
  - **Description:** Optional. Enter a description for the topic. Avoid entering confidential information.
4. Click **Create**.

### To delete a topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. For the topic you want to delete, click the Actions icon (three dots), and then click **Delete**.
3. Confirm when prompted.

### To update the description for a topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Click the name of the topic you want to update.
3. On the topic detail page, next to **Description**, click the edit icon.
4. Edit the description.
5. Click the save icon.

### To move a topic to a different compartment

Associated subscriptions remain in their current compartments. For more information, see [Moving Topics and Subscriptions to a Different Compartment](#).



### Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Notifications resources, see [Details for the Notifications Service](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. In the **Scope** section, select a compartment.
3. Find the topic in the list, click the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### To create a subscription



### Note

While new subscriptions must be created in the same compartment as the topic, you can [move](#) them to different compartments after creating them.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Click the name of the topic that you want to add the subscription to.

3. On the topic detail page, click **Create Subscription**.
4. In the **Create Subscription** dialog box, configure your subscription for the protocol you want:

### Email subscription

Sends an email message when you publish a message to the subscription's parent topic.

- **Protocol:** Select **Email**.
- **Email:** Type an email address.

### HTTPS (Custom URL) subscription

Sends specified information when you publish a message to the subscription's parent topic.

Endpoint format (URL using HTTPS protocol):

```
https://<anyvalidURL>
```

Basic access authentication is supported, allowing you to specify a username and password in the URL, as in `https://user:password@domain.com` or `https://user@domain.com`. The username and password are encrypted over the SSL connection established when using HTTPS. For more information about Basic Access Authentication, see [RFC-2617](#).

Query parameters are not allowed in URLs.

- **Protocol:** Select **HTTPS (Custom URL)**.
- **URL:** Type (or copy and paste) the URL you want to use as the endpoint.

### PagerDuty subscription

Creates a PagerDuty incident by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://events.pagerduty.com/integration/<integrationkey>/enqueue
```

Query parameters are not allowed in URLs.

To create an endpoint for a PagerDuty subscription (set up and retrieve an integration key), see [the PagerDuty documentation](#).

- **Protocol:** Select **PagerDuty**.
- **URL:** Type (or copy and paste) the *integration key* portion of the URL for your PagerDuty subscription. (The other portions of the URL are hard-coded.)

### Slack subscription



#### Note

See the following [known issue](#) for up-to-date information about creating Slack subscriptions.

Sends a message to the specified Slack channel by default when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The *<webhook-token>* portion of the URL contains two slashes (/).

Query parameters are not allowed in URLs.

To create an endpoint for a Slack subscription (using a webhook for your Slack channel), see [the Slack documentation](#).

- **Protocol:** Select **Slack**.
- **URL:** Type (or copy and paste) the Slack endpoint, including your webhook token.

### 5. Click **Create**.

The subscription has been created and a subscription confirmation URL will be sent. The subscription remains in "Pending" status until it has been confirmed.

## To confirm a subscription

Navigate to the confirmation URL that is sent to the subscription's endpoint.

Some protocols provide confirmation URLs in unique ways:

- **HTTPS (Custom URL):** You can find the confirmation URL in the request header or body of the subscription confirmation message (request of content-type: "application/json") that is sent to the endpoint.
  - In the request header, see the value of the X-OCI-NS-ConfirmationURL field.

### Example request header:

```
"X-OCI-NS-SignatureVersion": "1.0"
"X-OCI-NS-Signature": "<example-signature>"
"X-OCI-NS-SigningCertURL": "<example-url>"
"X-OCI-NS-TopicOcid": "ocid.compartment.oc1..<unique_ID>"
"X-OCI-NS-Timestamp": "2019-04-19T21:26:00.310+0000"
"X-OCI-NS-MessageId": "<unique_ID>"
"X-OCI-NS-TopicName": "mytopic"
"X-OCI-NS-MessageType": "SubscriptionConfirmation"
"X-OCI-NS-ConfirmationURL": "<exampleConfirmationURL>"
"X-OCI-NS-SubscriptionId": "ocid1.onssubscription.oc1.phx.<unique_ID>"
"X-OCI-NS-State": "Pending"
```

- In the request body, see the value of the ConfirmationURL key.

Example ConfirmationURL key and value (request body):

```
"ConfirmationURL": "<exampleConfirmationURL>"
```

- **PagerDuty:** Incident titled "Oracle Notification Service Subscription Confirmation". For more information, see [the PagerDuty documentation for Oracle Cloud Infrastructure](#).
- **Slack:** Message sent to Slack channel containing the text "To confirm the subscription".

### To resend a subscription confirmation

The ability to resend subscription confirmations is only applicable for pending subscriptions.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. For the subscription you want to resend the confirmation for, click the Actions icon (three dots), and then click **Resend Confirmation**.

### To update the retry duration for a subscription

The retry duration is part of the delivery policy for the subscription. By default, Notifications retries delivery of a message for up to two hours.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. For the subscription you want to update, click the Actions icon (three dots), and then click **Update Delivery Policy**.
3. In the **Update Delivery Policy** dialog box, click the edit icon for **Max Retry Duration in Minutes**, type the new value, and then click the save icon.

### To move a subscription to a different compartment

The associated topic remains in its current compartment. For more information, see [Moving Topics and Subscriptions to a Different Compartment](#).



### Note

To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Notifications resources, see [Details for the Notifications Service](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. In the **Scope** section, select a compartment.
3. Find the subscription in the list, click the Actions icon (three dots), and then click **Move Resource**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### To delete a subscription (unsubscribe)



### Note

Every message sent out as email contains a link to *unsubscribe* from the related topic.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. For the subscription you want to delete, click the Actions icon (three dots), and then

click **Delete**.

3. Confirm when prompted.

### Managing Tags for a Topic or Subscription

You can apply tags to your resources, such as topics and subscriptions, to help you organize them according to your business needs. You can apply tags at the time you create a topic or subscription, or you can update the topic or subscription later with the tags you want. For general information about applying tags, see [Resource Tags](#).

#### To manage tags for a topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Choose the **Compartment** that contains the topic you want to tag, and then click the topic's name.
3. Click the **Tags** tab to view or edit existing tags, or click **Add Tags** to add new ones.

For more information, see [Resource Tags](#).

#### To manage tags for a subscription

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Choose the **Compartment** that contains the subscription you want to tag, and then click the name of the topic that has the subscription.
3. For the subscription you want to tag, click the Actions icon (three dots), and then click **Add Tags**.  
To view or edit existing tags, click the Actions icon (three dots), and then click **View Tags**.

For more information, see [Resource Tags](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage topics:

- [CreateTopic](#)
- [GetTopic](#)
- [ListTopics](#)
- [UpdateTopic](#)
- [ChangeTopicCompartment](#)
- [DeleteTopic](#)

Use these API operations to manage subscriptions:

- [CreateSubscription](#)
- [GetSubscription](#)
- [ListSubscriptions](#)
- [UpdateSubscription](#)
- [ChangeSubscriptionCompartment](#)
- [GetConfirmSubscription](#)
- [ResendSubscriptionConfirmation](#)
- [GetUnsubscription](#)
- [DeleteSubscription](#)

## Publishing Messages

This topic describes how to publish messages directly using the Notifications service. You can manually enter the message content or allow a service or app to programmatically define the message content.

Each message is broadcast to all subscriptions in the specified topic. Every message sent out as email contains a link to *unsubscribe* from the related topic.

Message delivery rate limits per endpoint: 60 messages per minute for HTTP-based protocols. (HTTP-based protocols use URL endpoints that begin with "http:" or "https:".) 10 messages per minute for Email protocol.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Prerequisites

- IAM policies: To use Notifications, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you get a response that you don't have permission or are unauthorized, check with your administrator. You may not have the required type of access in the current compartment. For more information on user authorizations, see [Notifications Overview](#).
- Before you can publish a message, you need a topic with at least one subscription. See [Managing Topics and Subscriptions](#).

## Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

## Using the Console

### To publish a message

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. On the **Topics** page, for the topic you want, click the Actions icon (three dots), and then click **Publish Message**.
3. In the **Publish Message** dialog box, fill in the fields:
  - **Title**: Enter the title you want to send.

### Rendering of the title by protocol

Protocol	Rendering of the title
<b>Email</b>	Subject line of the email message.
<b>HTTPS (Custom URL)</b>	Not rendered.
<b>PagerDuty</b>	Title field of the published message.
<b>Slack</b>	Not rendered.

- **Message**: Enter the content you want to send.



### Note

Message size limit per request: 64KB.

4. Click **Publish**.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to publish messages:

[PublishMessage](#)

## Notifications Metrics

You can monitor the health, capacity, and performance of your messages by using metrics, alarms, and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_notification` (the Notifications service).

Resources: Not applicable. Measures data for messages, which are not resources.

## Overview of the Notifications Service Metrics

The Notifications service metrics help you measure the number and size of [messages](#) that are in initial requests, are delivered, and are not delivered.

### Prerequisites

- **IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics

The metrics listed in the following table are automatically available for messages you publish to topics. You do not need to enable monitoring on any resources to get these metrics.

Each metric includes a subset of the following dimensions:

**AVAILABILITYDOMAIN**

The availability domain in which the associated topic resides.

**ENDPOINTTYPE**

The [subscription protocol](#) of the endpoint used for the delivery attempt.

**REGION**

The region in which the associated topic resides.

**RESOURCEID**

The OCID of the resource to which the metric applies.

**TOPICDISPLAYNAME**

The friendly name of the associated topic.

## CHAPTER 24 Notifications

Metric	Metric Display Name	Unit	Description	Dimensions
PublishedMessagesSize	<b>Published Messages Size (Bytes)</b>	bytes	Size of messages in request.	availabilityDomain region resourceId
PublishedMessagesCount	<b>Published Messages Count</b>	count	Count of messages in request.	topicDisplayName
DeliveredMessagesSize	<b>Delivered Messages Size (Bytes)</b>	bytes	Size of messages successfully delivered to endpoints.	availabilityDomain endpointType region resourceId
FailedMessagesSize	<b>Failed Messages Sizes (Bytes)</b>	bytes	Size of messages that did not get delivered to endpoints.	topicDisplayName
DeliveredMessagesCount	<b>Delivered Messages Count</b>	count	Count of messages successfully delivered to endpoints.	
FailedMessagesCount	<b>Failed Messages Count</b>	count	Count of messages that did not get delivered to endpoints.	

### Using the Console

#### To view default metric charts for a single topic

1. Open the navigation menu. Under **Solutions and Platform**, go to **Application Integration** and click **Notifications**.
2. Choose the **Compartment** that contains the topic you want to view, and then click the topic's name.
3. In the **Resources** menu, click **Metrics** (if necessary).  
The **Metrics** page displays a default set of charts for the current topic.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).  
For information about notifications for alarms, see [Notifications Overview](#).

#### To view default metric charts for multiple topics

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.

2. For **Metric Namespace**, select **oci\_notification**.

The **Service Metrics** page displays a default set of charts for the selected metric namespace. For more information about the emitted metrics, see the foregoing table. You can also use the Monitoring service to create [custom queries](#).



### Tip

- [Filter metrics](#) by dimension, such as a selected topic, by clicking **Add** above the charts (to the right of **Dimensions**).
- [Aggregate data](#) across all topics (show a single line in the chart) by selecting the check box for **Aggregate Metric Streams** on the right.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#). For information about notifications for alarms, see [Notifications Overview](#).

## Using the API

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

# CHAPTER 25 Object Storage

This chapter explains how to upload, manage, and access data using Object Storage.

## Overview of Object Storage

Oracle Cloud Infrastructure offers two distinct storage class tiers to address the need for both performant, frequently accessed "hot" storage, and less frequently accessed "cold" storage. Storage tiers help you maximize performance where appropriate and minimize costs where possible.

- Use **Object Storage** for data to which you need fast, immediate, and frequent access. Data accessibility and performance justifies a higher price point to store data in the Object Storage tier.
- Use **Archive Storage** for data to which you seldom or rarely access, but that must be retained and preserved for long periods of time. The cost efficiency of the Archive Storage tier offsets the long lead time required to access the data. For more information, see [Overview of Archive Storage](#).

## About Object Storage

The Oracle Cloud Infrastructure Object Storage service is an internet-scale, high-performance storage platform that offers reliable and cost-efficient data durability. The Object Storage service can store an unlimited amount of unstructured data of any content type, including analytic data and rich content, like images and videos.

With Object Storage, you can safely and securely store or retrieve data directly from the internet or from within the cloud platform. Object Storage offers multiple [management interfaces](#) that let you easily manage storage at scale. The elasticity of the platform lets you start small and scale seamlessly, without experiencing any degradation in performance or service reliability.

Object Storage is a regional service and is not tied to any specific compute instance. You can access data from anywhere inside or outside the context of the Oracle Cloud Infrastructure, as

long you have internet connectivity and can access one of the [Object Storage endpoints](#). Authorization and resource limits are discussed later in this topic.

Object Storage also supports private access from Oracle Cloud Infrastructure resources in a VCN through a *service gateway*. A service gateway allows connectivity to the Object Storage public endpoints from private IP addresses in private subnets. For example, you can back up DB systems to an Object Storage bucket over the Oracle Cloud Infrastructure backbone instead of over the internet. You can optionally use IAM policies to control which VCNs or ranges of IP addresses can access Object Storage. See [Access to Oracle Services: Service Gateway](#) for details.

Object Storage is Always Free eligible. For more information about Always Free resources, including additional capabilities and limitations, see [Oracle Cloud Infrastructure's Free Tier](#).

The following list summarizes some of the ways that you can use Object Storage.

### **HADOOP/BIG DATA SUPPORT**

You can use Object Storage as the primary data repository for big data. Object Storage offers a scalable storage platform that lets you store large datasets and operate seamlessly on those datasets. The [HDFS connector](#) provides connectivity to various big data analytic engines like Apache Spark and MapReduce. This connectivity enables the analytics engines to work directly with data stored in Object Storage. For more information, see [Hadoop Support](#).

### **BACKUP/ARCHIVE**

You can use Object Storage to preserve backup and archive data that must be stored for an extended duration to adhere to various compliance mandates.

### **CONTENT REPOSITORY**

You can use Object Storage as your primary content repository for data, images, logs, and video. You can reliably store and preserve this data for a long time, and serve this content directly from Object Storage. The storage scales as your data storage needs scale.

### **LOG DATA**

You can use Object Storage to preserve application log data so that you can retroactively analyze this data to determine usage pattern and debug issues.

### **LARGE DATASETS**

You can use Object Storage to store generated application data that needs to be preserved for future use. Pharmaceutical trials data, genome data, and Internet of Things (IoT) data are examples of generated application data that you can preserve using Object Storage.

## Object Storage Resources

The following summarizes the Object Storage resources. Authorization and resource limits are discussed later in this topic.

### **OBJECT**

Any type of data, regardless of content type, is stored as an object. The object is composed of the object itself and metadata about the object. Each object is stored in a bucket.

### **BUCKET**

A logical container for storing objects. Users or systems create buckets as needed [within a region](#). A bucket is associated with a single compartment that has policies that determine what actions a user can perform on a bucket and on all the objects in the bucket.

### **NAMESPACE**

A logical entity that serves as a top-level container for all buckets and objects, allowing you to control bucket naming within your tenancy. Each Oracle Cloud Infrastructure tenant is assigned one unique and uneditable Object Storage [namespace](#) that spans all compartments within a region. Bucket names must be unique within each region. Within an Object Storage namespace, buckets and objects exist in flat hierarchy, but you can simulate a directory structure to help navigate a large set of objects (for example, `guitars/fender/stratocaster.jpg`, `guitars/gibson/lespaul.jpg`).



### Tip

If your namespace was created based on your tenancy name, your namespace uses all lower-case letters (regardless of the presence of capital letters in your tenancy name). When using the [API](#), [CLI](#), or [SDKs](#), do not use capital letters in your namespace string.

### COMPARTMENT

Primary building block used to organize your cloud resources. When your tenancy is provisioned, a root compartment is created for you. You can then create compartments under your root compartment to organize your resources. You control access by creating policies that specify what actions groups of users can take on the resources in those compartments. An Object Storage bucket can only exist in one compartment.

## Object Storage Features

Object Storage provides the following features:

### STRONG CONSISTENCY

When a read request is made, Object Storage always serves the most recent copy of the data that was written to the system.

### DURABILITY

Object Storage is a regional service. Data is stored redundantly across multiple storage servers. Object Storage actively monitors data integrity using checksums and automatically detects and repairs corrupt data. Object Storage actively monitors and ensures data redundancy. If a redundancy loss is detected, Object Storage automatically creates more data copies. For more details about Object Storage durability, see the [Oracle Cloud Infrastructure Object Storage FAQ](#).

### **CUSTOM METADATA**

You can define your own extensive metadata as key-value pairs for any purpose. For example, you can create descriptive tags for objects, retrieve those tags, and sort through the data. You can assign custom metadata to objects and buckets using the Oracle Cloud Infrastructure CLI or SDK. See [Software Development Kits and Command Line Interface](#) for details.

### **ENCRYPTION**

Object Storage employs 256-bit Advanced Encryption Standard (AES-256) to encrypt object data on the server. Each object is encrypted with its own data encryption key. Data encryption keys are always encrypted with a master encryption key that is assigned to the bucket. Encryption is enabled by default and cannot be turned off. By default, Oracle manages the master encryption key. However, you can optionally configure a bucket so that it's assigned an Oracle Cloud Infrastructure Key Management master encryption key that you control and rotate on your own schedule.

## Ways to Access Object Storage

You can access Object Storage using any of the following options, based on your preference and its suitability for the task you want to complete:

- The Console is an easy-to-use, browser-based interface. To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported. You are prompted to enter your cloud tenant, your user name, and your password.
- The command line interface (CLI) provides both quick access and full functionality without the need for programming. For more information, see [Using the CLI](#).
- The REST API provides the most functionality, but requires programming expertise. [API Reference and Endpoints](#) provides endpoint details and links to the available API reference documents. For general information about using the API, see [REST APIs](#). Object Storage is accessible with the following APIs:

- Object Storage Service API
- Amazon S3 Compatibility API
- Swift API (for use with Oracle RMAN)
- Oracle Cloud Infrastructure provides SDKs that interact with Object Storage without you having to create a framework. For general information about using the SDKs, see [Software Development Kits and Command Line Interface](#).

### Using Object Storage

If you are ready to use Object Storage, you can find more information in the following topics:

- For instructions on how to create a bucket and store an object in the bucket, see "Putting Data into Object Storage" in the *Oracle Cloud Infrastructure Getting Started Guide*.
- For task documentation related to buckets, see [Managing Buckets](#).
- For task documentation related to objects, see [Managing Objects](#) and [Copying Objects](#).
- For task documentation related to lifecycle management, see [Using Object Lifecycle Management](#).
- For API reference documentation, see [Object Storage Service API](#).
- For SDK and CLI information, see [Software Development Kits and Command Line Interface](#).
- For more information about using Archive Storage, see [Overview of Archive Storage](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API). IAM also manages user credentials for things like API signing keys, auth tokens, and customer secret keys for Amazon S3 Compatibility API. See [User Credentials](#) for details.

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For

example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see the [Policy Reference](#). For specific details about writing policies for Object Storage, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Object Storage IP Addresses

The Oracle Cloud Infrastructure Object Storage service uses the CIDR block IP range 134.70.0.0/17 for all regions.

### Limits on Object Storage Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

Other limits include:

- Number of Object Storage namespaces per root compartment: 1
- Maximum object size: 10 TiB
- Maximum object part size in a multipart upload: 50 GiB
- Maximum number of parts in a multipart upload: 10,000
- Maximum size of object metadata: 2 K

### Understanding Object Storage Namespaces

Each Oracle Cloud Infrastructure tenant is assigned an Object Storage namespace that spans all compartments within a region. The namespace is a unique and uneditable system-

## CHAPTER 25 Object Storage

---

generated string assigned during account creation and applies to all regions. For some older tenancies, the namespace string might be the tenancy name in all lower-case letters.

The Object Storage namespace serves as a container for all of your buckets and objects. You control bucket names within your namespace, however, bucket names must be unique within each region. You can have a bucket named **MyBucket** in US West (Phoenix) and a bucket named **MyBucket** in Germany Central (Frankfurt).

The namespace metadata stores the default compartment assignments for the Amazon S3 Compatibility API and the Swift API. For more information, see [Viewing and Specifying Designated Compartments](#).

### Using the Console

To view your Object Storage namespace string:

Open the **Profile** menu () and click **Tenancy: <your\_tenancy\_name>**. Your namespace string is listed under **Object Storage Settings**.



#### Note

While the Object Storage namespace string is displayed under **Object Storage Settings**, you cannot edit the namespace string. The namespace string appears here for information only.

### Using the Command Line Interface (CLI)

Open a command prompt and run the following command to get your Object Storage namespace:

```
oci os ns get
```

Your Object Storage namespace is returned:

```
{
 "data": "MyNamespace"
}
```



### Tip

You can use `-ns`, `--namespace`, or `--namespace-name` for options that require you to specify the Object Storage namespace string.

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [GetNamespace](#) operation to get your Object Storage namespace. If you have the `OBJECTSTORAGE_NAMESPACE_READ` permission and supply the compartment or tenancy OCID in the optional `compartmentId` parameter, you can also get the namespace of a different tenancy's Object Storage namespace.

## Managing Buckets

In the Oracle Cloud Infrastructure Object Storage service, a bucket is a container for storing objects in a compartment within an Object Storage namespace. A bucket is associated with a single compartment. The compartment has policies that indicate what actions you can perform on a bucket and all the objects in the bucket.

You cannot nest buckets—a bucket cannot contain other buckets.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you are new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For administrators:

- The policy [Let Object Storage admins manage buckets and objects](#) lets the specified group do everything with buckets and the associated objects.
- If you must write more restrictive policies for buckets, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

### Pre-Authenticated Requests

Pre-authenticated requests provide a way to let you access a bucket or an object without having your own credentials. For example, you can create a request that lets you upload backups to a bucket without owning API keys. See [Using Pre-Authenticated Requests](#) for details.

### Object Lifecycle Policies

Using object lifecycle policies applied at the bucket level, you can automatically manage the archiving and deletion of objects according to a pre-defined schedule. See [Using Object Lifecycle Management](#) for information on this feature.

### Tagging Resources

You can add tags to your resources to help you organize them according to your business needs. You can add tags at the time you create a resource, or you can update the resource

later with the desired tags. For general information about applying tags, see [Resource Tags](#).

Object Storage currently supports adding tags to buckets.

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For more information about monitoring buckets, see [Object Storage Metrics](#).

### Usage Reports

A usage report is a comma-separated value (CSV) file that can be used to get a detailed breakdown of resources in Oracle Cloud Infrastructure for audit or invoice reconciliation. A usage report is generated daily and stored in an Object Storage bucket. For more information, see [Usage Reports Overview](#) and [Accessing Usage Reports](#).

### Creating Automation for Buckets and Objects Using the Events Service

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

Buckets emit events for bucket state changes by default. Events for objects are handled differently than other resources. Objects do not emit events by default. Use the [Console](#), [CLI](#), or [API](#) to enable a bucket to emit events for object state changes. You can enable events for object state changes during or after bucket creation.

### Bucket Names

Bucket names are system generated by default, but you can overwrite the default with a name you specify.

### System-Generated Bucket Names

When a bucket is created, the system generates a default name for that bucket, for example **bucket-20190306-1359**. This bucket name identifies the current year, month, and day that the bucket was created. You can use that system-generated name for your new bucket or you can specify a different name for it.

### User-Specified Bucket Names

If you change this default bucket name or the name of any bucket, observe the following:

- Use from 1 to 256 characters.
- Valid characters are letters (upper or lower case), numbers, hyphens, underscores, and periods.



#### Important

Bucket names and object names are case-sensitive. Object Storage handles accounts-payable and Accounts-Payable as separate buckets.

- Do not include confidential information.
- Make the name unique within your tenancy's Object Storage namespace.

### Storage Tiers

When you create a bucket, you also decide which tier is appropriate for storing objects:

- Use the standard Object Storage tier for data to which you need fast, immediate, and frequent access. For more information, see [Overview of Object Storage](#).
- Use the Archive Storage tier for data to which you seldom or rarely access, but that must be retained and preserved for long periods of time. For more information, see [Overview of Archive Storage](#).



### Important

You cannot change the storage tier in which a bucket resides.

## Public Buckets

When you create a bucket, the bucket is considered a private bucket and the access to the bucket and its contents requires authentication and authorization. However, Object Storage supports anonymous, unauthenticated access to a bucket. You make a bucket *public* by enabling read access to the bucket.



### Important

Carefully assess the business requirement for public access to a bucket. When you enable anonymous access to a bucket, any user can obtain object metadata, download bucket objects, and optionally list bucket contents.

## Required Permissions

The following permissions are required to configure a public bucket:

- To enable public access when creating a bucket, use permission `BUCKET_CREATE`.
- To enable public access for an existing bucket, use permission `BUCKET_UPDATE`.

## Options

When creating a public bucket, you have the following options:

- You can configure the access to allow listing and downloading objects. List and download access is the default.
- You can configure the access to allow downloading objects only. A user would not be able to list bucket contents.

### Scope and Constraints

Understand the following scope and constraints regarding public access:

- Changing the type of access is bi-directional. You can change a bucket's access from public to private or from private to public.
- Changing the type of access doesn't affect existing pre-authenticated requests. Existing pre-authenticated requests still work.

You can enable anonymous public access for new or existing buckets using the Console, CLI, or an SDK to access the API.

## Using the Console

### To get a list of buckets

Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.

A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

### To create a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Select a compartment from the **Compartment** list on the left side of the page.  
A list of existing buckets is displayed.
3. Click **Create Bucket**.
4. In the **Create Bucket** dialog box, specify the attributes of the bucket:
  - **Bucket Name:** The system generates a default bucket name that reflects the current year, month, day, and time, for example **bucket-20190306-1359**. If you change this default to any other bucket name, use letters, numbers, dashes, underscores, and periods. Do not include any confidential information.
  - **Storage Tier:** Select the tier in which you want to store your data. Available tiers include:
    - **Standard** is the primary, default Object Storage tier for storing frequently accessed data that requires fast and immediate access.
    - **Archive** is a special tier for storing infrequently accessed data that requires long retention periods. Access to data in the **Archive** tier is not immediate. You must restore archived data before it's accessible. For more information, see "[Overview of Archive Storage](#)" in the *Oracle Cloud Infrastructure User Guide*.
  - **Object Events:** Select **Emit Object Events** if you want to enable the bucket to emit events for object state changes. For more information about events, see [Overview of Events](#).
  - **Encryption:** Buckets are encrypted with keys managed by Oracle by default, but you can optionally encrypt the data in this bucket using your own Key Management encryption key. To use Key Management for your encryption needs, select **Encrypt Using Customer-Managed Keys**. Then, select the **Vault Compartment** and **Vault** that contain the master encryption key you want to use. Also select the **Master Encryption Key Compartment** and **Master Encryption Key**. For more information about encryption, see [Overview of Key Management](#). For details on how to create a vault, see [Managing Vaults](#).

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click **Create Bucket**.

The bucket is created immediately and you can add objects to it. Objects added to archive buckets are immediately archived and must be [restored](#) before they are available for download.

### To view bucket details

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains your buckets.  
A list of buckets is displayed.
3. Click the Actions icon (three dots) to the right of the bucket name, and then click **View Bucket Details**.

Object Storage displays bucket details including the following:

- [Visibility](#)
- [Encryption Key](#)
- Namespace
- Created
- [Storage tier](#)
- Compartment
- [Approximate Count](#)
- [Approximate Size](#)

- [ETag](#) (entity tag)
- [Emit Object Events](#)

### To change the visibility of a bucket

A bucket is either private (the default) or public. See [Public Buckets](#) for more information.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the bucket name to see the bucket details.  
**Visibility:** shows the current bucket setting, which is **Private** by default.
3. Click **Edit Visibility**.
4. In the **Edit Visibility** dialog box, edit the visibility settings:
  - **Visibility**
    - **Public**
    - **Private**
  - If you select **Public** to enable public access, decide whether or not you want to let users list the bucket contents. Click **Allow users to list objects from this bucket** to set the visibility of bucket object lists.
5. Click **Save Changes**.

### To move a bucket to a different compartment



#### Important

When moving buckets resources between compartments, ensure that the resource users have sufficient access permissions for the compartment to which the resource is being moved.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. In the **Scope** section, select a compartment.
3. Find the bucket in the list, click the the Actions icon (three dots), and then click **Move Resource**.  
Alternatively, you can choose a bucket, and then click **Move Resource** on the bucket details page.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### To manage tags for a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the bucket name.
3. You can manage tags in the following ways:
  - To view the tags associated with the bucket, click the **Tags** tab, located to the right of the **Bucket Information** tab.
  - To add one or more tags, click **Add Tags**.

- To rename a tag, click the pencil icon to the left of a tag name, edit the name, and save.
- To delete a tag, click the pencil icon to the left of a tag name and click **Remove Tag**.

For more information, see [Resource Tags](#).

### To delete a bucket

You can permanently delete an *empty* bucket. The bucket cannot contain any objects. For information on deleting objects from a bucket, see [To delete objects from a bucket](#). In addition, you cannot delete a bucket that has a multipart upload in progress or a pre-authenticated request associated with that bucket.



#### **Warning**

You cannot recover a deleted bucket.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Find the bucket that you want to delete.
3. Click the Actions icon (three dots), and then click **Delete**.
4. Confirm deletion when prompted.

### To assign a Key Management master encryption key to a bucket

Buckets are encrypted with keys managed by Oracle by default, but you can optionally encrypt the data encryption keys that encrypt the objects in a bucket using your own Key Management master encryption key.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the name of the bucket that you want to encrypt.
3. Next to **Encryption Key**, do one of the following:
  - If the bucket is encrypted with a key managed by Oracle, click the **Assign** link.
  - If the bucket already has a Key Management master encryption key assigned, to assign a different key, click the **Edit** link.
4. In the dialog box, provide or edit the following information:
  - **Vault Compartment** and **Vault** that contain the master encryption key you want to use. The current compartment is displayed by default.
  - **Master Encryption Key Compartment** and **Master Encryption Key**. The current compartment is displayed by default.
5. When you are finished, click **Assign** or **Edit**.

See [Overview of Key Management](#) for more details.

### To remove a Key Management master encryption key from a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the name of the bucket for which you want to remove a Key Management key assignment.
3. Next to **Encryption Key**, click the **Unassign** link.
4. In the **Confirm** dialog box, click **OK** to remove the key assignment from the bucket.

### To re-encrypt a bucket's data encryption keys

If you've rotated a master encryption key since the time you assigned it to a bucket, you might want to re-encrypt the bucket. Until you explicitly re-encrypt a bucket, the key version associated with the bucket when an object was inserted into the bucket continues to decrypt all data encryption keys. To encrypt and decrypt all data encryption keys with the same, most recent version of the assigned master encryption key, re-encrypt the bucket.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the name of the bucket for which you want to re-encrypt all data encryption keys.
3. Click **Re-encrypt**. (If the button is not enabled, that's because the bucket is using a master encryption key managed by Oracle rather than a Key Management master encryption. Or, the bucket does not contain any objects.)
4. In the confirmation dialog box, click **Re-encrypt** to generate a work request to re-encrypt all data encryption keys associated with the bucket.

The **Work Requests Details** dialog box that displays tells you about the work request, including the percentage completed and the work request ID. You can copy the work request ID to monitor its status later.

### To view the approximate bucket size and number of objects in the bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains your buckets.  
A list of buckets is displayed.
3. Click the Actions icon (three dots) to the right of the bucket name, and then click **View Bucket Details**.

- **Approximate Count** is the approximate number of objects in the bucket. Count statistics are reported periodically. A lag can occur between what is displayed and the actual object count.
- **Approximate Size** is the approximate total size of all objects in the bucket. Size statistics are reported periodically. A lag can occur between what is displayed and the actual size of the bucket.

### To enable or disable emitting events for object state changes

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains your buckets.  
A list of buckets is displayed.
3. Click the Actions icon (three dots) to the right of the bucket name, and then click **View Bucket Details**.
4. Next to **Emit Object Events**, click **Edit**.
5. In the dialog box, select (to enable) or deselect (to disable) **Emit Object Events**.
6. Click **Save Changes**.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).



### Note

The examples in this section use the full syntax for all parameters, for example `--namespace` and `--compartment-id`. In some cases, there are shortened parameter terms that you can use instead of the full ones, for example `-ns` and `-c`. See the CLI online help for instances of a shortened parameter associated with a command.

## To get a list of buckets

```
oci os bucket list --namespace <object_storage_namespace> --compartment-id <target_compartment_id>
```

For example:

```
oci os bucket list --namespace MyNamespace --compartment-id ocid.compartment.oc1..exampleuniqueID

{
 "data": [
 {
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocidl.user.oc1..exampleuniqueID",
 "defined-tags": null,
 "etag": "c8889cd1-8414-41fb-84b7-3738c39e62c5",
 "freeform-tags": null,
 "name": "MyBucket1",
 "namespace": "MyNamespace",
 "time-created": "2019-10-22T19:22:25.032000+00:00"
 },
 {
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocidl.user.oc1..exampleuniqueID",
 "defined-tags": null,
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e",
 "freeform-tags": null,
 "name": "MyBucket2",
```

## CHAPTER 25 Object Storage

---

```
 "namespace": "MyNamespace ",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 }
]
}
```

By default, getting a list of buckets returns up to the first 1,000 buckets in the compartment.

To list all buckets in a compartment, use the `--all` option:

```
oci os bucket list --namespace <object_storage_namespace> --compartment-id <target_compartment_id> --all
```

To include [resource tag](#) data, use the `--fields tags` option:

```
oci os bucket list --namespace <object_storage_namespace> --compartment-id <target_compartment_id> --fields tags
```

For example:

```
oci os bucket list --namespace MyNamespace --compartment-id ocid1.compartment.oc1..exampleuniqueID --fields tags
```

```
{
 "data": [
 {
 "compartment-id": "ocid1.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {
 "example_tag_namespace_Financials": {
 "production": "Unit 5"
 },
 "example_tag_namespace_Operations": {
 "costcenter": "85"
 }
 },
 "etag": "48af18cf-1edd-4b05-9f36-a629d5032260",
 "freeform-tags": {
 "Project": "prototype 3"
 },
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "time-created": "2019-02-27T18:52:16.951000+00:00"
 }
]
}
```

## CHAPTER 25 Object Storage

---

```
]
}
```



### Note

If you do not specify the `--fields tags` option when listing buckets, `null` is returned as the value for both freeform and defined tags.

## To create a standard Object Storage tier bucket

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id <target_compartment_id>
```

For example:

```
oci os bucket create --namespace MyNamespace --name MyBucket --compartment-id
ocid.compartment.oc1..exampleuniqueID
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 },
}
```

## CHAPTER 25 Object Storage

---

```
"etag": "7b7c3dc1-713f-4996-b176-a938345cae8e"
}
```

By default, a bucket is created in the standard Object Storage tier. You do not need to explicitly set `--storage-tier`.

A Standard tier bucket is created immediately and you can add objects to it.

### To create an Archive tier bucket

To create an Archive tier bucket, you must explicitly set `--storage-tier Archive`.

```
oci os bucket create --namespace <object_storage_namespace> --name <archivebucket_name> --compartment-id <target_compartment_id> --storage-tier Archive
```

For example:

```
oci os bucket create -ns MyNamespace --name MyArchiveBucket --compartment-id
ocid.compartment.oc1..exampleuniqueID --storage-tier Archive
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "c8889cd1-8414-41fb-84b7-3738c39e62c5",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyArchiveBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Archive",
 "time-created": "2019-10-22T19:22:25.032000+00:00"
 },
 "etag": "c8889cd1-8414-41fb-84b7-3738c39e62c5"
}
```

## CHAPTER 25 Object Storage

---

An Archive Storage bucket is created and you can add objects to it. Objects added to Archive Storage buckets are immediately archived and must be [restored](#) before they are available for download.

### To create a public bucket that allows listing and downloading bucket objects

To create a public bucket that allows listing and downloading bucket objects, you must explicitly set `--public-access-type ObjectRead`.

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id <target_compartment_id> --public-access-type ObjectRead
```

For example:

```
oci os bucket create --namespace MyNamespace --name MyPublicObjectReadBucket --compartment-id ocid.compartment.oc1..exampleuniqueID --public-access-type ObjectRead
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "01096e0b-659a-4d9d-a806-d57568cf1b22",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyPublicObjectReadBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "ObjectRead",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:47:11.649000+00:00"
 },
 "etag": "01096e0b-659a-4d9d-a806-d57568cf1b22"
}
```

### To create a public bucket that allows downloading bucket objects only

To create a public bucket that allows downloading bucket objects only, you must explicitly set `--public-access-type ObjectReadWithoutList`.

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id <target_compartment_id> --public-access-type ObjectReadWithoutList
```

For example:

```
oci os bucket create -ns MyNamespace --name MyPublicObjectReadBucket --compartment-id ocid.compartment.oc1..exampleuniqueID --public-access-type ObjectReadWithoutList{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "ec20c59a-f5ba-4a6d-8a7e-b69bb9bb76ad",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyPublicObjectReadWithoutListBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "ObjectReadWithoutList",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T20:18:29.203000+00:00"
 },
 "etag": "ec20c59a-f5ba-4a6d-8a7e-b69bb9bb76ad"
}
```

### To create a bucket with resource tags

You can create standard Object Storage tier or [Archive](#) tier buckets with [resource tags](#).

To add resource tags when creating a bucket, set one or both of the `--defined-tags` and `--freeform-tags` options.



### Tip

The `--defined-tags` and `--freeform-tags` options require that the input to be a complex type formatted in valid JSON. See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for information JSON formatting.

The following example syntax creates a standard Object Storage tier bucket with a defined tag:

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id <target_compartment_id> --defined-tags '<JSON_formatted_defined_tag>'
```

Examples of defined tag formatting:

```
'{"Operations": {"CostCenter": "42"}}'
```

```
'{"Logistics": {"Procurement": "Madrid Center"},"Financials":{"Production": "Unit 5"}}'
```



### Note

If you are running the CLI on a Windows computer, you might need to use the backslash (\) character to escape the strings containing the tag values. For example, a single defined tag is formatted as follows:

```
'{"\Logistics\": {"\Procurement\": \"Madrid Center\"}}'
```

For example:

```
oci os bucket create --namespace MyNamespace --name MyBucketDefined --compartment-id ocid.compartment.oc1..exampleuniqueID --defined-tags {"Operations": {"CostCenter": "42"}}
{
 "data": {
 "approximate-count": null,
```

## CHAPTER 25 Object Storage

```
"approximate-size": null,
"compartment-id": "ocid.compartment.oc1..exampleuniqueID",
"created-by": "ocid1.user.oc1..exampleuniqueID",
"defined-tags": {
 "operations": {
 "costcenter": "42"
 }
},
"etag": "ea88f444-842c-462d-965e-d3540b3b54f6",
"freeform-tags": {},
"kms-key-id": null,
"metadata": {},
"name": "MyBucketDefined",
"namespace": "MyNamespace",
"object-events-enabled": false,
"object-lifecycle-policy-etag": null,
"public-access-type": "NoPublicAccess",
"storage-tier": "Standard",
"time-created": "2019-10-23T19:47:51.362000+00:00"
},
"etag": "ea88f444-842c-462d-965e-d3540b3b54f6"
}
```

The following example syntax creates a Standard tier bucket with a free-form tag:

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id
<target_compartment_id> --freeform-tags <JSON_formatted_free-form_tag>
```

Examples of free-form tag formatting:

```
'{"Chicago_Team": "marketing_videos"}'
```

```
'{"Project": "prototype 3", "Manager": "Meadows"}'
```



### Note

If you are running the CLI on a Windows computer, you might need to use the backslash (\) character to escape the strings containing the tag values. For example, a single free-form tag is formatted as:

```
'{"Chicago_Team\": {\"marketing_videos\"}}'
```

## CHAPTER 25 Object Storage

---

For example:

```
oci os bucket create --namespace MyNamespace --name MyBucketFreeform --compartment-id
ocid.compartment.oc1..exampleuniqueID --freeform-tags {"Chicago_Team": "marketing_videos"}
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocidl.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "6f4bdal0-fc8b-462e-8563-875639fd7294",
 "freeform-tags": {
 "Chicago_Team": "marketing_videos"
 },
 "kms-key-id": null,
 "metadata": {},
 "name": "MyBucketFreeform",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-23T20:51:16.260000+00:00"
 },
 "etag": "6f4bdal0-fc8b-462e-8563-875639fd7294"
}
```

To view bucket details

```
oci os bucket get --name <bucket_name>
```

Object Storage displays bucket details including the following:

- [Visibility](#)
- [Encryption Key](#)
- Namespace
- Created

## CHAPTER 25 Object Storage

---

- [Storage tier](#)
- Compartment
- [Approximate Count](#)
- [Approximate Size](#)
- [ETag](#) (entity tag)
- [Emit Object Events](#)

For example:

```
oci os bucket get --name MyBucket
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 },
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e"
}
```

### To view bucket metadata

```
oci os bucket get --namespace <object_storage_namespace> --name <bucket_name>
```

For example:

## CHAPTER 25 Object Storage

---

```
oci os bucket get --namespace MyNamespace --name MyBucket

{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 },
 "etag": "7b7c3dc1-713f-4996-b176-a938345cae8e"
}
```

### To add custom metadata key-value pairs to a bucket

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --metadata <JSON-formatted_key-value_pair>
```

*<JSON-formatted\_key-value\_pair>* is a key-value pair input as valid formatted JSON. See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for more information about JSON formatting.

For example:

```
oci os bucket update --namespace MyNamespace --name MyBucket --metadata '{"Department": "Finance"}'
```

```
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
```

## CHAPTER 25 Object Storage

```
"compartment-id": "ocid.compartment.oc1..exampleuniqueID",
"created-by": "ocidl.user.oc1..exampleuniqueID",
"defined-tags": {},
"etag": "4b09d7b9-a8bf-42f6-8d67-bb357694f92d",
"freeform-tags": {},
"kms-key-id": null,
"metadata": {
 "department": "Finance"
},
"name": "MyBucket",
"namespace": "MyNamespace",
"object-events-enabled": false,
"object-lifecycle-policy-etag": null,
"public-access-type": "NoPublicAccess",
"storage-tier": "Standard",
"time-created": "2019-10-22T19:04:05.879000+00:00"
},
"etag": "4b09d7b9-a8bf-42f6-8d67-bb357694f92d"
}
```

### To make a bucket private or public

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --public-access-type
[NoPublicAccess|ObjectReadWithoutList|ObjectRead]
```

- **NoPublicAccess:** Allows only an authenticated caller to access the bucket and its contents. This is the default value.
- **ObjectReadWithoutList:** Allows public access for the `GetObject`, `HeadObject`, and `ListObjects` operations.
- **ObjectRead:** Allows public access for the `GetObject` and `HeadObject` operations.

### For example:

```
oci os bucket update --namespace MyNamespace --name MyBucket --public-access-type ObjectRead
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
```

## CHAPTER 25 Object Storage

---

```
"created-by": "ocid1.user.oc1..exampleuniqueID",
"defined-tags": {},
"etag": "09ab3193-a441-43cc-a8e2-e468e94c7c60",
"freeform-tags": {},
"kms-key-id": null,
"metadata": {
 "department": "Finance"
},
"name": "MyBucket",
"namespace": "MyNamespace",
"object-events-enabled": false,
"object-lifecycle-policy-etag": null,
"public-access-type": "ObjectRead",
"storage-tier": "Standard",
"time-created": "2019-10-22T19:04:05.879000+00:00"
},
"etag": "09ab3193-a441-43cc-a8e2-e468e94c7c60"
}
```

### To move a bucket to a different compartment

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --compartment-id
<new_target_compartment_id>
```

**<new\_target\_compartment\_id>** is the compartment ID associated with the compartment to which you are moving the bucket.

For example:

```
oci os bucket update --namespace MyNamespace --name MyBucket --compartment-id
ocid.compartment.oc1..exampleuniqueID
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "new_ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "fe4fb648-8ddd-42eb-9732-d431aafac354",
 "freeform-tags": {},
 "kms-key-id": null,
```

## CHAPTER 25 Object Storage

---

```
"metadata": {
 "department": "Finance"
},
"name": "MyBucket",
"namespace": "MyNamespace",
"object-events-enabled": false,
"object-lifecycle-policy-etag": null,
"public-access-type": "ObjectRead",
"storage-tier": "Standard",
"time-created": "2019-10-22T19:04:05.879000+00:00"
},
"etag": "fe4fb648-8ddd-42eb-9732-d431aafac354"
}
```

### To add resource tags to a bucket

To add defined [resource tags](#) to a bucket, open a command prompt and run `oci os bucket update` with the `--defined-tags` option:

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --defined-tags <JSON_
formatted_defined_tag>
```

For example:

```
oci os bucket update --namespace MyNamespace --name MyBucket --defined-tags '{"Operations":
{"CostCenter": "42"}}'
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {
 "operations": {
 "costcenter": "42"
 }
 },
 "etag": "0a26b47d-c43f-4ef8-9c26-02bb8d69fa34",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {
 "department": "Finance"
 },
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "ObjectRead",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 }
}
```

## CHAPTER 25 Object Storage

---

```
 },
 "etag": "0a26b47d-c43f-4ef8-9c26-02bb8d69fa34"
 }
}
```

To add free-form resource tags to a bucket, open a command prompt and run `oci os bucket update` with the `--freeform-tags` option:

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --freeform-tags
<JSON_formatted_free-form_tag>
```

For example:

```
oci os bucket update --namespace MyNamespace --name MyBucket --freeform-tags '{"Chicago_Team":
"marketing_videos"}'
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocidl.user.oc1..exampleuniqueID",
 "defined-tags": {
 "operations": {
 "costcenter": "42"
 }
 },
 "etag": "856a3c73-0194-4c02-8c6b-1b20be3c9a48",
 "freeform-tags": {
 "Chicago_Team": "marketing_videos"
 },
 "kms-key-id": null,
 "metadata": {
 "department": "Finance"
 },
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "ObjectRead",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 },
 "etag": "856a3c73-0194-4c02-8c6b-1b20be3c9a48"
}
```



### Tip

The `--defined-tags` and `--freeform-tags` options require that you provide key-value pair input as valid formatted JSON. For examples of JSON-formatted resource tags, see [To create a Standard or Archive tier bucket with resource tags](#). See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for more information about JSON formatting.

### To delete a bucket

You can permanently delete an empty bucket. The bucket cannot contain any objects. For information on deleting objects, see [To delete objects from a bucket](#). You also cannot delete a bucket that has a multipart upload in progress or a pre-authenticated request associated with that bucket.

```
oci os bucket delete --namespace <object_storage_namespace> --name <bucket_name>
```

For example:

```
oci os bucket delete --namespace MyNamespace --name MyDeletedBucket
```

```
Are you sure you want to delete this resource? [y/N]:
```

Select `y` and press `Enter`. The bucket is deleted with no further prompting.



### Warning

You cannot recover a deleted bucket.

### To assign a Key Management key to a bucket

```
oci os bucket create --namespace <object_storage_namespace> --name <bucket_name> --compartment-id
```

## CHAPTER 25 Object Storage

---

```
<target_compartment_id> --kms-key-id <target_key_id>
```

*<target\_key\_id>* is the ID of the key versions that contain the cryptographic material used to encrypt and decrypt data, protecting the data where it is stored.

For example:

```
oci os bucket create --namespace MyNamespace --name MyKeyBucket --compartment-id
ocid.compartment.oc1..exampleuniqueID --kms-key-id ocid1.key.region1.sea..exampleuniqueID
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "e7f29fdd-b5f5-42e5-a98b-80883f9f2f32",
 "freeform-tags": {},
 "kms-key-id": "ocid1.key.region1.sea..exampleuniqueID",
 "metadata": {},
 "name": "MyKeyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-29T23:00:35.490000+00:00"
 },
 "etag": "e7f29fdd-b5f5-42e5-a98b-80883f9f2f32"
}
```

See [Overview of Key Management](#) for more details.

### To update the Key Management key assigned to a bucket

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --kms-key-id <target_
key_id>
```

*<target\_key\_id>* is the ID of the key versions that contain the cryptographic material used to encrypt and decrypt data, protecting the data where it is stored.

## CHAPTER 25 Object Storage

---

For example:

```
oci os bucket update --namespace MyNamespace --name MyKeyBucket --kms-key-id
ocid1.key.region1.sea.exampleuniqueID_updated
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "e7f29fdd-b5f5-42e5-a98b-80883f9f2f32",
 "freeform-tags": {},
 "kms-key-id": "ocid1.key.region1.sea..exampleuniqueID_updated",
 "metadata": {},
 "name": "MyKeyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2019-10-29T23:00:35.490000+00:00"
 },
 "etag": "e7f29fdd-b5f5-42e5-a98b-80883f9f2f32"
}
```

To remove the Key Management key assigned to a bucket

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --kms-key-id ""
```

For example:

```
oci os bucket update --namespace MyNamespace --name MyKeyBucket --kms-key-id ""
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1.user.oc1..exampleuniqueID",
 "defined-tags": {},
 "etag": "10a50818-e495-45a9-b1ce-cc815f7b39ad",
 }
}
```

## CHAPTER 25 Object Storage

---

```
"freeform-tags": {},
"kms-key-id": null,
"metadata": {},
"name": "MyKeyBucket",
"namespace": "MyNamespace",
"object-events-enabled": false,
"object-lifecycle-policy-etag": null,
"public-access-type": "NoPublicAccess",
"storage-tier": "Standard",
"time-created": "2019-10-29T23:00:35.490000+00:00"
},
"etag": "10a50818-e495-45a9-b1ce-cc815f7b39ad"
}
```

### To re-encrypt a bucket's data encryption keys

If you rotated a master encryption key since the time you assigned it to a bucket, you might want to re-encrypt the bucket. Until you explicitly re-encrypt a bucket, the key version associated with the bucket when an object was inserted into the bucket continues to decrypt all data encryption keys. To encrypt and decrypt all data encryption keys with the same current version of the assigned master encryption key, re-encrypt the bucket.

```
oci os bucket reencrypt --name <bucket_name>
```

For example:

```
oci os bucket reencrypt --name MyBucket
```

### To view the approximate bucket size and number of objects in the bucket

```
oci os bucket get --name <bucket_name> --fields approximateCount --fields approximateSize
```

- `approximateCount` is the approximate number of objects in the bucket. Count statistics are reported periodically. You will see a lag between what is displayed and the actual object count.

## CHAPTER 25 Object Storage

---

- `approximateSize` is the approximate total size of all objects in the bucket. Size statistics are reported periodically. You will see a lag between what is displayed and the actual size of the bucket.

For example:

```
oci os bucket get --name MyBucket --fields approximateCount --fields approximateSize
{
 "data": {
 "approximate-count": 7,
 "approximate-size": 8075918,
 "compartment-id": "ocid1.compartment.oc1..exampleuniqueID",
 "created-by": "ocid1:user:oc1:phx:1458751937789:exampleuniqueID",
 "defined-tags": {},
 "etag": "218f201f-28a4-434d-9591-f05b6223c67a",
 "freeform-tags": {},
 "kms-key-id": null,
 "metadata": {},
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": false,
 "object-level-audit-mode": "Disabled",
 "object-lifecycle-policy-etag": null,
 "public-access-type": "NoPublicAccess",
 "storage-tier": "Standard",
 "time-created": "2017-10-19T04:11:32.040000+00:00"
 },
 "etag": "218f201f-28a4-434d-9591-f05b6223c67a"
}
```

### To enable or disable emitting events for object state changes

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

Open a command prompt and run `oci os bucket update` to enable or disable:

```
oci os bucket update --namespace <object_storage_namespace> --name <bucket_name> --object-events-enabled [true|false]
```

## CHAPTER 25 Object Storage

---

For example, to enable emitting events for all objects in the bucket named `MyBucket`:

```
oci os bucket update --namespace example_namespace --name MyBucket --object-events-enabled true
{
 "data": {
 "approximate-count": null,
 "approximate-size": null,
 "compartment-id": "ocidl.compartment.oc1..exampleuniqueID",
 "created-by": "ocidl:user:oc1:phx:1458751937789:exampleuniqueID",
 "defined-tags": {
 "operations": {
 "costcenter": "42"
 }
 },
 "etag": "39d1db02-27d0-4263-b3ff-5e6450495457",
 "freeform-tags": {
 "Chicago_Team": "marketing_videos"
 },
 "kms-key-id": null,
 "metadata": {
 "department": "Finance"
 },
 "name": "MyBucket",
 "namespace": "MyNamespace",
 "object-events-enabled": true,
 "object-lifecycle-policy-etag": null,
 "public-access-type": "ObjectRead",
 "storage-tier": "Standard",
 "time-created": "2019-10-22T19:04:05.879000+00:00"
 },
 "etag": "39d1db02-27d0-4263-b3ff-5e6450495457"
}
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

When accessing the Object Storage API, the bucket name is used with the Object Storage namespace name to form the request URL:

```
n/<object_storage_namespace>/b/<bucket>
```

Use the following operations to manage buckets:

- [CreateBucket](#)
- [DeleteBucket](#)
- [GetBucket](#)
- [HeadBucket](#)
- [ListBuckets](#)
- [ReencryptBucket](#)
- [UpdateBucket](#)



### Note

There are two key properties worthy of mention in the [CreateBucket](#) and [UpdateBucket](#) APIs:

- `publicAccessType` property controls whether the bucket is private or public and limits the capability to list public bucket contents.
- `objectEventsEnabled` property controls if events are emitted for the objects in this bucket.

## Managing Objects

In the Oracle Cloud Infrastructure Object Storage service, an object is a file or unstructured data you upload to a bucket within a compartment within an Object Storage [namespace](#). The object can be any type of data, for example, multimedia files, data backups, static web content, or logs. You can store objects that are up to 10 TiB. Objects are processed as a single entity. You can't edit or append data to an object, but you can replace the entire object.

This topic describes how to manage objects within a single bucket. For information on copying an object to another bucket, see [Copying Objects](#).

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For administrators:

- The policy [Let Object Storage admins manage buckets and objects](#) lets the specified group do everything with buckets and objects. Objects always reside in the same compartment as the bucket.
- If you need to write a more restrictive policy for objects, the `inspect objects` lets you list all the objects in a bucket and do a HEAD operation for a particular object. In comparison, `read objects` lets you download the object itself. See [Details for Object Storage, Archive Storage, and Data Transfer](#).

### Pre-Authenticated Requests

Pre-authenticated requests provide a way to let users access a bucket or object without having their own credentials. For example, you can create a request that lets a user upload backups to a bucket without owning API keys. See [Using Pre-Authenticated Requests](#) for details.

### Object Names

Unlike other resources, objects do not have Oracle Cloud Identifiers (OCIDs). Instead, users define an object name when they upload an object.

Use the following guidelines when naming an object:

- Use from 1 to 1024 characters.
- Valid characters are letters (upper or lower case), numbers, and characters other than linefeed, newline, and NULL.



### Important

Bucket names and object names are case-sensitive. Object Storage handles q3-field-assets.xlsx and Q3-Field-Assets.XSLX as separate objects.

- Use only Unicode characters for which the UTF-8 encoding does not exceed 1024 bytes. Clients are responsible for URL-encoding characters.
- Do not include confidential information.
- Make the name unique within the bucket. Do not use the name of an existing object within the bucket when naming an object unless you intend to overwrite the existing object with the contents of the new or renamed object.



### Tip

Object names can include one or more forward slash (/) characters in the name. See [Object Naming Using Prefixes and Hierarchies](#) for more information on using the forward slash in object names to create hierarchies.

## Object Naming Using Prefixes and Hierarchies

Within an Object Storage namespace, buckets and objects exist in a flat hierarchy, but you can simulate a directory structure using a prefix string that includes the forward slash (/) to add hierarchy to an object name. Doing so lets you list one directory at a time, which is helpful when navigating a large set of objects.

For example:

## CHAPTER 25 Object Storage

---

```
marathon/finish_line.jpg
marathon/participants/p_21.jpg
```

If you added hierarchy to object names, you can use the CLI or API to perform bulk downloads and bulk deletes of all objects at a specified level of the hierarchy. Bulk downloads and bulk deletes a specified level of the hierarchy do not affect objects in any level above.

When naming objects, you can also use prefix strings without a delimiter. No delimiters would allow certain bulk operations in the CL or API to match on the prefix portion of the object name. For example, in the object names below, the string `gloves_27_` can serve as a prefix for matching purposes when performing bulk downloads or deletions:

```
gloves_27_dark_green.jpg
gloves_27_light_blue.jpg
```

When you perform bulk uploads with the CLI or API, you can also prepend a prefix string to the names of the files you are uploading.

### Object Lifecycle Management

Using Object Lifecycle Management feature, you can automatically manage the archiving and deletion of objects according to a pre-defined schedule. See [Using Object Lifecycle Management](#) for information on this feature.

### Multipart Uploading and Downloading

The Oracle Cloud Infrastructure Object Storage service supports multipart uploading and downloading for objects. See [Using Multipart Uploads](#) for more information. This page includes links to API documentation for this functionality. For CLI information on multipart downloading, see the procedure [for downloading an object using multipart download](#). For API documentation related to multipart downloading, see the [GetObject](#) API call and its **range** parameter.

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For more information about monitoring objects, see [Object Storage Metrics](#).

### Creating Automation for Objects Using the Events Service

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

Events for objects are handled differently than other resources. Objects do not emit events by default. Use the [Console](#), [CLI](#), or [API](#) to enable a bucket to emit events for object state changes. You can enable events for object state changes during or after bucket creation.

### Using Storage Gateway to Upload and Download Objects

Storage Gateway is another way you can upload objects to and download objects from Oracle Cloud Infrastructure Object Storage.

Storage Gateway is installed in an Oracle Cloud Infrastructure compute instance or as a Linux Docker instance on one or more hosts in your on-premises data center. Applications store and retrieve objects from Oracle Cloud Infrastructure Object Storage through *file systems* that you create in Storage Gateway. Storage Gateway exposes an NFS mount point that can be mounted to any host that supports an NFSv4 client. The Storage Gateway mount point maps to an Object Storage bucket to upload and download objects.

See [Overview of Storage Gateway](#) for details.

### Using the Console

#### To upload objects to a bucket

1. From the Object Storage Details screen, click the bucket name to view its details.
2. Click **Objects** under **Resources**.
3. In the **Objects** table, click **Upload Objects**.
4. Optionally, specify an **Object Name Prefix**. If provided, this prefix is prepended to each one of the files you upload. The prefix lets you simulate hierarchy and perform bulk operations. See [Object Naming Using Prefixes and Hierarchies](#) for details.
5. In the **Upload Objects** dialog box, select the objects that you want to upload in one of two ways:
  - Drag and drop one or more files from your computer.
  - Click the **select files** link and use the **File Upload** dialog box.

The files you select to upload are displayed in a list. If you decide that you do not want to upload a particular file, click the **X** to the right of the file name.

If the files you select to upload are already stored in the bucket with the same name, the Console displays messages warning you of an overwrite.

6. Click **Upload Objects**.

The selected objects are uploaded and displayed in the list of objects in the bucket.

#### To download an object from a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains the bucket that contains your object.

A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.

A list of objects in the bucket is displayed.

5. For the object you want to download, click the Actions icon (three dots), and then click **Download**.

### To view object details

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains the bucket that contains your object.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. Choose the object for which you want details.
6. Click the Actions icon (three dots), and then click **View Object Details**. The following object details are displayed:
  - Name
  - URL Path (URI)
  - Storage Tier
  - Size
  - Content Type
  - [ETag](#) (entity tag)
  - Last Modified
7. Optionally, click **Download** to download the selected object.

### To rename an object

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains the bucket that contains your object.  
A list of buckets is displayed.

3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. For the object you want to rename, click the Actions icon (three dots), and then click **Rename**.
6. In the **Rename Object** dialog box, provide the new name for the object and an optional delimited directory structure prefix. For example, `p_94.jpg` or `/marathon/participants/p_94.jpg`.  
Avoid entering confidential information in the object name.



### Warning

Buckets cannot store two objects that use identical names (case-sensitive). If you choose to rename an object using the name of another object in the same bucket, the object that originally used the name is overwritten.

7. Click **Save Changes**.

### To restore objects from Archive Storage

Depending on the size of the object, it can take four or more hours to restore an object from Archive Storage. You cannot download an item until the item is fully restored.



### Tip

You need `OBJECT_RESTORE` permissions to restore Archive Storage objects.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment your bucket is in.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. To restore a single object, click the Actions icon (three dots) to the right of the object you want to restore, and then click **Restore**. To restore multiple objects, select the check boxes to the left of each object you want to restore, then click **Restore**.
6. Optionally, specify the **Time Available for Download in Hours**.  
By default, you have 24 hours to download an object after restoration. However you can alternatively specify a download time of from 1 to 240 hours. You can find out how much download time is remaining by looking at **Available for Download** in object **Details** or by looking at the Actions icon (three dots) menu to the right of **Download**. Refresh the browser to obtain up-to-date remaining download time information.  
After the allotted download time expires, the object returns to Archive Storage.
7. Click **Restore Objects**.  
Error messages are generated if there is a problem with restoring the selected objects. You can optionally click **Retry failed restore option**.

### To check the status of an Archive Storage object restoration

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment your bucket is in.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. Click the Actions icon (three dots) to the right of the object you want to check the

restoration or download status of, then click **Details**.

6. Check the **Status**.

**Status** displays one of the following:

- Archived
- Restoring
- Restored

### To delete objects from a bucket

You can permanently delete an object from a bucket. You cannot, however, recover a deleted object.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains the bucket that contains the object or objects you want to delete.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. To delete a single object, click the Actions icon (three dots) to the right of the object you want to delete, and then click **Delete**. To delete multiple objects, select the check boxes to the left of each object you want to delete, and then click **Delete**.
6. Confirm when prompted.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

## CHAPTER 25 Object Storage

---

### To list objects in a bucket

Open a command prompt and run `oci os object list` to get a list of the objects in a bucket:

```
oci os object list -ns <object_storage_namespace> -bn <bucket_name>
```

By default, the following details are displayed for each object:

- Name
- Object size
- "Last Modified" timestamp
- MD5 hash

### To get object details

Open a command prompt and run `oci os object head` to get [object details](#):

```
oci os object head -ns <object_storage_namespace> -bn <bucket_name> --name <object_name>
```

The system output includes the following object details:

- [ETag](#) (entity tag)
- Content length (object body size)
- Custom metadata key-value pairs
- Storage tier
- MD5 hash
- [Archival state](#)

### To upload an object to a bucket

Open a command prompt and run `oci os object put` to upload an object:

```
oci os object put -ns <object_storage_namespace> -bn <bucket_name> --file <file_location> --name <object_name> --no-multipart
```

Where *<file\_location>* refers to a directory path like `C:\workspace\myfile.txt`. If you want to use the filename as the uploaded object's name, you can omit the `--name` option. The resulting object name does not include the path information (for example, `C:\workspace\`).

To add custom metadata key-value pairs, use the `--metadata` option:

```
oci os object put -ns <object_storage_namespace> -bn <bucket_name> --file <file_location> --name <object_name> --metadata <json_formatted_key-value_pairs> --no-multipart
```



### Tip

The `--metadata` option requires that you provide complex type key-value pair input in valid JSON. See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for more information about JSON formatting.

An object can be uploaded as a single part or as multiple parts. Here we describe a single part upload. For information on multipart uploads, see [Using Multipart Uploads](#).

## To bulk upload objects to a bucket

Open a command prompt and run `oci os object bulk-upload` to upload all files in a given directory (including files in subdirectories):

```
oci os object bulk-upload -ns <object_storage_namespace> -bn <bucket_name> --src-dir <source_directory_location> --no-multipart
```

Where *<source\_directory\_location>* refers to a directory path like `C:\workspace\files_to_upload\`. If your source directory has subdirectories, the subdirectory names are prepended to the names of the files stored in those subdirectories, delimited with a forward slash (/) character. For example, if a file named `maple.jpg` is stored in the subdirectory `trees`, when the file is uploaded, Object Storage assigns the name `trees/maple.jpg` to the resulting object.

## CHAPTER 25 Object Storage

---

To append a [prefix string](#) to the object names created by your uploads, use the `--object-prefix` option:

```
oci os object bulk-upload -ns <object_storage_namespace> -bn <bucket_name> --src-dir <source_directory_location> --object-prefix <object_name_prefix_string> --no-multipart
```

For example:

```
oci os object bulk-upload -ns ansh8tvru7zp -bn apparel --src-dir C:\workspace\new_items\bicycling\gloves\ --object-prefix /bicycling/gloves/ --no-multipart
```

To add custom metadata key-value pairs, use the `--metadata` option:

```
oci os object bulk-upload -ns <object_storage_namespace> -bn <bucket_name> --src-dir <source_directory_location> --metadata <json_formatted_key-value_pairs> --no-multipart
```



### Tip

The `--metadata` option requires that you provide complex type key-value pair input in valid JSON format. See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for more information about JSON formatting.

### To download an object from a bucket

Open a command prompt and run `oci os object get` to download an object:

```
oci os object get -ns <object_storage_namespace> -bn <bucket_name> --name <object_name> --file <file_location>
```

Where *<file\_location>* refers to a directory path like `C:\workspace\myfile.txt`.

### To download an object using multipart download

Multipart object downloading is available using the byte-range request standard defined in [RFC 7233, section 2.1](#).

## CHAPTER 25 Object Storage

---

To initiate a multipart download, open a command prompt and run `oci os object get` with the `--range` option and the `bytes=<byte_range>` byte-range specifier:

```
oci os object get -ns <object_storage_namespace> -bn <bucket_name> --name <object_name> --file <file_location> --range bytes=<byte_range>
```

For example:

```
oci os object get -ns ansh8lvru1zp -bn my_bucket --name my_object.mp4 --file /Users/me/my_object.mp4 --range bytes=0-499
```

### To bulk download all objects within a bucket

Open a command prompt and run `oci os object bulk-download` to download all the objects in a bucket:

```
oci os object bulk-download -ns <object_storage_namespace> -bn <bucket_name> --download-dir <download_directory_location>
```

Where `<download_directory_location>` refers to a directory path like

`C:\workspace\objects\` where downloaded objects are saved. If the directory does not exist, Object Storage creates the directory.

For a complete list of object bulk download options, see [CLI Help](#).

### To bulk download objects by object name prefix string

If you have named your objects with [prefix](#) strings, you can bulk download those objects in a bucket that match a specified prefix string. Open a command prompt and run `oci os object bulk-download` command with the `--prefix` option:

```
oci os object bulk-download -ns <object_storage_namespace> -bn <bucket_name> --download-dir <download_directory_location> --prefix <prefix_string>
```

Where `<download_directory_location>` refers to a directory path like

`C:\workspace\objects\` where downloaded objects are saved. If the directory does not exist, Object Storage creates the directory.

For example:

## CHAPTER 25 Object Storage

---

```
oci os object bulk-download -ns ansh8tvru7zp -bn apparel --download-dir C:\objects\ --prefix gloves_27
```

In the example above, an object named `gloves_27_A.jpg` is downloaded, while an object named `gloves_31_A.jpg` is not downloaded.

If you named your objects so that they exist in Object Storage in a [hierarchy](#), you can download objects at a specified level and below. Specify the prefix that matches the level of your choosing:

```
oci os object bulk-download -ns <object_storage_namespace> -bn <bucket_name> --download-dir <download_directory_location> --prefix <level_1/level_2/>
```

The preceding command downloads the following objects:

- `<level_1/level_2/object_name>`
- `<level_1/level_2/level_3/object_name>`
- `<level_1/level_2/level_3/level_4/object_name>`

To download only those objects at a given hierarchy level (and not objects in levels above or below), see [To bulk download objects at a specified hierarchy level](#).

### To bulk download objects at a specified hierarchy level

If you named your objects so that they exist in Object Storage in a [hierarchy](#), you can bulk download all objects at a specified hierarchy level.

Open a command prompt and run `oci os object bulk-download` command with the `--prefix` and `--delimiter` flags:

```
oci os object bulk-download -ns <object_storage_namespace> -bn <bucket_name> --download-dir <download_directory_location> --prefix <level_1/level_2/> --delimiter /
```

Where `<download_directory_location>` refers to a directory path like

`C:\workspace\objects\` where files downloaded objects are saved. If the directory you specify does not exist, Object Storage creates this directory.



### Note

Currently, only the forward slash (/) is the supported delimiter for the `--delimiter` option.

The preceding command downloads objects only at `<level_2>` of the hierarchy. For example, the following object is downloaded:

`<level_1/level_2/object_name>`

The preceding command does not download objects in levels *above* or *below* `<level_2>`. For example, the preceding command does not download the following objects:

- `<level_1/object_name>`
- `<level_1/level_2/level_3/object_name>`
- `<level_1/level_2/level_3/level_4/object_name>`

To download objects at a given hierarchy level along with all objects in the hierarchy sublevels, see [To bulk download objects by object name prefix string](#).

### To rename an object

Open a command prompt and run `oci os object rename` to rename an object:

```
oci os object rename -ns <object_storage_namespace> -bn <bucket_name> --name <object_original_name> --new-name <object_new_name>
```



### Warning

Avoid entering confidential information in object name.

For example:

## CHAPTER 25 Object Storage

---

```
oci os object rename -ns ansh8tvru7zp -bn photo_collection --name /marathon/participants/p_93.jpg --new-name /marathon/participants/p_94.jpg
```

To make the rename operation dependent on the object having a specific entity tag, use the `--src-obj-if-match-e-tag` option:

```
oci os object rename -ns <object_storage_namespace> -bn <bucket_name> --name <object_original_name> --new-name <object_new_name> --src-obj-if-match-e-tag <etag_required_for_object_rename>
```

For example:

```
oci os object rename -ns ansh8lvru7zp -bn my_bucket --name my_object.jpg --new-name my_renamed_object.jpg --src-obj-if-match-e-tag 6672BECB67CCFFBCE0530292F20ZBACE
```

For rename operations where you intend to overwrite one object in a bucket with another, you can make the renaming dependent on having a specific entity tag. To do so, use the `--new-obj-if-match-e-tag` option:

```
oci os object rename -ns <object_storage_namespace> -bn <bucket_name> --name <source_object_name> --new-name <name_of_object_to_be_overwritten> --new-obj-if-match-e-tag <etag_of_object_to_be_overwritten>
```

For example:

```
oci os object rename -ns ansh8lvru7zp -bn my_bucket --name my_object.jpg --new-name my_renamed_object.jpg --new-obj-if-match-e-tag 6672BECB67CCFFBCE0530292F20ZBACE
```

When renaming an object, you can prevent the system from overwriting another object in the same bucket by using the `--new-obj-if-none-match-e-tag *` option. This option prevents the renaming operation from completing if an object exists with the `--new-name` value specified and the same entity tag of the source object.

```
oci os object rename -ns <object_storage_namespace> -bn <bucket_name> --name <source_object_name> --new-name <new_name_for_object> --new-obj-if-none-match-e-tag *
```

For example:

```
oci os object rename -ns ansh8lvru7zp -bn my_bucket --name my_object.jpg --new-name my_renamed_object.jpg --new-obj-if-none-match-e-tag *
```

### To restore an Archive Storage tier object



#### Tip

You need OBJECT\_RESTORE permissions to restore Archive Storage objects.

Open a command prompt and run `oci os object restore` to restore an object from Archive Storage:

```
oci os object restore -ns <object_storage_namespace> -bn <archive_bucket_name> --name <archived_object_name> [--hours <#_of_hours>]
```

By default, you have 24 hours to download an object after restoration. However, you can optionally specify `--hours` with an integer value of download time of from 1 to 240 hours.

### To check the status of an Archive Storage object restoration

Open a command prompt and run `oci os object restore-status` to check the status of restoring an object from Archive Storage:

```
oci os object restore-status -ns <object_storage_namespace> -bn <archive_bucket_name> --name <archived_object_name>
```

### To delete an object

You can permanently delete an object. Open a command prompt and run `oci os object delete` to delete an object:

```
oci os object delete -ns <object_storage_namespace> -bn <bucket_name> --name <object_name>
```

### To bulk delete all objects within a bucket

Open a command prompt and run `oci os object bulk-delete` to delete all the objects in a bucket:

## CHAPTER 25 Object Storage

```
oci os object bulk-delete -ns <object_storage_namespace> -bn <bucket_name>
```



### Tip

To see a list of the files that will be deleted by a bulk delete command without actually deleting the files, use the `--dry-run` option.

### To bulk delete objects by object name prefix string

If you named your objects with [prefix](#) strings, you can bulk delete objects in a given bucket by providing a prefix to match. Open a command prompt and run `oci os object bulk-delete` command with the `--prefix` option:

```
oci os object bulk-delete -ns <object_storage_namespace> -bn <bucket_name> --prefix <prefix_string>
```

For example:

```
oci os object bulk-delete -ns ansh8tvru7zp -bn apparel --prefix gloves_27
```

The preceding command deletes an object named `gloves_27_A.jpg`, but does not delete an object named `gloves_31_A.jpg`.

If you named your objects so that they exist in a [hierarchy](#), you can bulk delete objects at a given level and below by specifying a prefix to match:

```
oci os object bulk-delete -ns <object_storage_namespace> -bn <bucket_name> --prefix <level_1/level_2/>
```

The preceding command deletes the following files:

- `<level_1/level_2/object_name>`
- `<level_1/level_2/level_3/object_name>`
- `<level_1/level_2/level_3/level_4/object_name>`

To delete only those objects at a given hierarchy level (and not objects in levels above or below), see [To bulk delete objects at a specified hierarchy level](#).



### Tip

To see a list of the files that will be deleted by a bulk delete command without actually deleting the files, use the `--dry-run` option.

### To bulk delete objects at a specified hierarchy level

If you named your objects so that they exist in a [hierarchy](#), you can bulk delete only those objects at a given hierarchy level (and not objects in levels above or below). Open a command prompt and run the `oci os object bulk-delete` command with the `--prefix` and `--delimiter` flags:

```
oci os object bulk-delete -ns <object_storage_namespace> -bn <bucket_name> --prefix <level_1/level_2/> --delimiter /
```



### Note

Currently, only the forward slash (/) is the supported delimiter for the `--delimiter` option.

The preceding bulk delete command deletes the following object:

*<level\_1/level\_2/>object\_name*

The preceding command does not bulk delete objects in levels above or below *<level\_2>*. For example, the command would not delete the following objects:

- *<level\_2/object\_name>*
- *<level\_1/level\_2/level\_3/object\_name>*
- *<level\_1/level\_2/level\_3/level\_4/object\_name>*

To delete objects at a given hierarchy level along with all objects in the hierarchy sublevels, see [To bulk delete objects by object name prefix string](#).



### Tip

To see a list of the files that will be deleted by a bulk delete command without actually deleting the files, use the `--dry-run` option.

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Object Storage prepends the Object Storage namespace string and bucket name to the object name when constructing a URL for use with the API. Everything :

```
/n/<object_storage_namespace>/b/<bucket>/o/<object_name>
```

The object name is everything after the `/o/`, which could include hierarchy levels and prefix strings.

Use the following operations to manage objects:

- [DeleteObject](#)
- [GetObject](#)
- [HeadObject](#)
- [ListObjects](#)
- [PutObject](#) (see [Special Instructions for Object Storage PUT](#) for signing request requirements)
- [RenameObject](#)
- [RestoreObjects](#)

# Copying Objects

This topic describes how to copy objects in Object Storage. You can copy objects to other buckets in the same region and to buckets in other regions.

## Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).



### Warning

Object copy does not work if you do not authorize the Object Storage service to copy objects on your behalf. See [Service Permissions](#) for more information.

## User Permissions

You must have the required access to both the source and destination buckets when performing an object copy. You must also have permissions to manage objects in the source and destination buckets.

For administrators:

- You can create a policy that lets the specified IAM group manage Object Storage namespaces, buckets, and their associated objects in all compartments in the tenancy:

```
Allow group <IAM_group_name> to manage object-family in tenancy
```

- Alternatively, you can create policies that reduce the scope of access. For example, to

## CHAPTER 25 Object Storage

---

let the specified group manage only buckets and objects in a particular compartment in the tenancy:

```
Allow group <IAM_group_name> to manage buckets in compartment <compartment_name>
```

For more information about other alternatives for writing policies, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

### Service Permissions

Because Object Storage is a regional service, you must authorize the Object Storage service for each region carrying out copy operations on your behalf. For example, you might authorize the Object Storage service in region US East (Ashburn) to manage objects on your behalf. Once you authorize the Object Storage service and ensure that you have the required user permissions, you can copy an object stored in a US East (Ashburn) bucket to a bucket in another region.

To determine the region name value of an Oracle Cloud Infrastructure region, see [Regions and Availability Domains](#).

For administrators:

To enable object copy, you must authorize the service to manage objects on your behalf:

- You can create a policy that authorizes the service in the specified region to manage Object Storage namespaces, buckets, and their associated objects in all compartments in the tenancy:

```
Allow service objectstorage-<region_name> to manage object-family in tenancy
```

- Rather than use the [policy verb](#) `manage`, you can create a policy that reduces the scope of access by instead using one of the following statements:

```
Allow service objectstorage-<region_name> to {OBJECT_READ, OBJECT_INSPECT, OBJECT_CREATE, OBJECT_OVERWRITE, OBJECT_DELETE} in tenancy
```

```
Allow service objectstorage-<region_name> to {OBJECT_READ, OBJECT_INSPECT, OBJECT_CREATE, OBJECT_OVERWRITE, OBJECT_DELETE} in compartment <compartment_name>
```

### Copy Object Work Requests

The Object Storage service handles copy requests asynchronously. The service creates a queue for copy requests, and then processes the requests when system resources become available. To provide visibility for in-progress copy operations, Object Storage creates a [work request](#). You can track the progress of the copy operation by monitoring the status of the work request.

The work request statuses are:

**ACCEPTED**

The copy request is in the work request queue to be processed.

**IN PROGRESS**

The object copy is in progress.

**SUCCEEDED**

The copy operation has successfully completed.

**CANCELING**

The copy request is in the process of being canceled.

**CANCELED**

The copy request has been canceled.

**FAILED**

The copy operation has failed. Work requests that do not complete because of overwrite rules or insufficient user authorizations are assigned the failed status.

### Copy Object Overwrite Rules

You can use overwrite rules to control the copying of objects based on their entity tag (ETag) values.

- **Overwrite destination object:** Use this option when you do not want to limit a copy operation by an ETag value. This option is the default. This option can be used for any copy operation, regardless of whether it involves overwriting an existing object.
- **Do not overwrite any destination object:** Use this option to prevent the overwriting an existing copy of an object in the destination location, regardless of the destination object's ETag value.
- **Overwrite destination object only if it matches the specified ETag:** Use this option to prevent the accidental overwriting of an object in the destination location that does not have the specified ETag. When you use this option, the copy operation only succeeds if the ETag you supply when initiating the copy request matches the ETag of the destination object.
- **Copy object only if the source matches the specified ETag:** Use this option if you want the copy operation successful only if the ETag you supply when initiating the copy request matches the ETag of the source object. For objects that are intentionally updated and overwritten as part of data management activity, this option ensures that only the specified *version* of the object (as indicated by the ETag) is allowed to be copied. If the object's ETag value changes after the copy work request is created, but before the copy operation is executed, the copy operation will not complete.



### Warning

If you overwrite an object, the operation cannot be undone.

## Scope and Constraints

- Objects cannot be copied directly from [Archive Storage](#). To copy objects that are currently in Archive Storage, you must first [restore](#) the object to the standard Object Storage tier. Objects can be copied directly to Archive Storage tier buckets from the standard Object Storage tier. When you copy objects into an Archive Storage bucket,

the copy of the object is immediately archived.

- Specify an existing target bucket for the copy request. The copy operation does not automatically create buckets.
- When an object is copied, the destination object receives a new ETag value.
- If you rename, overwrite, or delete a source object during a copy operation, the copy operation fails and the destination object is not created or overwritten.
- Bulk copying is not supported. Identify a single object in the copy request.

### Using the Console

The Console consumes the REST API and is subject to the same considerations as any Oracle Cloud Infrastructure client.

### To make a copy of an object

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment that contains the bucket that contains your object.
3. Select the bucket containing the object that you want to copy.
4. Click **Objects** under **Resources** to display a list of objects in the bucket.
5. For the object you want to copy (the source object), click the Actions icon (three dots), and then click **Copy**.
6. In the **Copy** dialog, enter the following:
  - **Destination Namespace:** The [Object Storage Namespace](#) of the destination bucket for your copied object. The namespace string of your tenancy is supplied as the default value.
  - **Destination Region:** The Oracle Cloud Infrastructure region containing the destination bucket for your copied object. Your tenancy must be [subscribed to a region](#) in order for you to copy an object to a bucket in that region.

- **Destination Bucket:** The name of the destination bucket for your copied object. Specify an existing target bucket. The copy operation does not automatically create buckets.
  - **Destination Object Name:** Optionally, you can specify a different destination object name. By default, the **Destination Object Name** is the same name as the object you are copying.
  - **Overwrite Rule:** Select the overwrite rule appropriate for your copy request. See [Copy Object Overwrite Rules](#) for information on the overwrite rule options.
7. Click **Copy Object**.  
A dialog confirms that your copy request was submitted successfully.

### To monitor the status of an object copy work request

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment containing your bucket.
3. Click the bucket name of the bucket containing the object being copied.
4. Click **Work Requests** under **Resources**.  
A list of work requests is displayed. The [status](#) of the request and details including object name, request ID, and the destination bucket's name, region, and [namespace](#) is also displayed.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To make a copy of an object

Open a command prompt and run `oci os object copy` to copy an object:

## CHAPTER 25 Object Storage

---

```
oci os object copy --namespace-name <object_storage_namespace> --bucket-name <source_bucket_name> --
source-object-name <source_object> --destination-namespace <destination_namespace_string> --
destination-region <destination_region> --destination-bucket <destination_bucket_name> --destination-
object-name <destination_object_name>
```

For example:

```
oci os object copy --namespace-name ansh8lvru1zp --bucket-name photos --source-object-name
hummingbird.jpg --destination-namespace ansh8lvru1zp --destination-region uk-london-1 --destination-
bucket UK_photos --destination-object-name hummingbird.jpg
```

For a complete list of object copy options, see [CLI Help](#).

### To get the status of an object copy work request

Use the `work-request get` command to get the [status](#) of an object copy work request using the work request ID. If you do not have the work request ID, you can get a list of work requests, including the request IDs, for a specified compartment using the [work-request list command](#).

Open a command prompt and run `oci os work-request get` to get the status of a work request:

```
oci os work-request get --work-request-id <request_id>
```

For a complete list of work request options, see [CLI Help](#).

### To get a list of work requests for a compartment

Open a command prompt and run `oci os work-request list` to get a list of a work requests for a specified compartment:

```
oci os work-request list --compartment-id <compartment_id>
```

For a complete list of work request options, see [CLI Help](#).

### To cancel a copy object work request

Open a command prompt and run `oci os work-request cancel` to cancel a work request:

## CHAPTER 25 Object Storage

---

```
oci os work-request cancel --work-request-id <request_id>
```

For a complete list of work request options, see [CLI Help](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these operations to view and manage work requests for copy object operations:

- [CopyObject](#)
- [ListWorkRequests](#)
- [GetWorkRequest](#)
- [CancelWorkRequest](#)

### Using Pre-Authenticated Requests

Pre-authenticated requests provide a way to let users access a bucket or an object without having their own credentials, as long as the request creator has permissions to access those objects. For example, you can create a request that lets an operations support user upload backups to a bucket without owning API keys. Or, you can create a request that lets a business partner update shared data in a bucket without owning API keys.

When you create a pre-authenticated request, a unique URL is generated. Anyone you provide this URL to can access the Object Storage resources identified in the pre-authenticated request, using standard HTTP tools like curl and wget.



### Important

Assess the business requirement for and the security ramifications of pre-authenticated access to a bucket or objects.

A pre-authenticated request URL gives anyone who has the URL access to the targets identified in the request. Carefully manage the distribution of the URL.

## Required Permissions

### To Create a Pre-Authenticated Request

To create or manage pre-authenticated requests, you need `PAR_MANAGE` permission to the target bucket or object.

While you only need `PAR_MANAGE` permission to create a pre-authenticated request, you must also have the appropriate permissions for the access type that you are granting. For example:

- If you are creating a pre-authenticated request for uploading objects to a bucket, you need `OBJECT_CREATE` and `OBJECT_OVERWRITE` permissions in addition to `PAR_MANAGE`.
- If you are creating a pre-authenticated request for read/write access to objects in a bucket, you need `OBJECT_READ`, `OBJECT_CREATE`, and `OBJECT_OVERWRITE` permissions in addition to `PAR_MANAGE` to grant user read/write access to objects.



### Important

If the creator of a pre-authenticated request is deleted or loses the required permissions after they created the request, the request will no longer work.

### To Use a Pre-Authenticated Request

Permissions of the pre-authenticated request creator are checked each time you use a pre-authenticated request. The pre-authenticated request no longer works if any of the following occurs:

- Permissions of the pre-authenticated request creator change
- User who created the pre-authenticated request is deleted
- Federated user who created the pre-authenticated request has lost the user capabilities that they had when they created the request
- Pre-authenticated request has expired

### Options

When creating a pre-authenticated request, you have the following options:

- You can specify the name of a bucket that a pre-authenticated request user has write access to and can upload one or more objects to.
- You can specify the name of an object that a pre-authenticated request user can read from, write to, or read from and write to.

### Scope and Constraints

Understand the following scope and constraints regarding pre-authenticated requests:

- Users can't list bucket contents.
- You can create an unlimited number of pre-authenticated requests.
- There is no time limit to the expiration date that you can set.
- You can't edit a pre-authenticated request. If you want to change user access options in response to changing requirements, you must create a new pre-authenticated request.
- The target and actions for a pre-authenticated request are based on the creator's permissions. The request is not, however, bound to the creator's account login

credentials. If the creator's login credentials change, a pre-authenticated request is not affected.

- You cannot delete a bucket that has a pre-authenticated request associated with that bucket or with an object in that bucket.

### Working with Pre-Authenticated Requests

You can create, delete, or list pre-authenticated requests using the Console, using the CLI, or by using an SDK to access the API.



#### Important

The unique URL provided by the system when you create a pre-authenticated request is the only way a user can access the bucket or object specified as the request target. Copy the URL to durable storage. The URL is displayed only at the time of creation and cannot be retrieved later.

After creating a pre-authenticated request, you can use a tool like curl to read and write data using the pre-authenticated request.

#### To put an object

```
$ curl -X PUT <unique-PAR-URL>
```

For example:

```
$ curl -X PUT https://objectstorage.us-phoenix-1.oraclecloud.com/p/j3DoSvqQHbUaw6ADzHkDlnaqMuXWef_
lhTxCiS9ngCw/n/docs/b/par-bucket/o/using-dita-guide.pdf
```

#### To get an object

```
$ curl -X GET <unique-PAR-URL>
```

For example:

```
$ curl -X GET https://objectstorage.us-phoenix-1.oraclecloud.com/p/j3DoSvqQHbUaw6ADzHkDlnaqMuXWef_
lhTxCiS9ngCw/n/namespace/b/par-bucket/o/using-dita-guide.pdf
```

## Using the Console

To create a pre-authenticated request for a bucket

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment where the bucket is.
3. Click the bucket name.
4. Click **Pre-Authenticated Requests** under **Resources** to display the list of pre-authenticated requests.
5. Click **Create Pre-Authenticated Request**.
6. Provide the following information:
  - **Name:** The system automatically generates a default, pre-authenticated request name that reflects the current year, month, day, and time, for example **par-bucket-20191101-1327**.  
If you change this default or any other pre-authenticated request name, use letters, numbers, dashes, underscores, and periods. Avoid entering confidential information.
  - **Pre-Authenticated Request Target:** Select **Bucket**.
  - **Expiration:** Accept the one week, system-generated expiration date or use the date and time editor to use a different expiration date and time.
7. Click **Create Pre-Authenticated Request**.  
After a request is created, the **Pre-Authenticated Request Details** dialog box displays the URL used to access the bucket.

8. Click **Copy** to copy the URL for future reference.



### Note

The unique URL provided by the system when you create a pre-authenticated request is the only way a user can access the bucket or object specified as the request target. Copy the URL to durable storage. The URL is displayed only at the time of creation and cannot be retrieved later.

9. Click **Close**.

### To create a pre-authenticated request for an object

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment where the bucket is.
3. Click the bucket name.
4. Click **Objects** under **Resources** to display the list of objects.
5. For the object you want to create a pre-authenticated request, click the Actions icon (three dots), and then click **Create Pre-Authenticated Request**.
6. Provide the following information:
  - **Name:** The system generates a default, pre-authenticated request name that reflects the current year, month, day, and time, for example **par-object-object-name-20191101-1429**.  
If you change this default or any other pre-authenticated request name, use letters, numbers, dashes, underscores, and periods. Avoid entering confidential information.
  - **Pre-Authenticated Request Target:** Select **Object**.

- **Object Name:** The name of the object that you want authenticated by this rule.
  - **Access Type:** Select one of the following.
    - Permit read on the object
    - Permit writes to the object
    - Permit reads on and writes to the object
  - **Expiration:** Accept the one week, system-generated expiration date or use the date and time editor to a different expiration date and time.
7. Click **Create Pre-Authenticated Request**.  
After a request is created, the **Pre-Authenticated Request Details** dialog displays the URL used to access the object.
  8. Click **Copy** to copy the URL for future reference.



### Note

The unique URL provided by the system when you create a pre-authenticated request is the only way a user can access the bucket or object specified as the request target. Copy the URL to durable storage. The URL is displayed only at the time of creation and cannot be retrieved later.

9. Click **Close**.

## To copy a pre-authenticated request ID

To copy the ID for a pre-authenticated request to the clipboard, do the following:

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment where the bucket is.
3. Click the bucket name.

4. Click **Pre-Authenticated Requests** under **Resources** to display the list of pre-authenticated requests.
5. For the pre-authenticated request ID that you want to copy, click the Actions icon (three dots), and then click **Copy Pre-Authenticated Request ID**.

The ID for the selected pre-authentication request is copied to the clipboard.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To create a pre-authenticated request for a bucket

```
oci os preauth-request create -ns <object_storage_namespace> -bn <bucket_name> --name <preauthenticated_request_name> --access-type AnyObjectWrite --time-expires <timestamp>
```

Note the following:

- To create a pre-authenticated request for a bucket, use the `AnyObjectWrite` enum value with the `--access-type` flag. Pre-authenticated requests for buckets permit writes to the bucket by default.
- The `<timestamp>` is required and must be an RFC 3339 timestamp. For example: `2017-09-01T00:09:51.000+02:00`.



#### Note

The unique URL provided by the system when you create a pre-authenticated request is the only way a user can access the bucket or object specified as the request target. Copy the URL to durable storage. The URL is displayed only at the time of creation and cannot be retrieved later.

### To create a pre-authenticated request for an object

```
oci os preauth-request create -ns <object_storage_namespace> -bn <bucket_name> --name <preauthenticated_request_name> --access-type <enum_value> --time-expires <timestamp> -on <object_name_or_null>
```

The *<enum\_value>* for `--access-type` is one of the following:

- `ObjectRead` (permits read on the object)
- `ObjectWrite` (permits writes to the object)
- `ObjectReadWrite` (permits reads on and writes to the object)

The *<timestamp>* is required and must be an RFC 3339 timestamp. For example: 2017-09-01T00:09:51.000+02:00.

Avoid entering confidential information in the pre-authenticated request name.



#### Note

The unique URL provided by the system when you create a pre-authenticated request is the only way a user can access the bucket or object specified as the request target. Copy the URL to durable storage. The URL is displayed only at the time of creation and cannot be retrieved later.

### To list a pre-authenticated request

```
oci os preauth-request list -ns <object_storage_namespace> -bn <bucket_name>
```

### To get a pre-authenticated request

```
oci os preauth-request get -ns <object_storage_namespace> -bn <bucket_name> --par-id <preauthenticated_request_id>
```

### To delete a pre-authenticated request

```
oci os preauth-request delete -ns <object_storage_namespace> -bn <bucket_name> --par-id <preauthenticated_request_id>
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to work with pre-authenticated requests:

- [CreatePreauthenticatedRequest](#)
- [DeletePreauthenticatedRequest](#)
- [GetPreauthenticatedRequest](#)
- [ListPreauthenticatedRequests](#)

### Using Multipart Uploads

The Oracle Cloud Infrastructure Object Storage service supports multipart uploads for more efficient and resilient uploads, especially for large objects. You can perform multipart uploads using the [API](#), the [Software Development Kits and Command Line Interface](#), or the [Command Line Interface \(CLI\)](#). With multipart uploads, individual parts of an object can be uploaded in parallel to reduce the amount of time you spend uploading. Multipart uploads performed through the API can also minimize the impact of network failures by letting you retry a failed part upload instead of requiring you to retry an entire object upload.

Multipart uploads can accommodate objects that are too large for a single upload operation. Oracle recommends that you perform a multipart upload to upload objects larger than 100 MiB. The maximum size for an uploaded object is 10 TiB. Object parts must be no larger than 50 GiB. For large uploads performed through the API, you have the flexibility of pausing between the uploads of individual parts, and resuming the upload when your schedule and resources allow.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you are new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For administrators:

- The policy in [Let Object Storage admins manage buckets and objects](#) lets the specified group do everything with buckets and objects.
- If you need to write more restrictive policies for buckets, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For more information about monitoring multipart uploads, see [Object Storage Metrics](#).

### Using the Multipart Upload API

A multipart upload performed using the API consists of the following steps:

1. Initiating an upload
2. Uploading object parts
3. Committing the upload

Before you use the multipart upload API, you are responsible for creating the parts to upload. Object Storage provides API operations for the remaining steps. The service also provides API

operations for listing in-progress multipart uploads, listing the object parts in an in-progress multipart upload, and aborting in-progress multipart uploads initiated through the API.

Here we provide a high-level overview of the API steps, but you can refer to the [API Reference](#) for specifics about supported API calls.

### Creating Object Parts

With multipart upload, you split the object you want to upload into individual parts. Individual parts can be as large as 50 GiB or as small as 10 MiB. (Object Storage waives the minimum part size restriction for the last uploaded part.) Decide what part number you want to use for each part. Part numbers can range from 1 to 10,000. You do not need to assign contiguous numbers, but Object Storage constructs the object by ordering part numbers in ascending order.

### Initiating an Upload

After you finish creating object parts, initiate a multipart upload by making a `CreateMultipartUpload` REST API call. Provide the object name and any object metadata. Object Storage Responds with a unique upload ID that you must include in any requests related to this multipart upload. Object Storage also marks the upload as active. The upload remains active until you explicitly commit it or abort it.

### Uploading Object Parts

Make an `UploadPart` request for each object part upload. In the request parameters, provide the Object Storage [namespace](#), bucket name, upload ID, and part number. In the request body, include the object part. Object parts can be uploaded in parallel and in any order. When you commit the upload, Object Storage uses the part numbers to sequence object parts. Part numbers do not have to be contiguous. If multiple object parts are uploaded using the same upload ID and part number, the last upload overwrites the part and is committed when you call the `CommitMultipartUpload` API.

Object Storage returns an ETag (entity tag) value for each part uploaded. You need both the part number and corresponding ETag value for each part when you commit the upload.

If you have network issues, you can restart a failed upload for an individual part. You do not need to restart the entire upload. If for some reason, you cannot perform an upload all at once, multipart upload lets you continue uploading parts at your own pace. While a multipart upload is still active, you can keep adding parts as long as the total number is less than 10,000.

You can check on an active multipart upload by listing all parts that have been uploaded. (You cannot list information for an individual object part in an active multipart upload.) The `ListMultipartUploadParts` operation requires the Object Storage namespace, bucket name, and upload ID. Object Storage responds with information about the parts associated with the specified upload ID. Parts information includes the part number, ETag value, MD5 hash, and part size (in bytes).

Similarly, if you have multiple multipart uploads occurring simultaneously, you can see what uploads are in-progress. Make an `ListMultipartUploads` API call to list active multipart uploads in the specified Object Storage namespace and bucket.

Charges for parts storage begin accruing when you upload data.

### **Committing the Upload**

When you have uploaded all object parts, commit the upload. Use the `CommitMultipartUpload` request parameters to specify the Object Storage namespace, bucket name, and upload ID. Include the part number and corresponding ETag value for each part in the body of the request. When you commit the upload, Object Storage constructs the object from its constituent parts. The object is stored in the specified bucket and Object Storage namespace. You can treat it like you would any other object. Garbage collection releases storage space occupied by any part numbers you uploaded, but did not include in the `CommitMultipartUpload` request.

You cannot list or retrieve parts from a completed upload. You cannot append or remove parts from the completed upload either. If you want, you can replace the object by initiating a new upload.

If you decide to abort a multipart upload instead of committing it, wait for in-progress part uploads to complete and then use the `AbortMultipartUpload` operation. If you abort an upload

while part uploads are still in progress anyway, Object Storage cleans up both completed and in-progress parts. Upload IDs from aborted multipart uploads cannot be reused.

### API Documentation

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage multipart uploads:

- [AbortMultipartUpload](#)
- [CommitMultipartUpload](#)
- [CreateMultipartUpload](#)
- [ListMultipartUploadParts](#)
- [ListMultipartUploads](#)
- [UploadPart](#) (see [Special Instructions for Object Storage PUT](#) for signing request requirements)

### Using the CLI

When you perform a multipart upload using the CLI, you do not need to split the object into parts as you are required to do by the API. Instead, you specify the part size of your choice, and Object Storage splits the object into parts and performs the upload of all parts automatically. You can choose to set the maximum number of parts that can be uploaded in parallel. By default, the CLI limits the number of parts that can be uploaded in parallel to three. When using the CLI, you do not have to perform a commit when the upload is complete.

You can also use the CLI to list in-progress multipart uploads, and to abort multipart uploads initiated through the API.

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To perform a multipart upload using the CLI

To upload an object, open a command prompt and run `oci os object put` with the `--part-size` flag. The `--part-size` value represents the size of each part in mebibytes (MiBs). Object Storage waives the minimum part size restriction for the last uploaded part. The `--part-size` value must be an integer.

Optionally, you can use the `--parallel-upload-count` flag to set the maximum number of parallel uploads allowed.

```
oci os object put --namespace <object_storage_namespace> -bn <bucket_name> --file <file_location> --name <object_name> --part-size <upload_part_size_in_MB> --parallel-upload-count <maximum_number_parallel_uploads>
```

For example:

```
oci os object put --namespace MyNamespace -bn MyBucket --file ~/path/to/file --name MyObject --parallel-upload-count 10 --part-size 500
Upload ID: 277ffff5-e1b5-e81d-5f81-c374a8f33998
Split file into 12 parts for upload.
Uploading object ##### 100%
{ "etag": "861c8341-74d8-4142-8da4-28e1ce7783ba", "last-modified": "Wed, 25 Sep 2019 19:59:15 GMT", "opc-multipart-md5": "9Qnleyou2yMiy009Bc7o1A==12" }
```

For more information on the `oci os object put` command, see [To upload an object to a bucket](#).

### To list the parts of an unfinished or failed multipart upload

```
oci os multipart list -ns <object_storage_namespace> -bn <bucket_name>
```

For example:

```
oci os multipart list --bucket-name MyBucket{
 "data": [
 {
 "bucket": "MyBucket",
 "namespace": "MyNamespace",
 "object": "MyObject",
 "time-created": "2019-07-25T21:55:21.973000+00:00",
 "upload-id": "0b7abd48-9ff2-9d5f-2034-63a02fdd7afa"
 },
 {
 "bucket": "MyBucket",
 "namespace": "MyNamespace",

```

## CHAPTER 25 Object Storage

```
"object": "MyObject",
"time-created": "2019-07-25T21:53:09.246000+00:00",
"upload-id": "1293ac9d-83f8-e055-a5a7-d1e13277b5c0"
},
{
 "bucket": "MyBucket",
 "namespace": "MyNamespace",
 "object": "MyObject",
 "time-created": "2019-07-25T21:46:34.981000+00:00",
 "upload-id": "33e7a875-9e94-c3bc-6577-2ee5d8226b53"
}
...

```



### Tip

See [CLI Help](#) for command options to control the pagination of the list output.

## To remove a part of an unfinished or failed multipart upload

```
oci os multipart abort -ns <object_storage_namespace> -bn <bucket_name> --object-name <object_name> --
upload-id <upload_ID>
```

For example:

```
oci os multipart abort --bucket-name MyBucket --object-name MyObject --upload-id 0b7abd48-9ff2-9d5f-
2034-63a02fdd7afa
WARNING: Are you sure you want to permanently remove this incomplete upload? [y/N]: y
```



### Tip

The CLI interface asks you to confirm the deletion request. To abort without the confirmation prompt, use the `--force` flag.

## To remove all parts of an unfinished or failed multipart upload

Use the following script with the `--force` flag to remove all parts:

```
#!/bin/bash
BUCKET=$1
oci os multipart list --bucket-name $BUCKET | \
jq -c '.data | map({'o': .object, 'i': ."upload-id"}) | .[]' | \
while read JSON; do
 OBJECTNAME=$(echo $JSON | jq '.o' | sed -e 's/\\/"/g;')
 UPLOADID=$(echo $JSON | jq '.i' | sed -e 's/\\/"/g;')
 echo Removing Object name $OBJECTNAME, ID $UPLOADID
 oci os multipart abort --bucket-name $BUCKET \
 --object-name $OBJECTNAME \
 --upload-id $UPLOADID \
 --force
done
```

## Using Object Lifecycle Management

Object Lifecycle Management lets you automatically manage the archiving and deletion of objects. By using Object Lifecycle Management to manage your [Object Storage](#) and [Archive Storage](#) data, you can reduce your storage costs and the amount of time you spend managing data.

Object Lifecycle Management works by defining rules that instruct Object Storage to archive or delete objects on your behalf within a given bucket. A bucket's lifecycle rules are collectively known as an object lifecycle policy. For example, you could have Object Storage automatically move objects to Archive Storage 30 days after creation, and then automatically delete the archived objects 120 days after creation.

Each Object Storage or Archive Storage bucket can have a single lifecycle policy consisting of up to 1,000 rules. Rules can have object name [prefix](#) and [pattern matching](#) conditions. You can create, edit, delete, enable, and disable individual rules in the Console as needed. To update a lifecycle policy using the CLI or API, overwrite the entire policy with a new policy that is inclusive of all the policy rules that you want to apply to the bucket.

### Required IAM Policy



#### Important

You cannot use Object Lifecycle Management until you authorize the Object Storage service to archive and delete objects on your behalf. See [Service Permissions](#) for more information.

### User Permissions

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

The policy [Let Object Storage admins manage buckets and objects](#) lets the specified group do everything with buckets and objects, including adding and managing lifecycle policies. See [Details for Object Storage, Archive Storage, and Data Transfer](#) for more information on Object Storage user permissions.

### Service Permissions

To execute object lifecycle policies, you must authorize the service to archive and delete objects on your behalf. To do so, create the following policy in the root compartment of your tenancy:

```
Allow service objectstorage-<region_name> to manage object-family in compartment <compartment_name>
```

Because Object Storage is a regional service, you must authorize the Object Storage service in each region you use lifecycle policies. Object Storage ensures that your data is not read from any unauthorized region.

## CHAPTER 25 Object Storage

---

If you don't have permissions to write policies for the root compartment of your tenancy, contact your Oracle Cloud Infrastructure administrator. To determine the region name value of an Oracle Cloud Infrastructure region, see [Regions and Availability Domains](#).

If you want to grant individual permissions to the service rather than use the [policy verb](#) `manage`, you can use the following syntax:

```
Allow service objectstorage-<region_name> to {BUCKET_INSPECT, BUCKET_READ, OBJECT_INSPECT, OBJECT_CREATE, OBJECT_DELETE} in compartment <compartment_name>
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

### Options

When creating object lifecycle policy rules, you have the following options:

- When a lifecycle rule is created, the system generates a default name for that rule, for example **lifecycle-rule-20190321-1559**. This rule name identifies the current year, month, day, and time that the rule was created. You can use that system-generated name for your new rule or you can specify a different name for it.
- You can use a rule to either archive or delete objects and specify the number of days until the specified action is taken.
- You can apply a rule to all objects in a bucket. Alternatively, you can use object name filters to specify which objects the lifecycle rule applies to. You can select objects using both object name prefixes and pattern matching. See [Using Object Name Filters](#) for details.
- You can decide whether a new rule is enabled or disabled upon creation.

### Using Object Name Filters

Use object name filters to specify which objects the lifecycle rule applies to.



### Important

If you want the rule to apply to all objects in the bucket, do not specify any object name filters.

You can add object filters in any order. Object Lifecycle Management evaluates the precedence of the rules as follows:

1. Pattern exclusions
2. Pattern inclusions
3. Prefix inclusions

### Using Prefix Matching to Filter Objects

When naming objects, you can use prefix strings without a delimiter so that certain bulk operations can be performed by matching on the prefix portion of the object name. For example, in the object names below, the string `gloves_27_` serves as a prefix for matching purposes when performing lifecycle management archive or deletions:

```
gloves_27_dark_green.jpg
gloves_27_light_blue.jpg
gloves_27_deep_purple.jpg
gloves_27_bright_orange.jpg
```

See [Object Naming Using Prefixes and Hierarchies](#) for complete object naming details.

### Using Pattern Matching to Filter Objects

Object Storage supports the following pattern matching characters to either include or exclude objects:

## CHAPTER 25 Object Storage

Character	Description	Pattern Examples	Matches	Doesn't Match
*	Matches 0 or more characters	*.tmp	foo.tmp foo/bar/baz.tmp	tmp Atmp
		*.xls	.xls /home/user/file.xlsx	xls .xl
		/archive/*	/archive/sub/dir/ /archive/1/2/3/4/foo.txt	/src/archive/a archive/b
?	Matches any one character	X?Z	XyZ X_Z	XZ XYYZ
\	Escapes the next character	\\dir\\sub\\*	\\dir\\sub\\ABC \\dir\\sub\\	dir\\sub\\abc dirsub

Character	Description	Pattern Examples	Matches	Doesn't Match
[...]	<p>Matches a group of characters, which can be:</p> <ul style="list-style-type: none"> <li>• A set of characters, for example: [Zafg9@]. Matches any character in the brackets.</li> <li>• A range of characters, for example: [a-f]. Matches any character in the range: <ul style="list-style-type: none"> <li>◦ [a-f] is equivalent to [abcdef].</li> <li>◦ For character ranges only the CHARACTER R-CHARACTER pattern is</li> </ul> </li> </ul>	[-ab3]	- a b 3	-a -ab 3b

Character	Description	Pattern Examples	Matches	Doesn't Match
	supported:			
	<ul style="list-style-type: none"> <li>▪ [ab- yz] is not valid</li> <li>▪ [a- mn- z] is not valid</li> </ul>	backup.tar.g z.[0-9]	backup.tar.gz.0 backup.tar.gz.5 backup.tar.gz.9	backup.tar.gz 10 backup.tar.gz
	<ul style="list-style-type: none"> <li>○ Character ranges cannot start with ^ or :</li> <li>○ To include a hyphen (-) in the range, make it the first or last character.</li> </ul>	page-[0-9]*	page-0 page-2 page-22 page-2X	page- page-A1
		\[a-z\]	[a-z]	a z [a-z

Patterns are limited to 1024 characters. The following are examples of invalid patterns:

- \
- [^a-z]

- [z-a]
- [:isalpha:]

### Scope and Constraints

Understand the following scope and constraints regarding object lifecycle policies:

- When you create a lifecycle policy for a bucket, Object Storage applies that policy to all objects that exist in the bucket unless you add object name filters.
- A rule that deletes an object always takes priority over a rule that would archive that same object.
- When creating a lifecycle policy rule that deletes objects from Archive Storage, Archive Storage has a [minimum retention requirement](#) of 90 days. Objects deleted from Archive Storage that have not met the 90-day retention minimum are billed for 90 days of storage.
- You can create up to 1,000 lifecycle rules per bucket.

### Working with Object Lifecycle Management Policies

You can create, delete, edit, or disable lifecycle policy rules using the Console, the [Command Line Interface \(CLI\)](#), an SDK, or the API.



#### Warning

Objects deleted on your behalf by lifecycle policies cannot be recovered. Be sure when creating and editing your lifecycle policies that you are not unintentionally deleting data you want to retain. Oracle recommends that you test your lifecycle policy on development data before using the policy in production.

### Using the Console

#### To create a lifecycle policy rule

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment containing bucket for which you want to create a lifecycle rule.
3. Click the bucket name.
4. Click **Lifecycle Policy Rules** under **Resources** to access the lifecycle policy rule list.
5. Click **Create Rule**.
6. Provide the following information:
  - **Name:** Required. The system generates a default rule name that reflects the current year, month, day, and time, for example **lifecycle-rule-20190321-1559**. If you change this default to any other rule name, use letters, numbers, dashes, underscores, and periods. Do not include any confidential information.
  - **Lifecycle Action:** Select rule type **Archive** or **Delete**.
  - **Number of Days:** The number of days until the specified action is taken.
7. Optionally, you can add one or more **Object Name Filters** to specify which objects the lifecycle rule applies to. You can choose objects using [prefixes](#) and [pattern matching](#). If no object name filters are specified, the rule applies to all objects in the bucket.

To create an object name filter:

  - a. Click **Add Filter**.
  - b. Select the **Filter Type**.
  - c. Enter the **Filter Value**.
  - d. Click **Add Another Filter** to add as many filters as you need for this rule.
8. Select whether the rule is enabled or disabled upon creation using the **State** selector.
9. Click **Create**.

### To edit a lifecycle policy rule

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment where the bucket is.
3. Click the bucket name.
4. Click **Lifecycle Policy Rules** under **Resources** to access the rule list.
5. For the rule you want to edit, click the Actions icon (three dots), and then click **Edit**.
6. In the **Edit Lifecycle Rule** dialog box, edit the following as needed for each rule you want to change:
  - **Name:** A user-friendly name for the rule. Avoid entering confidential information.
  - **Lifecycle Action:** Rule type **Archive** or **Delete**.
  - **Number of Days:** The number of days until the specified action is taken.
  - **Object Name Filters:** Edit, delete, or add a [prefix](#) or [pattern](#) filter.
7. Click **Save Changes**.

### To enable, disable, or delete a lifecycle policy rule

You can disable and enable a rule on demand using the Console. The system stops the execution of disabled or deleted rules immediately.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment where the bucket is.
3. Click the bucket name.
4. Click **Lifecycle Policy Rules** under **Resources** to access the rule list.
5. For the rule you want to manage, click the Actions icon (three dots), and then click one of the following:
  - **Enable** (only displays if the rule is disabled)
  - **Disable** (only displays if the rule is enabled)

- **Delete**

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To create or replace a lifecycle policy for a bucket

Open a command prompt and run `oci os object-lifecycle-policy put` to create or replace the object lifecycle policy for a bucket. To edit individual rules, replace the bucket's existing policy with a new version of the policy that includes the changes to your rules.

```
oci os object-lifecycle-policy put -ns <object_storage_namespace> -bn <bucket_name> --items <json_formatted_lifecycle_policy>
```



#### Tip

The `--items` option requires that you provide key-value pair input as valid formatted JSON. See [Passing Complex Input](#) and [Using a JSON File for Complex Input](#) for information on JSON formatting.

For example, the following lifecycle policy archives objects after 30 days and deletes them after 180 days:

```
oci os object-lifecycle-policy put -ns MyNamespace -bn MyBucket --items '[
{
 "action": "ARCHIVE",
 "is-enabled": true,
 "name": "ArchiveAfter30Days",
 "object-name-filter": {
 "exclusion-patterns": [
 "*.jpg"
],
 "inclusion-patterns": [
 "*.doc"
],
 "inclusion-prefixes": [
 "documents/"
]
 }
}]'
```

```
],
 "time-amount": 30,
 "time-unit": "DAYS"
 },
 {
 "action": "DELETE",
 "is-enabled": true,
 "name": "DeleteAfter180Days",
 "object-name-filter": {
 "exclusion-patterns": null,
 "inclusion-patterns": null,
 "inclusion-prefixes": null
 },
 "time-amount": 180,
 "time-unit": "DAYS"
 }
]
```

On Windows, to pass complex input to the CLI as a JSON string, you must enclose the entire block in double quotes. Inside the block, each double quote for the key and value strings must be escaped with a backslash (\) character.

For example:

```
oci os object-lifecycle-policy put -ns MyNamespace -bn MyBucket --items "[{"action":"ARCHIVE","is-enabled":true,"name":"Archive After 30 Days","object-name-filter":{"exclusion-patterns":["*.jpg"],"inclusion-patterns":["*.doc"],"inclusion-prefixes":["documents/"]},"time-amount":30,"time-unit":"DAYS"},{"action":"DELETE","is-enabled":true,"name":"DeleteAfter180Days","object-name-filter":{"exclusion-patterns":null,"inclusion-patterns":null,"inclusion-prefixes":null},"time-amount":180,"time-unit":"DAYS"}]"
```

### To delete a bucket's lifecycle policy

Open a command prompt and run `oci os object-lifecycle-policy delete` to delete a bucket's object lifecycle policy.

```
oci os object-lifecycle-policy delete -ns <object_storage_namespace> -bn <bucket_name>
```

### To get a bucket's lifecycle policy

Open a command prompt and run `oci os object-lifecycle-policy get` to get a bucket's object lifecycle policy.

```
oci os object-lifecycle-policy get -ns <object_storage_namespace> -bn <bucket_name>
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to manage object lifecycle policies:

- [PutObjectLifecyclePolicy](#)
- [GetObjectLifecyclePolicy](#)
- [DeleteObjectLifecyclePolicy](#)

### Object Storage Metrics

You can monitor the health, capacity, and performance of your buckets and objects by using [metrics](#), [alarms](#), and [notifications](#).

This topic describes the metrics emitted by the metric namespace `oci_objectstorage` (the Object Storage service).

Resources include buckets and objects.

### Overview of the Object Storage Service Metrics

Object Storage can store an unlimited amount of unstructured data of any content type, including analytic data and rich content, like images and videos. The Object Storage service metrics help you measure the amount of storage you're using. You can also use these metrics to monitor the performance of requests in terms of latency, and utilization as measured by counts of various types of requests made per bucket.

### Required IAM Policy

To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources

being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see the Authentication and Authorization section for the related service: [Monitoring](#) or [Notifications](#).

### Available Metrics: oci\_ objectstorage

The metrics listed in the following tables are automatically available for any buckets you create. You do not need to enable monitoring on the resource to get these metrics. However, you must have an object stored in a bucket to get *any* metrics. Buckets with no objects emit *no* metric data.

Each metric includes the following dimensions:

**RESOURCEID**

The OCID of the bucket to which the metric applies.

**RESOURCEDISPLAYNAME**

The name of the bucket.

**TIER**

The current storage tier of the object: standard or archive.

### Default Metrics

Default metrics are available on buckets in default charts.

Metric	Metric Display Name	Unit	Description	Dimensions
StoredBytes	<b>Bucket Size</b>	bytes	The size of the bucket, excluding any multipart upload parts that have not been discarded (aborted) or committed.	resourceID resourceDisplayName tier
ObjectCount	<b>Number of Objects</b>	count	The count of objects in the bucket, excluding any multipart upload parts that have not been discarded (aborted) or committed.	

### Custom Query Metrics

The following metrics are only available from a custom query. See [To view Object Storage metrics with a custom query](#).

## CHAPTER 25 Object Storage

---

Metric	Metric Display Name	Unit	Description	Dimensions
UncommittedParts	<b>Incomplete MultiPart Upload Size</b>	bytes	The size of any multipart upload parts that have not been discarded (aborted) or committed.	resourceID resourceDisplayName tier

Metric	Metric Display Name	Unit	Description	Dimensions
GetRequests	<b>GetObject Request Count</b>	count	The total number of GetObject requests made in a bucket.	resourceID resourceDisplayName
HeadRequests	<b>HeadObject Request Count</b>	count	The total number of HeadObject requests made in a bucket.	
DeleteRequests	<b>DeleteObject Request Count</b>	count	The total number of DeleteObject requests made in a bucket.	
PutRequests	<b>PutObject Request Count</b>	count	The total number of PutObject requests made in a bucket.	
ListRequests	<b>ListObjects Request Count</b>	count	The total number of ListObjects requests made in a bucket.	

## CHAPTER 25 Object Storage

Metric	Metric Display Name	Unit	Description	Dimensions
RenameRequests	<b>RenameObject Request Count</b>	count	The total number of RenameObject requests made in a bucket.	
PostRequests	<b>Post Object Request Count</b>	count	The total number of HTTP Post requests made in a bucket.	
ClientErrors	<b>Client Side Error Count</b>	count	The total number of 4xx errors for requests made in a bucket.	
TotalRequestLatency	<b>Overall Latency Time</b>	time (ms)	The per-request time from the first byte received by Object Storage to the last byte sent from Object Storage.	

## CHAPTER 25 Object Storage

---

Metric	Metric Display Name	Unit	Description	Dimensions
FirstByteLatency	<b>First Byte Latency Time</b>	time (ms)	The per-request time measured from the time Object Storage receives the complete request to when Object Storage returns the first byte of the response.	
AllRequests	<b>All Request Count</b>	count	The total number of all HTTP requests made in a bucket.	

Metric	Metric Display Name	Unit	Description	Dimensions
CopyRequests	<b>Copy Object Request Count</b>	count	The total number of CopyObject requests made in a bucket.	
ArchiveRequests	<b>Archive Object Request Count</b>	count	The total number of ArchiveObject requests made in a bucket. To archive objects, you must configure an Object Lifecycle policy. See <a href="#">Using Object Lifecycle Management</a> .	

## Using the Console

### To view default metric charts for a bucket

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the **Compartment** that contains the bucket you want to view, and then click the bucket's name.

3. In the **Resources** menu, click **Metrics**.

The **Metrics** page displays a default set of charts for the current bucket.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

### To view default metric charts by dimension

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Service Metrics**.
2. For **Metric Namespace**, select **oci\_objectstorage**.
3. For **Dimensions**, click **Add**.
4. For **Dimension Name**, select a dimension and then select a **Dimension Value**. Add more dimensions as needed.
5. Click **Done**.

The **Service Metrics** page displays a default set of charts for the selected metric namespace and dimension. For more information about the emitted metrics, see the foregoing table. You can also use the Monitoring service to create [custom queries](#).

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

### To view Object Storage metrics with a custom query

1. Open the navigation menu. Under **Solutions and Platform**, go to **Monitoring** and click **Metrics Explorer**.  
The **Metrics Explorer** page displays an empty chart with fields to build a query.
2. Select a compartment.
3. From **Metric Namespace**, select **oci\_objectstorage**.
4. From **Metric Name**, select a metric.

5. (Optional) Refine your query.

For instructions, see [To create a query](#).

6. Click **Update Chart**.

The chart shows the results of your new query. You can optionally add more queries by clicking **Add Query** below the chart.

For more information about monitoring metrics and using alarms, see [Monitoring Overview](#).

For information about notifications for alarms, see [Notifications Overview](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following APIs for monitoring:

- [Monitoring API](#) for metrics and alarms
- [Notifications API](#) for notifications (used with alarms)

### Hadoop Support

Using the HDFS connector, you can run Hadoop or Spark jobs against data in the Oracle Cloud Infrastructure Object Storage service. The connector has the following features:

- Supports read and write data stored in Object Storage
- Is compatible with existing buckets of data
- Is compatible with Hadoop 2.7.2

For information about downloading, configuring, and using the HDFS connector, see [HDFS Connector for Object Storage](#).

## Designating Compartments for the Amazon S3 Compatibility and Swift APIs

In the Oracle Cloud Infrastructure Object Storage service, a bucket is a container for storing objects in a compartment within an Object Storage [namespace](#). A bucket is associated with a single compartment and data is stored as objects in buckets.

In addition to the native Object Storage APIs, Object Storage provides API support for both Amazon S3 Compatibility API and Swift API. However these APIs do not understand the Oracle Cloud Infrastructure concept of a compartment. By default, buckets created using the Amazon S3 Compatibility API or the Swift API are created in the root compartment of the Oracle Cloud Infrastructure tenancy. Instead, you can [designate a different compartment](#) for the Amazon S3 Compatibility API or Swift API to create buckets in.

When you designate a different compartment to use for the Amazon S3 Compatibility API or Swift API, any new buckets you create using the Amazon S3 Compatibility API or the Swift API are created in this newly designated compartment. Buckets previously created in a different compartment are not automatically moved to the newly designated compartment. See [Managing Buckets](#) if you want to move previously created buckets to this newly designated compartment.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

Compartments have policies that indicate what actions a user can perform on a bucket and all the objects in the bucket.

For administrators:

- To change the default compartments for Amazon S3 Compatibility API and Swift API, a user must belong to a group with `NAMESPACE_UPDATE` permissions.

- To see the current default compartments for Amazon S3 Compatibility API and Swift API, a user must belong to a group with `NAMESPACE_READ` permissions.
- To move a bucket to a different compartment, a user must belong to a group with `BUCKET_UPDATE` and `BUCKET_CREATE` permissions in the source compartment, and `BUCKET_CREATE` permissions in the target compartment.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for buckets and objects, see [Details for Object Storage, Archive Storage, and Data Transfer](#).

### Viewing and Specifying Designated Compartments

You can view the current default compartment designations for Amazon S3 Compatibility API and Swift API data. If your permissions allow, you can also change the Amazon S3 Compatibility API and Swift API compartment designations.

Designated compartment names:

- Must be unique across all the compartments in your tenancy.
- Can be from 1 to 100 characters in length.
- Must not contain confidential information.
- Valid are letters (upper or lower case), numbers, hyphens, and underscore.

### Using the Console

To view your Amazon S3 Compatibility API and Swift API compartment designations

Open the **Profile** menu () and click **Tenancy:** *<your\_tenancy\_name>*.

Your default compartment designations for the APIs are listed under **Object Storage Settings**.

### To edit your tenancy's Amazon S3 Compatibility API and Swift API compartment designations

1. Open the **Profile** menu (👤) and click **Tenancy:** *<your\_tenancy\_name>*.
2. Click **Edit Object Storage Settings**.
3. In the **Edit Object Storage Settings** dialog:
  - Select the compartment that you want for the **Amazon S3 Compatibility API Designated Compartment** from the drop-down menu.
  - Select the compartment that you want for the **Swift API Designated Compartment** from the drop-down menu.
4. Click **Save**.  
The new **Object Storage Settings** are displayed.

### Using the Command Line Interface (CLI)

For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To get your tenancy's Amazon S3 Compatibility API and Swift API compartment designations

Use this CLI command to display metadata associated with the Amazon S3 and Swift compartments for the specified namespace in your tenancy.

```
oci os ns get-metadata --namespace <object_storage_namespace>
```

For example:

```
oci os ns get-metadata --namespace MyNamespace
{
 "data": {
 "default-s3-compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "default-swift-compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "namespace": "MyNamespace"
 }
}
```

## CHAPTER 25 Object Storage

---

```
}
}
```

### To update your tenancy's Amazon S3 Compatibility API compartment designation

Use this CLI command to specify the default Amazon S3 compartment for the specified namespace in your tenancy.

```
oci os ns update-metadata --namespace <object_storage_namespace> --default-s3-compartment-id <your_oci_compartment_id>
```

*<your\_oci\_compartment\_id>* specifies a compartment that is not the root compartment of your tenancy.

For example:

```
oci os ns update-metadata --namespace MyNamespace --default-s3-compartment-id
ocid.compartment.oc1..exampleuniqueID
{
 "data": {
 "default-s3-compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "default-swift-compartment-id": null,
 "namespace": null
 }
}
```

### To update your tenancy's Swift API compartment designations

Use this CLI command to specify the default Swift compartment for the specified namespace in your tenancy.

```
oci os ns update-metadata --namespace <object_storage_namespace> --default-swift-compartment-id <your_oci_compartment_id>
```

*<your\_oci\_compartment\_id>* specifies a compartment that is not the root compartment of your tenancy.

For example:

## CHAPTER 25 Object Storage

---

```
oci os ns update-metadata --namespace MyNamespace --default-swift-compartment-id
ocid.compartment.oc1..exampleuniqueID
{
 "data": {
 "default-s3-compartment-id": null,
 "default-swift-compartment-id": "ocid.compartment.oc1..exampleuniqueID",
 "namespace": null
 }
}
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operation to get your default Amazon S3 Compatibility API and Swift API compartment designations, and change those compartment designations:

- [GetNamespaceMetadata](#)
- [UpdateNamespaceMetadata](#)

### Amazon S3 Compatibility API

Using the Amazon S3 Compatibility API, customers can continue to use their existing Amazon S3 tools (for example, SDK clients) and partners can make minimal changes to their applications to work with Object Storage. The Amazon S3 Compatibility API and Object Storage datasets are congruent. If data is written to the Object Storage using the Amazon S3 Compatibility API, the data can be read back using the native Object Storage API and conversely.

### Differences between the Object Storage API and the Amazon S3 Compatibility API

The Object Storage Service provided by Oracle Cloud Infrastructure and Amazon S3 use similar concepts and terminology. In both cases, data is stored as objects in buckets. The differences are in the implementation of features and tools for working with objects.

The following highlights the differences between the two storage technologies:

- **Compartments**

Amazon S3 doesn't use compartments. By default, buckets created using the Amazon S3 Compatibility API or the Swift API are created in the root compartment of the Oracle Cloud Infrastructure tenancy. Instead, you can [designate a different compartment](#) for the Amazon S3 Compatibility API or Swift API to create buckets in.

- **Global bucket namespace**

Object Storage doesn't use a global bucket namespace. Each tenant is associated with one default namespace that spans all compartments within a region. The namespace serves as a container for all of your buckets and objects. You control bucket names within your namespace, however, bucket names must be unique within each region. You can have a bucket named **MyBucket** in US West (Phoenix) and a bucket named **MyBucket** in Germany Central (Frankfurt).

- **Encryption**

The Oracle Cloud Infrastructure Object Storage service encrypts all data at rest by default. Encryption can't be turned on or off using the API.

- **Object Level Access Control Lists (ACLs)**

Oracle Cloud Infrastructure does not use ACLs for objects. Instead, IAM policies are used to manage access to compartments, buckets, and objects.

For more information, see [Overview of the Object Storage service](#).

### Amazon S3 Compatibility API Prerequisites

To enable application access from Amazon S3 to Object Storage, you need to set up access to

Oracle Cloud Infrastructure and modify your application.

### Setting up access to Oracle Cloud Infrastructure:

- [Sign Up for Oracle Cloud Infrastructure](#) and obtain a unique namespace.
- Create an [Amazon S3 Compatibility API key](#). An Amazon S3 Compatibility API key consists of an Access Key/Secret key pair.

### Modifying your application:

- Configure a new endpoint for the application that includes . For example:  
`mynamespace.compat.objectstorage.us-phoenix-1.oraclecloud.com`.
- Set the target region as one of the Oracle Cloud Infrastructure regions.



#### Important

If your application does not support setting the region name to the correct Oracle Cloud Infrastructure region name, you must either set the region to `us-east-1` or leave it blank. Using this configuration, you can only use the Amazon S3 Compatibility API in your Oracle Cloud Infrastructure home region.

If you can manually set the region, you can use the application against any Oracle Cloud Infrastructure region.

- Configure the application to use the [Amazon S3 Compatibility API key](#).
- The application must use path -based access. Virtual host-style access (accessing a bucket as `bucketname.namespace.compat.objectstorage.region.oraclecloud.com`) is not supported.

You can now use the Amazon S3 Compatibility API to access Object Storage in Oracle Cloud Infrastructure.

### Amazon S3 Compatibility API Support

Amazon S3 Compatibility API support is provided at the bucket level and object level.

#### Bucket APIs

The following bucket APIs are supported:

- [DeleteBucket](#)
- [GetLocation](#)
- [HeadBucket](#)
- [GetService](#) (list all my buckets)
- [ListObjects](#)
- [PutBucket](#)

#### Object APIs

The following object APIs are supported:

- [BulkDelete](#)
- [DeleteObject](#)
- [GetObject](#)
- [HeadObject](#)
- [PutObject](#)
- [RestoreObjects](#)

#### Multipart Upload APIs

The following multipart upload APIs are supported:

- [AbortMultipartUpload](#)
- [CompleteMultipartUpload](#)
- [InitiateMultipartUpload](#)
- [ListParts](#)
- [ListUploads](#)
- [UploadPart](#)

### Tagging APIs

The following tagging APIs are supported:

- [DeleteBucketTagging](#)
- [GetBucketTagging](#)
- [PutBucketTagging](#)

### Supported Amazon S3 Clients

Here are some examples of configuring various client applications to talk to Object Storage's Amazon S3-compatible endpoints.

#### AWS SDK for Java

The following is an example of configuring [AWS SDK for Java](#).

```
// Get S3 credentials from the console and put them here
AWSCredentialsProvider credentials = new AWSStaticCredentialsProvider(new BasicAWSCredentials(
 "ocid1.credential.oc1..anEXAMPLE",
 "anEXAMPLE="));

// The name of your tenancy
String tenancy = "tenancy";

// The region to connect to
String region = "us-ashburn-1";

// Create an S3 client pointing at the region
String endpoint = String.format("%s.compat.objectstorage.%s.oraclecloud.com", tenancy, region);
```

## CHAPTER 25 Object Storage

---

```
AwsClientBuilder.EndpointConfiguration endpointConfiguration = new
AwsClientBuilder.EndpointConfiguration(endpoint, region);
AmazonS3 client = AmazonS3Client.builder()
 .standard()
 .withCredentials(credentials)
 .withEndpointConfiguration(endpointConfiguration)
 .disableChunkedEncoding()
 .enablePathStyleAccess()
 .build();
```

### AWS SDK for Javascript

The following is an example of configuring [AWS SDK for Javascript](#).

```
s3 = new AWS.S3({
 region: 'us-ashburn-1',
 endpoint: 'https://' + namespace + '.compat.objectstorage.us-ashburn-1.oraclecloud.com',
 accessKeyId: 'ocidl.credential.ocl..something',
 secretAccessKey: 'something=',
 s3ForcePathStyle: true,
 signatureVersion: 'v4',
});
```

### AWS SDK for Python (Boto 3)

The following is an example of configuring [AWS SDK for Python \(Boto 3\)](#).

```
import boto3

s3 = boto3.resource(
 's3',
 aws_access_key_id="ocidl.credential.ocl..something", # Put your ocid for the secret key here
 aws_secret_access_key="something=", # Put your secret key here
 region_name="us-phoenix-1", # Set the region to match the endpoint
 endpoint_url="https://namespace.compat.objectstorage.us-phoenix-1.oraclecloud.com" # Endpoint url
 has namespace ahead of compat
)

Print out bucket names
for bucket in s3.buckets.all():
 print bucket.name
```

### AWS Command Line Interface (CLI)

The following is an example of configuring [AWS Command Line Interface \(CLI\)](#).

```
pip3 install awscli --upgrade --user

Configure AWS Cli
aws configure --profile mynamespace
AWS Access Key ID [None]: ocid1.credential.oc1..something
AWS Secret Access Key [None]: secretkey=
Default region name [None]: us-phoenix-1
Default output format [None]:
Install this plugin to override the url or you need to pass it on command line
pip3 install awscli-plugin-endpoint --upgrade --user

Enable the plugin or it won't work
aws configure set plugins.endpoint_url awscli_plugin_endpoint

Let the plugin set the endpoint URL
aws configure --profile mynamespace set s3.endpoint_url https://mynamespace.compat.objectstorage.us-phoenix-1.oraclecloud.com

Test out the command!
aws s3 ls --profile mynamespace
```

# CHAPTER 26 Registry

This chapter explains how to store, share, and manage development artifacts like Docker images in an Oracle-managed registry.

## Overview of Registry

Oracle Cloud Infrastructure Registry is an Oracle-managed registry that enables you to simplify your development to production workflow. Oracle Cloud Infrastructure Registry makes it easy for you as a developer to store, share, and manage development artifacts like Docker images. And the highly available and scalable architecture of Oracle Cloud Infrastructure ensures you can reliably deploy your applications. So you don't have to worry about operational issues, or scaling the underlying infrastructure.

You can use Oracle Cloud Infrastructure Registry as a private Docker registry for internal use, pushing and pulling Docker images to and from the Registry using the [Docker V2 API](#) and the standard Docker command line interface (CLI). You can also use Oracle Cloud Infrastructure Registry as a public Docker registry, enabling any user with internet access and knowledge of the appropriate URL to pull images from public repositories in Oracle Cloud Infrastructure Registry.

Oracle Cloud Infrastructure Registry supports private access from other Oracle Cloud Infrastructure resources in a virtual cloud network (VCN) in the same region through a service gateway. Setting up and using a service gateway on a VCN lets resources (such as worker nodes in clusters managed by Container Engine for Kubernetes) access Oracle Cloud Infrastructure services such as Oracle Cloud Infrastructure Registry without exposing them to the public internet. No internet gateway is required and resources can be in a private subnet and use only private IP addresses. For more information, see [Access to Oracle Services: Service Gateway](#).

Oracle Cloud Infrastructure Registry is integrated with IAM, which provides easy authentication with native Oracle Cloud Infrastructure identity.

For an introductory tutorial, see [Pushing an Image to Oracle Cloud Infrastructure Registry](#).



### Note

Registry is not available in Oracle Cloud Infrastructure Government Cloudrealms.

## Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

Note that Oracle Cloud Infrastructure Registry fully implements a Docker protocol that enables you to use the Docker Registry HTTP API (rather than the Oracle Cloud Infrastructure API) to manage images. See the [Docker documentation](#) for information about using the Docker Registry HTTP API.

## Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

### Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Registry Capabilities and Limits

In each region that is enabled for your tenancy, you can create up to 500 repositories in Oracle Cloud Infrastructure Registry. Each repository can hold up to 500 images. See [Service Limits](#).

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#).

For more details about policies for Oracle Cloud Infrastructure Registry, see:

- [Policies to Control Repository Access](#)
- [Details for Registry](#)

## Preparing for Registry

Before you can push and pull Docker images to and from Oracle Cloud Infrastructure Registry:

- You must have access to an Oracle Cloud Infrastructure tenancy. The tenancy must be subscribed to one or more of the regions in which Registry is available (see [Availability by Region Name and Region Code](#)).
- You must have access to the Docker CLI (for example, to push and pull images on a local machine, you'll need to have installed Docker on the local machine).
- You must either belong to a group to which a policy grants the appropriate permissions, or belong to the tenancy's Administrators group. See [Policies to Control Repository Access](#).
- You must have an Oracle Cloud Infrastructure auth token. If you don't have an auth token already, see [Getting an Auth Token](#).

## Availability by Region Name and Region Code

Registry is available in the following regions. Note that you have to use the region code in some commands. In some cases, you might have to use shortened versions of availability domain names.

## CHAPTER 26 Registry

Region Name	Region Code	Shortened Availability Domain Names
US East (Ashburn)	iad	<ul style="list-style-type: none"><li>• US-ASHBURN-AD-1</li><li>• US-ASHBURN-AD-2</li><li>• US-ASHBURN-AD-3</li></ul>
Germany Central (Frankfurt)	fra	<ul style="list-style-type: none"><li>• EU-FRANKFURT-1-AD-1</li><li>• EU-FRANKFURT-1-AD-2</li><li>• EU-FRANKFURT-1-AD-3</li></ul>
UK South (London)	lhr	<ul style="list-style-type: none"><li>• UK-LONDON-1-AD-1</li><li>• UK-LONDON-1-AD-2</li><li>• UK-LONDON-1-AD-3</li></ul>
India West (Mumbai)	bom	<ul style="list-style-type: none"><li>• AP-MUMBAI-1-AD-1</li></ul>
US West (Phoenix)	phx	<ul style="list-style-type: none"><li>• PHX-AD-1</li><li>• PHX-AD-2</li><li>• PHX-AD-3</li></ul>
Brazil East (Sao Paulo)	gru	<ul style="list-style-type: none"><li>• SA-SAOPAULO-1-AD-1</li></ul>
South Korea Central (Seoul)	icn	<ul style="list-style-type: none"><li>• AP-SEOUL-1-AD-1</li></ul>
Australia East (Sydney)	syd	<ul style="list-style-type: none"><li>• AP-SYDNEY-1-AD-1</li></ul>
Japan East (Tokyo)	nrt	<ul style="list-style-type: none"><li>• AP-TOKYO-1-AD-1</li></ul>
Canada Southeast (Toronto)	yyz	<ul style="list-style-type: none"><li>• CA-TORONTO-1-AD-1</li></ul>
Switzerland North (Zurich)	zrh	<ul style="list-style-type: none"><li>• EU-ZURICH-1-AD-1</li></ul>

### About Images

You can store, share, and manage Docker images in Oracle Cloud Infrastructure Registry. A Docker image is a read-only template with instructions for creating a Docker container. A Docker image holds the application that you want Docker to run as a container, along with any dependencies. To create a Docker image, you first create a Dockerfile to describe that application. You then build the Docker image from the Dockerfile. Having created a Docker image, you store it in a Docker registry such as Oracle Cloud Infrastructure Registry.

### About Repositories

Related images in Oracle Cloud Infrastructure Registry can be grouped into meaningfully named repositories for convenience.

Repositories can be private or public. Any user with internet access and knowledge of the appropriate URL can pull images from a public repository in Oracle Cloud Infrastructure Registry.

A repository exists within a particular region and tenancy. When referring to the tenancy that owns a repository, you specify the tenancy's namespace. The tenancy namespace is an auto-generated random string of alphanumeric characters. For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`). To find out the tenancy namespace of the current tenancy, open the **Profile** menu () and click **Tenancy:**.

You must belong to the tenancy's Administrators group or have been granted the `REPOSITORY_MANAGE` permission to:

- create a new public repository
- change an existing repository into a public repository
- change an existing public repository into a private repository

If you make a repository private, you (along with users belonging to the tenancy's Administrators group) will be able to perform any operation on the repository. You can use

identity policies to allow other users to perform other operations on repositories (both public and private) that you create.

Typically, the images in a repository are all different versions of the same source image (for example 'acme-web-app'), with each version identified by a tag (for example, 'acme-web-app:4.6.3').

For example, for convenience you might want to group together multiple versions of the acme-web-app image in the acme-dev tenancy in the Ashburn region into a repository called project01. You do this by including the name of the repository in the image name when you push the image, in the format `<region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>`. For example, `iad.ocir.io/ansh81vrulzp/project01/acme-web-app:4.6.3`. Subsequently, when you use the `docker push` command, the presence of the repository in the image's name ensures the image is pushed to the intended repository.

If you push an image and include the name of a repository that doesn't already exist, a new private repository is created automatically. For example, if you enter a command like `docker push iad.ocir.io/ansh81vrulzp/project02/acme-web-app:7.5.2` and the `project02` repository doesn't exist, a private repository called `project02` is created automatically.

If you push an image and don't include a repository name, the image's name is used as the name of the repository. For example, if you enter a command like `docker push iad.ocir.io/ansh81vrulzp/acme-web-app:7.5.2` that doesn't contain a repository name, the image's name (`acme-web-app`) is used as the name of a private repository.

Alternatively, you can use the Console to create an empty repository and give it a name. If you belong to the tenancy's Administrators group or have been granted the `REPOSITORY_MANAGE` permission, you can also specify whether the repository is to be private or public. Any images you subsequently push to Oracle Cloud Infrastructure Registry that include the repository in the image name are pushed to that repository.

## Creating a Repository

Using the Console, you can create an empty repository in Oracle Cloud Infrastructure Registry and give it a name. Any images you subsequently push to the registry that include the repository in the image name are grouped into that repository.

## CHAPTER 26 Registry

---

Having created the new repository, you can push an image to the repository using the Docker CLI (see [Pushing Images Using the Docker CLI](#)).

Note that although creating an empty repository in advance can be a convenient placeholder, it is not strictly necessary. When you push an image, you use a command in the format `docker push <region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>`. However:

- If you push an image and the command includes the name of a repository that doesn't already exist, a new private repository is created automatically. For example, if you enter a command like `docker push iad.ocir.io/ansh81vrulzp/project02/acme-web-app:7.5.2` and the `project02` repository doesn't exist, a private repository called `project02` is created automatically.
- If you push an image and the command doesn't include a repository name, the image's name is used as the name of the repository. For example, if you enter a command like `docker push iad.ocir.io/ansh81vrulzp/acme-web-app:7.5.2` that doesn't contain a repository name, the image's name (`acme-web-app`) is used as the name of a private repository.

## Using the Console

To create a repository in Oracle Cloud Infrastructure Registry:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the region in which to create the repository.
3. Click **Create Repository**.
4. In the **Add Repository** dialog box, specify details for the new repository:
  - **Repository Name:** A name of your choice for the new repository. Avoid entering confidential information.
  - **Public:** Whether the new repository will be a public repository or a private repository. You can only make the new repository public if you belong to the tenancy's Administrators group or have been granted the `REPOSITORY_MANAGE`

permission. If you make the new repository public, any user with internet access and knowledge of the appropriate URL will be able to pull images from the repository. If you make the repository private, you (along with users belonging to the tenancy's Administrators group) will be able to perform any operation on the repository.

5. Click **Submit**.

## Pushing Images Using the Docker CLI

You use the Docker CLI to push images to Oracle Cloud Infrastructure Registry.

To push an image, you first use the `docker tag` command to create a copy of the local source image as a new image (the new image is actually just a reference to the existing source image). As a name for the new image, you specify the fully qualified path to the target location in Oracle Cloud Registry where you want to push the image, optionally including the name of a repository.

For example, assume you have a local image named `acme-web-app:latest`. Let's say you want to push this image to Oracle Cloud Infrastructure Registry with a name of `acme-web-app:version2.0.test` into a repository called `project01` in the Ashburn region of the `acme-dev` tenancy. When you use the `docker tag` command, you'd name the new image with the fully qualified path to its destination, in the format `<region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>`. So in this case, you'd name the new image `iad.ocir.io/ansh81vrulzp/project01/acme-web-app:version2.0.test`. Subsequently, when you use the `docker push` command, the image's name ensures it is pushed to the correct destination.

Your permissions control the images you can push to Oracle Cloud Infrastructure Registry. You can push images to repositories you've created, and to repositories that the groups to which you belong have been granted access by appropriate identity policies. If you belong to the Administrators group, you can push images to any repository in the tenancy.

To push images to Oracle Cloud Infrastructure Registry using the Docker CLI:

1. If you already have an auth token, go to the next step. Otherwise:
  - a. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
  - b. On the **Auth Tokens** page, click **Generate Token**.
  - c. Enter a friendly description for the auth token. Avoid entering confidential information.
  - d. Click **Generate Token**. The new auth token is displayed.
  - e. Copy the auth token immediately to a secure location from where you can retrieve it later, because you won't see the auth token again in the Console.
  - f. Close the Generate Token dialog.
2. In a terminal window on the client machine running Docker, log in to Oracle Cloud Infrastructure Registry by entering `docker login <region-code>.ocir.io`, where `<region-code>` corresponds to the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `docker login iad.ocir.io`. See [Availability by Region Name and Region Code](#) for the list of region codes.
3. When prompted, enter your username in the format `<tenancy-namespace>/<username>`, where `<tenancy-namespace>` is the auto-generated Object Storage namespace string of your tenancy (as shown on the **Tenancy Information** page). For example, `ansh81vrulzp/jdoe@acme.com`. If your tenancy is federated with Oracle Identity Cloud Service, use the format `<tenancy-namespace>/oracleidentitycloudservice/<username>`.
4. When prompted, enter the auth token you copied earlier.
5. Locate the image on the client machine that you want to push:
  - a. In a terminal window on your client machine, enter `docker images` to list the available images.

For example:

```
$ docker images
```

## CHAPTER 26 Registry

REPOSITORY	TAG	IMAGE ID	CREATED	SIZE
acme-web-app	latest	8e0506e14874	2 hours ago	162.6 MB
acme-web-app	version1.0	7d9495d03763	2 hours ago	162.6 MB
<none>	<none>	6ebd328f833d	5 hours ago	162.6 MB
hello-world	latest	80b84820d442	5 weeks ago	890 B

- b. Find the image on the client machine that you want to push to Oracle Cloud Infrastructure Registry.

In the output of the `docker images` command, look for the specific image that you want to push. You'll need to uniquely identify this image later, in one of the following ways:

- using its id
- using its name and tag, separated by a colon

For example, you might have an image named `acme-web-app` on the client machine. In the output of the `docker images` command, look for the specific `acme-web-app` image that you want to push. You can uniquely identify that particular image in one of the following ways:

- using its id (for example, `8e0506e14874`)
- using its name and tag, separated by a colon (for example `acme-web-app:latest`)

- c. Give a tag to the image that you're going to push to Oracle Cloud Infrastructure Registry by entering:

```
docker tag <image-identifier> <target-tag>
```

where:

- `<image-identifier>` uniquely identifies the image, either using the image's id (for example, `8e0506e14874`), or the image's name and tag separated by a colon (for example, `acme-web-app:latest`).

## CHAPTER 26 Registry

- `<target-tag>` is in the format `<region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>` where:
  - `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `iad`. See [Availability by Region Name and Region Code](#) for the list of region codes.
  - `ocir.io` is the Oracle Cloud Infrastructure Registry name.
  - `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy that owns the repository to which you want to push the image (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`). Note also that your user must have access to the tenancy.
  - `<repo-name>` (if specified) is the name of a repository to which you want to push the image (for example, `project01`). Note that specifying a repository is optional (see [About Repositories](#)).
  - `<image-name>` is the name you want to give the image in Oracle Cloud Infrastructure Registry (for example, `acme-web-app`).
  - `<tag>` is an image tag you want to give the image in Oracle Cloud Infrastructure Registry (for example, `version2.0.test`).

For example, combining the previous examples, you might enter:

```
docker tag 8e0506e14874 iad.ocir.io/ansh81vrulzp/project01/acme-web-app:version2.0.test
```

6. Confirm that the Docker image has been correctly tagged on the client machine by entering `docker images` and verifying that the list of images includes an image with the tag you specified.

For example:

```
$ docker images
REPOSITORY TAG IMAGE ID CREATED
SIZE
```

## CHAPTER 26 Registry

```
iad.ocir.io/ansh81vrulzp/project01/acme-web-app version2.0.test 8e0506e14874 1 minute ago
162.6 MB
acme-web-app latest 8e0506e14874 2 hours ago
162.6 MB
acme-web-app version1.0 7d9495d03763 2 hours ago
162.6 MB
<none> <none> 6ebd328f833d 5 hours ago
162.6 MB
hello-world latest 80b84820d442 5 weeks ago
890 B
```

7. Push the Docker image from the client machine to Oracle Cloud Infrastructure Registry by entering:

```
docker push <target-tag>
```

where `<target-tag>` is in the format `<region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>` where:

- `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `iad`. See [Availability by Region Name and Region Code](#) for the list of region codes.
- `ocir.io` is the Oracle Cloud Infrastructure Registry name.
- `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy that owns the repository to which you want to push the image (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`). Note also that your user must have access to the tenancy.
- `<repo-name>` (if specified) is the name of a repository to which you want to push the image (for example, `project01`). Note that specifying a repository is optional (see [About Repositories](#)).
- `<image-name>` is the name you want to give the image in Oracle Cloud Infrastructure Registry (for example, `acme-web-app`).

- `<tag>` is an image tag you want to give the image in Oracle Cloud Infrastructure Registry (for example, `version2.0.test`).

For example:

```
docker push iad.ocir.io/ansh81vrulzp/project01/acme-web-app:version2.0.test
```

## Pulling Images Using the Docker CLI

You use the Docker CLI to pull images from Oracle Cloud Infrastructure Registry.

Your permissions control the images you can pull from Oracle Cloud Infrastructure Registry. You can pull images from repositories you've created, from public repositories, and from repositories that the groups to which you belong have been granted access by identity policies. If you belong to the Administrators group, you can pull images from any repository in the tenancy.

To pull images from Oracle Cloud Infrastructure Registry using the Docker CLI:

1. If you already have an auth token, go to the next step. Otherwise:
  - a. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
  - b. On the **Auth Tokens** page, click **Generate Token**.
  - c. Enter a friendly description for the auth token. Avoid entering confidential information.
  - d. Click **Generate Token**. The new auth token is displayed.
  - e. Copy the auth token immediately to a secure location from where you can retrieve it later, because you won't see the auth token again in the Console.
  - f. Close the Generate Token dialog.
2. In a terminal window on the client machine running Docker, log in to Oracle Cloud Infrastructure Registry by entering `docker login <region-code>.ocir.io`, where `<region-code>` corresponds to the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `docker login iad.ocir.io`. See [Availability by Region Name and Region Code](#) for the list of region codes.

3. When prompted, enter your username in the format `<tenancy-namespace>/<username>`, where `<tenancy-namespace>` is the auto-generated Object Storage namespace string of your tenancy (as shown on the **Tenancy Information** page). For example, `ansh81vrulzp/jdoe@acme.com`. If your tenancy is federated with Oracle Identity Cloud Service, use the format `<tenancy-namespace>/oracleidentitycloudservice/<username>`.
4. When prompted, enter the auth token you copied earlier.
5. Pull the Docker image from Oracle Cloud Infrastructure Registry to the client machine by entering:

```
docker pull <region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>
```

where:

- `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `iad`. See [Availability by Region Name and Region Code](#) for the list of region codes.
- `ocir.io` is the Oracle Cloud Infrastructure Registry name.
- `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy that owns the repository from which you want to pull the image (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`). Note also that your user must have access to the tenancy.
- `<repo-name>` (optional) is the name of a repository from which you want to pull the image (for example, `project01`). Note that your user must have access to the repository. Omit this argument if the image does not exist within a repository (see [About Repositories](#)).
- `<image-name>` is the name of the image that you want to pull from Oracle Cloud Infrastructure Registry (for example, `acme-web-app`)

- `<tag>` is the tag of the image that you want to pull from Oracle Cloud Infrastructure Registry (for example, `version2.0.test`)

For example:

```
docker pull iad.ocir.io/ansh81vrulzp/project01/acme-web-app:version2.0.test
```

Note that if you don't specify a `<tag>` in the `docker pull` command, Docker pulls the image that has the `latest` tag.

6. Confirm that the image has been pulled from Oracle Cloud Infrastructure Registry by entering `docker images` and verifying that the list of images on the client machine now includes the image you just pulled.

For example:

```
$ docker images
REPOSITORY TAG IMAGE ID CREATED
SIZE
iad.ocir.io/ansh81vrulzp/project01/acme-web-app version2.0.test 8e0506e14874 1 minute ago
162.6 MB
acme-web-app latest 8e0506e14874 2 hours ago
162.6 MB
acme-web-app version1.0 7d9495d03763 2 hours ago
162.6 MB
<none> <none> 6ebd328f833d 5 hours ago
162.6 MB
hello-world latest 80b84820d442 5 weeks ago
890 B
```

## Pulling Images from Registry during Kubernetes Deployment

During the deployment of an application to a Kubernetes cluster, you'll typically want one or more images to be pulled from a Docker registry. In the application's manifest file you specify the images to pull, the registry to pull them from, and the credentials to use when pulling the images. The manifest file is commonly also referred to as a pod spec, or as a `deployment.yaml` file (although other filenames are allowed).

## CHAPTER 26 Registry

---

If you want the application to pull images that reside in Oracle Cloud Infrastructure Registry, you have to perform two steps:

- You have to use `kubectl` to create a Docker registry secret. The secret contains the Oracle Cloud Infrastructure credentials to use when pulling the image. When creating secrets, Oracle strongly recommends you use the latest version of `kubectl` (see the [kubectl documentation](#)).
- You have to specify the image to pull from Oracle Cloud Infrastructure Registry, including the repository location and the Docker registry secret to use, in the application's manifest file.

To create a Docker registry secret:

1. If you haven't already done so, follow the steps to download the cluster's kubeconfig configuration file and set the `KUBECONFIG` environment variable to point to the file. Note that you must download your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user downloaded. See [Downloading a kubeconfig File to Enable Cluster Access](#).
2. In a terminal window, enter:

```
$ kubectl create secret docker-registry <secret-name> --docker-server=<region-code>.ocir.io --
docker-username='<tenancy-namespace>/<oci-username>' --docker-password='<oci-auth-token>' --
docker-email='<email-address>'
```

where:

- `<secret-name>` is a name of your choice, that you will use in the manifest file to refer to the secret. For example, `ocirsecret`
- `<region-code>` is the code for the Oracle Cloud Infrastructure Registry region you're using. For example, `iad`. See [Availability by Region Name and Region Code](#) for the list of region codes.
- `ocir.io` is the Oracle Cloud Infrastructure Registry name.
- `<tenancy-namespace>` is the auto-generated Object Storage namespace string of the tenancy containing the repository from which the application is to pull the image (as shown on the **Tenancy Information** page). For example, the namespace of the `acme-dev` tenancy might be `ansh81vrulzp`. Note that for some

older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, `acme-dev`).

- `<oci-username>` is the username to use when pulling the image. The username must have access to the tenancy specified by `<tenancy-namespace>`. For example, `jdope@acme.com`. If your tenancy is federated with Oracle Identity Cloud Service, use the format `oracleidentitycloudservice/<oci-username>`
- `<oci-auth-token>` is the auth token of the user specified by `<oci-username>`. For example, `k]j64r{1sJSSF-;)K8`
- `<email-address>` is an email address. An email address is required, but it doesn't matter what you specify. For example, `jdope@acme.com`

Note the use of single quotes around strings containing special characters. For example, combining the previous examples, you might enter:

```
$ kubectl create secret docker-registry ocirsecret --docker-server=phx.ocir.io --docker-username='ansh81vrulzp/jdope@acme.com' --docker-password='k]j64r{1sJSSF-;)K8' --docker-email='jdope@acme.com'
```

Having created the Docker secret, you can now refer to it in the application manifest file.

To specify the image to pull from Oracle Cloud Infrastructure Registry, along with the Docker secret to use, during deployment of an application to a cluster:

1. Open the application's manifest file in a text editor.
2. Add the following sections to the manifest file:
  - a. Add a `containers` section that specifies the name and location of the container you want to pull from Oracle Cloud Infrastructure Registry, along with other deployment details.
  - b. Add an `imagePullSecrets` section to the manifest file that specifies the name of the Docker secret you created to access the Oracle Cloud Infrastructure Registry.

Here's an example of what the manifest might look like when you've added the `containers` and `imagePullSecrets` sections:

```
apiVersion: v1
kind: Pod
metadata:
 name: nginx-image
spec:
 containers:
 - name: nginx
 image: phx.ocir.io/ansh81vrulzp/project01/nginx-lb:latest
 imagePullPolicy: Always
 ports:
 - name: nginx
 containerPort: 8080
 protocol: TCP
 imagePullSecrets:
 - name: ocirsecret
```

3. Save and close the manifest file.

## Viewing Images and Image Details

To make sure you pull the correct image or to identify images that you no longer need, you can find out detailed information about the images in Oracle Cloud Infrastructure Registry.

Your permissions control the images in Oracle Cloud Infrastructure Registry that you can view information about. You can view information about images in repositories you've created, and in repositories that the groups to which you belong have been granted access by identity policies. If you belong to the Administrators group, you can view information about images in any repository in the tenancy.

### Using the Console

To view images and image details:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.

2. Choose the registry's region. You see all the repositories in the registry to which you have access.
3. Click the name of the repository that contains the image you want to see detailed information about. You see all the different images in the repository, along with the tag of each image and when it was pushed to the registry. You can sort the different images by the date they were pushed or by their tag.
4. Click the image for which you want to see detailed information. The **Summary** page shows you the size of the image, when it was pushed and by which user, and the number of times the image has been pulled. Use the options on the **Summary** page as follows:
  - Display the **Layers** tab to see the SHA message digest of each layer in the selected image.
  - Display the **Associated Tags** tab to see the full path for the image with the tag you select. Note that if you select a different tag, the summary details change accordingly.
5. (Optional) If you want to pull an image, click the **Download** button beside the image name and copy the command shown. The command includes the image name in the format `<region-code>.ocir.io/<tenancy-namespace>/<repo-name>/<image-name>:<tag>`. For example, `docker pull iad.ocir.io/ansh81vrulzp/project01/acme-web-app:version2.0.test`. See [Pulling Images Using the Docker CLI](#).

## Deleting an Image

When you no longer need an old image or you simply want to clean up the list of image tags in a repository, you can delete images from Oracle Cloud Infrastructure Registry.

Your permissions control the images in Oracle Cloud Infrastructure Registry that you can delete. You can delete images from repositories you've created, and from repositories that the groups to which you belong have been granted access by identity policies. If you belong to the Administrators group, you can delete images from any repository in the tenancy.

Note that as well deleting individual images as described below, you can set up image retention policies to delete images automatically based on selection criteria you specify (see [Retaining and Deleting Images Using Retention Policies](#)).

### Using the Console

To delete an image from Oracle Cloud Infrastructure Registry:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories to which you have access.
3. Click the name of the repository from which to delete the image.
4. Click the name of the image that you want to delete.
5. Click **Delete** on the **Summary** page and confirm that you want to delete the image.

The image is permanently removed from Oracle Cloud Infrastructure Registry.

### Retaining and Deleting Images Using Retention Policies

You can set up image retention policies to automatically delete images that meet particular selection criteria, namely:

- images that have not been pulled for a certain number of days
- images that have not been tagged for a certain number of days
- images that have not been given particular Docker tags specified as exempt from automatic deletion

An hourly process checks images against the selection criteria, and any that meet the selection criteria are automatically deleted.

You'll often find image retention policies are a more convenient way to manage the images in a repository than manually deleting individual images (see [Deleting an Image](#)).

In each region in a tenancy, there's a global image retention policy. The global image retention policy's default selection criteria retain all images, so that no images are

automatically deleted. However, you can change the global image retention policy so that images are deleted if they meet the criteria you specify. A region's global image retention policy applies to all repositories in the region, unless it is explicitly overridden by one or more custom image retention policies.

You can set up custom image retention policies to override the global image retention policy with different criteria for specific repositories in a region. Having created a custom image retention policy, you apply the custom retention policy to a repository by adding the repository to the policy. The global image retention policy no longer applies to repositories that you add to a custom retention policy.

If you have `manage` permission on the tenancy, you can:

- modify each region's own global image retention policy
- create new custom image retention policies
- modify the criteria of existing custom image retention policies
- delete custom image retention policies

If you have `manage` permission on a repository, you can:

- add the repository to a custom image retention policy
- remove the repository from a custom image retention policy

Note the following:

- Only one custom image retention policy at a time can apply to a repository. If a repository has already been added to a custom retention policy and you want to add the repository to a different custom retention policy, you have to remove the policy from the first retention policy before adding it to the second.
- When you create or update an image retention policy, the hourly process that checks images for deletion will ignore the new or updated policy for several hours. This cooling-off period enables you to refine the policy criteria to select only the images you want to delete, and thus reduces the chance of images being deleted unexpectedly. After this period, the policy is included in the hourly process and images are checked and deleted accordingly.

- The global image retention policy (and any custom image retention policies you create) are specific to a particular region. To delete images consistently in different regions in your tenancy, set up image retention policies in each region with identical selection criteria .

### Using the Console to Edit the Global Image Retention Policy

Provided you have `manage` permission on the tenancy, you can edit the region's global image retention policy that applies to all repositories in a region (except for repositories that have been explicitly added to a custom image retention policy).

To edit the global image retention policy:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories to which you have access.
3. Click **Settings**, and then select **Image retention policies**.  
You see the current selection criteria of the region's global image retention policy, along with any custom image retention policies that override the global image retention policy for specific repositories.
4. Click **Edit Global Policy**.
5. In the **Global Image Retention Policy** dialog, specify new criteria for the global retention policy:
  - **Delete any images that haven't been pulled in  $n$  days:** Select this option if you want to delete images that have not been pulled for the number of days you specify.
  - **Delete any images that haven't been tagged in  $n$  days:** Select this option if you want to delete images that have not been tagged for the number of days you specify.
  - **Exempt Tags:** If you want to prevent images from being deleted on the basis of Docker tags they've been given, specify those tags as exempt in a comma-separated list. An image that has been given one of the exempt tags will not be

deleted, even if the image meets the other criteria. You can include the asterisk (\*) as a wildcard to represent none, one, or more characters. For example, you might specify `latest,prod-*,*-tail,*.100.*`.

### 6. Click **Save Settings**.

Going forward, the criteria you entered for the region's global image retention policy will apply to all repositories in the region, except for repositories that have been explicitly added to a custom image retention policy. Images in repositories that have not been added to a custom image retention policy will be deleted from Oracle Cloud Infrastructure Registry if they meet the criteria you specified in the global image retention policy.

When you create or update an image retention policy, the hourly process that checks images for deletion will ignore the new or updated policy for several hours. This cooling-off period enables you to refine the policy criteria to select only the images you want to delete, and thus reduces the chance of images being deleted unexpectedly. After this period, the policy is included in the hourly process and images are checked and deleted accordingly.

## Using the Console to Create a New Custom Image Retention Policy to Override the Global Policy

Provided you have `manage` permission on the tenancy, you can create a new custom image retention policy to override the region's global image retention policy for the repositories you specify. A custom image retention policy is specific to the region in which you create it.

To create a new custom image retention policy:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories to which you have access.
3. Click **Settings**, and then select **Image retention policies**.  
You see the current selection criteria of the region's global image retention policy, along with any existing custom image retention policies that override the global image retention policy for specific repositories.
4. Click **Create Policy**.

5. In the **Create Repository Image Retention Policy** dialog, specify criteria for the new retention policy:
  - **Policy Name:** A name of your choice for the policy. Avoid entering confidential information.
  - **Delete any images that haven't been pulled in *n* days:** Select this option if you want to delete images that have not been pulled for the number of days you specify.
  - **Delete any images that haven't been tagged in *n* days:** Select this option if you want to delete images that have not been tagged for the number of days you specify.
  - **Exempt Tags:** If you want to prevent images from being deleted on the basis of Docker tags they've been given, specify those tags as exempt in a comma-separated list. An image that has been given one of the exempt tags will not be deleted, even if the image meets the other criteria. You can include the asterisk (\*) as a wildcard to represent none, one, or more characters. For example, you might specify `latest,prod-*,*-tail,*.100.*`.
6. Click **Save Settings**.

You can now add repositories to the new custom retention policy.

### Using the Console to Remove a Repository from a Custom Image Retention Policy

Provided you have `manage` permission on a repository, you can remove a repository from a custom image retention policy to which it was previously added.

You might want to remove the repository from a custom image retention policy:

- if you want the region's global image retention policy to apply to the repository
- if you want a different custom image retention policy to apply to the repository (only one custom image retention policy at a time can apply to a repository)

To remove a repository from a custom image retention policy:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories to which you have access.
3. Click **Settings**, and then select **Image retention policies**.  
You see the current selection criteria of the region's global image retention policy, along with any existing custom image retention policies that override the global image retention policy for specific repositories.
4. Locate the custom image retention policy to which the repository has been added.
5. Click the delete icon beside the repository name to remove it from the custom image retention policy.

Going forward, the region's global image retention policy will apply to the repository (unless you add the repository to a different custom image retention policy). The images in the repository will be deleted from Oracle Cloud Infrastructure Registry if they meet the criteria specified in the global image retention policy.

When you create or update an image retention policy, the hourly process that checks images for deletion will ignore the new or updated policy for several hours. This cooling-off period enables you to refine the policy criteria to select only the images you want to delete, and thus reduces the chance of images being deleted unexpectedly. After this period, the policy is included in the hourly process and images are checked and deleted accordingly.

### Using the Console to Add a Repository to a Custom Image Retention Policy

Provided you have `manage` permission on a repository, you can add a repository to an existing custom image retention policy.

Note that if a custom image retention policy already applies to the repository, you'll have to remove the repository from the current policy before adding it to a different policy. Note also that a custom image retention policy is specific to the region in which it was created.

To add a repository to an existing custom image retention policy:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories to which you have access.
3. Click **Settings**, and then select **Image retention policies**.

You see the current selection criteria of the region's global image retention policy, along with the custom image retention policies that have been defined to override the global image retention policy for specific repositories.

4. Locate the custom image retention policy to which you want to add the repository.
5. Click **Add Repository** and select from the list the repository you want to add to the custom image retention policy.

Note that the repository list includes all repositories in the region, regardless of whether you have permission to add them to a retention policy. You can only add a repository to a retention policy if you have `manage` permission on that repository,

If a repository in the list has a policy name beside it, the repository has already been added to a policy. Before you can add the repository to a different policy, you'll have to remove it from the first policy.

Going forward, the custom retention policy to which you added the repository will override the region's global image retention policy. The images in the repository will be deleted from Oracle Cloud Infrastructure Registry if they meet the criteria specified in the custom retention policy.

When you create or update an image retention policy, the hourly process that checks images for deletion will ignore the new or updated policy for several hours. This cooling-off period enables you to refine the policy criteria to select only the images you want to delete, and thus reduces the chance of images being deleted unexpectedly. After this period, the policy is included in the hourly process and images are checked and deleted accordingly.

## Deleting a Repository

There is a limit to the number of repositories you can have in any given region in a tenancy. So when you no longer need a repository, it makes sense to delete it from Oracle Cloud Infrastructure Registry.

Your permissions control the repositories in Oracle Cloud Infrastructure Registry that you can delete. You can delete repositories you've created, and from repositories that the groups to which you belong have been granted access by identity policies. If you belong to the Administrators group, you can delete any repository in the tenancy.

## Using the Console

To delete a repository from Oracle Cloud Infrastructure Registry:

1. In the Console, open the navigation menu. Under **Solutions and Platform**, go to **Developer Services** and click **Registry**.
2. Choose the registry's region. You see all the repositories in the registry to which you have access.
3. Click the name of the repository that you want to delete.
4. Click **Delete** on the **Summary** page and confirm that you want to delete the repository.

The repository is permanently removed from Oracle Cloud Infrastructure Registry.

## Getting an Auth Token

Before you can push and pull Docker images to and from Oracle Cloud Infrastructure Registry, you must already have an Oracle Cloud Infrastructure username and an auth token. If you

haven't got an auth token, or you've forgotten it, or you're not sure, you can create a new auth token. You only see the auth token string when you create it, so be sure to copy the auth token to a secure location immediately.

**Tip:** Each user can have up to two auth tokens at a time. So if you do lose or forget the auth token, you can always create a second auth token.

To create a new auth token:

1. In the top-right corner of the Console, open the **Profile** menu () and then click **User Settings** to view the details.
2. On the **Auth Tokens** page, click **Generate Token**.
3. Enter a friendly description for the auth token. Avoid entering confidential information.
4. Click **Generate Token**. The new auth token is displayed.
5. Copy the auth token immediately to a secure location from where you can retrieve it later, because you won't see the auth token again in the Console.
6. Close the Generate Token dialog.

## Policies to Control Repository Access

You have fine-grained control over the operations that users are allowed to perform on repositories in Oracle Cloud Infrastructure Registry.

A user's permissions to access repositories comes from the groups to which they belong. The permissions for a group are defined by identity policies. Policies define which actions the members of a group can perform. Users access repositories and perform operations based on the policies set for the groups they are members of. Identity policies to control repository access must be set at the tenancy level. See [Details for Registry](#) .

Before you can control access to repositories, you must have already created users and already placed them in appropriate groups (see [Managing Users](#) and [Managing Groups](#)). You can then create policies and policy statements to control repository access (see [Managing Policies](#)).

Note that users in the tenancy's Administrators group can perform any operation on any repository in Oracle Cloud Infrastructure Registry that belongs to the tenancy.

### Common Policies



#### Note

The policies in this section use example group names, as follows:

- **acme-viewers:** A group that you want to limit to seeing a list of repositories in the tenancy.
- **acme-pullers:** A group that you want to limit to pulling images.
- **acme-pushers:** A group that you want to allow to push and pull images.
- **acme-managers:** A group that you want to allow to push and pull images, delete repositories, and edit repository metadata (for example, to make a private repository public).

Make sure to replace the example group names with your own group names.

### Enable users to view a list of all the repositories belonging to the tenancy

**Type of access:** Ability to see a list of all repositories in Oracle Cloud Infrastructure Registry belonging to the tenancy. Users will not be able to:

## CHAPTER 26 Registry

---

- view the images or layers in a repository
- push or pull images from or to a repository

Note that there is currently no way to restrict the repositories shown on the Registry page in the Console.

**Where to create the policy:** In the tenancy.

```
Allow group acme-viewers to inspect repos in tenancy
```

### Enable users to pull images from any repository belonging to the tenancy

**Type of access:** Ability to pull images (layers and manifests) from any repository in Oracle Cloud Infrastructure Registry that belongs to the tenancy.

**Where to create the policy:** In the tenancy.

```
Allow group acme-pullers to read repos in tenancy
```

### Enable users to pull images from specific repositories

**Type of access:** Ability to pull images (layers and manifests) from any repository in Oracle Cloud Infrastructure Registry that belongs to the tenancy and that has a name starting with "acme-web-app".

**Where to create the policy:** In the tenancy.

```
Allow group acme-pullers to read repos in tenancy where all { target.repo.name=/acme-web-app*/ }
```

### Enable users to push images to any repositories (and create new repositories if necessary)

**Type of access:** Ability to push images (layers and manifests) to any repository in Oracle Cloud Infrastructure Registry that belongs to the tenancy. If a repository with the same name as the image doesn't exist yet, the `REPOSITORY_CREATE` permission ensures users are able to

## CHAPTER 26 Registry

---

create the repository when they push the image.

**Where to create the policy:** In the tenancy.

```
Allow group acme-pushers to use repos in tenancy
```

```
Allow group acme-pushers to manage repos in tenancy where ANY {request.permission = 'REPOSITORY_CREATE', request.permission = 'REPOSITORY_UPDATE'}
```

Enable managers to perform any operation on any repository belonging to the tenancy

**Type of access:** Ability to perform any operation on any repository in Oracle Cloud Infrastructure Registry that belongs to the tenancy, including:

- pull an image from any repository
- push an image to any repository
- create a new repository (either an empty repository, or when pushing an image for which no repository exists yet)
- delete a repository
- change a public repository to a private repository, or a private repository to a public repository

**Where to create the policy:** In the tenancy.

```
Allow group acme-managers to manage repos in tenancy
```

# CHAPTER 27 Overview of Resource Manager

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. Using Terraform, Resource Manager helps you install, configure, and manage resources through the "infrastructure-as-code" model.



## Note

Resource Manager is not available in Oracle Cloud Infrastructure Government Cloudrealms.

A Terraform configuration codifies your infrastructure in declarative configuration files. Resource Manager allows you to share and manage infrastructure configurations and state files across multiple teams and platforms. This infrastructure management can't be done with local Terraform installations and Oracle Terraform modules alone. For more information about the Oracle Cloud Infrastructure Terraform provider, see [Terraform Provider](#). For a general introduction to Terraform and the "infrastructure-as-code" model, see <https://www.terraform.io>.

## Key Concepts

Following are brief descriptions of key concepts and the main components of Resource Manager.

### CONFIGURATION

A set of one or more Terraform configuration files that codify your infrastructure. Use your configuration to specify the Oracle Cloud Infrastructure resources in a given stack. For example, specify resource metadata, data source definitions, and variable declarations. Each Terraform configuration file is either HashiCorp Configuration Language (HCL) format or JSON format, as indicated by the file's extension (either `.tf` or `.tf.json`, respectively).

For example configuration files, see [Terraform provider examples](#). For more information, see [Terraform Configurations for Resource Manager](#) and [Writing Terraform Configurations](#); see also [Hashicorp: Configuration](#).

### JOB

Instructions to perform the actions defined in your configuration. Only one job at a time can run on a given stack; further, you can have only one set of Oracle Cloud Infrastructure resources on a given stack. To provision a different set of resources, you must create a separate stack and use a different configuration.

Resource Manager provides the following job types:

- **Plan:** Parses your Terraform configuration and creates an execution plan for the associated stack. The execution plan lists the sequence of specific actions planned to provision your Oracle Cloud Infrastructure resources. The execution plan is handed off to the apply job, which then executes the instructions.
- **Apply.** Applies the execution plan to the associated stack to create (or modify) your Oracle Cloud Infrastructure resources. Depending on the number and type of resources specified, a given apply job can take some time. You can check status while the job runs.
- **Destroy.** Releases resources associated with a stack. Released resources are not deleted. For example, terminates a Compute instance controlled by a stack. The stack's job history and state remain after running a destroy job. You can monitor the status and review the results of a destroy job by inspecting the stack's log files.
- **Import State.** Sets the provided Terraform state file as the current state of the stack. Use this job to migrate local Terraform environments to Resource Manager.

Jobs store history about their associated stack. For example, plan jobs store generated execution plans and apply jobs store configurations (snapshots) and state files. Jobs reside in the compartment that is occupied by the stack they are associated with. An OCID is assigned to each job.

### MODULE

A group of related resources. Use modules to create lightweight and reusable abstractions, so that you can describe your infrastructure in terms of its architecture. For more information, see [Creating Modules](#).

### STACK

The collection of Oracle Cloud Infrastructure resources corresponding to a given [Terraform configuration](#). Each stack resides in the compartment you specify, in a single region; however, resources on a given stack can be deployed across multiple regions. An OCID is assigned to each stack.

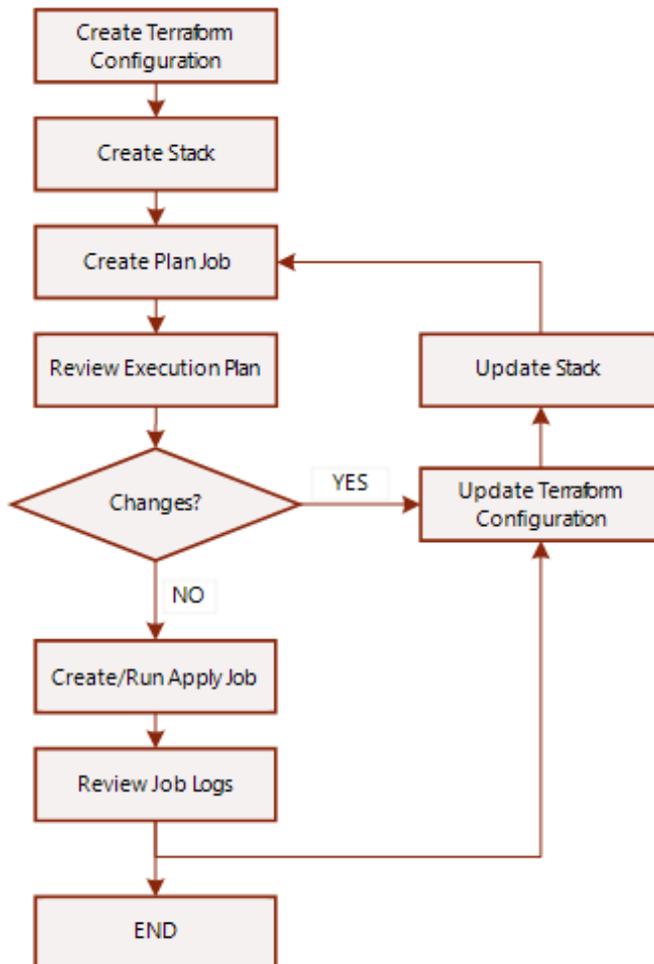
### STATE

The state of your resource configuration, stored in JSON format in a state file (.tfstate). The state file maps your stack's resources to your configuration and also maintains essential configuration metadata, such as resource dependencies. Resource Manager generates and updates state files automatically. You cannot edit the file manually.

Resource Manager supports state locking by allowing only one job at a time to run on a given stack. For more information about state files, see [Hashicorp: State](#).

## Generalized Workflow

The following image represents a generalized view of the Resource Manager workflow.



Links in the following steps reference Console instructions; however, you can do the same tasks using the [API](#) (through the [CLI](#) or other tool).

1. [Create a Terraform configuration.](#)
2. [Create a stack.](#)
3. [Run a plan job](#), which produces an execution plan.
4. [Review the execution plan.](#)

5. If changes are needed in the execution plan, [update the configuration](#) and run a plan job again.
6. [Run an apply job](#) to provision resources.
7. [Review state file and log files](#), as needed.
8. You can optionally reapply your configuration, with or without making changes, by running an apply job again.
9. Optionally, to release the resources running on a stack, [run a destroy job](#).

For a detailed walkthrough of the Resource Manager workflow, see [Sample: Creating a Compute Instance Using Resource Manager](#).

## Ways to Access Resource Manager

You can access the Resource Manager service using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

**Console:** To access Resource Manager using the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.

**API:** To access Resource Manager through APIs, use [Resource Manager API](#). To access this API using the Command Line Interface (CLI), use the `oci resource-manager` designation.

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

Administrators: For common policies that give groups access to stacks and jobs, see [Policies for Managing Stacks and Jobs](#). For a complete list of Resource Manager permissions, see [Details for Resource Manager](#). Policies for managing accessed resource types are also required.



### Important

Policies for managing Oracle Cloud Infrastructure resources are also required for Resource Manager operations that access resources. For example, running an apply job on a stack that includes Compute instances and subnets requires policies that grant you permissions for those resource types, in the compartments where you want to provision the resources. To see examples of policies for managing Oracle Cloud Infrastructure resources, see [Common Policies](#).

## Limits on Resource Manager Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

# Managing Stacks and Jobs

This topic describes how to create, edit, and delete stacks as well as work with jobs, including generating and applying execution plans.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Prerequisites

- **IAM policies:** To manage stacks and jobs, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. Administrators: For common policies that give groups access to stacks and jobs, see [Policies for Managing Stacks and Jobs](#). For a complete list of Resource Manager permissions, see [Details for Resource Manager](#). Policies for managing accessed resource types are also required.



### Important

Policies for managing Oracle Cloud Infrastructure resources are also required for Resource Manager operations that access resources. For example, running an apply job on a stack that includes Compute instances and subnets requires policies that grant you permissions for those resource types, in the compartments where you want to provision the resources. To see examples of policies for managing Oracle Cloud Infrastructure resources, see [Common Policies](#).

- Terraform configuration file: To create or update a stack, you must have a valid Terraform configuration file. See [Terraform Configurations for Resource Manager](#) and [Writing Terraform Configurations](#).

## Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

## Moving Resources to a Different Compartment

You can move stacks from one compartment to another. When you move a stack to a new compartment, its associated jobs move with it. After you move the stack to the new compartment, inherent policies apply immediately and affect access to the stack and associated jobs through the Console. For more information, see [Managing Compartments](#).

### Using the Console

#### To create a stack

Creating a stack involves uploading your Terraform configuration file, providing identifying information for the new stack, and optionally updating variables. You can always edit your stack later.



#### Important

Make sure your Terraform configuration file is valid. See [Writing Terraform Configurations](#) and [Terraform Configurations for Resource Manager](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click **Create Stack**.
4. In the **Create Stack** dialog, add your [Terraform configuration](#) (.zip) file. You can either drag and drop it onto the dialog's control or click **Browse** and navigate to the file location. The dialog box is populated with information contained in the configuration file.
5. Enter a **Name** for the new stack (or accept the default name provided).
6. Optionally enter a **Description**.
7. From the **Create in Compartment** drop-down, select the compartment where you want to create the stack. A compartment from the list scope is set by default.
8. Select a **Terraform Version**.



### Note

Terraform version 0.12.x is not backward-compatible.

9. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
10. Click **Next**.  
The **Configure Variables** panel displays variables auto-populated from the Terraform file you uploaded in step 1.
11. Review the variables and make changes as necessary.



### Important

Do not add your private key or other confidential information to configuration variables.

12. Click **Next**.
13. In the **Review** panel, verify your stack configuration.
14. Click **Create** to create your stack.

## To view stacks

You can view stack names, descriptions, states, and time created.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.

2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.

### To edit a stack

You can edit stacks. When editing a stack, you can upload a different .zip file and change its name, description, Terraform version, and variables.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the Actions icon (three dots), and then select **Edit**.  
You can also edit a stack from its detail page. Click the name of the stack to display its detail page and then click **Edit Stack**.
4. In the **Edit Stack** dialog, change the properties you want.



#### Note

Changing the Terraform version requires the stack's Resource Manager configuration (.zip) file to be compatible with the new version. Downgrading the Terraform version from 0.12.x to 0.11.x is only available before an apply job is run on the stack.

- To edit the values assigned to variables in a stack, click **Configure Variables**. You can also edit a stack from its detail page. Click the name of the stack to display the **Stack Details** page, click **Variables** (under **Resources**) and then click **Edit Variables**.



### Important

Do not add your private key or other confidential information to configuration variables.

If you want to add, reconfigure, or delete variables in a stack, update the Terraform configuration (.zip) file.

5. Click **Save Changes**.

### To manage tags for a stack

Tags are key/value pairs that you can attach to resources to help you organize and track your resources across compartments. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack you want. The **Stack Details** page lists the details about the selected job.
4. Click **Tags** to view or edit existing tags, or click **Add Tags** to add new ones.

### To move a stack to a different compartment

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the Actions icon (three dots), and then select **Move Stack**.  
You can also move a stack from its detail page. Click the name of the stack to display the **Stack Details** and then click **Move Stack**.
4. In the **Move Resource To A Different Compartment** dialog box, select the compartment that you want to move the stack to.
5. Click **Move Resource**.

### To delete a stack



#### Note

Associated resources persist after stack deletion. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by [running a destroy job](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.

2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the Actions icon (three dots), select **Delete**, and confirm the operation when prompted.



### Note

You cannot undo the delete stack operation.

You can also delete a stack from its detail page. Click the name of the stack to display the **Stack Details** and then click **Delete Stack**.

### To view jobs and job details

You can view name, type, status, and other key information about jobs for a given compartment or stack. You can view name, type, status, and other key information about a given job. You can also access the job's execution plan (represented by the job log), Terraform configuration, and Terraform state, as well as view the variables used in the job.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. To view job details, click the name of the job you want.  
The **Job Details** page lists the details about the selected job.
4. To view variables used in the job, click **Variables** under **Resources**.

### To manage tags for a job

Tags are key/value pairs that you can attach to resources to help you organize and track your resources across compartments. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the job you want.  
The **Job Details** page lists the details about the selected job.
4. Click **Tags** to view or edit existing tags, or click **Add Tags** to add new ones.

### To generate an execution plan (run a plan job)

Running a plan job parses your Terraform configuration (.zip) file and converts it into an execution plan listing resources and actions that will result when an apply job is run. We recommend generating an execution plan before running an apply job.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack that you want to use.

The **Stack Details** page is displayed.

4. Go to **Terraform Actions** and select **Plan**.
5. In the **Plan** dialog, review the plan job **Name** and update it if needed.
6. Click **Plan**.

The new plan job is listed under **Jobs**, with an initial state of "Accepted." Soon the status changes to "In Progress." When the job is complete, you can [review the execution plan](#) or [download the job information](#).

### To view the job log

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.

You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.

2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the plan job that you ran.
4. On the **Job Details** page, under **Resources**, click **Logs**.

For plan jobs, the log file is the execution plan. View the log file for the plan job and note the "message" fields in the sequence of log entries of the log file. These values represent the sequence of operations specified in your configuration.

You can also [download the job information](#).

### To update the configuration for a stack

Updating a stack involves uploading a new Terraform configuration (.zip) file, which overwrites the existing configuration. You might want to update a stack after reviewing the generated execution plan.



### Important

Make sure your Terraform configuration file is valid. See [Writing Terraform Configurations](#) and [Terraform Configurations for Resource Manager](#).

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack that you want to update. The **Stack Details** page is displayed.
4. In the **Stack Information** tab, next to **Terraform Configuration File (.zip)**, click **Upload New**.
5. In the **Edit Stack** dialog, add your revised [Terraform configuration](#) (.zip) file. You can either drag and drop it onto the dialog's control or click **Browse** and navigate to the file location. The dialog box is populated with information contained in the configuration file.
6. Click **Next** as needed and then click **Save Changes**. Now you can [generate a new execution plan](#) using your revised configuration.

### To download job information

You can download files associated with jobs: Terraform configurations, Terraform states, and logs.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.

## CHAPTER 27 Overview of Resource Manager

---

You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.

2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.

3. Click the name of the job you want.

The **Job Details** page is displayed.

You can view the log by clicking **Logs** under **Resources**.

You can view the state of your resources (for relevant jobs) by clicking **View State** under **Resources**.

4. Download the job information you want:

To download this job-associated file	Click
Terraform configuration (.zip file)	<b>Download Terraform Configuration</b>
Terraform state (.json file)	<b>Download Terraform State</b>
Logs (.txt file)	<b>Download Logs (Logs section under Resources)</b> <b>Note:</b> In cases of a long-running job, click the <b>Refresh</b> button to update the log.

### To run an apply job

When you run an apply job for a stack, Terraform creates the resources and executes the actions defined in your Terraform configuration (.zip) file. The time required to complete an apply job depends on the number and type of cloud resources to be created.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.

2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack that you want to use. The **Stack Details** page is displayed.
4. Go to **Terraform Actions** and select **Apply**.
5. In the **Apply** dialog, review the apply job **Name** and other settings and update it if needed.
6. Click **Apply**.  
The new apply job is listed under **Jobs**. Monitor its status: "Succeeded" indicates that the job has completed. While the job runs, or after it completes, you can [download its log file](#).
7. To view the Terraform state file (shows the state of your resources after running the job), click the name of the apply job and then click **View State** under **Resources**.

### To view the state of a job

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the job you want.
4. On the **Job Details** page, click **View State** under **Resources**.

### To import an existing Terraform state file (run an import job)

You can import state files for existing resources already managed by Terraform.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack that you want to use. The **Stack Details** page is displayed.
4. Go to **Terraform Actions** and select **Import State**.
5. In the **Import State File** dialog, review the job **Name** and update it if needed.
6. Add your Terraform state file, either by dragging and dropping it onto the dialog's control, or by clicking **Browse** and navigating to the file location.
7. Click **Import**.

### To release a stack's resources (run a destroy job)

Run a destroy job to tear down the resources and clean up the tenancy.



#### Note

We recommend running a destroy job before deleting a stack to release associated resources first. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by running a destroy job.

1. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
2. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
3. Click the name of the stack that you want to use. The **Stack Details** page is displayed.
4. Go to **Terraform Actions** and select **Destroy**.
5. Click **Destroy** again to confirm your action. You can monitor the status and review the results of a destroy job by viewing the state or the logs.
6. To view the Terraform state file (shows the state of your resources after running the job), click the name of the job to display the **Job Details** page, then click **View State** under **Resources**.
7. To view the logs for the job, click the name of the job to display the **Job Details** page, then click **Logs** under **Resources**.

### Using the CLI

This section provides basic sample CLI commands for managing stacks and jobs. For information about using the CLI, see [Command Line Interface \(CLI\)](#). For a complete list of flags and options available for CLI commands, see [CLI Help](#).

### To create a stack

On Windows, be sure the .zip file and variables.json files are in the same directory from which you're running the CLI. The CLI currently has a limitation on Windows that prevents correct handling of the files if either one is in a subdirectory.

Open a command prompt and run `oci resource-manager stack create` to create a stack:

## CHAPTER 27 Overview of Resource Manager

---

```
oci resource-manager stack create --compartment-id <compartment_OCID> --config-source <config_file_name> --variables <var_file_path> --display-name "<friendly_name>" --description "<description>" --working-directory ""
```



### Note

You can return later to update stack settings or add variables after you have created the stack.

## Options

For a complete list of flags and options available for CLI commands, see [CLI Help](#).

- `--compartment-id` is the OCID of the compartment where you want to create the stack.
- `--config-source` is the name of a .zip file that contains one or more Terraform configuration files.
- `--variables` is the path to the file specifying input variables for your resources. Optional.

The Oracle Cloud Infrastructure Terraform provider requires additional parameters when running Terraform locally (unless you are using instance principals). For more information on using variables in Terraform, see [Input Variables](#). See also [Input Variable Configuration](#).

- `--display-name` is the friendly name for the new stack. Optional.
- `--description` is the description for the new stack. Optional.
- `--working-directory` is the root configuration file in the directory. Optional. If not specified, or if null as in this example, then the service assumes that the top-level file in the directory is the root configuration file.

For example:

```
oci resource-manager stack create --compartment-id ocid1.tenancy.oc1..uniqueid --config-source vcn.zip --variables file://variables.json --display-name "My Example Stack" --description "My Tutorial to Create a VCN" --working-directory ""
```

### Example response

```
{
 "data": {
 "config-source": {
 "working-directory": null,
 "config-source-type": "ZIP_UPLOAD"
 },
 "defined-tags": {},
 "description": "My Tutorial to Create a VCN",
 "display-name": "My Example Stack",
 "freeform-tags": {},
 "id": "ocid1.ormstack.oc1..uniqueid",
 "lifecycle-state": "ACTIVE",
 "time-created": "2019-04-03T18:26:56.299000+00:00",
 "variables": {
 "compartment_ocid": "ocid1.compartment.oc1..uniqueid",
 "region": "us-phoenix-1"
 }
 }
}
```

### To list stacks in a compartment

Open a command prompt and run `oci resource-manager stack list` to list the stacks in a compartment:

```
oci resource-manager stack list --compartment-id <compartment_OCID>
```

### To list full details of a stack

Open a command prompt and run `oci resource-manager stack get` to list the details for the specified stack:

```
oci resource-manager stack get --stack-id <stack_OCID>
```

### To delete a stack



#### Note

Associated resources persist after stack deletion. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by running a destroy job.

Open a command prompt and run `oci resource-manager stack delete` to delete the specified stack:

```
oci resource-manager stack delete --stack-id <stack_OCID>
```

### To generate an execution plan (run a plan job)

Open a command prompt and run `oci resource-manager job create-plan-job` to run a plan job on the specified stack (`--display-name` is optional):

```
oci resource-manager job create-plan-job --stack-id <stack_OCID> --display-name "<friendly_name>"
```

Depending on the complexity of the configuration, the plan job can take several minutes to complete. When the job is complete, make sure you [review the generated execution plan](#) before running an apply job.

### To check the current state of the plan job

Open a command prompt and run `oci resource-manager job get` to retrieve information

## CHAPTER 27 Overview of Resource Manager

---

about the job:

```
oci resource-manager job get --job-id <plan_job_OCID>
```

### Lifecycle states

Possible values for `lifecycle-state`:

- **ACCEPTED**: The job is queued for execution.
- **IN\_PROGRESS**: The job is running.
- **FAILED**: The job has failed and stopped running.
- **SUCCEEDED**: The job has completed successfully.
- **CANCELING**: The job has been notified to cancel, but has not yet stopped running.
- **CANCELED**: The job was canceled and has stopped running.

### Example response

This example shows **ACCEPTED** for `lifecycle-state`.

```
{
 "data": {
 "compartment-id": " ocid1.compartment.oc1..uniqueid",
 "defined-tags": null,
 "display-name": "Example Plan Job",
 "freeform-tags": {},
 "id": "ocid1.ormjob.oc1..uniqueid",
 "lifecycle-state": "ACCEPTED",
 "operation": "PLAN",
 "jobOperationDetails": {
 "operation": "PLAN"
 },
 "stack-id": " ocid1.ormstack.oc1..uniqueid",
 "time-created": "2019-03-09T20:52:13.922000+00:00",
 "time-finished": null,
 }
}
```

## CHAPTER 27 Overview of Resource Manager

---

```
"variables": {
 "compartment_ocid": "ocid1.compartment.oc1..uniqueid",
 "region": "us-phoenix-1"
}
}
```

### To review an execution plan (view the log for a plan job)

Review the execution plan to ensure that it accurately reflects your intentions. View the log file and note the "message" fields in the sequence of log entries of the log file. These values represent the sequence of operations specified in your configuration.

Open a command prompt and run `oci resource-manager job get-job-logs` to view the log file for the specified job:

```
oci resource-manager job get-job-logs --job-id <plan_job_OCID>
```

If you see problems or errors and wish to make changes, then update the appropriate [configuration file](#) (.tf file), [update the stack](#) to use the revised configuration, [generate a new execution plan](#), and then review the new execution plan.

### Example response

The command returns JSON objects that describe log entries. Each object has a message member with a property that displays one line of the execution plan. In the example shown below, the plan creates a single virtual cloud network (VCN); the remaining members show details about the VCN.

```
...
{
 "level": "INFO",
 "message": "Terraform will perform the following actions:",
 "timestamp": "2018-05-24T00:57:14.170000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
```

## CHAPTER 27 Overview of Resource Manager

---

```
{
 "level": "INFO",
 "message": "",
 "timestamp": "2018-05-24T00:57:14.170000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "+ oci_core_virtual_network.vcn1",
 "timestamp": "2018-05-24T00:57:14.170000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "id: <computed>",
 "timestamp": "2018-05-24T00:57:14.172000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "cidr_block: \"10.0.0.0/16\"",
 "timestamp": "2018-05-24T00:57:14.172000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "compartment_id:
\"ocid1.tenancy.oc1..exampleaqpqpfqfmr6dw5gcew7yqpirvarueirj2mv4jzn5goejsxma\"",
 "timestamp": "2018-05-24T00:57:14.172000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "default_dhcp_options_id: <computed_value>",
 "timestamp": "2018-05-24T00:57:14.172000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": "default_route_table_id: <computed_value>",
```

## CHAPTER 27 Overview of Resource Manager

---

```
"timestamp": "2018-05-24T00:57:14.172000+00:00",
"type": "TERRAFORM_CONSOLE"
},
{
 "level": "INFO",
 "message": " default_security_list_id: <computed_value>",
 "timestamp": "2018-05-24T00:57:14.172000+00:00",
 "type": "TERRAFORM_CONSOLE"
},
...
```

To update an execution plan (update the configuration for a stack)



### Important

To update the execution plan after running the plan job, you must first update the configuration and recreate the configuration .zip file. Then, upload the new .zip file and rerun the plan job.

Open a command prompt and run `oci resource-manager stack update` with the option `--config-source` to update the Terraform configuration for the specified stack:

```
oci resource-manager stack update --stack-id <stack_OCID> --config-source <config_file_name>
```

After updating the stack, regenerate and review an execution plan (run a new plan job and then view the log file).

### To run an apply job

To check the current state of the apply job

## CHAPTER 27 Overview of Resource Manager

---

Open a command prompt and run `oci resource-manager job create-apply-job` with the relevant value for `--execution-plan-strategy` (examples use `--display-name`, which is optional):

- To specify a plan job ("apply" an execution plan), use `FROM_PLAN_JOB_ID`:

```
oci resource-manager job create-apply-job --stack-id <stack_OCID> --execution-plan-strategy
FROM_PLAN_JOB_ID --execution-plan-job-id <plan_job_OCID> --display-name "Example Apply Job"
```

Use this option to "apply" your confirmed execution plan to the stack, execute the instructions, and provision the stack with the specified resources.

- To automatically approve the apply job (no plan job specified), use `AUTO_APPROVED`:

```
oci resource-manager job create-apply-job --stack-id <stack_OCID> --execution-plan-strategy
AUTO_APPROVED --display-name "Example Apply Job"
```

Depending on the complexity of your execution plan, the operation can take some time. Periodically check the lifecycle state of your apply job to see when it switches from `IN_PROGRESS` to `SUCCEEDED`.

### To check the current state of the apply job

Open a command prompt and run `oci resource-manager job get` to retrieve information about the job:

```
oci resource-manager job get --job-id <apply_job_OCID>
```

### Lifecycle states

Possible values for `lifecycle-state`:

- `ACCEPTED`: The job is queued for execution.
- `IN_PROGRESS`: The job is running.
- `FAILED`: The job has failed and stopped running.
- `SUCCEEDED`: The job has completed successfully.

## CHAPTER 27 Overview of Resource Manager

---

- **CANCELING:** The job has been notified to cancel, but has not yet stopped running.
- **CANCELED:** The job was canceled and has stopped running.

To confirm existence of newly provisioned resources, [inspect resources in the compartment](#).

### To download or view job information

You can download Terraform configurations and Terraform states associated with jobs. You can also view logs associated with jobs.

### To download the configuration for a job

Open a command prompt and run `oci resource-manager job get-job-tf-config` to download the Terraform configuration of the specified job to the specified file:

```
oci resource-manager job get-job-tf-config -job-id <job_OCID> --file <output_file_name>
```

### To download the state file for a job

Open a command prompt and run `oci resource-manager job get-job-tf-state` to download the Terraform state of the specified job to the specified file:

```
oci resource-manager job get-job-tf-state --job-id <job_OCID> --file <output_file_name>
```

### Example response for an apply job

```
{
 "data": {
 "lineage": "57ef4f0c-c8cd-8a32-d45f-d2c40be7b915",
 "modules": [
 {
 "depends_on": [],
 "outputs": {},
 }
]
 }
}
```

## CHAPTER 27 Overview of Resource Manager

---

```
"path": [
 "root"
],
"resources": {
 "oci_core_virtual_network.vcn1": {
 "depends_on": [],
 "deposed": [],
 "primary": {
 "attributes": {
 "cidr_block": "10.0.0.0/16",
 "compartment_id": "ocidl.tenancy.oc1..uniqueid",
 "default_dhcp_options_id": "ocidl.dhcpoptions.oc1.phx.uniqueid",
 "default_route_table_id": "ocidl.routetable.oc1.phx.uniqueid",
 "default_security_list_id": "ocidl.securitylist.oc1.phx.uniqueid",
 "display_name": "My VCN display name",
 "dns_label": "myvcntest",
 "id": "ocidl.vcn.oc1.phx.uniqueid",
 "state": "AVAILABLE",
 "time_created": "2018-05-24 01:13:05.855 +0000 UTC",
 "vcn_domain_name": "myvcntest.oraclevcn.com"
 },
 "id": "ocidl.vcn.oc1.phx.uniqueid",
 "meta": {
 "e2bfb730-ecaa-11e6-8f88-34363bc7c4c0": {
 "create": 300000000000,
 "delete": 300000000000,
 "update": 300000000000
 }
 },
 "tainted": false
 },
 "provider": "provider.oci",
 "type": "oci_core_virtual_network"
 }
}
],
"serial": 4,
```

## CHAPTER 27 Overview of Resource Manager

---

```
"terraform_version": "0.11.7",
 "version": 3
}
}
```

### To view the log for a job

View the log file and note the "message" fields in the sequence of log entries of the log file. You can view the log file for the specified job as either a paged list of entries or in its raw form.

To view the log as a paged list of entries, open a command prompt and run `oci resource-manager job get-job-logs`:

```
oci resource-manager job get-job-logs --job-id <job_OCID>
```

To view the log in raw form, open a command prompt and run `oci resource-manager job get-job-logs-content`:

```
oci resource-manager job get-job-logs-content --job-id <job_OCID>
```

### To import an existing Terraform state file (run an import job)

Open a command prompt and run `oci resource-manager x` to import an existing state file for resources already managed by Terraform:

```
oci resource-manager job create-import-tf-state-job --stack-id <stack_id> --tf-state-file <state_file>
```

### To inspect resources in a compartment

Inspecting resources in a compartment allows you to confirm existence of a resource that you provisioned (by running an apply job) or absence of a resource that you released (by running a destroy job).

## CHAPTER 27 Overview of Resource Manager

---

Open a command prompt and run the CLI command corresponding to the resources you want to inspect.

For example, run `oci network vcn list` to inspect VCN resources in the specified compartment:

```
oci network vcn list --compartment-id <compartment_OCID>
```

To release a stack's resources (run a destroy job)



### Note

We recommend running a destroy job before deleting a stack to release associated resources first. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by running a destroy job.

Open a command prompt and run `oci resource-manager job create-destroy-job` to tear down and clean up the resources provisioned by the specified stack:

```
oci resource-manager job create-destroy-job --stack-id <stack_OCID> --execution-plan-strategy=AUTO_APPROVED
```

To confirm deletion of the resources, [inspect resources in the compartment](#).

## Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line](#).

### [Interface.](#)

Use these API operations to manage stacks:

- [CreateStack](#)
- [GetStack](#)
- [GetStackTfConfig](#)
- [ListStacks](#)
- [ListTerraformVersions](#)
- [UpdateStack](#)
- [ChangeStackCompartment](#)
- [DeleteStack](#)

Use these API operations to manage jobs:

- [CreateJob](#)
- [GetJob](#)
- [GetJobLogs](#)
- [GetJobLogsContent](#)
- [GetJobTfConfig](#)
- [GetJobTfState](#)
- [ListJobs](#)
- [UpdateJob](#)
- [CancelJob](#)

Use these API operations to manage work requests:

- [GetWorkRequest](#)
- [ListWorkRequestErrors](#)
- [ListWorkRequestLogs](#)
- [ListWorkRequests](#)

# Sample: Creating a Compute Instance Using Resource Manager

This sample provides an end-to-end walkthrough of the tasks required to create and deploy an Oracle Cloud Infrastructure Compute instance using Resource Manager. For a brief introduction to Resource Manager, see [Overview of Resource Manager](#).



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Highlights

This walkthrough provides samples that demonstrate how to use Resource Manager to create a Compute instance. Resource Manager uses Terraform to provision the resources that you've defined in a Terraform configuration. The resources are organized into stacks, which you create and provision using jobs.

The walkthrough covers the following tasks:

- Create a Terraform configuration using the HashiCorp configuration language (HCL). For more information, see [Configuration Syntax](#).
- Provision the infrastructure:
  - Create a stack in which to provision your infrastructure.
  - Run a plan job against your stack, which parses your configuration and creates an execution plan.
  - Review the generated execution plan.
  - Run an apply job against your stack, which provisions your resources. The apply

job follows the execution plan, which is based on your Terraform configuration.

- Review the resulting infrastructure.

### Before We Begin

Ensure that you have installed, obtained, or created the prerequisites:

- An Oracle Cloud Infrastructure tenancy
- The OCID for the compartment where you wish to create your stack.
- A user account that includes the following:
  - An API signing key. For guidance, see [Required Keys and OCIDs](#).
  - Required IAM permissions. For more information, see [How Policies Work](#) and [Details for Resource Manager](#).
- If you want to use the Oracle Cloud Infrastructure CLI, install and configure the CLI first. See [Quickstart](#) and [Configuration](#)

### Task 1: Create the Terraform Configuration

A Terraform configuration is a .zip file containing one or more files that codify your infrastructure. The configuration defines your Terraform provider, the resources you intend to provision, variables, and specific instructions for provisioning the resources. You need a Terraform configuration to get started with Resource Manager. For more information, see [Terraform Configuration](#).



#### Warning

Do not provide user credentials or other confidential information in your Terraform configuration.

In this example, our configuration uses several configuration files (.tf files) to direct Resource Manager to execute the following sequence of operations.

### Create an Oracle Cloud Infrastructure Provider

The following code sample creates a basic Oracle Cloud Infrastructure Terraform provider. You can provide values as variables that are defined either in a variables file or in the provider definition (.tf) file. For more information, see [Provider Configuration](#).

```
provider "oci" {
 region = "${var.region}"
}
```

### Define Variables

Define the variables you want to use when provisioning your resources. A best practice is to create a "variables" file in the configuration package that you upload. Following is an example from a configuration file that we've named `variables.tf`. For more information about using variables, see [Input Variables](#). See also [Configuring Input Variables](#).

```
variable "compartment_ocid" {
 default = "ocidl.compartment.oc1..uniqueid"
}

variable "region" {
 default = "us-phoenix-1"
}

variable "InstanceImageOCID" {
 type = "map"
 default = {
 // See https://docs.cloud.oracle.com/images/
 // Oracle-provided image "Oracle-Linux-7.5-2018.10.16-0"
 eu-frankfurt-1 = "ocidl.image.oc1.eu-frankfurt-
1.aaaaaaaaaitzn6tdyjer7j134h2ujz74jwy5nkbukbh55ekp6oyzwrta4zma"
 uk-london-1 = "ocidl.image.oc1.uk-london-
1.aaaaaaaa32voyikkzfxyo4xbdmac2dmvorfxxgdhpnk6dw64fa314jh7wa"
 us-ashburn-1 =
"ocidl.image.oc1.iad.aaaaaaaaageeenzyuxgia726xur4ztaoxbyjlxogdhreu3ngfj2gji3bayda"
 us-phoenix-1 =
"ocidl.image.oc1.phx.aaaaaaaaoj42sokaoh42176wsyhn3k2beuntrh5maj3gmgmzeyr55zzrwwa"
 }
}
```

```
variable "ssh_public_key" {
 default = "ssh-rsa <public_key_value>"
}

Defines the number of instances to deploy
variable "NumInstances" {
 default = "1"
}

variable "InstanceShape" {
 default = "VM.Standard2.1"
}

Specifies the Availability Domain
variable "localAD" {
 default = "<AD_name>"
}
```

For more information about variables declared in the preceding examples, see the following:

- InstanceImageOCID: [Oracle-Provided Images](#)
- InstanceShape: [Compute Shapes](#)
- region and localAD: [Regions and Availability Domains](#)

### Create a Virtual Cloud Network (VCN)

The following code sample creates an Oracle Cloud Infrastructure virtual cloud network (VCN) named "ExampleVCN."

```
resource "oci_core_virtual_network" "ExampleVCN" {
 cidr_block = "10.1.0.0/16"
 compartment_id = "${var.compartment_ocid}"
 display_name = "TFExampleVCN"
 dns_label = "tfexamplevcn"
}
```

### Create a Subnet in Your VCN

The following code sample creates a subnet named "ExampleSubnet" in the VCN defined in the

previous code sample.

```
resource "oci_core_subnet" "ExampleSubnet" {
 availability_domain = "${var.localAD}"
 cidr_block = "10.1.20.0/24"
 display_name = "TFExampleSubnet"
 dns_label = "tfexamplesubnet"
 security_list_ids = ["${oci_core_virtual_network.ExampleVCN.default_security_list_id}"]
 compartment_id = "${var.compartment_ocid}"
 vcn_id = "${oci_core_virtual_network.ExampleVCN.id}"
 route_table_id = "${oci_core_route_table.ExampleRT.id}"
 dhcp_options_id = "${oci_core_virtual_network.ExampleVCN.default_dhcp_options_id}"
}
```

### Create an Internet Gateway

The following code sample creates an internet gateway named "ExampleIG" in the VCN that we created.

```
resource "oci_core_internet_gateway" "ExampleIG" {
 compartment_id = "${var.compartment_ocid}"
 display_name = "TFExampleIG"
 vcn_id = "${oci_core_virtual_network.ExampleVCN.id}"
}
```

### Create a Core Route Table

The following code sample creates a Oracle Cloud Infrastructure core route table in the VCN and then applies two route rules.

```
resource "oci_core_route_table" "ExampleRT" {
 compartment_id = "${var.compartment_ocid}"
 vcn_id = "${oci_core_virtual_network.ExampleVCN.id}"
 display_name = "TFExampleRouteTable"
 route_rules {
 cidr_block = "0.0.0.0/0"
 network_entity_id = "${oci_core_internet_gateway.ExampleIG.id}"
 }
}
```

### Create a Compute Instance

The following extended code example creates an Oracle Cloud Infrastructure Compute instance. The code also references the image on which the Compute instance is created, sets boot volume size, adds essential metadata, and applies both free-form and defined tags.

## CHAPTER 27 Overview of Resource Manager

---

```
resource "oci_core_instance" "TFInstance" {
 count = "${var.NumInstances}"
 availability_domain = "${var.localAD}"
 compartment_id = "${var.compartment_ocid}"
 display_name = "TFInstance${count.index}"
 shape = "${var.InstanceShape}"

 create_vnic_details {
 subnet_id = "${oci_core_subnet.ExampleSubnet.id}"
 display_name = "primaryvnic"
 assign_public_ip = true
 hostname_label = "tfexampleinstance${count.index}"
 },

 source_details {
 source_type = "image"
 source_id = "${var.InstanceImageOCID[var.region]}"

 # Apply this to set the size of the boot volume that's created for this instance.
 # Otherwise, the default boot volume size of the image is used.
 # This should only be specified when source_type is set to "image".
 #boot_volume_size_in_gbs = "60"
 }

 # Apply the following flag only if you wish to preserve the attached boot volume upon destroying this
instance
 # Setting this and destroying the instance will result in a boot volume that should be managed outside
of this config.
 # When changing this value, make sure to run 'terraform apply' so that it takes effect before the
resource is destroyed.
 #preserve_boot_volume = true

 metadata {
 ssh_authorized_keys = "${var.ssh_public_key}"
 }

 timeouts {
 create = "60m"
 }
}
```

### Finalize the Configuration

Ensure that all of the configuration files are in a single directory. When using the CLI, you create the configuration .zip file, then specify it using the `--config-source` parameter. When using the Console, you create the configuration .zip file , then upload it.



#### Important

Make sure your Terraform configuration file is valid. See [Writing Terraform Configurations](#) and [Terraform Configurations for Resource Manager](#).

### Task 2: Provision the Infrastructure

Use your Terraform configuration to build and deploy your infrastructure by taking the following actions:

1. Create a stack in a tenancy compartment of your choosing.  
A stack is a collection of resources that you can act on as a group. All of the resources that you specify in your configuration are provisioned in the stack that you create.

#### To create a stack (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click **Create Stack**.
- d. In the **Create Stack** dialog, add your [Terraform configuration](#) (.zip) file.

You can either drag and drop it onto the dialog's control or click **Browse** and navigate to the file location.

The dialog box is populated with information contained in the configuration file.

- e. Enter a **Name** for the new stack (or accept the default name provided).
- f. Optionally enter a **Description**.
- g. From the **Create in Compartment** drop-down, select the compartment where you want to create the stack.  
A compartment from the list scope is set by default.
- h. Select a **Terraform Version**.



### Note

Terraform version 0.12.x is not backward-compatible.

- i. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- j. Click **Next**.  
The **Configure Variables** panel displays variables auto-populated from the Terraform file you uploaded in step 1.
- k. Review the variables and make changes as necessary.



### Important

Do not add your private key or other confidential information to configuration variables.

- l. Click **Next**.
- m. In the **Review** panel, verify your stack configuration.
- n. Click **Create** to create your stack.

### To create a stack (CLI)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click **Create Stack**.
- d. In the **Create Stack** dialog, add your [Terraform configuration](#) (.zip) file.  
You can either drag and drop it onto the dialog's control or click **Browse** and navigate to the file location.  
The dialog box is populated with information contained in the configuration file.
- e. Enter a **Name** for the new stack (or accept the default name provided).
- f. Optionally enter a **Description**.
- g. From the **Create in Compartment** drop-down, select the compartment where you want to create the stack.  
A compartment from the list scope is set by default.
- h. Select a **Terraform Version**.



### Note

Terraform version 0.12.x is not backward-compatible.

- i. Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
- j. Click **Next**.  
The **Configure Variables** panel displays variables auto-populated from the Terraform file you uploaded in step 1.
- k. Review the variables and make changes as necessary.



### Important

Do not add your private key or other confidential information to configuration variables.

- l. Click **Next**.
  - m. In the **Review** panel, verify your stack configuration.
  - n. Click **Create** to create your stack.
2. Generate an execution plan.  
The plan job parses your configuration to create an "execution plan," which is a step-by-step representation of the planned deployment in job log entries. Once the plan job has completed, you can evaluate the execution plan by viewing the job's log entries to confirm that it performs the expected operations, and in the intended sequence.

### To run a plan job (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the stack that you want to use.  
The **Stack Details** page is displayed.
- d. Go to **Terraform Actions** and select **Plan**.
- e. In the **Plan** dialog, review the plan job **Name** and update it if needed.
- f. Click **Plan**.

The new plan job is listed under **Jobs**, with an initial state of "Accepted." Soon the status changes to "In Progress." When the job is complete, you can [review the execution plan](#) or [download the job information](#).

### To run a plan job (CLI)

Open a command prompt and run `oci resource-manager job create-plan-job` to run a plan job on the specified stack (`--display-name` is optional):

```
oci resource-manager job create-plan-job --stack-id <stack_OCID> --display-name "<friendly_name>"
```

Depending on the complexity of the configuration, the plan job can take several minutes to complete. When the job is complete, make sure you [review the generated execution plan](#) before running an apply job.

### To check the current state of the plan job

Open a command prompt and run `oci resource-manager job get` to retrieve information about the job:

```
oci resource-manager job get --job-id <plan_job_OCID>
```

### Lifecycle states

Possible values for `lifecycle-state`:

- **ACCEPTED**: The job is queued for execution.
- **IN\_PROGRESS**: The job is running.
- **FAILED**: The job has failed and stopped running.
- **SUCCEEDED**: The job has completed successfully.
- **CANCELING**: The job has been notified to cancel, but has not yet stopped running.
- **CANCELED**: The job was canceled and has stopped running.

### Example response

This example shows **ACCEPTED** for `lifecycle-state`.

```
{
 "data": {
 "compartment-id": " ocid1.compartment.oc1..uniqueid",
 "defined-tags": null,
 "display-name": "Example Plan Job",
 "freeform-tags": {},
 "id": "ocid1.ormjob.oc1..uniqueid",
 "lifecycle-state": "ACCEPTED",
 "operation": "PLAN",
 "jobOperationDetails": {
 "operation": "PLAN"
 },
 "stack-id": " ocid1.ormstack.oc1..uniqueid",
 "time-created": "2019-03-09T20:52:13.922000+00:00",
 "time-finished": null,
 "variables": {
```

## CHAPTER 27 Overview of Resource Manager

---

```
"compartment_ocid": "ocid1.compartment.oc1..uniqueid",
 "region": "us-phoenix-1"
}
}
}
```

3. Review the execution plan to confirm that it represents your intentions. The execution plan is represented in the log for the plan job you ran previously.

### To review an execution plan (the log for the plan job) (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the plan job that you ran.
- d. On the **Job Details** page, under **Resources**, click **Logs**.  
For plan jobs, the log file is the execution plan. View the log file for the plan job and note the "message" fields in the sequence of log entries of the log file. These values represent the sequence of operations specified in your configuration.  
You can also [download the job information](#).

If changes are needed, revise the configuration, update your stack, then rerun the plan job to obtain an updated execution plan.

### To update a stack

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource**

**Manager** and click **Stacks**.

- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the stack that you want to update.  
The **Stack Details** page is displayed.
- d. In the **Stack Information** tab, next to **Terraform Configuration File (.zip)**, click **Upload New**.
- e. In the **Edit Stack** dialog, add your revised [Terraform configuration](#) (.zip) file. You can either drag and drop it onto the dialog's control or click **Browse** and navigate to the file location.  
The dialog box is populated with information contained in the configuration file.
- f. Click **Next** as needed and then click **Save Changes**.  
Now you can [generate a new execution plan](#) using your revised configuration.

### To review an execution plan (the log for the plan job) (CLI)

Review the execution plan to ensure that it accurately reflects your intentions. View the log file and note the "message" fields in the sequence of log entries of the log file. These values represent the sequence of operations specified in your configuration.

Open a command prompt and run `oci resource-manager job get-job-logs` to view the log file for the specified job:

```
oci resource-manager job get-job-logs --job-id <plan_job_OCID>
```

If you see problems or errors and wish to make changes, then update the appropriate [configuration file](#) (.tf file), [update the stack](#) to use the revised configuration, [generate a new execution plan](#), and then review the new execution plan.

If changes are needed, revise the configuration, update your stack, then rerun the plan job to obtain an updated execution plan.

### To update a stack



#### Important

To update the execution plan after running the plan job, you must first update the configuration and recreate the configuration .zip file. Then, upload the new .zip file and rerun the plan job.

Open a command prompt and run `oci resource-manager stack update` with the option `--config-source` to update the Terraform configuration for the specified stack:

```
oci resource-manager stack update --stack-id <stack_OCID> --config-source <config_file_name>
```

After updating the stack, regenerate and review an execution plan (run a new plan job and then view the log file).

4. Provision your resources by running an apply job against the execution plan.  
When satisfied with the execution plan, we're ready to do the work of provisioning the stack with the resources that we've defined. The apply job takes the execution plan and "applies" it to the stack. The result is a fully provisioned stack.

### To run an apply job (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the stack that you want to use.  
The **Stack Details** page is displayed.

- d. Go to **Terraform Actions** and select **Apply**.
- e. In the **Apply** dialog, review the apply job **Name** and other settings and update it if needed.
- f. Click **Apply**.  
The new apply job is listed under **Jobs**. Monitor its status: "Succeeded" indicates that the job has completed. While the job runs, or after it completes, you can [download its log file](#).
- g. To view the Terraform state file (shows the state of your resources after running the job), click the name of the apply job and then click **View State** under **Resources**.

### To run an apply job (CLI)

To check the current state of the apply job

Open a command prompt and run `oci resource-manager job create-apply-job` with the relevant value for `--execution-plan-strategy` (examples use `--display-name`, which is optional):

- To specify a plan job ("apply" an execution plan), use `FROM_PLAN_JOB_ID`:

```
oci resource-manager job create-apply-job --stack-id <stack_OCID> --execution-plan-strategy FROM_PLAN_JOB_ID --execution-plan-job-id <plan_job_OCID> --display-name "Example Apply Job"
```

Use this option to "apply" your confirmed execution plan to the stack, execute the instructions, and provision the stack with the specified resources.

- To automatically approve the apply job (no plan job specified), use `AUTO_APPROVED`:

```
oci resource-manager job create-apply-job --stack-id <stack_OCID> --execution-plan-strategy AUTO_APPROVED --display-name "Example Apply Job"
```

Depending on the complexity of your execution plan, the operation can take some time. Periodically check the lifecycle state of your apply job to see when it switches from `IN_PROGRESS` to `SUCCEEDED`.

### To check the current state of the apply job

Open a command prompt and run `oci resource-manager job get` to retrieve information about the job:

```
oci resource-manager job get --job-id <apply_job_OCID>
```

### Lifecycle states

Possible values for `lifecycle-state`:

- **ACCEPTED**: The job is queued for execution.
- **IN\_PROGRESS**: The job is running.
- **FAILED**: The job has failed and stopped running.
- **SUCCEEDED**: The job has completed successfully.
- **CANCELING**: The job has been notified to cancel, but has not yet stopped running.
- **CANCELED**: The job was canceled and has stopped running.

To confirm existence of newly provisioned resources, [inspect resources in the compartment](#).

5. Review the log entries and state file for the apply job you just ran.
  - See the entries in the job log for more details about the job.

### To view the job log (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that

compartment. If you're not sure which compartment to use, contact an administrator.

- c. Click the name of the plan job that you ran.
- d. On the **Job Details** page, under **Resources**, click **Logs**.

For plan jobs, the log file is the execution plan. View the log file for the plan job and note the "message" fields in the sequence of log entries of the log file. These values represent the sequence of operations specified in your configuration.

You can also [download the job information](#).

### To view the job log (CLI)

View the log file and note the "message" fields in the sequence of log entries of the log file. You can view the log file for the specified job as either a paged list of entries or in its raw form.

To view the log as a paged list of entries, open a command prompt and run `oci resource-manager job get-job-logs`:

```
oci resource-manager job get-job-logs --job-id <job_OCID>
```

To view the log in raw form, open a command prompt and run `oci resource-manager job get-job-logs-content`:

```
oci resource-manager job get-job-logs-content --job-id <job_OCID>
```

- The job state file represents the job's output in JSON format. The state file maps your stack's resources to your configuration and also maintains essential configuration metadata, such as resource dependencies. Resource Manager generates and updates state files automatically when you run jobs. The Resource Manager supports state locking by allowing only one job at a time to run on a given stack. For more information about state files, see [Hashicorp: State](#).

### To view the state of the job (Console)

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Jobs**.  
You can also access jobs from a stack detail page. Click **Stacks** and then click the name of the stack you want.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the job you want.
- d. On the **Job Details** page, click **View State** under **Resources**.

### To view the state of the job (CLI)

Open a command prompt and run `oci resource-manager job get-job-tf-state` to download the Terraform state of the specified job to the specified file:

```
oci resource-manager job get-job-tf-state --job-id <job_OCID> --file <output_file_name>
```

### Example response for an apply job

```
{
 "data": {
 "lineage": "57ef4f0c-c8cd-8a32-d45f-d2c40be7b915",
 "modules": [
 {
 "depends_on": [],
 "outputs": {},
 "path": [
 "root"
],
 "resources": {
```

## CHAPTER 27 Overview of Resource Manager

---

```
"oci_core_virtual_network.vcn1": {
 "depends_on": [],
 "deposed": [],
 "primary": {
 "attributes": {
 "cidr_block": "10.0.0.0/16",
 "compartment_id": "ocidl.tenancy.oc1..uniqueid",
 "default_dhcp_options_id": "ocidl.dhcptoptions.oc1.phx.uniqueid",
 "default_route_table_id": "ocidl.routetable.oc1.phx.uniqueid",
 "default_security_list_id": "ocidl.securitylist.oc1.phx.uniqueid",
 "display_name": "My VCN display name",
 "dns_label": "myvcntest",
 "id": "ocidl.vcn.oc1.phx.uniqueid",
 "state": "AVAILABLE",
 "time_created": "2018-05-24 01:13:05.855 +0000 UTC",
 "vcn_domain_name": "myvcntest.oraclevcn.com"
 },
 "id": "ocidl.vcn.oc1.phx.uniqueid",
 "meta": {
 "e2bfb730-ecaa-11e6-8f88-34363bc7c4c0": {
 "create": 300000000000,
 "delete": 300000000000,
 "update": 300000000000
 }
 },
 "tainted": false
 },
 "provider": "provider.oci",
 "type": "oci_core_virtual_network"
}
}
}
"serial": 4,
"terraform_version": "0.11.7",
"version": 3
}
```

```
}
```



### Note

You can also [import state files](#) for resources already managed by Terraform.

6. When you need to release the resources that you provisioned, run a destroy job on the stack.

A destroy job tears down the stack that you created and then cleans up associated resources without deleting them. For example, the destroy job terminates Compute instances associated with the stack.

### To run a destroy job (Console)



### Note

We recommend running a destroy job before deleting a stack to release associated resources first. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by running a destroy job.

- a. Open the navigation menu. Under **Solutions and Platform**, go to **Resource Manager** and click **Stacks**.
- b. Choose a compartment you have permission to work in (on the left side of the page). The page updates to display only the resources in that compartment. If you're not sure which compartment to use, contact an administrator.
- c. Click the name of the stack that you want to use.  
The **Stack Details** page is displayed.
- d. Go to **Terraform Actions** and select **Destroy**.
- e. Click **Destroy** again to confirm your action.  
You can monitor the status and review the results of a destroy job by viewing the state or the logs.
- f. To view the Terraform state file (shows the state of your resources after running the job), click the name of the job to display the **Job Details** page, then click **View State** under **Resources**.
- g. To view the logs for the job, click the name of the job to display the **Job Details** page, then click **Logs** under **Resources**.

### To run a destroy job (CLI)



#### Note

We recommend running a destroy job before deleting a stack to release associated resources first. When you delete a stack, its associated state file is also deleted; therefore, you lose track of the state of its associated resources. Cleaning up resources associated with a deleted stack can be difficult without the state file, especially when those



resources are spread across multiple compartments. To avoid difficult cleanup later, we recommend that you release associated resources first by running a destroy job.

Open a command prompt and run `oci resource-manager job create-destroy-job` to tear down and clean up the resources provisioned by the specified stack:

```
oci resource-manager job create-destroy-job --stack-id <stack_OCID> --execution-plan-strategy=AUTO_APPROVED
```

To confirm deletion of the resources, [inspect resources in the compartment](#).

## Using Remote Exec

With Resource Manager, you can use [Terraform's remote exec functionality](#) to execute scripts or commands on a remote computer. You can also use this technique for other provisioners that require access to the remote resource.



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Prerequisites

- The location where the script is remotely executed must be an Oracle Cloud Infrastructure resource that has a public IP and supports remote login.
- On Windows, WinRM must be enabled. On Linux or Unix, SSH must be enabled.

- A key pair used for signing API requests, with the public key uploaded to Oracle. For more information on generating and uploading keys, see [Required Keys and OCIDs](#).

### Authenticating

We recommend using one of the following approaches, depending on whether you have access to the Key Management service. For more information, see [Overview of Key Management](#).

#### With Key Management

First, use Key Management to encrypt your private key. For more information, see [Managing Keys](#) and [Using Keys](#).

Next, provide the encrypted private key to Resource Manager. You can use the decrypt data source to decrypt it.

The following code sample demonstrates this process.

```
data "oci_kms_decrypted_data" "private_key_decrypted" {
 #Required
 ciphertext = "${file(var.encrypted_private_key_path)}"
 crypto_endpoint = "${var.decrypted_data_crypto_endpoint}"
 key_id = "${var.kms_encryption_key_id}"
}

resource "oci_core_instance" "TFInstance1" {
 availability_domain = "${lookup(data.oci_identity_availability_domains.ADs.availability_domains
[var.availability_domain - 1], "name")}"
 compartment_id = "${var.compartment_ocid}"
 display_name = "TFInstance"
 hostname_label = "instance3"
 shape = "${var.instance_shape}"
 subnet_id = "${oci_core_subnet.ExampleSubnet.id}"

 source_details {
 source_type = "image"
 source_id = "${var.instance_image_ocid[var.region]}"
 }
}
```

```
extended_metadata {
 ssh_authorized_keys = "${var.ssh_public_key}"
}

resource "null_resource" "remote-exec" {
 connection {
 agent = false
 timeout = "30m"
 host = "${oci_core_instance.TFInstance1.public_ip}"
 user = "${var.opc_user_name}"
 private_key = "${data.oci_kms_decrypted_data.test_decrypted_data.plaintext}"
 }

 inline = [
 "touch ~/IMadeAFile.Right.Here"
]
}
```

### Without Key Management

If you do not have access to the Key Management service, you can dynamically generate a key pair and store them in the state file.

1. Generate a key pair using a TLS resource.
2. When you launch the Compute instance, use the public key from the TLS resource.
3. When you establish the SSH connection, provide the private key.



#### Warning

You should not save your private key in your Terraform configuration file because that is not a secure location.

The following sample demonstrates how to use the TLS private key resource to provision a Compute instance, then perform a remote execution on that instance.

## CHAPTER 27 Overview of Resource Manager

---

```
resource "tls_private_key" "public_private_key_pair" {
 algorithm = "RSA"
}

resource "oci_core_instance" "TFInstance1" {
 availability_domain = "${lookup(data.oci_identity_availability_domains.ADs.availability_domains
[var.availability_domain - 1], "name")}"
 compartment_id = "${var.compartment_ocid}"
 display_name = "TFInstance"
 hostname_label = "instance3"
 shape = "${var.instance_shape}"
 subnet_id = "${oci_core_subnet.ExampleSubnet.id}"

 source_details {
 source_type = "image"
 source_id = "${var.instance_image_ocid[var.region]}"
 }

 extended_metadata {
 ssh_authorized_keys = "${tls_private_key.public_private_key_pair.public_key_openssh}"
 }
}

resource "null_resource" "remote-exec" {
 depends_on = ["oci_core_instance.TFInstance1"]

 provisioner "remote-exec" {
 connection {
 agent = false
 timeout = "30m"
 host = "${oci_core_instance.TFInstance1.public_ip}"
 user = "${var.opc_user_name}"
 private_key = "${tls_private_key.public_private_key_pair.private_key_pem}"
 }

 inline = [
 "touch ~/IMadeAFile.Right.Here"
]
 }
}
```

### Connection Construct

This example demonstrates how to use a `connection` construct for remote exec. Terraform uses a number of defaults when connecting to a resource, but these can be overridden using a `connection` block in either a `resource` or `provisioner`. For more information, see [Provisioner Connections](#).

## Terraform Configurations for Resource Manager

This topic describes requirements and recommendations for Terraform configurations used with Resource Manager. For basic information about Terraform configurations, see [Writing Terraform Configurations](#). For instructions on using configurations with stacks and jobs, see [Managing Stacks and Jobs](#).



### Warning

Do not provide user credentials or other confidential information in your Terraform configurations.

### Terraform Provider

When using Resource Manager, the `region` field in the `provider "oci"` block is the only required field. All other fields, such as `userid` or `fingerprint`, are optional.

### File Structure

Resource Manager requires the following Terraform configuration file (`.zip`) structure for the configuration:

- The working directory must contain at least one `.tf` file. The working directory cannot contain a `.terraform` directory.

The working directory is the path from which to run Terraform. By default, the working directory is the root of the `.zip` file. When using the API, you can specify a different location for the working directory by setting the `workingDirectory` parameter.

- The configuration must follow guidelines specified in [Writing Terraform Configurations](#).
- No Terraform state files (`.tfstate`) can exist in the configuration.

Modules: We recommend that all modules used by a Resource Manager stack are included locally in the configuration and referenced using a local path. To use a module for Oracle Cloud Infrastructure from the [Terraform Module Registry](#), do the following.

- Download the source from GitHub and include the relevant portion in a subdirectory in your `.zip` file.
- Reference the module using a local path. For more information, see [Local Paths](#).

## Variables

Resource Manager supports the native Terraform behavior for handling variables. You can include a `terraform.tfvars` file and files with the `.auto.tfvars` extension in the configuration `.zip` file. For more information about defining variables, see [Define Variables](#) and [Input Variables](#).

### **Adding a Schema Document to Facilitate Variable Entry**

You can make it easier for your users to enter variables when they create or update a stack in the Console using your Terraform configuration. A Resource Manager schema document clarifies the meaning and use of variables in the Console. The schema document allows you to:

- Validate and constrain variables to communicate permitted values more clearly.
- Name, group, and order variables.
- Require variables.
- Control visibility of variables.
- Dynamically prepopulate variable values.

## CHAPTER 27 Overview of Resource Manager

---

### SCHEMA DOCUMENT REQUIREMENTS

The Resource Manager schema document has the following requirements:

- YAML format.
- Placement under the root folder of the Resource Manager Terraform configuration (.zip) file. (By default, the schema document assumes that the root folder is the working folder.)

For example schema documents and more information, see the [example schema](#) and [meta schema](#).

### VALIDATING AND CONSTRAINING VARIABLE ENTRIES

Terraform's configuration language supports a default value, a description, and three types for input variables: strings, lists, and maps. Because these three types are all strings, users can have difficulty understanding what values are permitted for each variable.

Use a schema document to define validations and constraints for your variables. These validations and constraints help users determine permitted values when creating and editing stacks using the Console.

**Example 1:** Enum type for validating and constraining entry in a variable field. The Console renders the `num_nodes` field as a selectable list, with available selections of 1 and 2. The default value is 2.

```
num_nodes:
 type: enum
 enum:
 - "1"
 - "2"
 default: "2"
```

**Example 2:** Pattern for validating entry in a variable field. The Console renders the `url` field as a free text entry field, validating the entered value against the provided URL pattern (and the type).

```
url:
 type: string
 pattern: ^https?:\\\/(www\\.)?[-a-zA-Z0-9@:%._\\+~#={2,256}\\.[a-z]{2,4}\\b([-a-zA-Z0-9@:%_\\+.~#?&//=]*)$
```

### NAMING AND REQUIRING VARIABLES

In your schema document, add a `groupings` node. In this node, add a `title` and list the variables you want in the group in the order you want. You can specify friendly names and descriptions for the variables and explicitly requiring variables. When a variable is required, the user must specify a value.

The following example defines a group titled "Create A VCN" that lists the `region` and `compartment_ocid`. Defining a variable by type (such as `oci:identity:region:name` for `region`) makes the fields selectable instead of free text input. Each field also includes a friendly name and description. Both variables are required.

```
schemaVersion: 1.0.0
version: "20190612"
locale: "en"
groupings:
 - title: "Create A VCN"
 variables:
 - region
 - compartment_ocid

variables:
 region:
 type: oci:identity:region:name
 title: Region
 description: The region in which to create all resources
 required: true

 compartment_ocid:
 type: oci:identity:compartment:id
 title: Target Compartment
 description: The target compartment for all of the provisioned resources
 required: true
```

The Console renders the "Create A VCN" variable group as follows.

### Create A VCN

**REGION**

us-ashburn-1 

The region in which to create all resources

**TARGET COMPARTMENT**

mycompartment 

The target compartment for all of the provisioned resources

#### CONTROLLING VISIBILITY OF VARIABLES

In more complex stack launches, you sometimes specify input variables for certain configurations of the stack only. Use the schema to control visibility of groups and variables based on the values of other variables. Supported operations include "eq", "and", "or", and "not".

In the following example, the "Equality Conditional Section" is visible if the value of the variable "objectStorageTier" is equal to "standard":

```
- title: "Equality Conditional Section"
 variables:
 - ${myVcn}
 visible:
 eq:
 - ${objectStorageTier}
 - standard
```

In the next example, the "mySubnet" variable is visible if the value of the variable "useExistingVcn" is true:

```
mySubnet:
```

## CHAPTER 27 Overview of Resource Manager

---

```
type: oci:core:subnet:id
dependsOn:
 compartmentId: ${subnetCompartment}
 vcnId: ${myVcn}
visible: ${useExistingVcn}
```

Groups have higher priority than the groups' constituent variables. For example, if a variable is visible within a group that is not visible, then the entire group is not visible.

### DYNAMICALLY PREPOPULATING VARIABLE VALUES

Use the schema to define variables that can be dynamically prepopulated in the Console. Dynamic prepopulation has two requirements to find the elements to prepopulate with: type and dependencies. Type describes the entity you're listing. Dependencies describe what is needed to be able to find the list.

The following example shows variables that can be prepopulated based on the VCN and compartment IDs:

```
vcnCompartment:
 type: oci:identity:compartment:id

myVcn:
 type: oci:core:vcn:id
 dependsOn:
 compartmentId: ${vcnCompartment}

subnetCompartment:
 type: oci:identity:compartment:id

mySubnet:
 type: oci:core:subnet:id
 dependsOn:
 compartmentId: ${subnetCompartment}
 vcnId: ${myVcn}
```

The Console renders the dynamically prepopulated variables from this schema as follows.

## CHAPTER 27 Overview of Resource Manager

Configure networking

Virtual cloud network compartment

Sandbox ⌵

marketplacedev (root)/Sandbox

Virtual cloud network

vcn20190115053808 ⌵

Subnet compartment

Sandbox ⌵

marketplacedev (root)/Sandbox

Subnet ⓘ

Public Subnet qYRb:US-ASHBURN-AD-1 ⌵

### Example Configuration

The following example shows a Terraform configuration that is contained in a single file. This basic sample defines just one Terraform provider, one Oracle Cloud Infrastructure resource, and a set of variables.

```
variable "compartment_ocid" {}
variable "region" {}

provider "oci" {
 region = "${var.region}"
}

resource "oci_core_virtual_network" "vcn1" {
 cidr_block = "10.0.0.0/16"
 dns_label = "vcn1"
 compartment_id = "${var.compartment_ocid}"
 display_name = "vcn1"
}
```

## CHAPTER 27 Overview of Resource Manager

---

More often, Terraform configurations consist of two or more files bundled together and uploaded in a .zip file. To see more complex, multi-file Terraform configurations, explore the examples at the Oracle Cloud Infrastructure GitHub: [terraform-provider-oci/docs/examples](https://github.com/oracle/terraform-provider-oci/docs/examples).

# CHAPTER 28 Search

This chapter explains how to search for resources across compartments.

## Overview of Search

Oracle Cloud Infrastructure Search lets you find resources in your tenancy without requiring you to navigate through different services and compartments. You do not need to know the compartment or availability domain where a resource exists in order to locate and view its details. Rather, with a query, you can use as little as a single piece of information, such as the creation date or other supported attribute, to find a resource. Querying also helps you avoid the latency associated with loading a long list of results onto a single page or the inconvenience of viewing a long list that spans multiple pages.

You might find it helpful to use Search to find related resources when creating or deleting another resource. For example, you might want to find what compartments already exist before creating a new one because compartments cannot be deleted. Or, if you want to delete a volume, you can use a query to verify that a backup exists.

Another benefit of Search is that you can find resources that require action. For example, you might want to delete terminated block volumes because you no longer need them and don't want them to count against your service limits. Or, you can search for all resources that match a specific naming scheme, in case you want to act on a category of associated resources. Sometimes, resources in a specific lifecycle state, such as databases in a failed state, require troubleshooting. With Search, you can quickly identify those resources and resolve problems.



### Note

Search is not available in Oracle Cloud Infrastructure Government Cloud realms.

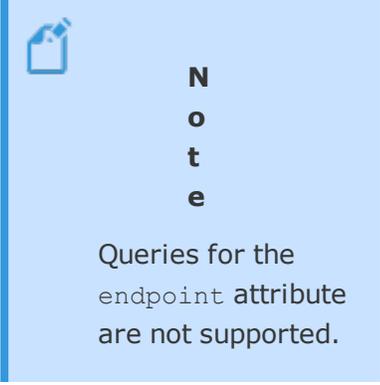
## Supported Resources

Search supports queries for the following Oracle Cloud Infrastructure services and resources. This table will be updated as query support is added for more resources.

Service	Resource Type	Attributes
Block Volume	bootvolume	See <a href="#">BootVolume Reference</a> .
Block Volume	bootvolumebackup	See <a href="#">BootVolumeBackup Reference</a> .
Block Volume	volume	See <a href="#">Volume Reference</a> .
Block Volume	volumebackup	See <a href="#">VolumeBackup Reference</a> .
Compute	autoscalingconfiguration	See <a href="#">AutoScalingConfiguration Reference</a> . <div data-bbox="954 932 1334 1318" style="border: 1px solid #0070c0; background-color: #e1f5fe; padding: 10px; margin-top: 10px;">  <p style="text-align: center;"><b>N o t e</b></p> <p>Queries for the policies attribute are not supported.</p> </div>
Compute	consolehistory	See <a href="#">ConsoleHistory Reference</a> .
Compute	image	See <a href="#">Image Reference</a> .
Compute	instance	See <a href="#">Instance Reference</a> .
Database	autonomousdatabase	See <a href="#">AutonomousDatabase Reference</a> .
Database	database	See <a href="#">Database Reference</a> .

## CHAPTER 28 Search

Service	Resource Type	Attributes
Database	dbsystem	See <a href="#">DbSystem Reference</a> .
Events	eventrule	See <a href="#">Rule Reference</a> .
Functions	functionsapplication	See <a href="#">Application Reference</a> .
Functions	functionsfunction	See <a href="#">Function Reference</a> .
IAM	compartment	See <a href="#">Compartment Reference</a> .
IAM	group	See <a href="#">Group Reference</a> .
IAM	identityprovider	See <a href="#">IdentityProvider Reference</a> .
IAM	user	See <a href="#">User Reference</a> .
Key Management	key	See <a href="#">Key Reference</a> .
Key Management	vault	See <a href="#">Vault Reference</a> .
Monitoring	alarm	See <a href="#">Search-Supported Attributes for Alarms</a> .
Networking	routetable	See <a href="#">RouteTable Reference</a> .
Networking	securitylist	See <a href="#">SecurityList Reference</a> .
Networking	subnet	See <a href="#">Subnet Reference</a> .
Networking	vcn	See <a href="#">Vcn Reference</a> .

Service	Resource Type	Attributes
Notifications	onssubscription	See <a href="#">Subscription Reference</a> .  <b>N o t e</b> Queries for the <code>endpoint</code> attribute are not supported.
Notifications	onstopic	See <a href="#">NotificationTopic Reference</a> .
Object Storage	bucket	See <a href="#">Bucket Reference</a> .
Resource Manager	ormjob	See <a href="#">Job Reference</a> .
Resource Manager	ormstack	See <a href="#">Stack Reference</a> .
WAF	waascertificate	See <a href="#">WaasCertificate Reference</a> .
WAF	waaspolicy	See <a href="#">WaasPolicy Reference</a> .

Although you can use the query language to search fields and values for any supported attribute, query results only provide information about the following resource attributes:

- Resource type
- Oracle Cloud Identifier (OCID)

- Compartment
- Availability domain
- Display name
- Creation date and time
- Lifecycle state
- Tags (visible in the API only)

The preceding attributes are common to most Oracle Cloud Infrastructure resources. Their meaning is consistent across resource types. Query results do not contain information specific to any resource type. For example, you can query for volumes of a certain size, but search results will not display the **Size** attribute. You must view the details of a resource to see resource-specific information.



### Tip

If you use the Console, neither query results nor resource details will include either defined tags or free-form tags, due to display constraints. Any given resource might contain hundreds of tags. If you want to see tags, use the API to view resource details.

## Required IAM Permissions

The resources that you see in query results depend on the permissions you have in place for the resource type. You do not necessarily see results for everything in the compartment or tenancy. For example, if your user account is not associated with a policy that grants you the ability to, at a minimum, `inspect` the `dbssystem` resource type, then you can't query for DB systems. (The verb `inspect` lets you list and get resources.) Instead, Search will show no results for queries of DB system resources. For more information about policies, see [How Policies Work](#). For information about the specific permissions required for the list

API operation for your desired resource type, see the [Policy Reference](#) for the appropriate service.

## Search Language Syntax

This topic describes the basics of the query language for Search, including an explanation of syntax and rules so you can create your own queries. Queries apply search conditions to specific resource types and let you sort results. If you want to search across all supported resource types and resource attributes and do not need ordered search results, you do not need to construct a query. Instead, you can search for a partial or exact match of free-form text without applying query language syntax to your search.

When you are ready to run a query, see [Querying Resources](#) for instructions.

### Query Basics

The following examples show the basic syntax of a query:

```
query <resourceType> resources where <conditions> sorted by <fieldName> <order>
```

Or:

```
query <resourceType> resources matching <keywords>
```

Search ignores white space, indentation, and line breaks. Sample queries include indentation to improve readability. For the purposes of demonstrating syntax only, angle brackets (<>) and italicized text indicate variables, which can consist of one or more keywords.

In a query, clauses include the following:

- `query` - (Required) Selects which resources to return based on subsequent clauses. Query statements always begin with the word `query`.
- `where` - Matches resources to the specified `conditions`.
- `matching` - Matches resources to the specified text regardless of whether the text matches exactly, matches the resource type, or appears in an indexed resource attribute.

## CHAPTER 28 Search

---

- `sorted by` - Orders resources according to `fieldName` in the order specified by `order`. If you do not include this clause, Search lists results by creation date in descending order, with the newest resources listed first.

Clauses are optional unless indicated otherwise. For matching purposes, you can use the `where` clause and the `matching` clause either separately or together.

In the `query` clause, you specify the following information:

- `resourceType` - (Required) Specifies the resource type to which the subsequent clauses apply when you run the query. You can specify either the resource type name (for example, `database` or `group`) or `all`. If you specify `all`, Search searches for the conditions against all resource types. You can query for individual resource types, but not family types. For a list of supported resource types, see the [Supported Resources](#) section of [Overview of Search](#).
- `resources` - (Required) Specifies that this is a resource query.

## Conditions

The `where` clause applies conditions that act as a filter to restrict the results returned by Search. You can specify one or more condition statements. For more information about multiple conditions, see [Grouping Conditions](#).

In a query, conditions consist of the following:

```
<fieldName> <operation> <value>
```

The `fieldName` keyword is the resource attribute against which the `operation` and desired `value` of that attribute are evaluated. Each field is associated with a field type. The type of `operation` you can use in your conditions filter depends on the field type. The API documentation includes a reference for each supported resource type that specifies attributes, their field types, and any restrictions. For more information, see the [Supported Resources](#) section of [Overview of Search](#).

In query conditions, an `operation` is a comparison operator that applies to the `value` in the statement. The `value` keyword refers to the value of the `fieldName` you specified. Search evaluates whether the specified attribute of the desired resource type matches or does not

## CHAPTER 28 Search

match the desired `value`, according to the operation. You must enclose a value in opening and closing straight single quotes (`'`).

The following table describes supported operations for resource queries:

Operation	Description	Supported Field Types	Case-sensitive?	Example
=	Equals, or exact matching for strings	String, integer, rational, Boolean, date-time	No	If the <code>value</code> was <code>'backUp'</code> , it would match <code>"backup"</code> , <code>"BACKUP"</code> , <code>"BackUp"</code> , <code>"backUp"</code> , or any other variation in casing.
!=	Does not equal	String, integer, rational, Boolean, date-time	No	If the <code>value</code> was <code>'backUp'</code> , it would match anything that does not equal <code>"backUp"</code> , <code>"backup"</code> , or any other variation in casing. It also would match anything that does not contain the characters <code>'backup'</code> in that order.
==	Strictly equals	String	Yes	If the <code>value</code> was <code>'backUp'</code> , it would only match <code>"backUp"</code> and no other variation in casing.
!==	Strictly does not equal	String	Yes	If the <code>value</code> was <code>'backUp'</code> , it would match <code>"backup"</code> , <code>"BACKup"</code> , or anything except <code>"backUp"</code> , with that exact casing.

Operation	Description	Supported Field Types	Case-sensitive?	Example
=~	Contains	String	No	If the <code>value</code> was ' <b>backUp</b> ', it would match anything that equals "backup", "BACKUP", "BackUp", "backUp", or any other variation in casing, or contains those characters in that order, alongside other characters.
>=	Greater than or equal to	Integer, rational, date-time	Not applicable	For a query where you have <code>size &gt;= 5</code> as the condition, all results have a value of <b>5</b> or greater in the field named <b>size</b> .
>	Greater than	Integer, rational, date-time	Not applicable	For a query where you have <code>size &gt; 5</code> as the condition, all results have a value of greater than <b>5</b> in the field named <b>size</b> .
<=	Less than or equal to	Integer, rational, date-time	Not applicable	For a query where you have <code>size &lt;= 5</code> as the condition, all results have a value of <b>5</b> or less in the field named <b>size</b> .
<	Less than	Integer, rational, date-time	Not applicable	For a query where you have <code>size &lt; 5</code> as the condition, all results have a value of <b>5</b> or less in the field named <b>size</b> .

## Date and Time Values

You can specify date and time values by using any of the following pattern string formats:

## CHAPTER 28 Search

Format	Examples	Comments
<code>&lt;yyyy&gt;-&lt;MM&gt;-&lt;dd&gt;</code> <code>&lt;HH&gt;:&lt;mm&gt;:&lt;ss&gt;</code> <code>&lt;TimeZone&gt;</code>	'2018-06-19 16:15:41 PDT', '2018-06-19 16:15:41 -08:00'	TimeZone is optional. If TimeZone is omitted, UTC is used.
<code>&lt;EEE&gt;, &lt;d&gt; &lt;MMM&gt; &lt;yyyy&gt;</code> <code>&lt;HH&gt;:&lt;mm&gt;:&lt;ss&gt;</code> <code>&lt;TimeZone&gt;</code>	'Tue, 19 Jun 2018 16:15:41 +0300', '19 June 2018 16:15:41'	TimeZone is optional. If TimeZone is omitted, UTC is used.
<code>&lt;yyyy&gt;-&lt;MM&gt;-</code> <code>&lt;dd&gt;T&lt;HH&gt;:&lt;mm&gt;:&lt;ss&gt;Z</code>	2018-06-19T16:15:41Z	Time in UTC. 'T' and 'Z' are case-sensitive.

You must observe spacing. Interpret dashes, colons, commas, and the characters 'T' and 'Z' literally. To interpret placeholder values in the preceding table, you can refer to the following pattern syntax:

Letter	Date or Time Component	Presentation
Y	Year	Year
M	Month in year	Month
d	Day in month	Day
H	Hour in day (from 00-23)	Number
m	Minute in hour	Number
s	Seconds in minute	Number
E	Day in week	Text

Repeating pattern letters indicate their exact presentation. For example, 'HH' means you must use '00' and not '0' to represent midnight. Similarly, 'EEE' means 'Tue' and not 'Tuesday'. Likewise, 'MM' requires '09' instead of '9' to represent the month of September.

`TimeZone` is optional, but in your chosen format, you can specify `TimeZone` in any of the following ways:

- **Name.** You can specify a time zone by its name, such as **GMT** or **PDT**. Values are case-insensitive.
- **GMT offset value.** You can specify a time zone according to its GMT offset. For example, **GMT-08:00**. Values are case-insensitive.
- **ISO 8601 time zone.** You can specify a time zone according to ISO 8601 standards. For example, **-08**, **-0800**, or **-08:00**.

Instead of using one of the preceding formats, you can also specify a date-time value as the constant `now`. The constant `now` represents the current time to the level of granularity of seconds in a minute.

Lastly, you can add or subtract time intervals from any date-time values. For example, you can query for resources that were created within five minutes of a specific time. Search supports the following time intervals:

Letter	Date or Time Component
s	Seconds
m	Minutes
h	Hours
d	Days
w	Weeks

To specify a time interval in relation to a date-time value, use one of the following formats:

- **now - 3h**
- **2018-06-19 16:15:41 PDT + 1h**

### Matching

For matching purposes, instead of or in addition to using a `where` clause with `conditions`, you might want to use the `matching` clause. The `matching` clause obviates the need to specify `conditions` (that contain a field name, operation, and value). A `matching` clause effectively queries all indexed fields by applying the `=` (equals) operator along with the text you specify. However, it does so without strictly requiring an exact match. For example, the following query uses a `matching` clause to behave the same way as a free text search: `query all resources matching 'instance'`. The query produces results that match all resources and resource attributes that contain the word "instance".

### Sorting

The last clause of a resource query is the `sorted by` clause and is optional. The `sorted by` clause orders the results returned by Search based on the field name and lists them according to the `order` you specify. By default, if you do not specify sort order, results are always sorted by date-time created in descending order.

In the `sorted by` clause, you can specify the following:

- `fieldName` - The field that Search uses to sort results. You can specify any field of any resource. Resources that do not contain the field you specify are listed after the resources that do.
- `order` - You can specify either **asc** or **desc**. Specifying **asc** lists results in ascending order. Specifying **desc** lists results in descending order.

### Grouping Conditions

You can group multiple conditions by using either the logical operators `&&` (ampersands, to indicate a logical AND) or `||` (vertical bars, to indicate a logical OR). For example:

```
licenseModel = 'LICENSE_INCLUDED' && dataStoragePercentage > 40 && lifecycleState != 'FAILED'
```

You cannot combine two different logical operators in the same query unless you wrap parentheses around one group of predicates. (Multiple conditions can only use the same logical operator otherwise.) For example:

## CHAPTER 28 Search

---

```
(licenseModel = 'LICENSE_INCLUDED' && dataStoragePercentage > 40) || lifecycleState != 'FAILED'
```

In the preceding example, all results returned will have either "LICENSE\_INCLUDED" as the value in the field named "licenseModel" and a value greater than 40 for the field named "dataStoragePercentage" or the value of their "lifecycleState" field name is anything other than "FAILED".

The following group is also acceptable:

```
licenseModel = 'LICENSE_INCLUDED' && (dataStoragePercentage > 40 || lifecycleState != 'FAILED')
```

In the preceding example, all results returned will have "LICENSE\_INCLUDED" as the value in the field named "licenseModel" and either a value greater than 40 as the value for the field named "dataStoragePercentage" or anything that is not "FAILED" for the value of the field named "lifecycleState".

Search does not perform left-to-right evaluation to reduce ambiguity or clarify intent.

### Querying Multiple Resource Types

You can query multiple resource types at once by joining queries. Each query can have its own conditional clause. If the queries that you want to join have different "where" conditions, then the syntax is different from when you have queries for multiple resource types that share the same "where" condition.

The basic syntax for a query for multiple resource types is as follows:

```
query <resourceType>, <resourceType> resources
```

For example:

```
query group, user resources
```

The preceding example query returns all groups and all users in the tenancy.

The following shows the syntax for a query for multiple resource types with conditions, but where the conditions are the same for all resource types:

```
query <resourceType>, <resourceType> resources where <conditions>
```

For example:

```
query group, user resources where displayName = 'administrator'
```

## CHAPTER 28 Search

---

The preceding example query returns all groups with the display name "administrator" and all users with the display name "administrator," with any variation in casing.

If you need to apply differing conditions to any resource type, you must use a `union` keyword instead of comma separation between the joined queries. The following shows the syntax for a query for multiple resource types where some of the resource types share conditions while others do not:

```
query <resourceType>, <resourceType> resources where <conditions> union <resourceType> resources
```

For example:

```
query group, user resources where displayName = 'administrator' union compartment resources
```

The preceding example returns all groups with the display name "administrator" and all users with the display name "administrator," with any variation in casing, and all compartment resources.

Or, for example:

```
query group resources union user resources where displayName = 'administrator' union compartment resources
```

The preceding example returns all groups and all compartments. It also returns all users with the display name "administrator," with any variation in casing.

Optionally, you can add the `sorted by` clause to the end of the query to order all results in ascending or descending order.

## Sample Queries

This topic provides an explanation of sample queries, including what results to expect from a given sample query. For more information about the syntax for constructing a query, see [Search Language Syntax](#).

### *Example Values*

Sample queries show example values for resource attributes. Replace those examples with values from your own tenancy.

Search provides the following sample queries in the Console:

- Query for everything
- Query for everything, sorted by time created
- Query for volumes and users
- Query for volumes and users, sorted by time created
- Query for volumes and users that have any indexed field matching "production," sorted by time created
- Query for all resources that have a specific freeform tag
- Query for all resources that have one of two specific defined tags
- Query for instances in a "Running" state
- Query for instances in either a "Terminated" or "Terminating" state
- Query for all resources in a specific compartment
- Query for all instances due for a maintenance reboot

### Query All Resources

**Query name:** Query for everything

**Expected results:** Returns all supported resources in the tenancy across all compartments. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
 all resources
```

### Query All Resources, Sort by Time Created

**Query name:** Query for everything, sorted by timeCreated

**Expected results:** Returns all supported resources in the tenancy across all compartments, listed in order of time created, from newest to oldest.

**Sample query language:**

```
query
 all resources
sorted by timeCreated desc
```

### Query Volumes and Users

**Query name:** Query for volumes and users

**Expected results:** Returns all block volumes and users in the tenancy. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
 volume, user resources
```

### Query All Volumes and Users, Sort by Time Created

**Query name:** Query for volumes and users, sorted by timeCreated

**Expected results:** Returns all block volumes and users in the tenancy, listed in order of time created, from newest to oldest

**Sample query language:**

```
query
 volume, user resources
sorted by timeCreated desc
```

### Query Volumes and Users Matching "Production," Sorted by Time Created

**Query name:** Query for volumes and users, with anything matching production, sorted by timeCreated

**Expected results:** Returns all block volumes and users in the tenancy that have any indexed fields that exactly or partially match the search string "production", irrespective of casing. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
 volume, user resources
matching 'production'
sorted by timeCreated desc
```

### Query All Resources With Specific Freeform Tags

**Query name:** Query for all resources that have specific freeform tags

**Expected results:** Returns all resources in the tenancy that have a freeform tag of "costcenter" with a value of "1234."

**Sample query language:**

```
query
 all resources
 where
 (freeformTags.key = 'costcenter' && freeformTags.value = '1234')
```

### Query All Resources According to Defined Tags

**Query name:** Query for all resources that have one of two specific defined tags

**Expected results:** Returns all resources in the tenancy that have either a tag with the key "region" and value "phx" in the tag namespace "categorization," or all resources in the tenancy that have a tag with the key "region" and value "iad" in the namespace "categorization." Ignores casing for all keys and values.

**Sample query language:**

```
query
 all resources
 where
 (definedTags.namespace = 'categorization' && definedTags.key = 'region' && definedTags.value = 'phx')
 ||
 (definedTags.namespace = 'categorization' && definedTags.key = 'region' && definedTags.value = 'iad')
```

### Query Instances According to Specific Lifecycle State

**Query name:** Query for running instances

**Expected results:** Returns all instances in the tenancy in a "Running" state. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
instance resources
 where lifeCycleState = 'RUNNING'
```

### Query Instances According to One of Two Lifecycle States

**Query name:** Query for instances terminated or terminating

**Expected results:** Returns all instances in the tenancy in either a "Terminated" or "Terminating" state. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
instance resources
 where lifeCycleState = 'TERMINATED' || lifeCycleState = 'TERMINATING'
```

### Query All Resources According to Compartment ID

**Query name:** Query for all resources in a compartment

**Expected results:** Returns all resources in the tenancy with a specific compartment ID. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
all resources
 where compartmentId = 'compartmentOcid'
```

### Query All Instances Due for Maintenance Reboot

**Query name:** Query for all instances which have an upcoming scheduled maintenance reboot

**Expected results:** Returns all instances in the tenancy with a scheduled maintenance reboot time value of "now." Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
instance resources
 where timeMaintenanceRebootDue = 'now'
```

### Query All Always Free Resources

**Query name:** Query for all resources that are Always Free

**Expected results:** Returns all resources in the tenancy that are free of charge for the life of the account. Lists results, by default, in order of time created, from newest to oldest.

**Sample query language:**

```
query
all resources
 where
 systemTags.namespace = 'orcl-cloud' &&
 systemTags.key = 'free-tier-retained' &&
 systemTags.value = 'true'
```

## Querying Resources

This topic describes how to find Oracle Cloud Infrastructure resources in your tenancy by performing a free text search or running a query. Queries let you find resources according to specific fields and conditions while free text searches locate resources with the desired text anywhere in the resource metadata.



### Note

#### *Supported Resources and Query Language Syntax*

Search bases search results on supported resources. To see what Oracle Cloud Infrastructure services and resources Search supports, see the [Supported Resources](#) section of [Overview of Search](#).

Furthermore, if free text search results do not produce the results you want, you might need to run a query using query language syntax. For more information about syntax for queries, see [Search Language Syntax](#).

## Using the Console

You can find resources by doing one of the following:

- typing free-form text for a free text search
- typing a query (based on resource query language syntax)
- modifying a sample query

By default, text entered into the **Search** box is interpreted as a free text search.

### To perform a free text search

1. Open the Console, and then, in the top navigation bar, click **Search**.
2. Type the free-form text you want to search for, and then press ENTER.
3. In the left-hand navigation bar, under **Resource Type**, click the type of supported resource to filter query results.
4. (Optional) If you do not see the results that you expect, you can refine your search with

a query by clicking **Advanced Search**. Then, follow the instructions in [To run a custom, free-form query](#) or [To run a sample query](#).

Except for volume backups, compartments, and subnets, you can click the display name of the resource to view details. Search results are eventually consistent, but might not immediately include resources that you created very recently.

### To run a custom, free-form query

1. In the Console, if you are not already there, append "/a/query" to the end of your base Console URL. For example, <https://console.us-ashburn-1.oraclecloud.com/a/query>.
2. In the query text box, type a query using query language syntax, and then click **Search**.
3. In the left-hand navigation bar, under **Resource Type**, click the type of supported resource to filter query results.

Except for volume backups, compartments, and subnets, you can click the display name of the resource to view details. Query results are eventually consistent, but might not immediately include resources that you created very recently.

### To run a sample query

1. In the Console, if you are not already there, append "/a/query" to the end of your base Console URL. For example, <https://console.us-ashburn-1.oraclecloud.com/a/query>.
2. Click **Select Sample Query**, and then click one of the listed sample queries. For an explanation of sample queries, see [Sample Queries](#).
3. Verify that the query language in the query text box satisfies your needs. Change all example values. Add, delete, or modify clauses, as appropriate, and then click **Search**.
4. In the left-hand navigation bar, under **Resource Type**, click the type of supported resource to filter query results.

Except for volume backups, compartments, and subnets, you can click the display name of the resource to view details. Query results are eventually consistent, but might not immediately include resources that you created very recently.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to search for resources or find out what resources you can search for:

- [SearchResources](#)
- [ListResourceTypes](#)

#### **Example: Finding Instance Resources With a Specific Defined Tag**

This section describes how to use the API to query for a specific type of resource based on the resource's defined tags.

The following query will find instances with a defined tag within the namespace "rqs", where the tag's key is "costcenter" and the key's value is "1234".

```
query
instance resources
where
 (definedTags.namespace = 'rqs' && definedTags.key = 'costcenter' && definedTags.value = '1234')
```

When you use the [SearchResources](#) operation to issue the query, the request will look similar to the following. (This example purposefully omits the authorization header and other headers.)

```
POST /20180409/resources
Host: query.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
 "type": "Structured",
 "query": "query instance resources where (definedTags.namespace = 'rqs' && definedTags.key =
```

## CHAPTER 28 Search

---

```
'costcenter' && definedTags.value = '1234)'),
 "matchingContextType": "HIGHLIGHTS"
}
```

If your query produces results, the response will list the resources that match the resource type and tag that you specified. The response will look similar to the following:

```
{
 "items" : [{
 "resourceType" : "Instance",
 "identifier" :
"ocidl.instance.oc1.phx.exampleawcbfhncvbh3siw2svbpgr3bopovy6hgnywfauxqo37ckdmr6hjya",
 "compartmentId" : "ocidl.tenancy.oc1..examplea46vssm7l5wsk5qa7cvbl63ctajep4bh6lv4vaifauxz6ec7jzg4q",
 "timeCreated" : "2018-10-31T22:48:47.855Z",
 "displayName" : "service-pkgs",
 "availabilityDomain" : "ABCd:PHX-AD-1",
 "lifecycleState" : "RUNNING",
 "freeformTags" : { },
 "definedTags" : {
 "rqs" : {
 "costcenter" : "1234"
 }
 },
 "searchContext" : null
 }, {
 "resourceType" : "Instance",
 "identifier" :
"ocidl.instance.oc1.phx.exampleanb3poce6z4omcvbzw66epp3pvbbww6hq7e2jfaux2lxvi3daxhra",
 "compartmentId" :
"ocidl.compartment.oc1..examplea43m3udlwrzwmbevbk5hm3umk2khgfhjcgdtawjlfauuxqwsjiya",
 "timeCreated" : "2018-10-09T23:35:30.167Z",
 "displayName" : "prod-test",
 "availabilityDomain" : "ABCd:PHX-AD-2",
 "lifecycleState" : "RUNNING",
 "freeformTags" : { },
 "definedTags" : {
 "rqs" : {
 "costcenter" : "1234"
 }
 },
 "searchContext" : null
 }, {
 "resourceType" : "Instance",
```

## CHAPTER 28 Search

---

```
"identifier" :
"ocidl.instance.oc1.phx.examples7cz4z6b5hpdly2cvb56obhaiy4gvh2hdpz4akq4fauxpakvlqgya",
 "compartmentId" : "ocidl.tenancy.oc1..examplea46vssm715wsk5qa7cvb163ctajep4bh6lv4fauxf4iz6ec7jzg4q",
 "timeCreated" : "2018-06-12T19:45:24.945Z",
 "displayName" : "BackupTest",
 "availabilityDomain" : "ABCd:PHX-AD-3",
 "lifecycleState" : "STOPPED",
 "freeformTags" : { },
 "definedTags" : {
 "rqs" : {
 "costcenter" : "1234"
 }
 },
 "searchContext" : null
}, {
 "resourceType" : "Instance",
 "identifier" :
"ocidl.compartment.oc1..exampleaexfjsiad7gbi6r4hvmcvbk3a5hgkvutlswf54ulfauxks4p2jasq",
 "compartmentId" : "ocidl.tenancy.oc1..examplea46vssm715cvb5qa7gg5163ctajep4bh6lv4fauxf4iz6ec7jzg4q",
 "timeCreated" : "2018-06-12T19:25:16.942Z",
 "displayName" : "personal_abc",
 "availabilityDomain" : "ABCd:PHX-AD-2",
 "lifecycleState" : "TERMINATED",
 "freeformTags" : { },
 "definedTags" : {
 "rqs" : {
 "costcenter" : "1234"
 }
 },
 "searchContext" : null
}, {
 "resourceType" : "Instance",
 "identifier" :
"ocidl.compartment.oc1..examplealrskzczjqmrb3cvbj4yxdvqxahhffauxtu24tk5dhikoff4uliha",
 "compartmentId" : "ocidl.tenancy.oc1..examplea46vssm715wsk5qa7gg5163cvbjep4bh6lv4fauxf4iz6ec7jzg4q",
 "timeCreated" : "2018-11-29T23:40:29.005Z",
 "displayName" : "test_unused",
 "availabilityDomain" : null,
 "lifecycleState" : "AVAILABLE",
 "freeformTags" : { },
 "definedTags" : {
 "rqs" : {
```

```
 "costcenter" : "1234"
 }
},
"searchContext" : null
}]
}
```

With these results, you can take additional action, if needed. For more information about a resource type, such as its attributes, see its reference page in the API Reference Guide. For the reference pages of resource types that have been indexed for Search, see [Supported Resources](#).

## Troubleshooting Search

This topic covers common issues related to Search and how you can address them:

- [Query or Search Results are Not as Expected](#)

### Query or Search Results are Not as Expected

There are several reasons why you might not see results that you expect from a search or query.

Not all resource types have been indexed for Search. For a list of currently supported resource types, see [Supported Resources](#).

You might not have the required permissions for the resource type that you want to view in search or query results. If there's no policy that grants you the permissions you need, then an administrator must create one for you or add you to a group that's already named in a policy. For more information, see [Details for Search](#).

The query syntax you used might need adjustment. Verify that the conditions in your query language haven't restricted the results to a narrower set than you intended.

If you recently created a resource, it might not show up in search results immediately. Similarly, if you recently updated a resource, your changes might not immediately appear. At times, you might see a resource in a list view before you can see it in search results. The Search service is eventually consistent. Wait, and then try again.



# CHAPTER 29 Security Guide and Announcements

This section of the Oracle Cloud Infrastructure documentation provides a guide to help you securely configure services and resources, and timely announcements relevant to emerging security issues.

- [Oracle Cloud Infrastructure Security Guide](#)
- [Oracle Cloud Security Response to Intel L1TF Vulnerabilities](#)
- [Oracle Cloud Security Response to Intel Microarchitectural Data Sampling \(MDS\) Vulnerabilities](#)

## Oracle Cloud Infrastructure Security Guide

Oracle Cloud Infrastructure enables enterprises to maximize the number of mission-critical workloads that they can migrate to the cloud while continuing to maintain their desired security posture and reduce the overhead of building and operating data-center infrastructure. With Oracle Cloud Infrastructure, enterprise customers get unparalleled control of and transparency into their applications running in the cloud, including:

- Customer isolation that allows you to deploy your application and data assets in an environment that commits full isolation from other tenants and Oracle's staff, as well as between the same tenant's workloads.
- Always-on encryption that protects customer data at-rest and HTTPS-only public APIs.
- Easy-to-use security policy that allows you to constrain access to your services and segregate operational responsibilities to reduce risk associated with malicious and accidental user actions.
- Comprehensive log data that allows you to audit and monitor actions on your resources, helping you to meet your audit requirements while reducing security and operational risk.
- Identity federation that allows you to use your existing users and groups in the cloud.

- Support for bringing in third-party software solutions for protecting customer data and resources in the cloud.
- Fault-independent data centers that enable high availability scale-out architectures and are resilient against network attacks, ensuring constant uptime in the face of disaster and security attack.
- Rigorous internal processes and use of effective security controls in all phases of cloud service development and operation.
- Adherence to Oracle's strict security standards through third-party audits, certifications, and attestations. Oracle helps customers demonstrate compliance readiness to internal security and compliance teams, their customers, auditors, and regulators.

All of the Oracle Cloud Infrastructure security capabilities have been designed with one goal in mind: allowing you to run your mission-critical workloads in the cloud with complete control and confidence. Oracle continues to invest in the above areas and more to offer unmatched security and assurance to enterprise customers.

For an overview of Oracle Cloud Infrastructure's security, see [Security Overview](#).

For service-specific best practices and policy examples, see [Security Best Practices](#).

### Security Overview

Oracle's mission is to build cloud infrastructure and platform services for your business to have effective and manageable security to run your mission-critical workloads and store your data with confidence.

Oracle Cloud Infrastructure's security approach is based on seven core pillars. Each pillar has multiple solutions designed to maximize the security and compliance of the platform.

#### **CUSTOMER ISOLATION**

Allow customers to deploy their application and data assets in an environment that commits full isolation from other tenants and Oracle's staff.

### **DATA ENCRYPTION**

Protect customer data at-rest and in-transit in a way that allows customers to meet their security and compliance requirements for cryptographic algorithms and key management.

### **SECURITY CONTROLS**

Offer customers effective and easy-to-use security management solutions that allow them to constrain access to their services and segregate operational responsibilities to reduce risk associated with malicious and accidental user actions.

### **VISIBILITY**

Offer customers comprehensive log data and security analytics that they can use to audit and monitor actions on their resources, allowing them to meet their audit requirements and reduce security and operational risk.

### **SECURE HYBRID CLOUD**

Enable customers to use their existing security assets, such as user accounts and policies, as well as third-party security solutions when accessing their cloud resources and securing their data and application assets in the cloud.

### **HIGH AVAILABILITY**

Offer fault-independent data centers that enable high availability scale-out architectures and are resilient against network attacks, ensuring constant uptime in the face of disaster and security attack.

### **VERIFIABLY SECURE INFRASTRUCTURE**

Follow rigorous processes and use effective security controls in all phases of cloud service development and operation. Demonstrate adherence to Oracle's strict security standards through third-party audits, certifications, and attestations. Help customers demonstrate compliance readiness to internal security and compliance teams, their customers, auditors, and regulators.

Also, Oracle employs some of the world's foremost security experts in information, database, application, infrastructure, and network security. By using Oracle Cloud Infrastructure, our customers directly benefit from Oracle's deep expertise and continuous investments in security.

### Basic Security Considerations

The following principles are fundamental to using any application securely:

- Keep software up-to-date. This includes the latest product release and any patches that apply to it.
- Limit privileges as much as possible. Users should be given only the access necessary to perform their work. User privileges should be reviewed periodically to determine relevance to current work requirements.
- Monitor system activity. Establish who should access which system components, and how often, and monitor those components.
- Learn about and use the Oracle Cloud Infrastructure security features. For more information, see [Security Services and Features](#).
- Use secure best practices. For more information, see [Security Best Practices](#).
- Keep up-to-date on security information. Oracle regularly issues security-related patch updates and security alerts. Install all security patches as soon as possible. See the [Critical Patch Updates and Security Alerts](#) website.

### Understanding the Oracle Cloud Infrastructure Environment

When planning your Oracle Cloud Infrastructure deployment, consider the following:

#### WHICH RESOURCES MUST BE PROTECTED?

- Protect customer data, such as credit card numbers.
- Protect internal data, such as proprietary source code.
- Protect system components from being disabled by external attacks or intentional system overloads.

#### WHO ARE YOU PROTECTING DATA FROM?

For example, you must protect your subscribers' data from other subscribers, but someone in your organization needs to access that data to manage it. Analyze your workflows to determine who needs access to the data. Consider carefully how much access to give a

system administrator; it is possible that a system administrator can manage your system components without needing to access the system data.

### **WHAT WILL HAPPEN IF PROTECTIONS ON A STRATEGIC RESOURCE FAIL?**

Sometimes, a fault in your security scheme is nothing more than an inconvenience. In other cases, a fault might cause great damage to you or your customers. Understanding the security ramifications of each resource will help you protect it properly.

### **Shared Security Model**

Oracle Cloud Infrastructure offers best-in-class security technology and operational processes to secure its enterprise cloud services. However, for you to securely run your workloads in Oracle Cloud Infrastructure, you must be aware of your security and compliance responsibilities. By design, Oracle provides security of cloud infrastructure and operations (cloud operator access controls, infrastructure security patching, and so on), and you are responsible for securely configuring your cloud resources. Security in the cloud is a shared responsibility between you and Oracle.

In a shared, multi-tenant compute environment, Oracle is responsible for the security of the underlying cloud infrastructure (such as data-center facilities, and hardware and software systems) and you are responsible for securing your workloads and configuring your services (such as compute, network, storage, and database) securely.

In a fully isolated, single-tenant, bare metal server with no Oracle software on it, your responsibility increases as you bring the entire software stack (operating systems and above) on which you deploy your applications. In this environment, you are responsible for securing your workloads, and configuring your services (compute, network, storage, database) securely, and ensuring that the software components that you run on the bare metal servers are configured, deployed, and managed securely.

More specifically, your and Oracle's responsibilities can be divided into the following areas:

- **Identity and Access Management (IAM):** As with all Oracle cloud services, you should protect your cloud access credentials and set up individual user accounts. You are responsible for managing and reviewing access for your own employee accounts

and for all activities that occur under your tenancy. Oracle is responsible for providing effective IAM services such as identity management, authentication, authorization, and auditing.

- **Workload Security:** You are responsible for protecting and securing the operating system and application layers of your compute instances from attacks and compromises. This protection includes patching applications and operating systems, operating system configuration, and protection against malware and network attacks. Oracle is responsible for providing secure images that are hardened and have the latest patches. Also, Oracle makes it simple for you to bring the same third-party security solutions that you use today.
- **Data Classification and Compliance:** You are responsible for correctly classifying and labeling your data and meeting any compliance obligations. Also, you are responsible for auditing your solutions to ensure that they meet your compliance obligations.
- **Host Infrastructure Security:** You are responsible for securely configuring and managing your compute (virtual hosts, containers), storage (object, local storage, block volumes), and platform (database configuration) services. Oracle has a shared responsibility with you to ensure that the service is optimally configured and secured. This responsibility includes hypervisor security and the configuration of the permissions and network access controls required to ensure that hosts can communicate correctly and that devices are able to attach or mount the correct storage devices.
- **Network Security:** You are responsible for securely configuring network elements such as virtual networking, load balancing, DNS, and gateways. Oracle is responsible for providing a secure network infrastructure.
- **Client and Endpoint Protection:** Your enterprise uses various hardware and software systems, such as mobile devices and browsers, to access your cloud resources. You are responsible for securing all clients and endpoints that you allow to access Oracle Cloud Infrastructure services.
- **Physical Security:** Oracle is responsible for protecting the global infrastructure that runs all of the services offered in Oracle Cloud Infrastructure. This infrastructure

consists of the hardware, software, networking, and facilities that run Oracle Cloud Infrastructure services.

For information about using security credentials to access Oracle Cloud Infrastructure, see [Security Credentials](#).

### Infrastructure Security

Our security model is built around people, process, tooling, and a common security “platform” of methodologies and approaches from which we build our products. We apply this model to our core security components of Security Culture, Security Design and Controls, Secure Software Development, Personnel Security, Physical Security, and Security Operations that we use to protect and secure our customers and business.

#### SECURITY CULTURE

We believe that a dynamic security-first culture is vital to building a successful security-minded organization. We have cultivated a holistic approach to security culture in which all our team members internalize the role that security plays in our business and are actively engaged in managing and improving our products' security posture. We have also implemented mechanisms that assist us in creating and maintaining a security-aware culture.

- **Security-minded leadership:** Our senior leadership is actively involved in our security planning, monitoring and management. We define and measure ourselves against security metrics and include security as a component of our team evaluation processes.
- **Embedded expertise:** To help with driving security practices within our team, we have an embedded security-engineering model with security team members sitting and working with our product development teams. This approach enables our security organization to build deep understanding of the product-development processes and system architectures. We are also able to better assist teams in solving security challenges in real time and drive security initiatives more effectively.
- **Common security standards:** We actively work to integrate security into our products and operations. One way we have done this is to establish a security standards baseline. Our objective in creating this baseline is to provide a single security point of

reference for business that establishes clear and actionable guidelines. The security baseline is updated frequently to incorporate learned lessons and reflect emerging business factors. We have also created a series of support materials to assist our teams in implementing security controls including reference architectures, implementation guides, and access to security experts.

- **Values of openness, constructive debate, and encouraged escalation:** Security issues can be addressed only when the people who can fix them are aware of them. We believe that openness and transparency, constructive debate, and encouraged escalation make us stronger. We encourage escalation, and we work to create an environment where raising issues early and often is rewarded.
- **Security training awareness:** We maintain robust security and awareness training programs that raise awareness and reinforce our security culture. We require in-depth security training sessions for all new employees as well as annual refresher trainings, and we provide security training that is tailored to our employees' specific job roles. All our software developers undergo a secure development training that establishes baseline security requirements for product development and provides best practices. We also work to provide engaging and innovative forms of security awareness training such as guest speakers and interactive forums (and we're not above providing food, drinks, or swag to drive attendance).

### SECURITY DESIGNS AND CONTROLS

Security is integrated into our products and operations through our Oracle Cloud Infrastructure Methodology. This centralized methodology defines our approach for the core security areas that form the security foundation from which we build our products. This approach lends itself to agility and helps us apply best practices and lessons learned from one product across the business, thus raising the security of all our products.

- **User authentication and access control:** Least-privilege access is used to grant access to production systems, and the approved lists of service team members are periodically reviewed to revoke access when there is no justifiable need. Access to production environments requires multi-factor authentication (MFA). The MFA tokens are granted by the security team, and tokens of inactive members are disabled. All access to production systems is logged, and the logs are stored for security analysis.

- **Change management:** Oracle Cloud Infrastructure follows a defined and rigorous change management and deployment process that uses purpose-built proprietary testing and deployment tools. All changes deployed into our production environment follow a testing and approval process prior to release. This process is designed to ensure that changes operate as intended, and can otherwise be rolled back to a previous known good state to recover gracefully from unforeseen bugs or operational issues. We also track the integrity of critical system configurations to ensure that they align with expected state.
- **Vulnerability management:** We use both internal penetration testing teams and external industry experts to help us identify potential vulnerabilities in our products. These exercises help us improve the security of our products, and we work to incorporate the lessons that we learn into our future development work. Oracle Cloud Infrastructure hosts undergo periodic vulnerability scanning using industry-standard scanners. Scan results are triaged to validate applicability of findings to the Oracle Cloud Infrastructure environment, and that applicable findings are patched by our product teams.
- **Incident response:** We have developed strong processes and mechanisms to enable us to respond to and address incidents as they arise. We maintain 24/7 incident response teams ready to detect and respond to events. Our critical staff members carry paging devices that enable us to call on the expertise needed to bring issues to resolution. We have also built a process to help us learn from our incidents. We perform root cause analysis through our Corrective Action/Preventative Action (CAPA) process. CAPAs are intended to discover process gaps and changes that should be made by the business after an incident. CAPAs act as a common language that we can use to reflect on an issue and capture concrete steps to improve future operational readiness. CAPAs capture the root cause of an issue, what is required to contain or fix the issue, and what steps we must take to ensure that the issue does not recur. Our leadership team reviews all CAPAs, looks for cross-organizational applications for learned lessons, and ensures that actions are implemented in a timely manner.
- **Security logging and monitoring:** We have created automated mechanisms to log various security-relevant events (for example, API calls and network events) in the

infrastructure, and monitor the logs for anomalous behavior. Alerts generated by monitoring mechanisms are tracked and triaged by the security team.

- **Network security:** By default, customer communications with Oracle Cloud Infrastructure services are done using the latest TLS ciphers and configuration to secure customer data in transit, and hinder any man-in-the-middle attacks. As a further defense in depth, customer commands to the services are digitally signed using public keys, to prevent any tampering. The services also deploy proven, industry-leading tools and mechanisms to mitigate distributed denial of service (DDoS) attacks and maintain high availability.
- **Control plane security:** Oracle Cloud Infrastructure back-end (control plane) hosts are securely isolated from customer instances by using network ACLs. Provisioning and management of customer instances are done by software agents that must interact with the backend hosts. Only authenticated and authorized software agents can successfully interact with Oracle Cloud Infrastructure back-end hosts. For back-end hosts, pre-production environments (for example, dev, test, and integ) are separated from production environments so that any development and test activities do not have any impact on production systems.
- **Server security and media management:** Oracle has a long history of enterprise-class secure hardware development. Our Hardware Security team is responsible for designing and testing the security of the hardware used to deliver Oracle Cloud Infrastructure services. This team works with our supply chain and tests hardware components to validate them against rigorous Oracle Cloud Infrastructure hardware security standards. This team also works closely with our product development functions to ensure that hardware can be returned to a pristine, safe state after being released by customers.
- **Secure host wipe and media destruction:** Oracle Cloud Infrastructure instances are securely wiped after hardware is released by customers. This secure wipe restores hardware to a pristine state. We have re-engineered the platform with proprietary hardware components that allow us to wipe and reinitialize the hardware in a secure manner. When the underlying hardware has reached end-of-life, it is securely destroyed. Before leaving our data centers, drives are rendered unusable by using industry-leading media destruction devices.

### **SECURE SOFTWARE DEVELOPMENT**

Secure product development requires consistently applied methodologies that conform to clear security objectives and principles. We build security practices into every element of our product development life cycle. Oracle employs formal secure product development standards that are a roadmap and guide for developers. These standards discuss general security knowledge areas such as design principles and common vulnerabilities, and provide specific guidance on topics such as data validation, data privacy, and user management.

Oracle secure product development standards have evolved and expanded over time to address the common issues affecting code, new threats as they are discovered, and new use cases by Oracle customers. The standards incorporate insights and learned lessons; they do not live in a vacuum, nor are they an “after the fact” addendum to software development. They are integral to language-specific standards such as C/C++, Java, PL/SQL, and others, and are a cornerstone to Oracle's secure development programs and processes.

Security assurance analysis and testing verify security qualities of Oracle products against various types of attacks. There are two broad categories of tests employed for testing Oracle products: static and dynamic analysis. These tests fit differently in the product development lifecycle and tend to find different categories of issues, so they are used together by Oracle product teams.

### **PERSONNEL SECURITY**

Our people make our business. We strive to hire the best, and we invest in and continue to develop our employees. We value training, and we require not only baseline security training for all our employees but also specialized training to keep our teams abreast of the latest security technologies, exploits, and methodologies. In addition to standard annual corporate training programs that cover our information security and privacy programs (among many others), we engage with a broad spectrum of industry groups and send our employees to specialist conferences to collaborate with other industry experts on emerging challenges. The objectives of our security training programs are to help our employees better protect our customers and products, to enable employees to grow in their knowledge areas around security, and to further our mission to attract and retain the best talent.

We work to recruit the best talent for our team as we grow, and we hire people with strong ethics and good judgment. All our employees undergo pre-employment screening as

permitted by law, including criminal background checks and prior-employment validation. We also maintain performance evaluation processes to recognize good performance and help our teams and employees identify opportunities for growth. We maintain both team and employee evaluation processes, and we use security as a component of our team evaluation processes. This approach provides our teams and leadership visibility into how our teams are performing against our security standards and enables us to identify best practices and improvement areas for critical security processes.

### **PHYSICAL SECURITY**

Oracle Cloud Infrastructure data centers are designed for security and availability of customer data. This approach begins with our site selection process. Candidate build sites and provider locations undergo an extensive risk evaluation process that considers environmental threats, power availability and stability, vendor reputation and history, neighboring facility functions (for example, high-risk manufacturing or high-threat targets), and geopolitical considerations, among other criteria.

Oracle Cloud Infrastructure data centers align with Uptime Institute and Telecommunications Industry Association (TIA) ANSI/TIA-942-A Tier 3 or Tier 4 standards and follow a N2 redundancy methodology for critical equipment operation. Data centers housing Oracle Cloud Infrastructure services use redundant power sources and maintain generator backups in case of widespread electrical outage. Server rooms are closely monitored for air temperature and humidity, and fire suppression systems are in place. Data center staff are trained in incident response and escalation procedures to address security or availability events that may arise.

We take a layered approach to physical security that starts with the site build. Oracle Cloud Infrastructure data center facilities are durably built with steel, concrete, or comparable materials and are designed to withstand impact from a light vehicle strike. Our sites are staffed with security guards who are ready to respond to incidents 24 hours a day, 7 days a week, 365 days a year. The exterior of the sites is secured with perimeter barriers and vehicle checks are actively monitored by a guard force and cameras that cover the building perimeter.

All persons entering our data centers must first go through a layer of security at the site entrances, which are staffed with security guards. Persons without site-specific security badges entering the site must present government-issued identification and have an approved

access request granting them access to the data center building. All employees and visitors must wear visible, official identification badges at all times. There are additional security layers between the entrance and server rooms that vary depending on the site build and risk profile. Data center server rooms are built with additional security layers including cameras that cover server rooms, two-factor access control, and intrusion-detection mechanisms. Physical barriers are in place to create isolated security zones around server and networking racks that span from the floor (including below the raised floor where applicable) to the ceiling (including above ceiling tiles where applicable).

Access to Oracle Cloud Infrastructure data centers is carefully controlled and follows a least-privilege access approach. All access to server rooms must be approved by authorized personnel and is granted only for the necessary period. Access usage is audited, and access provisioned within the system is periodically reviewed by data-center leadership. Server rooms are isolated into secure zones that are managed on a zone-by-zone basis, and access is provisioned only for those zones required by personnel.

### **SECURITY OPERATIONS**

The Oracle Cloud Infrastructure Security Operations team is responsible for monitoring and securing the unique Oracle Cloud Infrastructure hosting and virtual networking technologies. The team works and trains directly with the Oracle engineers who develop these technologies to leverage the unique security and introspection capabilities they provide.

We monitor emerging internet security threats daily and implement appropriate response and defense plans to address risks to the business. When we determine that urgent changes are recommended that are within the scope of the customers' responsibilities, we issue security alert bulletins to those customers to ensure their protection.

In the case of a detected or reported security issue that affects Oracle Cloud Infrastructure servers or networks, Security Operations staff is available 24/7 to respond, escalate, or take required corrective action. When necessary, we will escalate and coordinate with external parties (including network and hosting service providers, hardware vendors, or law enforcement) to protect Oracle Cloud Infrastructure, our customers, and our network's security and reputation.

All actions performed in response to a security issue by the Security Operations team are done according to our documented process, and are logged in accordance with compliance

requirements. Care is always taken to protect the goals of service and data integrity, privacy, and business continuity.

### **Customer Data Protection**

#### **DATA RIGHTS AND OWNERSHIP**

Oracle Cloud Infrastructure customers retain all ownership and intellectual property rights in and to their content. Customer data protection is critically important, and we strive to be transparent with our data protection processes as well as law enforcement requests that we might receive.

#### **DATA PRIVACY**

Oracle complies with the EU-U.S. Privacy Shield Framework as set forth by the U.S. Department of Commerce regarding the collection, use, and retention of personal information transferred from the European Union to the United States. Oracle is also responsible for ensuring that third parties who act as an agent on our behalf do the same.

Oracle has certified to the Department of Commerce that it adheres to the Privacy Shield Principles. If there is any conflict between the terms in our privacy policy and the Privacy Shield Principles, the Privacy Shield Principles shall govern. To learn more about the Privacy Shield program, and to view our certification, visit <https://www.privacyshield.gov/list>.

For personal information received or transferred pursuant to the Privacy Shield Framework, Oracle is subject to the regulatory enforcement powers of the U.S. Federal Trade Commission.

Oracle continues to adhere to the underlying European privacy principles of the U.S.-Swiss Safe Harbor for the processing of Personal Information received from Switzerland. To learn more about the Safe Harbor program, and to view our certification, visit <https://safeharbor.export.gov/swisslist.aspx>.

#### **LAW ENFORCEMENT REQUESTS**

Except as otherwise required by law, Oracle will promptly notify customers of any subpoena, judicial, administrative or arbitral order of an executive or administrative agency or other governmental authority that it receives and which relates to the personal data Oracle is

processing on the customer's behalf. Upon customer request, Oracle will provide customers with reasonable information in its possession relevant to the law enforcement request and any assistance reasonably required for them to respond to the request in a timely manner.

### **COMPLIANCE**

Oracle Cloud Infrastructure is built for enterprises. We operate under practices aligned with the ISO/IEC 27002 Code of Practice for information security controls, from which we have identified a comprehensive set of security controls that apply to our business. Oracle Cloud Infrastructure is still a new product line, and we must operate for a period of time in order for these security controls and our operations to undergo external audit. As an enterprise cloud, we plan to pursue a broad suite of industry and government certifications, audits, and regulatory programs.

## Security Services and Features

A key objective of Oracle Cloud Infrastructure is to allow you to create a logical extension of your on-premises infrastructure and data centers in Oracle Cloud Infrastructure. You can gain the benefits of a modern public cloud without having to compromise or reinvent your existing security posture. This idea was central to the design of all our infrastructure and services.

### **Regions and Availability Domains**

To provide data availability and durability, Oracle Cloud Infrastructure enables you to select from infrastructure with distinct geographic and threat profiles. A region is the top-level component of the infrastructure. Each region is a separate geographic area with multiple, fault-isolated locations called availability domains. An availability domain is the subcomponent of a region and is independent and highly reliable. Each availability domain is built with fully independent infrastructure: buildings, power generators, cooling equipment, and network connectivity. With physical separation comes protection against natural and other disasters. Availability domains within the same region are connected by a secure, high-speed, low-latency network, which allows customers to build and run highly reliable applications and workloads with minimum impact to application latency and performance. All links between availability domains are encrypted. Each region has one or more availability domains, each allowing customers to deploy highly available applications.

### Identity and Access Management (IAM) Service

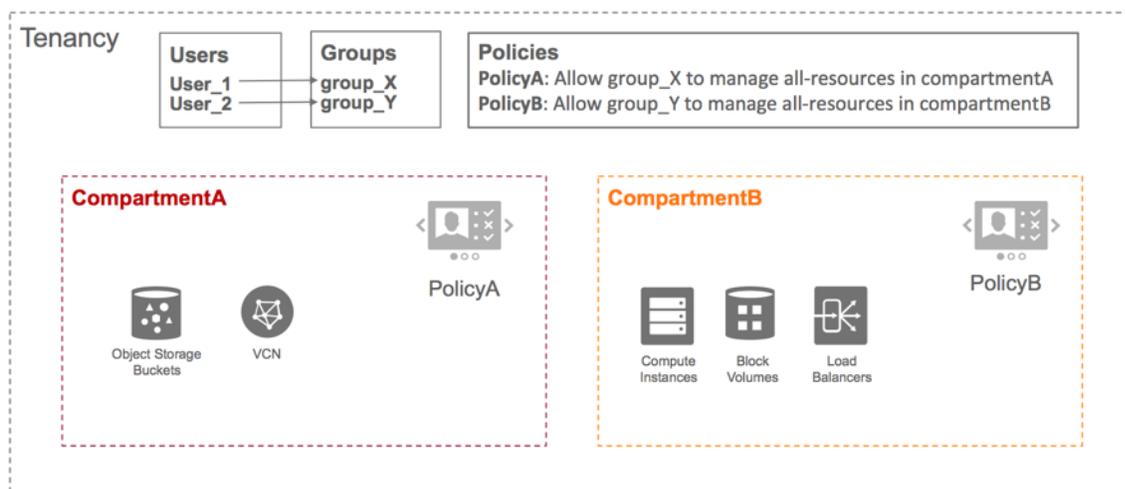
The Oracle Cloud Infrastructure Identity and Access Management (IAM) service is built to meet the requirements of enterprises, and it provides authentication and authorization for all their Oracle Cloud Infrastructure resources and services. An enterprise can use a single tenancy shared by various business units, teams, and individuals while maintaining security, isolation, and governance.

When a customer joins Oracle Cloud Infrastructure, a tenancy is created. A tenancy is a virtual construct that contains all of the Oracle Cloud Infrastructure resources that belong to the customer. The administrator of the tenancy can create users and groups and assign them least-privileged access to resources that are partitioned into compartments. A compartment is a group of resources that can be managed as a single logical unit, providing a streamlined way to manage large infrastructure. For example, a customer can create a compartment (`HR-Compartment`) to host a specific set of cloud network, compute instances, and storage volumes necessary to host its HR applications. Compartments are a fundamental component of Oracle Cloud Infrastructure for organizing and isolating cloud resources. Customers use them to clearly separate resources for the purposes of isolation (separating the resources for one project or business unit from another). A common approach is to create a compartment for each major part of an organization. Unlike most Oracle Cloud Infrastructure services that are regionally scoped, the IAM service resources are global. Customers can have a single tenancy across multiple regions.

The following are key IAM primitives:

- **Resource:** A cloud object that a company's employees create and use when interacting with Oracle Cloud Infrastructure services, for example, compute instances, block storage volumes, virtual cloud networks (VCNs), subnets, and route tables.
- **Policy:** A set of authorization rules that define access to resources within a tenancy.
- **Compartment:** A heterogeneous collection of resources for the purposes of security isolation and access control.
- **Tenancy:** The root compartment that contains all of an organization's resources. Within a tenancy, administrators can create one or more compartments, create more users and groups, and assign policies that grant groups the ability to use resources within a compartment.

- **User:** A human being or system that needs access to manage their resources. Users must be added to groups in order to access resources. Users have one or more credentials that must be used to authenticate to Oracle Cloud Infrastructure services. Federated users are also supported.
- **Group:** A collection of users who share a similar set of access privileges. Administrators can grant access policies that authorize a group to consume or manage resources within a tenancy. All users in a group inherit the same set of privileges.
- **Identity Provider:** A trusted relationship with a federated identity provider. Federated users who attempt to authenticate to the Oracle Cloud Infrastructure console are redirected to the configured identity provider. After successfully authenticating, federated users can manage Oracle Cloud Infrastructure resources in the console just like a native IAM user. Currently, Oracle Cloud Infrastructure supports the Oracle Identity Cloud Service and Microsoft Active Directory Federation Service (ADFS) as identity providers. Federated groups are mapped to native IAM groups to define the policies apply to a federated user.



All customer calls to access Oracle Cloud Infrastructure resources are first authenticated by the IAM service (or federated provider) and then authorized based on IAM policies. A customer can create a policy that gives a set of users permission to access the infrastructure resources (network, compute, storage, and so on) within a compartment in the tenancy.

## CHAPTER 29 Security Guide and Announcements

---

These policies are flexible and are written in a human-readable form that is easy to understand and audit. A policy contains one or more policy statements that follow this easy-to-understand syntax:

```
Allow group <group_name> to <verb> <resource-type> in compartment <compartment_name>
```

A verb defines the type of access covered. Oracle defines the following verbs that you can use in your policy statements:

- **inspect:** Provides the ability to list resources, without access to any confidential information or user-specified metadata that might be part of that resource.
- **read:** Includes inspect plus the ability to get user-specified metadata and the actual resource itself.
- **use:** Includes read plus the ability to work with existing resources (the actions vary by resource type). Includes the ability to update the resource, except for resource types where the update operation has the same effective impact as the create operation (for example, `UpdatePolicy` and `UpdateSecurityList`). In such cases, the update ability is available only with the manage verb. In general, this verb does not include the ability to create or delete that type of resource.
- **manage:** Includes all permissions for the resource.

The following example policy enables the `GroupAdmins` group to create, update, or delete any groups:

```
Allow group GroupAdmins to manage groups in tenancy
```

Each user has one or more of the following credentials to authenticate themselves to Oracle Cloud Infrastructure. Users can generate and rotate their own credentials. In addition, a tenancy security administrator can reset credentials for any user within their tenancy.

- **Console password:** Used to authenticate a user to the Oracle Cloud Infrastructure Console.
- **API key:** All API calls are signed using a user-specific 2048-bit RSA private key. The user creates a public key pair, and uploads the public key in the Console.
- **Auth token:** Auth tokens are Oracle-generated token strings that you can use to authenticate with third-party APIs that do not support Oracle Cloud Infrastructure's

signature-based authentication. For example, use an auth token to authenticate with a Swift client. To ensure sufficient complexity, the token is created by the IAM service and cannot be provided by a customer.

- **Customer secret key:** Used by Amazon S3 clients to access the Object Storage service's S3-compatible API. To ensure sufficient complexity, the password is created by the IAM service and cannot be provided by a customer.

### Audit Service

The Oracle Cloud Infrastructure Audit service records all API calls to resources in a customer's tenancy as well as login activity from the graphical management Console. Using the Audit service, customers can achieve their own security and compliance goals by monitoring all user activity within their tenancy. Because all Console, SDK, and command line (CLI) calls go through our APIs, all activity from those sources is included. Audit records are available through an authenticated, filterable query API or can be retrieved as batched files from Oracle Cloud Infrastructure Object Storage. Audit log contents include what activity occurred, the user that initiated it, the date and time of the request, as well as source IP, user agent, and HTTP headers of the request.

### Compute Service

Compute is a core component of Oracle Cloud Infrastructure and provides on-demand and elastic compute capabilities with enterprise-grade security and performance. Customers can provision thousands of compute instances and scale them up or down through an easy-to-use web-based management console. Programmatic support to do the same is available through feature-rich SDKs and command-line interfaces (CLIs). All compute instances are hosted in Oracle enterprise-grade data centers.

Compute instances are based on high-performance server hardware that uses latest-generation, multi-core server CPUs, large amounts of memory, and high-throughput NVMe local storage. Oracle Cloud Infrastructure provides bare metal (BM) and virtual machine (VM) instances. Customers can choose instances that fit their performance, cost, and software flexibility requirements.

- **Bare metal instances:** In bare metal instances, physical servers are dedicated to a single customer who has complete control over the server. There is no Oracle-managed hypervisor and Oracle personnel have no access to memory or local (NVMe) storage while the instance is running. All network virtualization is performed off-box and only the Oracle Integrated Lights Out Manager (ILOM) is accessible to the infrastructure (required in order to remotely reboot or reprovision instances). These bare metal instances offer consistent high performance and are immune to any noisy-neighbor issues. Customers have OS-level administrative privileges to the bare metal instance. After a customer terminates their bare metal instance, the server undergoes an automated disk and firmware-level wipe process to ensure isolation between customers.
- **Virtual machine (VM) instances:** Customers with flexibility requirements or those who don't need a dedicated bare metal instance can opt for VMs. Multi-tenant customer VMs in Oracle Cloud Infrastructure are managed by a security hardened hypervisor that provides strong isolation between customers.

Oracle Cloud Infrastructure instances use key-based SSH by default. Customers provide the SSH public keys to Oracle Cloud Infrastructure and securely use the SSH private keys for accessing the instances. Oracle recommends using key-based SSH to access Oracle Cloud Infrastructure instances. Password-based SSH could be susceptible to brute-forcing attacks, and are not recommended.

Oracle Linux images hardened with the latest security updates are available for you to run on Oracle Cloud Infrastructure instances. Oracle Linux images run the Unbreakable Enterprise Kernel (UEK) and support advanced security features such as Ksplice to apply security patches without booting, which allows enterprises to live-update their instances without any disruption. In addition to Oracle Linux, Oracle Cloud Infrastructure makes a growing list of other OS images available, including CentOS, Ubuntu, and Windows Server. You can also bring your own custom images. All Oracle-provided images come with secure defaults including OS-level firewalls turned on by default.

### Networking Service

High-throughput and reliable networking is fundamental to public-cloud infrastructure that delivers compute and storage services at scale. As a result, we invested significant innovation

in Oracle Cloud Infrastructure networking to support requirements of enterprise customers and their workloads. Oracle Cloud Infrastructure regions have been built with a state-of-the-art, non-blocking Clos network that is not over-subscribed and provides customers with a predictable, high-bandwidth, low latency network. The data centers in a region are networked to be highly available and have low-latency connectivity between them.

The Oracle Cloud Infrastructure Networking service offers a customizable private network (a VCN, or virtual cloud network) to customers, which enforces logical isolation of customer Oracle Cloud Infrastructure resources. As with their on-premises network in their data centers, customers can set up a VCN with hosts with private IP addresses, subnets, route tables and gateways using VCN. The VCN can be configured for internet connectivity, or connected to the customer's private data center through an IPSec VPN gateway or FastConnect. FastConnect offers a private connection between an existing network's edge router and Dynamic Routing Gateways (DRG). Traffic does not traverse the internet.

The Networking service also supports bi-directional, stateful and stateless firewalls that allow customers to initialize network security access controls. Firewalls and ACLs specified for a customer VCN are propagated throughout the network topology and control plane, ensuring a multi-tiered and defense-in-depth implementation. Each tenant (customer) can create multiple VCNs to implement logical grouping of their resources.

The following are key networking concepts associated with a VCN:

- **Subnets:** The primary subdivision of a VCN. Subnets have historically been specific to an availability domain, but can now be regional (covering all availability domains in the region). Subnets can be marked as private upon creation, which prevents instances launched in that subnet from having public IP addresses.
- **Internet gateway:** A virtual router that provides public internet connectivity from a VCN. By default, a newly created VCN has no internet connectivity.
- **Dynamic routing gateway:** A virtual router that provides a path for private traffic between a VCN and a data center's network. It is used with an IPSec VPN or Oracle Cloud Infrastructure FastConnect connection to establish private connectivity between a VCN and an on-premises or other cloud network.

- **NAT gateway:** A virtual router that gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections.
- **Service gateway:** A virtual router that gives cloud resources private access to Oracle services such as Object Storage without using an internet gateway or NAT gateway.
- **Routing tables:** Virtual routing tables that give the subnets access to the VCN's gateways (Internet Gateway and Dynamic Routing Gateway). Routes can also use private IPs as a target to implement network functionality such as NAT, firewalls, IDS, and so on.
- **Primary VNICs:** Subnets contain virtual network interface cards (VNICs) that attach to instances. The VNIC determines how the instance connects with endpoints inside and outside the VCN. Each instance has a primary VNIC that is created during instance launch and cannot be removed. During instance launch, the Networking service also assigns a public IP address. Customers can override that behavior during instance launch and request to have no public IP address assigned.
- **Secondary VNICs:** VNICs with public and private IP addresses that can be attached to an instance. In a bring-your-own-hypervisor (BYOH) scenario, where customers can run their hypervisor on a BM instance, a secondary VNIC can be assigned to a VM, to allow VCN networking for the VM. This is very useful for running virtual security appliances in a VCN.
- **IPSec VPN connection:** A secure VPN connection between a VCN and a data center.
- **Security lists:** Virtual firewall rules that define allowed ingress and egress to an instance at the packet level. Individual rules can be defined to be stateful or stateless.

Virtual firewalls are implemented by using VCN security lists. Customers can specify a set of firewall rules and associate them with one or more subnets. Associating a security list with a subnet applies those firewall rules to all instances running inside the subnet, at the packet level. There are two types of firewall rules:

- **Ingress rules:** Ingress rules specify the source (IP CIDR and port range), destination port range, and protocol to match on, and are applied to ingress network connections.

- **Egress rules:** Egress rules specify the destination (IP CIDR and port range), source port range, and protocol to match on, and are applied to egress network connections.

Every VCN has a default security list customers may optionally use that allows only SSH and certain types of important ICMP ingress traffic, and all egress traffic. Customers can associate multiple security lists with a subnet. The subnet uses the default security list if the customer doesn't specify another list for the subnet to use.

For further information about security in the Networking service, see:

- [Ways to Secure Your Network](#)
- [Access Control](#)
- [Security Lists](#)

### Storage Services

Oracle Cloud Infrastructure offers multiple storage solutions to meet the performance and durability requirements of customers:

- **Local Storage:** NVMe-backed storage on compute instances, offering extremely high IOPS.
- **Block Volumes:** Network-attached storage volumes, attachable to compute instances.
- **Object Storage:** Regional service for storing large amounts of data as objects, providing strong consistency and durability.

The Oracle Cloud Infrastructure Block Volumes service provides persistent storage that can be attached to compute instances using the iSCSI protocol. The volumes are stored in high-performance network storage and support automated backup and snapshot capabilities. Volumes and their backups are accessible only from within a customer's VCN and are encrypted at rest using unique keys. For additional security, iSCSI CHAP authentication can be required on a per-volume basis.

The Oracle Cloud Infrastructure Object Storage service provides highly scalable, strongly consistent, and durable storage for objects. API calls over HTTPS provide high-throughput access to data. All objects are encrypted at rest using unique keys. Objects are organized by

bucket, and, by default, access to buckets and objects within them requires authentication. Users can use IAM security policies to grant users and groups access privileges to buckets.

To allow bucket access by users who do not have IAM credentials, the bucket owner (or a user with necessary privileges) can create pre-authenticated requests that allow authorized actions on buckets or objects for a specified duration. Alternately, buckets can be made public, which allows unauthenticated and anonymous access. Given the security risk of inadvertent information disclosure, Oracle highly recommends carefully considering the business case for making buckets public. Object Storage enables you to verify that an object was not unintentionally corrupted by allowing an MD5 hash to be sent with the object (or with each part, in the case of multipart uploads) and returned upon successful upload. This hash can be used to validate the integrity of the object.

In addition to its native API, the Object Storage service supports Amazon S3 compatible APIs. Using the Amazon S3 Compatibility API, customers can continue to use the existing Amazon S3 tools (for example, SDK clients), and partners can modify their applications to work with Object Storage, with minimal changes to their applications. Their native API can co-exist with the Amazon S3 Compatibility API, which supports CRUD operations. Before customers can use the Amazon S3 Compatibility API, they must create an S3 Compatibility API key. After they've generated the necessary key, they can use the Amazon S3 Compatibility API to access Object Storage in Oracle Cloud Infrastructure.

### **Database Service**

Oracle Cloud Infrastructure makes it easy to run, scale, and secure your Oracle databases (DBs) in the cloud. The Oracle Cloud Infrastructure Database service offers three types of DB systems:

- Bare metal: Comprising 1-node DB and 2-node Real Application Cluster (RAC) systems, providing exceptional performance at cost-effective pricing.
- Exadata: Proven industry-leading Exadata DB systems in quarter, half, and full rack configurations.
- Virtual machine: Allows customers to create full-featured Oracle databases on VM shapes with various cores.

DB systems are accessible only from a customer's VCN, and customers can configure VCN security lists to control network access to their databases. The Database service is integrated with Oracle Cloud Infrastructure IAM for controlling which users can launch and manage DB systems. By default, data is encrypted at rest using Oracle TDE with master keys stored in an Oracle Wallet on each DB system. RMAN backups of DB systems are encrypted and stored in customer-owned buckets in the Object Storage service. Customers need to create a bucket for DB backups and configure the Oracle Database Cloud Backup module with an auth token (to use with the Swift API) and IAM permissions to access the bucket. Alternately, DB backups can be made to local NVMe storage on the DB system.

Each user automatically has the ability to create, update, and delete their own auth tokens in the Console or the API. An administrator does not need to create a policy to give a user those abilities. Administrators (or anyone with permission to the tenancy) also have the ability to manage auth tokens for other users. Any user of a Swift client that integrates with Object Storage needs permission to work with the service.

### **Load Balancing Service**

Oracle Cloud Infrastructure Load Balancing provides automated traffic distribution to compute instances in a customer's VCN. Load balancers (LBs) can be created as public (accepting traffic from the internet and directing it to private instances) or private (directing traffic between private instances). LBs can be configured for SSL termination using customer-provided certificates; end-to-end SSL, whereby the LB terminates the SSL connection and creates a new SSL connection to the backend; or SSL tunneling, in which the SSL connection is passed through to the backend (TCP load balances only). The Load Balancing service supports TLS 1.2 by default, and prioritizes the following forward-secrecy ciphers in the TLS cipher-suite:

- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-RSA-AES256-SHA384
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES128-SHA256
- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES256-SHA256

- DHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES128-SHA256

### **Managed Domain Name Servers (DNS) Service**

The Oracle Cloud Infrastructure DNS service provides dynamic, static, and recursive DNS solutions for enterprise customers. The service connects visitors to customer websites and applications with fast and secure services. The DNS service operates on a global anycast network with 18 points of presence (PoPs) on five continents and offers fully redundant DNS constellations and multiple Tier 1 transit providers per PoP. The solution provides a DNS-based Distributed Denial of Services (DDoS) protection and in-house security expertise that leverages a vast sensor network that collects and analyzes over 240 billion data points per day. The DNS service also fully supports the secondary DNS features to complement the customer's existing DNS service, providing resiliency at the DNS layer.

### Oracle Cloud Security Testing Policy

This section describes the Oracle Cloud Security Testing policy and how you can submit a request to schedule the tests of your Oracle Cloud services. The Oracle Cloud Security Testing Policy describes when and how you may conduct certain types of security testing of Oracle Cloud Services, including vulnerability and penetration tests, as well as tests involving data scraping tools. Notwithstanding anything to the contrary, any such testing of Oracle Cloud Services may be conducted only by customers who have an Oracle Account with the necessary privileges to file service maintenance requests, and who are signed-in to the environment that will be the subject of such testing.

#### **PENETRATION AND VULNERABILITY TESTING**

Oracle regularly performs penetration and vulnerability testing and security assessments against the Oracle cloud infrastructure, platforms, and applications. These tests are intended to validate and improve the overall security of Oracle Cloud Services.

However, Oracle does not assess or test any components (including, non-Oracle applications, non-Oracle databases or other non-Oracle software, code or data, as may be applicable) that you manage through or introduce into – including introduction through your development in or

creation in - the Oracle Cloud Services (the “**Customer Components**”). This policy does not address or provide any right to conduct testing of any third party materials included in the Customer Components.

Except as otherwise permitted or restricted in your Oracle Cloud Services agreements, your service administrator who has system level access to your Oracle Cloud Services may run penetration and vulnerability tests for the Customer Components included in certain of your Oracle Cloud Services in accordance with the following rules and restrictions.

### **Permitted Cloud Penetration and Vulnerability Testing**

The following explains where penetration and vulnerability testing of Customer Components is permitted:

- **IaaS:** Using your own monitoring and testing tools, you may conduct penetration and vulnerability tests of your acquired single-tenant Oracle Infrastructure as a Service (IaaS) offerings. You must notify Oracle prior to conducting any such penetration and vulnerability tests in accordance with the process set forth below. Pursuant to such penetration and vulnerability tests, you may assess the security of the Customer Components; however, you may not assess any other aspects or components of these Oracle Cloud Services including the facilities, hardware, software, and networks owned or managed by Oracle or its agents and licensors.
- **PaaS:** Using your own monitoring and testing tools, you may conduct penetration and vulnerability tests of your acquired single-tenant PaaS offerings. You must notify Oracle prior to conducting any such penetration and vulnerability tests in accordance with the process set forth below. Pursuant to such penetration and vulnerability tests, you may assess the security of the Customer Components; however, you may not assess any other aspects or components of these Oracle Cloud Services including the facilities, hardware, networks, applications, and software owned or managed by Oracle or its agents and licensors. To be clear, you may not assess any Oracle applications that are installed on top of the PaaS service.
- **SaaS:** Penetration and vulnerability testing is not permitted for Oracle Software as a Service (SaaS) offerings.

### **Rules of Engagement**

The following rules of engagement apply to cloud penetration and vulnerability testing:

- Your testing must not target any other subscription or any other Oracle Cloud customer resources, or any shared infrastructure components.
- You must not conduct any tests that will exceed the bandwidth quota or any other subscribed resource for your subscription.
- You are strictly prohibited from utilizing any tools or services in a manner that perform Denial-of-Service (DoS) attacks or simulations of such, or any “load testing” against any Oracle Cloud asset including yours.
- Any port scanning must be performed in a non-aggressive mode.
- You are responsible for independently validating that the tools or services employed during penetration and vulnerability testing do not perform DoS attacks, or simulations of such, prior to assessment of your instances. This responsibility includes ensuring any contracted third parties perform assessments in a manner that does not violate this policy.
- Social Engineering of Oracle employees and physical penetration and vulnerability testing of Oracle facilities is prohibited.
- You must not attempt to access another customer’s environment or data, or to break out of any container (for example, virtual machine).
- Your testing will continue to be subject to terms and conditions of the agreement(s) under which you purchased Oracle Cloud Services, and nothing in this policy shall be deemed to grant you additional rights or privileges with respect to such Cloud Services.
- If you believe you have discovered a potential security issue related to Oracle Cloud, you must report it to Oracle within 24 hours by conveying the relevant information to [My Oracle Support](#). You must create a service request within 24 hours and must not disclose this information publicly or to any third party. Note that some of the vulnerabilities and issues you may discover may be resolved by you by applying the most recent patches in your instances.
- In the event you inadvertently access another customer’s data, you must immediately terminate all testing and report it to Oracle within one hour by conveying the relevant information to [My Oracle Support](#).
- You are responsible for any damages to Oracle Cloud or other Oracle Cloud customers

that are caused by your testing activities by failing to abide by these rules of engagement.

### Notification Process

The process for notifying Oracle of Your election to conduct a penetration or vulnerability test as required by this policy can be found in [Submitting a Cloud Security Testing Notification](#).

### DATA SCRAPING TOOLS

Any use of data scraping tools or technologies with Oracle Cloud Services to collect data available through any Oracle user interface or via web service calls requires the express written permission of Oracle. Oracle reserves the right to require that your proposed data scraping tools are validated and tested by Oracle prior to use in production, and are subsequently re-validated and tested annually.

## Security Best Practices

This guide provides actionable guidance and recommendations to Oracle Cloud Infrastructure customers for securely configuring Oracle Cloud Infrastructure services and resources. Understanding Oracle Cloud Infrastructure services and their security features is an essential prerequisite before reading. As a prerequisite, Oracle recommends that you become familiar with security of services. For more information, see [Oracle Cloud Infrastructure Security Guide](#).

Security of an Oracle Cloud Infrastructure tenancy is based on a combination of factors, all of which must be thought through and securely configured. From a practical perspective, take a hierarchical view of Oracle Cloud Infrastructure tenancy security configuration, where we start with addressing the foundational security issues. The following steps provide a roadmap of high-level guidelines to follow when configuring security of a tenancy, where we provide a link to a section enumerating detailed security guidance related to each step.

- **User authentication and authorization:** The initial step in securely configuring a tenancy is to create mechanisms for authenticating users and authorizing users to access tenancy resources in a least-privilege manner. This step comprises creating Oracle Cloud Infrastructure Identity and Access Management (IAM) users, creating IAM

groups, formulating authentication mechanisms (for example, Console access using password, API access using API keys, and auth token for object store) for the IAM users created, grouping customer tenancy resources into logical groups using compartments, and formulating IAM security policies authorizing access of IAM groups to tenancy or compartment resources. For enterprises, federating their on-premises users and groups to their tenancy is an important consideration. IAM allows you to create users, groups, security polices, and federation mechanisms. Security recommendations for configuring IAM are provided in the [IAM section](#).

- **Network security architecture:** After formulating IAM user authentication and authorization, a next step is creating a network security architecture for securely running the customer applications and storing their data in a tenancy. All the customer's compute and storage resources are enclosed in a virtual cloud network (VCN) created for the customer. VCN is a software-defined network, resembling the on-premises physical network used by customers to run their workloads. Formulating a VCN security architecture includes tasks such as:
  - Creating VCN subnets for network segmentation
  - Formulating VCN and load balancer firewalls using VCN security lists
  - Using load balancing for high-availability and TLS
  - Determining type of VCN external connectivity whether internet, on-premises network, peered VCN, or combination of these
  - Using virtual network security appliances (for example, next-generation firewalls, IDs)
  - Creating DNS zones and mappings. An important security consideration in load balancers is using customer Transport Layer Security (TLS) certificates to configure TLS connections to customer's VCN. Security recommendations for VCN are provided in the Networking section.
- **Compute instances security configuration:** Within a customer VCN, the customer applications run on compute instances including Bare Metal (BM) instances, Virtual Machine (VM) instances and GPUs. Compute instances are the basic compute building blocks. Bare metal instances have no Oracle-managed software running on them, with

the result that the instances and data stored (in memory and local drives) are completely controlled by the customer. VM instances are architected with least privilege mechanisms, and with corporate industry-leading hypervisor security best-practices. Depending on their security and performance requirements, customers have a choice of using BM and VM instances, to run their application workloads in their tenancy. It is imperative to securely configure compute instances, to maintain security of customer applications running on them. Recommendations related to instance security configuration with respect to instance firewalls, instance credential management, entropy, security patching, and security logging and monitoring are provided in the [Compute section](#).

- **Data storage security configuration:** Depending on the type of data and access required, customers can store data in local drives (attached to compute instances), remote block volumes, object store buckets, databases, or file storage in their tenancy. To handle these data storage requirements, Oracle Cloud Infrastructure offers multiple data storage services such as Block Volume, Object Storage, Database, and File Storage. In order to meet their data security requirements, customers need to formulate a tenancy data storage architecture for storing their data in their tenancy, and securely configure the storage services used. Compliance and regulatory requirements are an important factor in determining an appropriate data storage security architecture. Recommendations related to storage services security configuration are available in [Block Volume](#), [Object Storage](#), [Database](#), and [File Storage](#) sections. In addition to these services, customers need to consider security of data stored ephemeraly in compute instance memory (DRAM) and local NVMe storage.

API Audit logs record calls to APIs (for example, through the Console, SDKs, CLIs, and custom clients using the APIs) as log events. The API Audit logs are always on by default and can't be turned off. These logs are available to customers for 90 days, with retention period configurable up to 365 days. Information in the API Audit logs show what time API activity occurred, the source of the activity, the target of the activity, what the action was, and what the response was. Oracle recommends that customers periodically review the OCI API Audit logs to ensure they are in accordance with actions they took on their tenancy resources.

For service-specific best practices, see the following topics:

- [Securing Block Volume](#)
- [Securing Compute](#)
- [Securing Data Transfer](#)
- [Securing Database](#)
- [Securing Email Delivery](#)
- [Securing File Storage](#)
- [Securing IAM](#)
- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Securing Object Storage](#)
- [Securing Resource Manager](#)

### Securing Block Volume

#### SECURITY RECOMMENDATIONS

- There are two types of volumes: block volumes and boot volumes. Block volumes allow instance storage capacity to be expanded dynamically. A boot volume contains the image used to boot the compute instance. The IAM service groups the family of related volume resource types into a combined resource type called `volume-family`.
- Assign least privilege access for IAM users and groups to resource types in `volume-family`. The resource types in `volume-family` are `volumes`, `volume-attachments`, and `volume-backups`. The `volume-family` resources are detachable block volume devices that allow dynamic expansion of instance storage capacity or contain the image for booting the instance. The `volume-attachments` resources are attachments between volumes and instances. The `volume-backups` resources are point-in-time copies of volumes that can be used to create block volumes or recover block volumes.

#### *DATA DURABILITY*

To minimize loss of data due to inadvertent deletes by an authorized user or malicious deletes, Oracle recommends to giving `VOLUME_DELETE`, `VOLUME_ATTACHMENT_DELETE` and

## CHAPTER 29 Security Guide and Announcements

---

`VOLUME_BACKUP_DELETE` permissions to a minimum possible set of IAM users and groups. `DELETE` permissions should be given only to tenancy and compartment administrators.

To minimize loss of data due to deletes or corruption, Oracle recommends that you make periodic backups of volumes. Oracle Cloud Infrastructure allows automated scheduled backups. For more information about scheduled backups, see [Policy-Based Backups](#).

### *DATA-AT-REST ENCRYPTION*

By default, volumes and their backups are encrypted at rest using AES-256. You can also encrypt your data volumes using tools like dm-crypt, veracrypt, and Bit-Locker. Instructions on dm-crypt encryption are presented in the next section.

### **SECURITY POLICY EXAMPLES**

#### *PREVENT DELETE OF VOLUMES*

The following example policy allows group `VolumeUsers` to perform all actions on volumes and backups, except deleting them.

```
Allow group VolumeUsers to manage volumes in tenancy
 where request.permission!='VOLUME_DELETE'
Allow group VolumeUsers to manage volume-backups in tenancy
 where request.permission!='VOLUME_BACKUP_DELETE'
```

If `VolumeUsers` can't detach volumes from instances, you can add the following policy to the previous example.

```
Allow group VolumeUsers to manage volume-attachments in tenancy
 where request.permission!='VOLUME_ATTACHMENT_DELETE'
```

### **SECURITY-RELATED TASKS**

#### *ENCRYPTING NON-ROOT VOLUMES WITH DM-CRYPT*

dm-crypt is a kernel-level encryption mechanism (part of Linux device mapper framework) to provide encrypted volumes. It encrypts data passed from the filesystem (for example, ext4 and NTFS ), and stores it on a storage device in Linux Unified Key Setup (LUKS ) format. The encrypted volumes can be stored on a complete disk, disk partition, logical volume, or a file-backed storage created using loopback devices. Cryptsetup is the user-level utility used to

manage dm-crypt, and used to encrypt partitions and files. dm-crypt uses the Linux crypto APIs for encryption routines.

1. Attach block storage volume to an instance (for example, `/dev/sdb`)
2. Format `/dev/sdb` for LUKS encryption. Enter LUKS passphrase when prompted. The passphrase is used to encrypt the LUKS master key used for encrypting the volume.

```
cryptsetup -y luksFormat /dev/sdb
```

3. Verify that the LUKS formatting is successful.

```
cryptsetup isLuks /dev/sdb && echo Success
```

4. Get encryption information about the device.

```
cryptsetup luksDump /dev/sdb
```

5. Get LUKS UUID of the device. The UUID value is used to configure the `/etc/crypttab`.

```
cryptsetup luksUUID /dev/sdb
```

6. Create a LUKS container with device name, `dev_name`. This also creates a device node, `/dev/mapper/<dev_name>`.

```
cryptsetup luksOpen /dev/sdb <dev_name>
```

7. Get information about the mapped device.

```
dmsetup info <dev_name>
```

8. Format the device node as ext4 filesystem.

```
sudo mkfs -t ext4 /dev/sdb
```

9. Mount the device node.

```
mount /dev/mapper/<dev_name> /home/encrypt_fs
```

10. Add an entry to `/etc/crypttab`.

```
<dev_name> UUID=<LUKS UUID of /dev/sdb> none
```

All the files copied to `/home/encrypt_fs` are encrypted by LUKS.

11. Add a keyfile to an available keyslot of the encrypted volume. This keyfile can be used to access the encrypted volume.

## CHAPTER 29 Security Guide and Announcements

---

```
dd if=/dev/urandom of=$HOME/keyfile bs=32 count=1
chmod 600 $HOME/keyfile
cryptsetup luksAddKey /dev/sdb ~/keyfile
```

### 12. Verify the encryption status of files.

```
cryptsetup status /home/encrypt_fs
```

### 13. Unmount after you're finished.

```
umount /home/encrypt_fs
cryptsetup luksClose <dev_name>
```

To access the encrypted volume:

```
cryptsetup luksOpen /dev/sdb <dev_name> --key-file=/home/opc/keyfile
mount /dev/mapper/<dev_name> /home/encrypt_fs
```

If you lose the keyfile, or if the keyfile or passphrase gets corrupted, you can't decrypt the encrypted volume. This results in permanent loss of data. Oracle recommends that you store durable copies of the keyfile on an on-premises host.

#### *REMOTE MOUNTING OF DM-CRYPT ENCRYPTED DATA VOLUMES*

The following steps assume that the keyfile is on an on-premises host (`SRC_IP`) and that `<OCI_SSH_KEY>` is the SSH private key of the instance.

#### 1. Copy keyfile from the on-premises host to an instance.

```
scp -i <OCI_SSH_KEY> keyfile opc@SRC_IP:/home/opc
```

#### 2. Open the encrypted volume.

```
ssh i <OCI_SSH_KEY> opc@SRC_IP "cryptsetup luksOpen /dev/sdb <dev_name> --key-
file=/home/opc/keyfile"
```

#### 3. Mount the volume.

```
ssh -i <OCI_SSH_KEY> opc@SRC_IP "mount /dev/mapper/<dev_name> /home/encrypt_fs"
```

#### 4. Perform operations on data in the mounted volume.

#### 5. Unmount the encrypted volume.

## CHAPTER 29 Security Guide and Announcements

---

```
ssh -i <OCI_SSH_KEY> opc@SRC_IP "umount /home/encrypt_fs"
ssh -i <OCI_SSH_KEY> opc@SRC_IP "cryptsetup luksClose <dev_name>"
```

### 6. Delete the keyfile from the instance.

```
ssh -i <OCI_SSH_KEY> opc@SRC_IP "\rm -f /home/opc/keyfile"
```

## Securing Compute

### SECURITY RECOMMENDATIONS

Oracle Cloud Infrastructure Compute provides both bare metal and virtual machine (VM) instances, architected and managed in accordance with industry-leading security best practices.

#### *MANAGING INSTANCES AND CREDENTIALS*

- To prevent inadvertent or malicious termination of critical instances (for example, production instances), Oracle recommends that you give `INSTANCE_DELETE` permission to a minimal set of groups. Give `DELETE` permissions only to tenancy and compartment admins.
- Instances can be authorized to access Oracle Cloud Infrastructure services (Compute, Block volume, Networking, Load balancing, Object storage), on behalf of an IAM user, by using Oracle Cloud Infrastructure instance principals feature. To use this feature, you create dynamic groups and grant them access to service APIs. Members of dynamic groups are instances that you define based on rules to match instances to the group. A short-lived private key to sign API calls, is delivered through instance metadata service (<http://169.254.169.254/opc/v1/identity/cert.pem>), and the key is rotated multiple times a day. For more information about accessing services from instances, see [Calling Services from an Instance](#).

#### *INSTANCE METADATA ACCESS CONTROL*

- Instance metadata (<http://169.254.169.254>) provides predefined instance information (for example, OCID and display name), along with custom fields. Short-lived credentials, such as dynamic group credentials, could be provided through the instance metadata. Oracle recommends that you limit instance metadata access to only

## CHAPTER 29 Security Guide and Announcements

privileged users on the instance. For example, iptables can be used to restrict instance metadata access only to privileged users such as root, using the following rule.

```
iptables -A OUTPUT -m owner ! --uid-owner root -d 169.254.169.254 -j DROP
```

- Instances use link local addresses to access instance metadata service (169.254.169.254:80), DNS (169.254.169.254:53), NTP (169.254.169.254:123), kernel updates (169.254.0.3), and iSCSI connections to boot volumes (169.254.0.2:3260, 169.254.2.0/24:3260). Host-based firewalls such as iptables can be used to authorize only root user to access these IPs. Ensure that these OS firewall rules are not altered.

### INSTANCE NETWORK ACCESS CONTROLS

- Harden secure shell (SSH) on all instances. Some SSH security recommendations are shown in the following table. SSH configuration options can be set in the `sshd_config` file (located at `/etc/ssh/sshd_config` in Linux).

Security Recommendation	Configuration sshd_config	Comments
Use public-key logins only	PubkeyAuthentication yes	Periodically review SSH public keys in <code>~/.ssh/authorized_keys</code> file
Disable password logins	PasswordAuthentication no	Mitigates password brute-force attacks
Disable root logins	PermitRootLogin no	Prevents root privileges for remote logins
Change SSH port to non-standard port	Port <port number>	This is optional. Ensure that this change does not break applications using port 22 for SSH

- Secure SSH private keys used to access instances and prevent any inadvertent disclosure. For more information about creating an SSH key pair and configuring an instance with the keys, see [Creating a Key Pair](#).

- Use VCN network security groups or security lists as a mechanism to allow instance access from authorized IP addresses. Fail2ban is an application that blacklists IP addresses involved in brute-force login attempts (that is, too many failed attempts to an instance). Fail2ban inspects SSH accesses by default, and can be configured for other protocols. For more information about Fail2ban, see [Fail2ban Main Page](#).
- In addition to VCN network security groups and security lists, host-based firewalls (such as iptables and firewalld) can be used to restrict network access to instances, in terms of ports, protocols, and packet types. These firewalls can be used to prevent potential network security attack reconnaissance (for example, port scanning) and intrusion attempts. Custom firewall rules can be configured, saved, and initialized on every instance boot. The following example shows commands for iptables.

```
save iptables rules after configuration
sudo iptables-save > /etc/iptables/iptables.rules
restore iptables rules on next reboot
sudo /sbin/iptables-restore < /etc/iptables.rules
restart iptables after restore
sudo service iptables restart
```

### *INSTANCE ENTROPY*

Both bare metal and VM instances provide high-quality and high-throughput entropy source. Instances have hardware random number generator support, whose output is fed into the entropy pools used by the OS to generate random numbers. In Linux instances, `/dev/random` is non-blocking, and recommended to be used for security applications requiring random numbers. You can use the following commands to check the throughput and quality of random numbers generated by `/dev/random`, before using its output in applications.

```
check sources of entropy
sudo rngd -v
enable rngd, if not already
sudo systemctl start rngd
verify rngd status
sudo systemctl status rngd
verify /dev/random throughput and quality using rngtest
cat /dev/random | rngtest -c 1000
```

## CHAPTER 29 Security Guide and Announcements

---

### HOST SECURITY HARDENING AND PATCHING

- Establish baseline for security hardening of OS images (Linux, Windows) running on instances. For more information about security hardening of Oracle Linux images, see [Tips for Hardening an Oracle Linux Server](#). The [Center for Internet Security Benchmarks](#) provides a comprehensive set of OS security hardening benchmarks for various distributions of Linux and Windows Server.
- Keep the instance software up-to-date with security patches. Oracle recommends that you periodically apply the latest available software updates to your instances. For Oracle Linux, you can run `sudo yum update` command. On Oracle Linux, you can get information about available and installed security patches using the `yum-security` plugin. Commands for `yum-security` are available below. For Oracle Linux instances launched after February 15, 2017, Ksplice support is available for applying patches without rebooting the instance. For more information about using Ksplice on Oracle Cloud Infrastructure instances, see [Installing and Running Oracle Ksplice](#).

```
Install yum-security plugin
yum install yum-plugin-security
Get list of security patches without installing them
yum updateinfo list security all
Get list of installed security patches
yum updateinfo list security all
```

### INSTANCE SECURITY LOGGING AND MONITORING

- Various security-related events are captured in log files. Oracle recommends that you periodically review these log files for detecting any security issues. In Oracle Linux, the log files are located in `/var/log`. Some security-relevant log files are listed in the table below.

Log File or Directory	Description
/var/log/secure	Auth log showing failed and successful logins
/var/log/audit	Auditd logs capturing system calls issued, sudo attempts, user logins, etc. ausearch and aureport are two tools used to query auditd logs
/var/log/yum.log	Lists various packages installed or updated on the instance with yum
/var/log/cloud-init.log	cloud-init can run user-provided scripts as privileged user, during instance boot. For example, SSH keys can be introduced using cloud-init. Oracle recommends that you review the cloud-init logs for any unrecognized commands.

- Host-based intrusion detection system (IDS) monitors instances for unauthorized accesses. OSSEC and Wazuh (OSSEC fork) are popular open-source IDS that can monitor for unauthorized access, malware, file modifications, and security misconfigurations. In the case of Wazuh, Wazuh server and ELK stack are deployed on an instance, and agents are deployed on other instances in the VCN to send logs to the Wazuh server. The resulting alerts are displayed on a Kibana dashboard. Wazuh IDS was prototyped on instances, and below are instructions for deploying a working Wazuh server on an instance (with ELK version 5.6.3). For more information about installing Wazuh agents and accessing the Kibana dashboard, see the [Wazuh documentation](#).

```
#!/bin/sh

echo "installing elasticsearch"
sudo add-apt-repository ppa:webupd8team/java;
sudo apt-get update;
sudo apt-get install oracle-java8-installer;

sudo apt-get install curl apt-transport-https
curl -s https://artifacts.elastic.co/GPG-KEY-elasticsearch | sudo apt-key add - ;
echo "deb https://artifacts.elastic.co/packages/5.x/apt stable main" | sudo tee
```

## CHAPTER 29 Security Guide and Announcements

---

```
/etc/apt/sources.list.d/elastic-5.x.list;
sudo apt-get update;

sudo apt-get install elasticsearch=5.6.3;

sudo systemctl daemon-reload
sudo systemctl enable elasticsearch.service
sudo systemctl start elasticsearch.service

curl https://raw.githubusercontent.com/wazuh/wazuh-kibana-app/2.1/server/startup/integration_
files/template_file.json | curl -XPUT 'http://localhost:9200/_template/wazuh' -H 'Content-Type:
application/json' -d @-

curl https://raw.githubusercontent.com/wazuh/wazuh-kibana-app/2.1/server/startup/integration_
files/alert_sample.json | curl -XPUT "http://localhost:9200/wazuh-a
lerts-``date +%Y.%m.%d`"/wazuh/sample" -H 'Content-Type: application/json' -d @-

echo "installing logstash"
sudo apt-get install logstash=1:5.6.3-1;
sudo curl -so /etc/logstash/conf.d/01-wazuh.conf
https://raw.githubusercontent.com/wazuh/wazuh/2.1/extensions/logstash/01-wazuh.conf;
sudo curl -so /etc/logstash/wazuh-elastic5-template.json
https://raw.githubusercontent.com/wazuh/wazuh/2.1/extensions/elasticsearch/wazuh-elastic5-
template.json;
sudo usermod -a -G ossec logstash;

sudo systemctl daemon-reload;
sudo systemctl enable logstash.service;
sudo systemctl start logstash.service;

echo "installing kibana"
sudo apt-get install kibana=5.6.3;
sudo -u kibana /usr/share/kibana/bin/kibana-plugin install
https://packages.wazuh.com/wazuhapp/wazuhapp-2.1.1_5.6.3.zip;

sudo systemctl daemon-reload;
sudo systemctl enable kibana.service;
sudo systemctl start kibana.service;
```

### SECURITY POLICY EXAMPLES

In all the following examples, the policies are scoped to a tenancy. However, by specifying a compartment name, they can be scoped down to specific compartment in a tenancy.

#### *RESTRICT USERS ABILITY TO DELETE INSTANCES*

The following example allows the `InstanceUsers` group to launch instances but not delete them

```
Allow group InstanceUsers to manage instance-family in tenancy
 where request.permission!='INSTANCE_DELETE'
Allow group InstanceUsers to use volume-family in tenancy
Allow group InstanceUsers to use virtual-network-family in tenancy
```

#### *RESTRICT ABILITY TO USE INSTANCE CONSOLE*

For security compliance reasons, some customers do not want to expose instance console to users in their tenancy. The following policy example restricts ability to create or read from consoles.

```
Allow group InstanceUsers to manage instance-console-connection in tenancy
 where all {request.permission!= INSTANCE_CONSOLE_CONNECTION_READ,
 request.permission!= INSTANCE_CONSOLE_CONNECTION_CREATE}
```

### Securing Data Transfer

Oracle offers offline data transfer solutions that let you migrate large amounts of data to buckets in a tenancy in Oracle Cloud Infrastructure. Data transfer solutions include:

- Disk-Based Data Transfer  
For more information about securely transferring data using this service, see [Secure Disk Data Transfer to Oracle Cloud Infrastructure](#)
- Appliance-Based Data Transfer  
For more information about securely transferring data using this service, see [Secure Disk Data Transfer to Oracle Cloud Infrastructure.](#)

### Securing Database

#### SECURITY RECOMMENDATIONS

This section lists security recommendations for managing Oracle Cloud Infrastructure Database instances. Recommendations for securely configuring Oracle databases are available in the Oracle Database Security Guide.

#### *DATABASE ACCESS CONTROL*

- Users authenticate to the database using their password. Oracle recommends that these passwords be strong. For guidelines on choosing Oracle database passwords, see [Guidelines for Securing Passwords](#). In addition, Oracle database provides a PL/SQL script to verify database password complexity. This script is located at `$ORACLE_HOME/rdbms/admin/UTLPWDMG.SQL`. For instructions on running UTLPWDMG.SQL script to verify password complexity, see [Enforcing Password Complexity Verification](#).
- In addition to the database password, you can use VCN network security groups or security lists to enforce network access control to database instances. Oracle recommends that you configure VCN network security groups or security lists to allow least privilege access to customer databases in Oracle Cloud Infrastructure Database.
- DB systems created within a public subnet can send outbound traffic directly to the Internet. DB systems created within a private subnet do not have internet connectivity, and internet traffic (both egress and ingress) cannot reach the instance directly. If you try to define a route to a DB system within a private subnet using an internet gateway, the route is ignored.

To perform OS patching and backup for a DB system on private subnet, you can use a service gateway or a NAT gateway to connect to your patching or backup endpoints.

In an virtual cloud network (VCN), you can use security rules along with a private subnet to restrict access to a DB system. In multi-tier deployments, a private subnet and VCN security rules can be used to restrict access to the DB system from the application tiers.

### DATA DURABILITY

- Oracle recommends that you give database delete permissions (`DATABASE_DELETE`, `DB_SYSTEM_DELETE`) to a minimum possible set of IAM users and groups. This minimizes loss of data due to inadvertent deletes by an authorized user or due to malicious deletes. Only give `DELETE` permissions to tenancy and compartment administrators.
- You can use RMAN to do periodic backups of Database databases, where encrypted backup copies are stored in local storage (block volumes, for example) or Oracle Cloud Infrastructure Object Storage. RMAN encrypts each backup of a database with a unique encryption key. In transparent mode, the encryption key is stored in the Oracle Wallet. RMAN backups to Object Storage require internet gateway (IGW), and VCN network security groups or security lists need to be configured to allow secure access to Object Storage. For information about setting up the VCN for backing up bare metal databases, see [Backing Up a Database to Oracle Cloud Infrastructure Object Storage](#). For information about backing up and Exadata databases, see [Managing Exadata Database Backups by Using bkup\\_api](#).

### DATABASE ENCRYPTION AND KEY MANAGEMENT

- All databases created in Oracle Cloud Infrastructure are encrypted using transparent data encryption (TDE). Note that if you migrate an unencrypted database from on-premise to Oracle Cloud Infrastructure using RMAN, the migrated database will not be encrypted. Oracle strongly recommends encrypting such databases after migrating them to the cloud.

To learn how to encrypt your database with minimum downtime during migration, see the Oracle Maximum Availability Architecture white paper [Converting to Transparent Data Encryption with Oracle Data Guard using Fast Offline Conversion](#).

Note that virtual machine DB systems use Oracle Cloud Infrastructure block storage instead of local storage. Block storage is encrypted by default.

- User-created tablespaces are encrypted by default in Oracle Cloud Infrastructure Database. In these databases, `ENCRYPT_NEW_TABLESPACES` parameter is set to `CLOUD_ONLY` where tablespaces created in a Database Cloud Service (DBCS) database are

transparently encrypted with the AES128 algorithm unless a different algorithm is specified.

- The Database administrator creates a local Oracle Wallet on a newly created database instance, and initializes the Transparent Data Encryption (TDE) master key. Then the Oracle Wallet is configured to be "auto-open". However, a customer can choose to set a password for the Oracle Wallet, and Oracle recommends that you set a strong password (eight characters or more, with at least one capital letter, one small letter, one number, and one special symbol).
- Oracle recommends that you periodically rotate the TDE master key. The recommended rotation period is 90 days or less. You can rotate the TDE master key by using native database commands ("administer key management" in 12c, for example) or dbaascli. All previous versions of TDE master key are maintained in the Oracle Wallet.
- Oracle Key Vault (OKV) is a key management appliance used for managing Oracle TDE master keys. OKV can store, rotate, and audit accesses to TDE master keys. For instructions about installing and configuring OKV in Oracle Cloud Infrastructure, see [Managing Oracle Database Encryption Keys in Oracle Cloud Infrastructure with Oracle Key Vault](#).

### *DATABASE PATCHING*

Applying Oracle database security patches (Oracle Critical Patch Updates) is imperative to mitigate known security issues, and Oracle recommends that you keep patches up-to-date. Patchsets and Patch Set Updates (PSUs) are released on a quarterly basis. These patch releases contain security fixes and additional high-impact/low-risk critical bug fixes.

For information about the latest known security issues and available fixes, see [Critical Patch Updates, Security Alerts and Bulletins](#). If your application does not support the latest patches and needs to use a DB system with older patches, you can provision a DB system with an older version of the Oracle Database edition you are using. In addition to reviewing the critical patch updates and security alerts for your Oracle Database, Oracle recommends that you analyze and patch the operating system provisioned with the DB system.

For information about applying patches to Oracle Cloud Infrastructure Database instances, see [Patching a DB System](#) and [Patching an Exadata DB System](#).

## CHAPTER 29 Security Guide and Announcements

---

### *DATABASE SECURITY CONFIGURATION CHECKING*

- The [Oracle Database Security Assessment Tool](#) (DBSAT) provides automated security configuration checks of Oracle databases in Oracle Cloud Infrastructure. DBSAT performs security checks for user privilege analysis, database authorization controls, auditing polices, database listener configuration, OS file permissions, and sensitive data stored. Oracle database images in Oracle Cloud Infrastructure Database are scanned with DBSAT before provisioning. After provisioning, Oracle recommends that you periodically scan databases with DBSAT, and remediate any issues found. DBSAT is available free of charge to Oracle customers.

### *DATABASE SECURITY AUDITING*

Oracle Audit Vault and Database Firewall (AVDF) monitors database audit logs and creates alerts. For instructions about installing and configuring AVDF in Oracle Cloud Infrastructure, see [Deploying Oracle Audit Vault and Database Firewall in Oracle Cloud Infrastructure](#).

### *DATABASE BACKUPS*

Oracle recommends using Managed backups (backups created using the Oracle Cloud Infrastructure Console or the API) whenever possible. When you use managed backups, Oracle manages the object store user and credentials, and rotates these credentials every 3 days. Oracle Cloud Infrastructure encrypts all managed backups in the object store. Oracle uses the Database Transparent Encryption feature by default for encrypting the backups.

If you are not using managed backups, Oracle recommends that you change the object store passwords at regular intervals.

### **SECURITY POLICY EXAMPLES**

#### *PREVENT DELETE OF DATABASE INSTANCES*

The following example policy allows the group `DBUsers` to perform all management actions except delete databases and any artifacts.

```
Allow group DBUsers to manage db-systems in tenancy
 where request.permission!='DB_SYSTEM_DELETE'
Allow group DBUsers to manage databases in tenancy
 where request.permission!='DATABASE_DELETE'
```

```
Allow group DBUsers to manage db-homes in tenancy
where request.permission!='DB_HOME_DELETE'
```

### Securing Email Delivery

The Email Delivery service offers an SMTP endpoint, secured by a password generated in the Console. The SMTP password is required for sending emails using Email Delivery. Oracle recommends that you create a separate IAM user for SMTP. This user must have manage permissions for `approved-senders` and `suppressions` resource types. Oracle recommends that you securely store the SMTP credential, and periodically rotate it. For more information about generating an SMTP credential for Email Delivery, see [Generate SMTP Credentials for a User](#).

For Email Delivery best practices, including managing your sender reputation and help for avoiding being blacklisted, see [Deliverability Best Practices](#).

### Securing File Storage

The File Storage Service exposes an NFSv3 endpoint as a mount target in each customer's VCN subnet. The mount target is identified by a DNS name and is mapped to an IP address. Oracle recommends that you use VCN security lists (of the mount target subnet) to configure network access to the mount target from only authorized IP addresses.

You can mount a file system using the Console or from a Linux command line using NFS utilities. You can authorize users to mount file systems using IAM security policies, but this applies to the console only.

For data durability, Oracle recommends that you take periodic snapshots of the file system. To minimize accidental deletion of data, constrain the set of users having privileges to delete mount targets, file-systems, and snapshots.

All file-system data is encrypted at rest.

Access to mounted NFS file systems from a remote host is determined by POSIX user and group permissions. Oracle recommends that you use well-known NFS security best practices such as the `all_squash` option to map all users to `nfsnobody`, and NFS ACLs to enforce access control to the mounted file system.

## CHAPTER 29 Security Guide and Announcements

### SECURITY POLICY EXAMPLES

#### *PREVENT MOUNT TARGET AND FILE SYSTEM DELETION*

The following example prevents group FileUsers from deleting mount targets and file-systems.

```
Allow group FileUsers to manage file-systems in tenancy
 where request.permission!='FILE_SYSTEM_DELETE'
Allow group FileUsers to manage mount-targets in tenancy
 where request.permission!='MOUNT_TARGET_DELETE'
Allow group FileUsers to manage export-sets in tenancy
 where request.permission!='EXPORT_SET_DELETE'
```

### Securing IAM

#### SECURITY RECOMMENDATIONS

Oracle Cloud Infrastructure Identity and Access Management (IAM) provides authentication of users, and authorization to access resources. Security-relevant IAM concepts include:

#### IAM concepts and descriptions

Concept	Description
Compartment	A compartment is a fundamental mechanism to aggregate resources into logical groups. They also provide isolation.
Tenancy	Oracle Cloud Infrastructure automatically creates a tenancy for an account created by an organization. It is the root compartment that contains all the organization's resources.
Users and groups	A group is an aggregation of users who need similar access to a group of resources.
Resource	A resource is an object created in Oracle Cloud Infrastructure services.

Concept	Description
Security policy	A security policy specifies the type of access IAM groups have to resources in a specified aggregation level. An aggregation level can be the tenancy, a compartment, or a service.
Dynamic groups	A dynamic group allows aggregating Compute instances as principal actors (similar to user groups), in order to authorize instances to make calls to Oracle Cloud Infrastructure APIs.
Tags	Tags allow you to organize resources across multiple compartments for reporting purposes or for taking bulk actions.
Federation	Mechanism to federate IAM with other identity providers (IdP) used by an organization to authenticate their users.

Oracle recommends that you periodically monitor Audit logs to review changes to IAM users, groups, and security policies.

### IAM TENANCY AND COMPARTMENTS

- Compartments are unique to IAM, and offer a mechanism that allows an enterprise customer to meet its central needs by having a single account or tenancy. This single account or tenancy provides full central control and visibility while also allowing the account or tenancy to be subdivided to meet the needs of constituent teams, projects, and initiatives.
- For security and governance reasons, users should only have access to resources they need. For example, enterprise users working on a project or belonging to a business unit should have access only to resources belonging to the project or business unit. Compartments provide an effective mechanism to group tenancy resources based on their access privileges and authorize groups of users to access the compartments on as needed basis. In the example above, a compartment can be created to include all resources belonging to a business unit, and authorize only members of the business unit to access the compartment. Similarly, a groups' access to a compartment can be revoked when they do not need it anymore.

- Keep the following in mind when you create a compartment and assign resources:
  - Every resource should belong to a compartment.
  - A resource can be reassigned to a different compartment after creation. See [Managing Compartments](#).
  - A compartment can be deleted after creation. See [Managing Compartments](#).
- Resource tags provide a way to logically aggregate resources distributed across multiple compartments. For example, tenancy resources can be tagged as `test` or `production` depending on their use. For more information about resource tags (free-form and defined tags), see [Resource Tags](#).
- Every tenancy comes with a default administrators group. This group can perform any action on all resources in a tenancy (that is, they have root access to the tenancy). Oracle recommends that you keep the group of tenancy administrators as small as possible. Some security recommendations on managing tenancy administrators:
  - Have security policies granting membership of tenancy administrator group strictly on a as-needed basis.
  - Tenancy administrators should use high-complexity passwords, along with MFA, and periodically rotate their passwords.
  - After account set up and configuration, Oracle recommends that you don't use the tenancy administrator account for day-to-day operations. Instead, create less privileged users and groups.
  - Though administrator accounts are not used for daily operations, they are still needed to address emergency scenarios impacting customer tenancy and operations. Specify secure and auditable "break-glass" procedures for using administrator accounts in such emergencies.
  - Disable tenancy administration access immediately when an employee leaves the organization.
  - Because the tenancy administrator group membership is restricted, Oracle recommends that you create security policies which prevent administrator account lock-out (for example, if the tenancy administrator leaves the company and no current employees have administrator privileges).

### IAM USERS AND GROUPS

- Create an IAM user for everyone in the customer organization who needs access to resources. Do not share IAM user accounts across multiple users, especially those with administrative accounts. Using distinct IAM users enables enforcing least privilege access for each user, and captures their actions in audit logs.
- The recommended unit of administration is IAM groups, which makes it easier to manage and keep track of security permissions (as opposed to individual users). Create IAM groups with permissions to do commonly needed tasks (for example, network administration, volume administration), and assign users to these groups on an as-needed basis. IAM permissions can be used to give a group access to resources across multiple compartments in a tenancy.
- Periodically review membership of IAM users in IAM groups, and remove IAM users from groups they do not need access to anymore. Using group membership to manage user access scales well with increasing number of users.
- Deactivate IAM users who do not need access to tenancy resources. Deleting an IAM user removes the user permanently. You can temporarily deactivate an IAM user by doing the following:
  - Rotate the user password and throw it away.
  - Remove all tenancy permissions of the user by removing membership from all groups.

### IAM CREDENTIALS

IAM user credentials (Console password, API signing key, auth tokens, and customer secret keys) grant access to resources. It is important to secure these credentials to prevent unauthorized access to Oracle Cloud Infrastructure resources. General guidelines for handling credentials include:

- Create a strong console password for each IAM user, with sufficient complexity. Oracle recommends the following for a complex password:
  - Password has a minimum length of 12 characters
  - Password contains at least one uppercase letter

- Password contains at least one lowercase letter
- Password contains at least one symbol
- Password contains at least one number
- Rotate IAM passwords and API keys regularly, every 90 days or less. In addition to a security engineering best practice, this is also a compliance requirement. For example, PCI-DSS Section 3.6.4 states, "Verify that key-management procedures include a defined cryptoperiod for each key type in use and define a process for key changes at the end of the defined crypto period(s)."
- Do not hard code sensitive IAM credentials directly in software or documents accessible to a wide audience. Examples include code uploaded to GitHub, presentations, or documents available on the internet. There have been known, highly publicized cases of hackers breaching customer cloud accounts, using credentials inadvertently disclosed on public sites. When software applications need to access Oracle Cloud Infrastructure resources, Oracle recommends that you use instance principals. If it is not feasible to use instance principals, other recommendations include using user environment variables to store credentials, and using locally stored credential files with API keys to be used by the Oracle Cloud Infrastructure SDK or CLI.
- Do not share IAM credentials between multiple users.
- By federating the Console login through Oracle Identity Cloud Service, customers can use multifactor authentication (MFA) for IAM users, especially administrators.

When rotating API keys, verify that the rotated keys work as expected before disabling older keys. For information about generating and uploading IAM API keys, see [Required Keys and OCIDs](#). The high-level steps in rotating an API key are:

1. Generate and upload a new API key.
2. Update the SDK and CLI configuration files with the new API key.
3. Verify that the SDK and CLI calls are working correctly with the new key.
4. Disable the old API key. Use [ListApiKeys](#) to list all active API keys.

### IAM SECURITY POLICIES

IAM policies are used to govern access of IAM groups to resources in compartments and in the tenancy. Oracle recommends that you assign least privilege access to IAM groups for accessing resources. The common format for IAM policies is shown in the following example.

```
Allow group <group_name> to <verb> <resource-type> in compartment <compartment_name>
Allow group <group_name> to <verb> <resource-type> in tenancy
```

IAM policies allow four predefined verbs: inspect, read, use and manage. Inspect allows least privilege and manage allows the maximum. The four verbs are shown in increasing order of privilege in the following table.

IAM policy verbs

Verb	Access Type	Example User
inspect	Should only show metadata. This usually results in ability to list resources only	third-party auditor
read	inspect plus ability to read resource and user metadata. This is the permission most users need to get work done.	internal auditors
use	read plus ability to work with resources (the actions vary by resource type). Excludes ability to create or delete resource	regular users (software developers, system engineers, dev managers, etc) setting up and configuring tenancy resources, and applications running on them
manage	All the permissions for all the resources	administrators, executives (for break-glass scenarios)

The resource types of Oracle Cloud Infrastructure resources are shown in the following table.

IAM resource families, descriptions, and resource types

Resource Type Family	Description	Resource Types
all-resources	All resource types	
No name by design	Resource types in IAM service	compartments, users, groups, dynamic-groups, policies, identity-providers, tenancy tag-namespaces, tag-definitions
instance-family	Resource types in compute service	console-histories, instance-console-connection, instance-images, instances, volume-attachments
volume-family	Resource types in block storage service	volumes, volume-attachments, volume-backups
virtual-network-family	Resource types in virtual networking service	vcns, subnets, route-tables, security-lists, dhcp-options, private-ips, public-ips, internet-gateways, local-peering-gatewaysdrgs, deg-attachments, cpes, ipsec-connections, cross-connects, cross-connect-groups, virtual-circuits, vnics, vnic-attachments
object-family	Resource types in object storage service	buckets, objects

## CHAPTER 29 Security Guide and Announcements

Resource Type Family	Description	Resource Types
database-family	Resource types in DbaaS service	db-systems, db-nodes, db-homes, databases, backups
load-balancers	Resources in Load Balancer service	load-balancers
file-family	Resources in file storage service	file-systems, mount-targets, export-sets
dns	Resources in DNS service	dns-zones, dns-records, dns-traffic
email-family	Resources in email delivery service	approved-senders, suppressions

For more information about IAM verbs and resource type permission mappings, see [Details for the Core Services](#).

IAM security policies can be made fine-grained through conditions. Access specified in the policy is allowed only if the condition statements evaluate to true. Conditions are specified using predefined variables. The variables use the key words `request` or `target`, depending on whether the variable is relevant to the request or the resource being acted on, respectively. For information about supported predefined variables, see [Policy Reference](#).

IAM dynamic groups are used to authorize Compute instances to access Oracle Cloud Infrastructure APIs. The instance principals feature can be used by applications, running on the instances, to programmatically access Oracle Cloud Infrastructure services. Customers create dynamic groups, which include instances as members, and authorize access to their tenancy resources using IAM security policies. All access by instances is captured in the audit logs available to customers.

### IAM FEDERATION

- Oracle recommends that you use federation to manage logins into the Console. Identity federation supports SAML 2.0 compliant identity providers, and can be used to federate on-premises users and groups to IAM users and groups. The enterprise administrator needs to set up a federation trust between the on-premises identity provider (IdP) and IAM, in addition to creating mapping between on-premises groups and IAM groups. Then, on-premises users can single sign-on (SSO) into the Console, and access resources based on authorization of IAM groups they belong to. For more information about federating to the Console, see [Federating with Identity Providers](#). Federation is especially important for enterprises using custom policies for user authentication (for example, multifactor authentication) . For more information about managing users and groups under federation, see [Federating with Identity Providers](#).
- When using federation, Oracle recommends that you create a federation administrators group that maps to the federated IdP administrator group. The federation administrators group will have administrative privileges to manage customer tenancy, while being governed by the same security policies as the federated IdP administrator group. In this scenario, it is a good idea to have access to the local tenancy administrator user (that is, member of the default tenancy administrator IAM group), to handle any break-glass type scenarios (for example, inability to access resources through federation). However, you must prevent any unauthorized use of this highly privileged local tenancy administrator user. Oracle recommends the following approach to securely managing the tenancy administrator user:
  1. Create a local user belonging to the default tenancy administrator group.
  2. Create a highly complex Console password or passphrase (18 characters or more, with at least one lowercase letter, one uppercase letter, one number, and one special character) for the local tenancy administrator user.

3. Securely escrow the local tenancy administrator user password in an on-premises location (for example, place the password in a sealed envelope in an on-premises physical safe).
4. Create security policies for accessing the escrowed password only under specific "break-glass" scenarios.
5. Have IAM security policy to prevent the federated administrators IAM group from adding or modifying membership of the default tenancy administrator group to prevent security by-passes.
6. Monitor audit logs for accesses by default tenancy administrator and changes to the administrator group, to alert on any unauthorized actions. For additional security, the local tenancy administrator user password can be rotated after every login, or periodically, based on a password policy.

For an example that shows the way various IAM components fit together, see [Example Scenario](#). Periodically monitor Audit logs to review changes to IAM users, groups, policies, compartments, and tags.

### SECURITY POLICY EXAMPLES

Common IAM security policy examples are available at [Common Policies](#). In all the examples that follow, the policies are scoped to a tenancy. However, by specifying a compartment name, you can scope down the policies to specific compartments in a tenancy.

#### *CREATE SERVICE-LEVEL ADMINS FOR LEAST PRIVILEGE*

To implement security principle of least privilege, you can create service-level admins in the tenancy to further scope down administrative access. This means that service-level administrators can only manage resources of a specific service. For instance, network administrators need administrative (`manage`) access only to VCN resources, and not to other resources. The following example shows how to create administrator groups for block storage (`VolumeAdmins`), VCN (`NetworkAdmins`), databases (`DBAdmins`), and object storage (`StorageAdmins`).

```
Allow group TenancyAdmins to manage all-resources in tenancy
Allow group VolumeAdmins to manage volume-family in tenancy
Allow group NetworkAdmins to manage virtual-network-family in tenancy
```

## CHAPTER 29 Security Guide and Announcements

---

```
Allow group StorageAdmins to manage object-family in tenancy
Allow group DBAdmins to manage database-family in tenancy
```

You can further constrain the security policies to a specific compartment. For example, the HR department in an enterprise can create group `HRAdmins` to manage resources within its compartment, `HR-compartment`. The `HRNetworkAdmins` group has administrative access to VCN resources only within the `HR-compartment` compartment.

```
Allow group HRAdmins to manage all-resources in compartment HR-compartment
Allow group HRNetworkAdmins to manage virtual-network-family in compartment HR-compartment
```

Compliance auditors are tasked with examining cloud resources and verifying for policy violations. The following policy allows group `InternalAuditors` to inspect (`list`) all resources in a tenancy.

```
Allow group InternalAuditors to inspect all-resources in tenancy
```

If you want to limit auditors to only inspect users and groups in a tenancy, you can create a group `UserAuditors` with the following policy:

```
Allow group UserAuditors to inspect users in tenancy
Allow group UserAuditors to inspect groups in tenancy
```

If you want to create an auditor group that can only inspect VCN firewalls in the tenancy, use the following policy:

```
Allow group FirewallAuditors to inspect security-lists in tenancy
```

In all the policy examples, you can constrain the policies to a compartment by specifying `Compartment <name>` (where `<name>` is the compartment name) in the policy.

### *RESTRICT ABILITY TO CHANGE TENANCY ADMINISTRATORS GROUP MEMBERSHIP*

Members in the group `Administrators` can manage all resources in a tenancy. Membership of the `Administrators` group is controlled by users in the group. Usually, it's convenient to have a group to create and add users in the tenancy, but restrict them from making changes to the `Administrators` group membership. The following example creates a group `UserAdmins` to do this.

```
Allow group UserAdmins to inspect users in tenancy
Allow group UserAdmins to inspect groups in tenancy
Allow group UserAdmins to use users in tenancy
 where target.group.name!='Administrators'
```

## CHAPTER 29 Security Guide and Announcements

---

```
Allow group UserAdmins to use groups in tenancy
where target.group.name!='Administrators'
```

Use verb with conditions (third and fourth policy statements) allows `UserAdmins` to add users and groups using APIs (`UpdateUser`, `UpdateGroup`) to all groups in the tenancy except the `Administrators` group. However, because `target.group.name!='Administrators'` is not related to the `list` and `get` APIs (`ListUsers`, `GetUser`, `ListGroups`, and `GetGroup`), these APIs will fail. So you must explicitly add the `inspect` verb (first and second policy statements) to allow `UserAdmins` to get user and group membership information.

### *PREVENT DELETE OR UPDATE OF SECURITY POLICIES*

The following example creates a group `PolicyAdmins` to be able to create and list security policies created by tenancy administrators, but not delete or update them.

```
Allow group PolicyAdmins to use policies in tenancy
Allow group PolicyAdmins to manage policies in tenancy
where request.permission='POLICY_CREATE'
```

This security policy statement explicitly only allows `POLICY_CREATE` permission, and not to `POLICY_DELETE` and `POLICY_UPDATE`.

### *PREVENT ADMINS FROM ACCESSING OR ALTERING USER CREDENTIALS*

Some compliance requirements need separation of duties, especially where user credential management functionality is separated from tenancy management. In this case, you can create two administration groups, `TenancyAdmins` and `CredentialAdmins` where `TenancyAdmins` can perform all tenancy management functions except user credential management, and `CredentialAdmins` can manage user credentials. `TenancyAdmins` can access all APIs except those that list, update, or delete user credentials. `CredentialAdmins` can only manage the user credentials.

```
Allow group TenancyAdmins to manage all resources in tenancy
where all {request.operation!='ListApiKeys',
 request.operation!='ListAuthTokens',
 request.operation!='ListCustomerSecretKeys',
 request.operation!='UploadApiKey',
 request.operation!='DeleteApiKey',
 request.operation!='UpdateAuthToken',
 request.operation!='CreateAuthToken',
 request.operation!='DeleteAuthToken',
```

## CHAPTER 29 Security Guide and Announcements

---

```
request.operation!='CreateSecretKey',
request.operation!='UpdateCustomerSecretKey',
request.operation!='DeleteCustomerSecretKey'}
Allow group CredentialAdmins to manage users in tenancy
where any {request.operation='ListApiKeys',
request.operation='ListAuthTokens',
request.operation='ListCustomerSecretKeys',
request.operation='UploadApiKey',
request.operation='DeleteApiKey',
request.operation='UpdateAuthToken',
request.operation='CreateAuthToken',
request.operation='DeleteAuthToken',
request.operation='CreateSecretKey',
request.operation='UpdateCustomerSecretKey',
request.operation='DeleteCustomerSecretKey'}
```

### USEFUL CLI COMMANDS

In all the following examples, environment variables `$T` and `$C` are set to tenancy OCID and compartment OCID, respectively.

#### *LIST COMPARTMENTS IN A TENANCY*

```
list all compartments (OCID, display name, description) in tenancy $T
oci iam compartment list -c $T
grep above command for important fields
oci iam compartment list -c $T | grep -E "name|description|\"id\""
```

#### *LIST IAM USERS*

```
lists all users (OCID, display name, description) in tenancy $T
oci iam user list -c $T
grep above command for important fields
oci iam user list -c $T | grep -E "name|description|\"id\""
```

#### *LIST IAM GROUPS*

```
lists all groups (OCID, display name, description) in tenancy $T.
oci iam group list -c $T
grep above command for important fields
oci iam group list -c $T | grep -E "name|description|\"id\""
```

## CHAPTER 29 Security Guide and Announcements

### LIST USERS IN A GROUP

The following command is helpful for listing users in groups, especially users with administrative privileges. This command requires the OCID of the group whose users are listed.

```
list users in group with OCID <GROUP_OCID>
oci iam group list-users -c $T --group-id <GROUP_OCID>
```

### LIST SECURITY POLICIES

```
lists all policies (OCID, name, statements) in tenancy $T. Remove pipe to grep to get entire
information
oci iam policy list -c $T
grep above command for important fields
oci iam policy list -c $T | grep -E "name|Allow|\"id\""
```

## Securing Networking: VCN, Load Balancers, and DNS

### SECURITY RECOMMENDATIONS

The Networking service has a collection of features for enforcing network access control and securing VCN traffic. These features are listed in the following table.

VCN Feature	Security Description
<a href="#">Public and private subnets</a>	Your VCN can be partitioned into subnets. Subnets have historically been specific to an availability domain, but can now be regional (covering all availability domains in the region). Instances inside private subnets cannot have public IP addresses. Instances inside public subnets can optionally have public IP addresses at your discretion.
<a href="#">Security rules</a>	Security rules provide stateful and stateless firewall capability to control network access to your instances. To implement security rules in your VCN, you can use <a href="#">network security groups (NSGs)</a> or <a href="#">security lists</a> . For more information, see <a href="#">Comparison of Security Lists and Network Security Groups</a> .

VCN Feature	Security Description
Gateways	<p>Gateways let resources in a VCN communicate with destinations outside the VCN. The gateways include:</p> <ul style="list-style-type: none"> <li>• <a href="#">Internet gateway</a>: for internet connectivity (for resources with public IP addresses in public subnets)</li> <li>• <a href="#">NAT gateway</a>: for internet connectivity without exposing the resources to incoming internet connections (for resources in private subnets)</li> <li>• <a href="#">Dynamic routing gateway (DRG)</a>: for connectivity to networks outside the VCN's region (for example, your on-premises network by way of an IPSec VPN or FastConnect, or a peered VCN in another region)</li> <li>• <a href="#">Service gateway</a>: for private connectivity to Oracle services such as Object Storage</li> <li>• <a href="#">Local peering gateway (LPG)</a>: for connectivity to a peered VCN in the same region</li> </ul>
<a href="#">Route table rules</a>	Route tables control how traffic is routed from your VCN's subnets to destinations outside the VCN. Routing targets can be VCN gateways or a private IP address in the VCN.
<a href="#">IAM policies for virtual-network-family</a>	IAM policies specify access and actions permitted by IAM groups to resources in a VCN. For example, IAM polices can give administrative privileges to network administrators who manage the VCNs, and scoped-down permissions to normal users.

Oracle recommends that you periodically monitor Oracle Cloud Infrastructure Audit logs to review changes to VCN network security groups, security lists, route table rules, and VCN gateways.

## CHAPTER 29 Security Guide and Announcements

---

### *NETWORK SEGMENTATION: VCN SUBNETS*

- Formulate a tiered subnet strategy for the VCN, to control network access. A common design pattern is to have the following subnet tiers:
  1. **DMZ subnet** for load balancers
  2. **Public subnet** for externally accessible hosts such as NAT instances, intrusion detection (IDS) instances, and web application servers
  3. **Private subnet** for internal hosts such as databases

No special routing is required for the instances in the different subnets to communicate. However, you can control the types of traffic between the different tiers by using the VCN's network security groups or security lists.

- Instances in the private subnet only have private IP addresses and can be reached only by other instances in the VCN. Oracle recommends that you place security-sensitive hosts (DB systems, for example) in a private subnet, and use security rules to control the type of connectivity to hosts in a public subnet. In addition to VCN security rules, configure host-based firewalls such as iptables, firewalld for network access control, as a defense-in-depth mechanism.
- You can add a service gateway to your VCN to enable DB systems in the private subnet to directly back up to Object Storage without the traffic traversing the internet. You must set up the routing and security rules to enable that traffic. For more information for bare metal or virtual machine DB systems, see [Network Setup for DB Systems](#). For more information for Exadata DB systems, see [Network Setup for Exadata DB Systems](#).

### *NETWORK ACCESS CONTROL: VCN SECURITY RULES*

- Use your VCN's security rules to restrict network access to instances. A security rule is stateful by default, but can also be configured to be stateless. A common practice is to use stateless rules for high-performance applications. In a case where network traffic matches both stateful and stateless security lists, the stateless rule takes precedence. For more information about configuring VCN security rules, see [Security Rules](#).
- To prevent unauthorized access or attacks on Compute instances, Oracle recommends that you use a VCN security rule to allow SSH or RDP access only from authorized CIDR blocks rather than leave them open to the internet (0.0.0.0/0). For additional security,

you can temporarily enable SSH (port 22) or RDP (port 3389) access on an as-needed basis using the VCN API `UpdateNetworkSecurityGroupSecurityRules` (if you're using network security groups) or `UpdateSecurityList` (if you're using security lists). For more information about enabling RDP access, see [To enable RDP access](#) in [Creating an Instance](#). For performing instance health checks, Oracle recommends that you configure VCN security rules to allow ICMP pings. For more information, see [Rules to Enable Ping](#).

- Oracle recommends bastion hosts as a way to control external access (for example, SSH) to VCN hosts. Usually bastion hosts in a VCN public subnet control access to VCN private subnet hosts. For more information about setting up an SSH bastion host in a VCN, see the white paper [Bastion Hosts: Protected Access for Virtual Cloud Networks](#).
- VCN network security groups (NSGs) and security lists enable security-critical network access control to Compute instances, and it is important to prevent any unintended or unauthorized changes to NSGs and security lists. To prevent unauthorized changes, Oracle recommends that you use IAM policies to allow only network administrators to make NSG and security list changes.

### *SECURE CONNECTIVITY: VCN GATEWAYS AND FASTCONNECT PEERING*

- VCN gateways provide external connectivity (internet, on-premises, or peered VCN) to VCN hosts. See the table earlier in this topic for a list of the type of gateways. Oracle recommends that you use an IAM policy to allow only network administrators to create or modify VCN gateways.
- Carefully consider allowing internet access to any instances. For example, you don't want to accidentally allow internet access to sensitive database instances. In order for an instance in a VCN to be *publicly accessible from the internet*, you must configure the following VCN options:
  - The instance must be in a VCN public subnet.
  - The VCN containing the instance must have an internet gateway enabled and configured to be the routing target for outbound traffic.
  - The instance must have a public IP address assigned to it.

- The VCN security list for the instance's subnet must be configured to allow inbound traffic from 0.0.0.0/0. Or if you're using network security groups (NSG), the instance must be in an NSG that allows that traffic.
- VPN IPsec provides connectivity between a customer's on-premises network and VCN. You can create two IPsec tunnels for high availability. For more information about creating VPN tunnels to connect VCN DRG to customer CPEs, see [VPN Connect](#).
- FastConnect peering allows you to connect your on-premises network to your VCN with a private circuit so that the traffic does not traverse the public internet. You can set up private peering (to connect to private IP addresses), or public peering (to connect to Oracle Cloud Infrastructure public endpoints, such as for Object Storage). For more information about FastConnect peering options, see [FastConnect](#).

### *VIRTUAL SECURITY APPLIANCES IN A VCN*

- The Networking service lets you implement network security functions such as intrusion detection, application-level firewalls, and NAT (although you can instead use a [NAT gateway](#) with your VCN). You can do this by routing all the subnet traffic to a network security host, using route table rules that use a local VCN private IP address as a target. For more information, see [Using a Private IP as a Route Target](#). For high availability, you can assign the gateway security host a secondary private IP address, which you can move to a VNIC on a standby host in case of primary host failure. For more information about setting up a NAT instance in a VCN, see the white paper [NAT Instance Configuration: Enabling Internet Access for Private Subnets](#). Full network packet capture or network flow logs can be captured on the NAT instances using tcpdump, and the logs can be uploaded periodically to an Object Storage bucket.
- Virtual security appliances can be run as virtual machines (VMs) on a bring-your-own-hypervisor (BYOH) model on a bare metal instance. Virtual security appliance VMs running on the BYOH bare metal instance each have their own secondary VNIC, giving direct connectivity to other instances and services in the VNIC's VCN. For information about enabling BYOH on a bare metal instance using an open-source KVM hypervisor, see [Installing and Configuring KVM on Bare Metal Instances with Multi-VNIC](#).
- Virtual security appliances can also be installed on Compute virtual machines (VMs) where VMDK or QCOW2 images of security appliances can be imported using the bring

your own image (BYOI) feature. However, due to infrastructure dependencies, the BYOI feature might not work for some appliances, in which case the BYOH model would be another option to use. For more information about importing appliance images into Oracle Cloud Infrastructure, see [Bring Your Own Image \(BYOI\)](#).

### *LOAD BALANCERS*

- Oracle Cloud Infrastructure load balancers enable end-to-end TLS connections between a client's applications and a customer's VCN. The TLS connection can be terminated at an HTTP load balancer, or on a back-end server by using a TCP load balancer. The load balancers use TLS1.2 by default. For information about configuring an HTTPS listener, see [Managing Load Balancer Listeners](#). You can also upload your own TLS certificates. For more information see [Managing SSL Certificates](#).
- You can configure network access to load balancers by using VCN network security groups or security lists. This method provides similar functionality to traditional load balancer firewalls. For public load balancers, Oracle recommends that you use a regional public subnet (for example, DMZ subnet) for instantiating the load balancers in a highly available configuration across two different availability domains. You can configure the load balancer firewall rules by setting up the load balancer's network security groups or the subnet's security lists. For more information about creating load balancer security lists, see [Update Load Balancer Security Lists and Allow Internet Traffic to the Listener](#). Similarly, you must configure the VCN network security groups or security lists for the backend servers to limit traffic only from the public load balancers. For more information about configuring backend server security lists, see [Update Rules to Limit Traffic to Backend Servers](#).

### *DNS ZONES AND RECORDS*

DNS zones and records are critical for accessibility of web properties. Incorrect updates or unauthorized deletions could result in outage of services, accessed through the DNS names. Oracle recommends that you limit IAM users who can modify DNS zones and records.

### SECURITY POLICY EXAMPLES

#### *ALLOW USERS TO ONLY VIEW SECURITY LISTS*

Your network administrators are the personnel who should have the ability to create and manage network security groups and security lists.

However, you may have network users who need to know what security rules are in a particular network security group (NSG) or security list.

The first line in the following example policy allows the NetworkUsers group to view security lists and their contents. This policy does not let the group create, attach, delete, or modify security lists.

The second line lets the NetworkUsers group view the security rules in NSGs, and also view what VNICs and parent resources are in NSGs. The second line does not let the NetworkUsers group change the security rules in NSGs.

```
Allow group NetworkUsers to inspect security-lists in tenancy
Allow group NetworkUsers to use network-security-groups in tenancy
```

#### *PREVENT USERS FROM CREATING EXTERNAL CONNECTION TO THE INTERNET*

In some cases, you might need to prevent users from creating external internet connectivity to their VCN. In the following example policy, the NetworkUsers group is prevented from creating an internet gateway.

```
Allow group NetworkUsers to manage internet-gateways in tenancy
 where request.permission!='INTERNET_GATEWAY_CREATE'
```

#### *PREVENT USERS FROM UPDATING DNS RECORDS AND ZONES*

In the following example policy, the NetworkUsers group is prevented from deleting and updating DNS zones and records

```
Allow group NetworkUsers to manage dns-records in tenancy
 where all {request.permission!='DNS_RECORD_DELETE',
 request.permission!='DNS_RECORD_UPDATE'}
Allow group NetworkUsers to manage dns-zones in tenancy
 where all {request.permission!='DNS_ZONE_DELETE',
 request.permission!='DNS_ZONE_UPDATE'}
```

## CHAPTER 29 Security Guide and Announcements

---

### USEFUL CLI COMMANDS

In all the following examples, the environment variables \$T , \$C and \$VCN are set to tenancy OCID, compartment OCID, and VCN OCID, respectively.

#### *LIST OPEN SECURITY LISTS IN A VCN*

```
list open (0.0.0.0/0) security lists in VCN $VCN in compartment $C
oci network security-list list -c $C --vcn-id $VCN | grep "source" | grep "\"0.0.0.0/0\""
```

#### *LIST GATEWAYS IN A VCN*

```
list all internet gateways in VCN $VCN in compartment $C
oci network internet-gateway list -c $C --vcn-id $VCN
list all DRGs in compartment $C
oci network drg list -c $C
list all local peering gateways in vcn $VCN in compartment $C
oci network local-peering-gateway list -c $C --vcn-id $VCN
```

#### *LIST ROUTE TABLE RULES IN A VCN*

```
list route table rules in VCN $VCN in compartment $C
oci network route-table list -c $C --vcn-id $VCN
```

## Securing Object Storage

### SECURITY RECOMMENDATIONS

Assign least privileged access for IAM users and groups to resource types in object-family (buckets and objects). For example, the inspect verb gives the least privilege. Inspect lets you check to see if a bucket exists (`HeadBucket`) and list the buckets in a compartment (`ListBucket`). The manage verb gives all permissions on the resource. You can create IAM security policies to give appropriate bucket and object access to various IAM groups. For more information about IAM verbs and permissions for Object Storage buckets and objects, see [Details for Object Storage, Archive Storage, and Data Transfer](#). For users without IAM credentials, we recommend that you use pre-authenticated requests (PARs) to give time-bound access to objects or buckets.

### *PUBLIC BUCKETS SECURITY CONTROLS*

- A public bucket allows unauthenticated and anonymous reads to all objects in the bucket. Carefully evaluate the intended use case for public buckets before you enable public buckets. We recommend that you use pre-authenticated requests (PARs) to give bucket or object access(read or write) to users without IAM credentials. By defaults, buckets are created with no public access (access type is set to `NoPublicAccess`).
- You can make existing buckets public by updating the bucket access type to `ObjectRead` or `ObjectReadWithoutList`. To minimize the possibility of existing buckets being made public inadvertently or maliciously, `BUCKET_UPDATE` permission should be restricted to a minimal set of IAM groups.

### *PRE-AUTHENTICATED REQUEST (PAR)*

- Pre-authenticated requests (PARs) provide a mechanism to provide access to objects stored in buckets, to users who do not have IAM user credentials. In a PAR, an IAM user who has appropriate privileges for accessing objects, can create URLs which grant time-bound read or write access to these objects. For more information about creating PARs, see [Using Pre-Authenticated Requests](#).
- The creator of a PAR must have `PAR_MANAGE` IAM permission. You can create the following PARs:
  - Bucket PAR to allow writes to a bucket
  - Object PAR for reading an object
  - Object PAR for writing an object
  - Object PAR to read or write an object
- A PAR cannot be used to list objects in a bucket.
- All PAR accesses to a bucket or object are logged in Audit logs.
- We recommend that you note down the PAR URL created. By design, it is not possible to retrieve a forgotten PAR URL. If you forget a PAR URL, you must create a new PAR.

## CHAPTER 29 Security Guide and Announcements

---

### DATA DURABILITY

- Minimize data loss because of inadvertent deletes by an authorized user or malicious deletes. We recommended giving `BUCKET_DELETE` and `OBJECT_DELETE` permissions to a minimum set of IAM users and groups. Grant `DELETE` permission only to tenancy and compartment admins.
- Write once read many (WORM) compliance requires that objects cannot be deleted or modified. WORM is achieved by granting `OBJECT_CREATE`, `OBJECT_READ`, and `OBJECT_INSPECT` permissions to an IAM group. Grant `OBJECT_OVERWRITE` permission to prevent modification to existing an object.

### DATA ENCRYPTION

- All data in Object Storage is encrypted at rest by using AES-256. Encryption is on by default and cannot be turned off. Each object is encrypted with its encryption key, and the object encryption keys are encrypted with a master encryption key. In addition, customers can use client-side encryption to encrypt objects with their encryption keys before storing them in Object Storage buckets. An available option for customers is to use the Amazon S3 Compatibility API, along with client-side object encryption support available in AWS SDK for Java. See [Amazon S3 Compatibility API](#) for more details about on this SDK.
- Data in transit between customer clients (for example, SDKs and CLIs) and Object Storage public endpoints is encrypted with TLS 1.2 by default. FastConnect public peering allows on-premises access to Object Storage to go over a private network, rather than the public internet.

### DATA INTEGRITY

- To verify object data integrity, a cryptographic hash using MD5 is provided for all objects uploaded to Object Storage. We recommend that you verify that the offline MD5 hash of an object matches the hash value returned by the Console or API after upload. Oracle Cloud Infrastructure provides the object hash value in base64 encoding. To covert the base64 encoded hash value to hexadecimal, use the following command:

```
python -c 'print "BASE64-ENCODED-MD5-VALUE".decode("base64").encode("hex")'
```

Linux provides an `md5sum` command line utility to compute MD5 hash value of an object in hexadecimal format.

- Object Storage service supports multipart uploads for more efficient and resilient uploads, especially for large objects. In a multipart upload, a large object is broken up into smaller parts by specifying a part size in MiB. Each part is uploaded separately. Object Storage then combines all the parts to create the original object. If any of the parts fail to upload, only those parts need to be retried for upload, and not the entire object. In a multipart upload, the MD5 hash values are computed for each part, and an MD5 hash computed over all the individual hash values to get the output MD5 value. To verify the MD5 value returned for a multipart upload, follow the same process for offline MD5 hash calculation. A sample script for offline calculation of MD5 hash value for a multipart upload to Object Storage is available here (link: <https://gist.github.com/itemir/f5bc9fded6483cd79c89ebf4ca1cfd30>).

### SECURITY POLICY EXAMPLES

In the following examples, the policies are scoped to a tenancy. However, specifying a compartment name reduces the scope to a specific compartment in a tenancy.

#### *RESTRICT GROUP ACCESS TO SPECIFIC BUCKETS*

You can restrict access by a group to a specific bucket by using the specific bucket name (`target.bucket.name`), regular expression matching (`/*name/`, `/name*`, `/*name*/`), or defined tags (`target.tag.definition.name`).

The following is an example of restricting access by groups `BucketUsers` to a specific bucket.

```
Allow group BucketUsers to use buckets in tenancy
where target.bucket.name='BucketFoo'.
```

You can modify this policy to restrict access by group `BucketUsers` to all buckets whose names are prefixed with `ProjectA_`.

```
Allow group BucketUsers to use buckets in tenancy
where target.bucket.name=/ProjectA_*/
```

You can also match for post-fix (`/*_ProjectA/`) or substring (`/*ProjectA*/`).

## CHAPTER 29 Security Guide and Announcements

---

### *RESTRICT GROUP ACCESS TO READ OR WRITE TO OBJECTS IN A SPECIFIC BUCKET*

The following example allows listing and reading objects by group `BucketUsers` from a specific bucket named `BucketFoo`.

```
Allow group BucketUsers to read buckets in tenancy
Allow group BucketUsers to manage objects in tenancy
 where all {target.bucket.name='BucketFoo',
 any {request.permission='OBJECT_INSPECT',
 request.permission='OBJECT_READ'}}
```

The following policy modifies the previous policy to allow listing and writing objects to `BucketFoo`.

```
Allow group BucketUsers to read buckets in tenancy
Allow group BucketUsers to manage objects in tenancy
 where all {target.bucket.name='BucketFoo',
 any {request.permission='OBJECT_INSPECT',
 request.permission='OBJECT_CREATE'}}
```

You can restrict this policy to read or write access to a set of buckets by using regular expressions or tags rather than a specific bucket.

### *PREVENT DELETE OF BUCKETS OR OBJECTS*

In the following example, the group `BucketUsers` can perform all actions on buckets and objects except delete.

```
Allow group BucketUsers to manage objects in tenancy
 where request.permission!='OBJECT_DELETE'
Allow group BucketUsers to manage buckets in tenancy
 where request.permission!='BUCKET_DELETE'
```

The following example further restricts object deletion from the specific bucket (`BucketFoo`).

```
Allow group BucketUsers to manage objects in tenancy
 where any {target.bucket.name!='BucketFoo',
 all {target.bucket.name='BucketFoo',
 request.permission!='OBJECT_DELETE'}}
```

### *ENABLE WORM COMPLIANCE FOR OBJECTS*

The following policy enables WORM compliance by removing permissions for group `BucketUsers` to delete or update objects.

## CHAPTER 29 Security Guide and Announcements

---

```
Allow group BucketUsers to manage objects in tenancy
where any {request.permission='OBJECT_INSPECT',
 request.permission='OBJECT_READ',
 request.permission='OBJECT_CREATE'}
```

The following policy allows for WORM compliance.

```
Allow group BucketUsers to manage buckets in tenancy
where any {request.permission='BUCKET_INSPECT',
 request.permission='BUCKET_READ',
 request.permission='BUCKET_CREATE',
 request.permission='PAR_MANAGE'}
```

### *PREVENT PUBLIC BUCKETS CONFIGURATION*

As mentioned in previous section, `BUCKET_CREATE` and `BUCKET_UPDATE` permissions are required to create buckets or make existing private buckets public. Removing these permissions prevents users from creating buckets or making existing buckets public.

```
Allow group BucketUsers to manage buckets in tenancy
where any {request.permission='BUCKET_INSPECT',
 request.permission='BUCKET_READ',
 request.permission='PAR_MANAGE'}
```

### **USEFUL CLI COMMANDS**

Here are some useful commands to determine if you have public buckets or PARS in your tenancy.

#### *LIST OF PUBLIC BUCKETS*

The following command returns the public-access-type assigned to a bucket.

```
"public-access-type" of 'NoPublicAccess' indicates a private bucket, and
anything else ('ObjectRead') indicates a public bucket
oci os bucket get -ns <your_namespace> --bucket-name <bucket_name> | grep "public-access-type"
```

#### *LIST OF BUCKET PRE-AUTHENTICATED REQUESTS (PARS)*

The following command returns a list of object PARS in a bucket.

```
list all PARS for objects in bucket $BUCKET_NAME
oci os preauth-request list -ns <your_namespace> -bn <bucket_name>
```

### Securing Resource Manager

Resource Manager allows you to automate installing and provisioning Oracle Cloud Infrastructure resources by committing the provisioning instructions to configuration files. These configuration files capture the step-by-step provisioning instructions using a declarative language that follows the "infrastructure-as-code" model. The provisioning instructions are executed as "jobs"; the Oracle Cloud Infrastructure resources that are provisioned when you run the jobs are organized into "stacks."

Executing jobs and provisioning stacks is gated using role-based access control (RBAC), which is enabled by Oracle Cloud Infrastructure Identity and Access Management (IAM). This gives administrators granular control over user access to Oracle Cloud Infrastructure resources and the actions that users can take on these resources.

The Resource Manager security scheme rests on three pillars:

- **Security groups.** Administrator-defined groups that have permission to perform specific operations on stacks and jobs. Individual users are assigned to security groups and can then perform operations that are allowed by that group.
- **Permission sets.** Sets of permissions that are specific to jobs and stacks. Permission sets for jobs and stacks are listed in Table 1.
- **Operations.** The operations (or actions) that are allowed and the permissions that are required to perform each one. These are listed in Table 2.

#### RESOURCE MANAGER OPERATIONS AND PERMISSIONS

The Resource Manager supports two permission sets: one for stack resources and another for job resources. Table 1 lists permission sets that are associated with each resource type.

TABLE 1. RESOURCE TYPES AND PERMISSION SETS

Resource Type	Permissions
Stacks ( <code>orm-stacks</code> )	inspect orm-stack read orm-stack use orm-stack create orm-stack update orm-stack delete orm-stack
Jobs ( <code>orm-jobs</code> )	inspect orm-job read orm-job manage orm-job

Each of the permissions in the preceding table are associated with specific Resource Manager operations. The table that follows lists Resource Manager operations, then shows which permissions are required to execute each of the operations. Notice that the `CreateJob` operation requires two permissions.

TABLE 2. RESOURCE OPERATIONS AND REQUIRED PERMISSIONS

Operation	Permission
Generate a list of stacks ( <code>ListStacks</code> )	inspect orm-stack
Create a stack ( <code>CreateStack</code> )	create orm-stack
Get a stack ( <code>GetStack</code> )	read orm-stack
Update a stack ( <code>UpdateStack</code> )	update orm-stack
Delete a stack ( <code>DeleteStack</code> )	delete orm-stack
Get a stack Terraform configuration ( <code>GetStackTfConfig</code> )	read orm-stack

## CHAPTER 29 Security Guide and Announcements

Operation	Permission
List jobs ( <code>ListJobs</code> )	inspect orm-job
Create a job ( <code>CreateJob</code> )	use orm-stack <i>and</i> manage orm-job
Get a job ( <code>GetJob</code> )	read orm-jobs
Update a job ( <code>UpdateJob</code> )	manage orm-job
Cancel a job ( <code>CancelJob</code> )	manage orm-job
Get a job Terraform state file ( <code>GetJobTfState</code> )	read orm-job
Get a job Terraform configuration ( <code>GetJobTfConfig</code> )	read orm-job
Get a job Terraform execution plan ( <code>GetJobTfExecutionPlan</code> )	read orm-job
Get job logs ( <code>GetJobLogs</code> )	read orm-job

### RECOMMENDED SECURITY POLICIES

Following are policy recommendations for securing Resource Manager. For more information about security policies, see [Getting Started with Policies](#). See also [How Policies Work](#) and [Policy Syntax](#).

#### PERMISSION TO MANAGE STACKS AND JOBS

The following policy grants permission to a specified group to manage both stacks and jobs, and also to manage resources on the tenancy stacks.

```
Allow group <group_name> to manage orm-stacks in compartment
Allow group <group_name> to manage orm-jobs in compartment
```

#### PREVENT USERS FROM RUNNING DESTROY JOBS

In addition to granting permission to perform specific operations, you can also create a policy that prevents certain actions. The following policy example explicitly prevents members of a

## CHAPTER 29 Security Guide and Announcements

---

specified group from running Destroy jobs on a stack.

```
Allow group <group_name> to use orm-stacks in compartment
Allow group <group_name> to read orm-jobs in compartment
Allow group <group_name> to manage orm-jobs in compartment where any {target.job.operation = 'PLAN',
target.job.operation = 'APPLY'}
```

Notice that you must include the new permission to read orm-jobs in compartment because the statement includes a `where` condition that references variables that are not relevant to listing or getting jobs.

### POTENTIAL SECURITY RISKS AND MITIGATIONS

#### TERRAFORM STATE FILES

Terraform state (.tfstate) can contain sensitive data, including resource IDs and in some cases sensitive user data like passwords. HashiCorp provides recommendations for handling Terraform state in the article [Sensitive Data in State](#).

To control access to the Terraform state file, you can create a security policy that limits access to reading jobs, such as the following:

```
Allow group <group_name> to read orm-jobs in compartment
```



#### Note

Because the permission **read orm-jobs** also affects other operations such as getting logs and Terraform configurations, you should segregate state files in a compartment on which a restrictive policy will not limit the ability to perform other operations.

#### TERRAFORM CONFIGURATION .ZIP FILES

The Resource Manager workflow includes uploading your Terraform configuration to the service as a .zip file. Because the Terraform configuration can be accessed using the ORM API (`GetJobTfConfig`), it is highly recommended that you do not include sensitive information in your configuration files.

### Addressing Basic Configuration Issues

This topic lists procedures to address common configuration issues that affect the security of your cloud resources.

#### Block Volume

##### Block volume detached from instance

**Issue:** Ensure that only Oracle Cloud Infrastructure administrators can detach block volumes from instances.

**Basics:** When you detach a block volume it decouples the volume from its associated instance, affecting the data available to the instance. This could impact data availability from business-critical data to the successful completion of scheduled volume backups. To minimize loss of data due to inadvertent volume detachments by an authorized user or malicious volume detachments you should restrict the `VOLUME_ATTACHMENT_DELETE` permission to administrators.

##### To prevent detachment of block volumes:

The following policy allows the group `VolumeUsers` to manage volumes and volume attachments except for detaching volumes:

```
Allow group VolumeUsers to manage volumes in tenancy
Allow group VolumeUsers to manage volume-attachments in tenancy
 where request.permission!='VOLUME_ATTACHMENT_DELETE'
```

This change prevents `VolumeUsers` from detaching volumes from instances.

##### More information:

- [Securing Block Volume](#)
- [Getting Started with Policies](#)

- [How Policies Work](#)
- [For volume-family Resource Types](#)

### Compute

#### Instance created based on unapproved custom image

**Issue:** An instance was created from a custom image that is unsupported in your environment.

**Basics:** When users create instances they can select from Oracle-provided images, boot volumes from terminated instances, or custom images. Custom images represent a wide variety of images which can include images that aren't approved for your environment. If you use tags in your Oracle Cloud Infrastructure tenancy to identify approved images, verify whether the image the instance is based on is an approved image and terminate the instance if necessary.

#### To verify the tags for the image the instance was created from:

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Click the instance you're interested in.
3. Click the **Image** link to view the source image.
4. Click the **Tags** tab to view the tags applied to this image.

If the custom image does not have an approved tag, and the instance needs to be terminated, see [Terminating an Instance](#)

#### More information:

- [Securing Compute](#)
- [Resource Tags](#)

- [Creating an Instance](#)
- [Image Import/Export](#)
- [Bring Your Own Image \(BYOI\)](#)

### IAM

#### Member of the Administrators group used API keys

**Issue:** A user who is a member of the Administrators group accessed resources using an API key.

**Basics:**

- API keys are credentials used to grant programmatic access to Oracle Cloud Infrastructure.
- For security and governance reasons, users should only have access to resources they need to interact with.
- For individuals who are members of the Administrators group who also need access to resources through the API, create another user in IAM to which you attach the API keys. Grant the user with the API keys permissions to only the resources they need to interact with programmatically.

**To create a user, group, and policy with limited permissions:**

The following set of procedures shows you how to set up an example user with limited permissions. In this example, the user needs to be able to launch instances in a specific compartment.

The following procedure is for the Oracle Cloud Infrastructure Console.

#### Create a User

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity**

and click **Users**.

2. Click **Create User**.
3. In the **New User** dialog:
  - **Name:** Enter a unique name or email address for the new user.  
The value will be the user's login to the Console and must be unique across all other users in your tenancy.
  - **Description:** Enter a description (required).
4. Click **Create User**.

### Create a Group

Next, create the group ("InstanceLaunchers" ) that you will create the policy for.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.
2. Click **Create Group**.
3. In the **Create Group** dialog:
  - **Name:** Enter a unique name for your group, for example, "InstanceLaunchers".  
Note that the name cannot contain spaces.
  - **Description:** Enter a description (required).
4. Click **Create Group**.

### Create a Policy

In this example, the policy grants members of the group InstanceLaunchers permissions to launch instances in a specific compartment (CompartmentA).

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.

2. Click **Create Policy**.
3. Enter a unique **Name** for your policy, for example, "InstanceLaunchersPolicy". Note that the name cannot contain spaces.
4. Enter a **Description** (required), for example, "Grants users permission to launch instances in CompartmentA".
5. Enter the following **Statement**:

```
Allow group InstanceLaunchers to manage instance-family in compartment CompartmentA
```

This statement grants members of the InstanceLaunchers group permissions to launch and manage instances in the compartment called CompartmentA.

6. Click **Create Policy**.

### Add the User to the Group

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**.
2. In the **Users** list, find the user and click the name.
3. On the user detail page, click **Groups** (on the left side of the page). The list of groups that the user belongs to is displayed.
4. Click **Add User to Group**.
5. From the **Groups** list, select InstanceLaunchers.
6. Click **Add**.

### Upload an API signing key for the user

The following procedure works for a regular user or an administrator. Administrators can upload an API key for either another user or themselves.

 **Important**

The API key must be an **RSA key in PEM format (minimum 2048 bits)**. The PEM format looks something like this:

```
-----BEGIN PUBLIC KEY-----
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAAoTFqF...
...
-----END PUBLIC KEY-----
```

For more information about generating a public PEM key, see [Required Keys and OCIDs](#).

1. View the user's details:
  - If you're uploading an API key for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator uploading an API key for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. Click **Add Public Key**.
3. Paste the key's value into the window and click **Add**.

The key is added and its fingerprint is displayed (example fingerprint: d1:b2:32:53:d3:5f:cf:68:2d:6f:8b:5f:77:8f:07:13).

### More information:

- [Securing IAM](#)
- [How Policies Work](#) and [Common Policies](#)
- [Managing User Credentials](#)

- [Managing Groups](#)
- [Managing Users](#)

### Policy grants broad permissions

**Issue:** A policy grants full management permissions for at least one service in a compartment or in the tenancy.

**Basics:**

- Access to resources is controlled through policies. A *policy* is a document that specifies who can access which Oracle Cloud Infrastructure resources that your company has, and how. A policy simply allows a group to work in certain ways with specific types of resources in a particular compartment.
- For security and governance reasons, users should only have access to resources they need.
- Consider carefully the access level a user needs. Policy language provides a default set of verbs (`manage`, `use`, `read`, `inspect`) that allow you to easily scope users' permissions to a set of common tasks. For example, if a user needs to be able to update resources, but does not need to create or delete them, grant them the `use` permission, rather than the `manage` permission.
- The policy language is designed to let you write simple statements involving only verbs and resource-types, without having to state the permissions in the statement. For more fine-grained access control, you can use conditions combined with permissions or API operations to reduce the scope of access granted by a particular verb.
- Wherever possible, scope access to the specific compartments a user needs access to, rather than scoping it to the full tenancy.

**Tips for writing least-privilege policies:**

### Scope the policy to a compartment instead of the tenancy

Each policy consists of one or more policy *statements* that follow a basic syntax. Where possible, scope policies to compartments, rather than to the tenancy. For example, update a policy like this:

```
Allow group <group_name> to <verb><resource-type> in tenancy
```

to include just the compartments needed:

```
Allow group <group_name> to <verb><resource-type> in compartment <compartment_name>
```

If the user needs access to multiple compartments, create a policy statement for each compartment. It is then easy to remove access to individual compartments, if necessary.

### Scope permissions to those required to perform a job function

Oracle defines the possible verbs you can use in your policies. Here's a summary of the verbs, from least amount of access to the most:

Verb	Types of Access Covered	Target User
inspect	Ability to list resources, without access to any confidential information or user-specified metadata that may be part of that resource.	Third-party auditors
read	Includes <i>inspect</i> plus the ability to get user-specified metadata and the actual resource itself.	Internal auditors
use	Includes <i>read</i> plus the ability to work with existing resources (the actions vary by resource type). In general, this verb does not include the ability to create or delete that type of resource.	Day-to-day end users of resources
manage	Includes all permissions for the resource.	Administrators

Users who don't need to create or delete resources generally don't need manage permissions. If you have a policy like

```
Allow group <group_name> to manage <resource-type> in compartment <compartment_name>
```

but the user will never create or delete the resource-type, consider rewriting the policy to

```
Allow group <group_name> to use <resource-type> in compartment <compartment_name>
```

The [Policy Reference](#) includes details of the specific resource-types for each service, and which verb + resource-type combination gives access to which API operations.

### Service-specific links

- [Details for the Audit Service](#)
- [Details for Container Engine for Kubernetes](#)
- [Details for the Core Services](#) (this includes Networking, Compute, and Block Volume)
- [Details for the Database Service](#)
- [Details for the DNS Service](#)
- [Details for the Email Service](#)
- [Details for the File Storage Service](#)
- [Details for IAM](#)
- [Details for Load Balancing](#)
- [Details for Object Storage, Archive Storage, and Data Transfer](#)
- [Details for Registry](#)
- [Details for the Search Service](#)

For fine-grained access control, scope access using conditions and API operations

In a policy statement, you can use [conditions](#) combined with permissions or API operations to

reduce the scope of access granted by a particular verb.

For example, let's say you want group XYZ to be able to list, get, create, or update groups (change their description), but not delete them. To list, get, create, and update groups, you need a policy with `manage groups` as the verb and resource-type. According to the table in [Details for Verbs + Resource-Type Combinations](#), the permissions covered are:

- GROUP\_INSPECT
- GROUP\_UPDATE
- GROUP\_CREATE
- GROUP\_DELETE

To restrict access to only the desired permissions, you could add a condition *that explicitly states the permissions you want to allow*:

```
Allow group XYZ to manage groups in tenancy

where any {request.permission='GROUP_INSPECT',
 request.permission='GROUP_CREATE',
 request.permission='GROUP_UPDATE'}
```

An alternative would be a policy that *allows all permissions except* GROUP\_DELETE:

```
Allow group XYZ to manage groups in tenancy where request.permission != 'GROUP_DELETE'
```

Another alternative would be to write a condition *based on the specific API operations*. Notice that according to the table in [Permissions Required for Each API Operation](#), both `ListGroups` and `GetGroup` require only the GROUP\_INSPECT permission. Here's the policy:

```
Allow group XYZ to manage groups in tenancy

where any {request.operation='ListGroups',
 request.operation='GetGroup',
 request.operation='CreateGroup',
 request.operation='UpdateGroup'}
```

It can be beneficial to use permissions instead of API operations in conditions. In the future, if a new API operation is added that requires one of the permissions listed in the permissions-

based policy above, that policy will already control XYZ group's access to that new API operation.

But notice that you can further scope a user's access to a permission by *also* specifying a condition based on API operation. For example, you could give a user access to GROUP\_INSPECT, but then only to ListGroups.

```
Allow group XYZ to manage groups in tenancy

where all {request.permission='GROUP_INSPECT',
 request.operation='ListGroups'}
```

### More information:

- [Securing IAM](#)
- [How Policies Work](#) and [Common Policies](#)
- [Advanced Policy Features](#)
- [Managing Policies](#)

### API signing keys over 90 days old

**Issue:** A user's API signing keys are older than 90 days. Oracle recommends that you rotate API keys at least every 90 days.

#### Basics:

- API keys are credentials used to grant programmatic access to Oracle Cloud Infrastructure.
- It is a security engineering best practice and compliance requirement to rotate API keys regularly, every 90 days or less.
- Ensure that you test the new keys before you deactivate the old keys.

#### To generate and upload new API keys:

The following procedure is for the Oracle Cloud Infrastructure Console.

### Generate new API keys

You can use the following [OpenSSL](#) commands to generate the key pair in the required PEM format. If you're using Windows, you'll need to install [Git Bash for Windows](#) and run the commands with that tool.

1. If you haven't already, create a `.oci` directory to store the credentials:

```
mkdir ~/.oci
```

2. Generate the private key with one of the following commands.

- Recommended: To generate the key, encrypted with a passphrase you provide when prompted:

```
openssl genrsa -out ~/.oci/oci_api_key.pem -aes128 2048
```

**Note:** For Windows, you may need to insert `-passout stdin` to be prompted for a passphrase. The prompt will just be the blinking cursor, with no text.

```
openssl genrsa -out ~/.oci/oci_api_key.pem -aes128 -passout stdin 2048
```

- To generate the key with no passphrase:

```
openssl genrsa -out ~/.oci/oci_api_key.pem 2048
```

3. Ensure that only you can read the private key file:

```
chmod go-rwx ~/.oci/oci_api_key.pem
```

4. Generate the public key:

```
openssl rsa -pubout -in ~/.oci/oci_api_key.pem -out ~/.oci/oci_api_key_public.pem
```

**Note:** For Windows, if you generated the private key with a passphrase, you may need to insert `-passin stdin` to be prompted for the passphrase. The prompt will just be the blinking cursor, with no text.

```
openssl rsa -pubout -in ~/.oci/oci_api_key.pem -out ~/.oci/oci_api_key_public.pem -passin stdin
```

5. Copy the contents of the public key to the clipboard using `pbcopy`, `xclip` or a similar tool (you'll need to paste the value into the Console later). For example:

```
cat ~/.oci/oci_api_key_public.pem | pbcopy
```

Your API requests will be signed with your private key, and Oracle will use the public key to verify the authenticity of the request. You must upload the public key to IAM (instructions below).

### Get the key's fingerprint

You can get the key's fingerprint with the following OpenSSL command. If you're using Windows, you'll need to install [Git Bash for Windows](#) and run the command with that tool.

```
openssl rsa -pubout -outform DER -in ~/.oci/oci_api_key.pem | openssl md5 -c
```

When you upload the public key in the Console, the fingerprint is also automatically displayed there. It looks something like this: 12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:cd:ef

### Upload the API signing key for the user

You can upload the PEM public key in the Console, located at <https://console.us-ashburn-1.oraclecloud.com>. If you don't have a login and password for the Console, contact an administrator.

1. Open the Console, and sign in.
2. View the details for the user who will be calling the API with the key pair:
  - If you're signed in as this user, click your username in the top-right corner of the Console, and then click **User Settings**.
  - If you're an administrator doing this for another user, instead click **Identity**, click **Users**, and then select the user from the list.
3. Click **Add Public Key**.
4. Paste the contents of the PEM public key in the dialog box and click **Add**.

The key's fingerprint is displayed (for example, 12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:cd:ef).

Notice that after you've uploaded your first public key, you can also use the [UploadApiKey](#) API operation to upload additional keys. You can have up to three API key pairs per user. In an API request, you specify the key's fingerprint to indicate which key you're using to sign the request.

### Test the new key

Test the key in a sample API call against Oracle Cloud Infrastructure.

### Delete the old key

The following procedure works for a regular user or an administrator. Administrators can delete an API key for either another user or themselves.

1. View the user's details:
  - If you're deleting an API key for *yourself*: Open the **Profile** menu () and click **User Settings**.
  - If you're an administrator deleting an API key for *another user*: In the Console, click **Identity**, and then click **Users**. Locate the user in the list, and then click the user's name to view the details.
2. For the API key you want to delete, click **Delete**.
3. Confirm when prompted.

The API key is no longer valid for sending API requests.

#### More information:

- [Securing IAM](#)
- [API Signing Key](#)

### Tenancy administrator privilege grant to an IAM group

**Issue:** A group other than the Administrators group has been granted administrator privileges.

**Basics:**

- Granting the tenancy administrator privilege (`manage all-resources in tenancy`) to a group enables the members to have full access to all resources in the tenancy.
- This high-privilege entitlement must be controlled and restricted to only the users who need it to perform their job function.
- Verify with the Oracle Cloud Infrastructure administrator that this entitlement grant was sanctioned and that the membership of the group remains valid after the grant of the administrator privilege.
- Rather than create an alternative group with administrator privileges, consider instead adding users needing administrator privileges to the default Administrators group.

**To resolve this issue:**

Add users who need administrator privileges to the Administrators group:

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Groups**.
2. In the **Groups** list, click **Administrators**.
3. Click **Add User to Group**.
4. In the **Add User to Group** dialog, select the user from the **User** list.
5. Click **Add User**.

Remove the policy or policy statement that grants the (non-Administrators) group administration privileges.

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Policies**.  
A list of the policies in the compartment you're viewing is displayed. If you don't see

the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).

2. Click the policy you want to update.  
The policy's details and statements are displayed.
3. Find the statement that grants administrator privileges to the group. This policy will look like:

```
Allow group <group_name> to manage all-resources in tenancy
```

Click the the Actions icon (three dots) and then click **Delete**.

4. If the policy has no other statements, you can delete the policy by clicking **Delete** next to the policy name.

### More information:

- [Securing IAM](#)
- [Managing Policies](#)

## Networking: VCN, Load Balancers, and DNS

### No ingress rules in security lists

**Issue:** A VCN's security lists have no ingress rules. This means that the instances in the VCN can't receive incoming traffic.

### Basics:

- Security lists provide stateful and stateless firewall capability to control network access to your instances.
- A security list is configured at the subnet level and enforced at the instance level.
- You can associate multiple security lists with a subnet. A packet is allowed if it matches any rule in any of the security lists used by the subnet.

- If there are no ingress (inbound) rules in any of the subnet's security lists, no traffic is allowed to the instances in that subnet.
- For defense in depth, ingress security list rules should state a specific known source and not an open source (0.0.0.0/0).
- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted security lists.

### To add an ingress rule to an existing security list:

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Confirm you're viewing the compartment that contains the cloud network you're interested in.
3. Click the cloud network you're interested in.
4. Click **Security Lists**.
5. Click the security list you're interested in.
6. Click **Edit All Rules**.
7. Add at least one ingress rule:
  - a. In the **Allow Rules for Ingress** section, click **+ Rule**.
  - b. Choose whether it's a stateful or stateless rule (see [Stateful Versus Stateless Rules](#)). By default, rules are stateful unless you specify otherwise.
  - c. Enter the source CIDR. Typical CIDRs you might specify in a rule are the CIDR block for your on-premises network or a particular subnet. If you're setting up a security list rule to allow traffic with a service gateway, instead see [Task 3: \(Optional\) Update security rules](#).
  - d. Select the protocol (for example, TCP, UDP, ICMP, "All protocols", and so on).

- e. Enter further details depending on the protocol:
  - If you chose TCP or UDP, enter a source port range and destination port range. You can enter "All" to cover all ports. If you want to allow a specific [port](#), enter the port number (for example, 22 for SSH or 3389 for RDP) or a port range (for example, 20-22).
  - If you chose ICMP, you can enter "All" to cover all types and codes. If you want to allow a specific [ICMP type](#), enter the type and an optional code separated by a comma (for example, 3,4). If the type has multiple codes you want to allow, create a separate rule for each code.
8. When you're done, click **Save Security List Rules**.

This change enables ingress access from the source CIDR block listed in the rule. Add additional rules if you want to allow ingress from other known sources.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Security Lists](#)
- [UpdateSecurityList](#)

### Security list allows traffic from any IP address (open source)

**Issue:** A security list has at least one rule with an open source (0.0.0.0/0). This means that traffic could come from any source and is not controlled.

### Basics:

- Security lists provide stateful and stateless firewall capability to control network access to your instances.
- A security list is configured at the subnet level and enforced at the instance level.
- You can associate multiple security lists with a subnet. A packet is allowed if it matches any rule in any of the security lists used by the subnet.

- If there are no ingress (inbound) rules in any of the subnet's security lists, no traffic is allowed to the instances in that subnet.
- For defense in depth, ingress security list rules should state a specific known source and not an open source (0.0.0.0/0).
- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted security lists.

### To change the source of a security list rule:

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Confirm you're viewing the compartment that contains the cloud network you're interested in.
3. Click the cloud network you're interested in.
4. Click **Security Lists**.
5. Click the security list you're interested in.
6. Click **Edit All Rules**.
7. Locate the rule that lists 0.0.0.0/0 as the source CIDR.
8. For that rule, change 0.0.0.0/0 to the CIDR block of a known source.
9. Click **Save Security List Rules**.

This change restricts ingress so packets are allowed only from a specific CIDR block. Add additional rules if you want to allow ingress from other known sources.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Security Lists](#)
- [UpdateSecurityList](#)

### Security list allows traffic to sensitive ports

**Issue:** A security list has at least one rule that enables access to a sensitive port.

**Basics:**

- Security lists provide stateful and stateless firewall capability to control network access to your instances.
- A security list is configured at the subnet level and enforced at the instance level.
- You can associate multiple security lists with a subnet. A packet is allowed if it matches any rule in any of the security lists used by the subnet.
- If there are no ingress (inbound) rules in any of the subnet's security lists, no traffic is allowed to the instances in that subnet.
- For defense in depth, ingress security list rules should state a specific known source and not an open source (0.0.0.0/0).
- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted security lists.

**Recommendation:** Update the subnet's security list to enable access to instances through SSH (TCP port 22) or RDP (TCP port 3389) on a temporary, as-needed basis, and only from authorized CIDR blocks (not 0.0.0.0/0). To perform instance health checks, update the security list to allow ICMP pings.

**To change an existing security list:**

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Confirm you're viewing the compartment that contains the cloud network you're interested in.
3. Click the cloud network you're interested in.
4. Click **Security Lists**.

5. Click the security list you're interested in.
6. Click **Edit All Rules**.
7. Make one or more of these changes:
  - Delete an existing rule by clicking the **X** next to the rule.
  - Change an existing rule in the list. For example: change the source from 0.0.0.0/0 to the CIDR block of a known source.
  - Add a new rule by clicking **+ Rule** and entering values for the new rule.
8. Click **Save Security List Rules**.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [To enable RDP access](#)
- [Security Lists](#)
- [UpdateSecurityList](#)

### Internet gateway attached to VCN

**Issue:** A VCN has an internet gateway. The gateway must be authorized to be attached to the VCN and must not unintentionally expose resources to the internet.

### Basics:

- Gateways provide external connectivity to hosts in a VCN. For example: an internet gateway enables direct internet connectivity for instances that are in a public subnet and have a public IP address. A dynamic routing gateway (DRG) enables connectivity with the on-premises network over an [IPSec VPN](#) or [FastConnect](#).
- To enable traffic through the internet gateway from a particular subnet in the VCN, there must be a rule in the subnet's route table that lists the internet gateway as a route target. To delete the internet gateway from the VCN, you must first delete any route rules that specify the internet gateway as the target.

- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted VCNs.

### To remove an internet gateway from a VCN:

Prerequisite: Ensure that there are no route rules that specify the internet gateway as a target.

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Confirm you're viewing the compartment that contains the cloud network you're interested in.
3. Click the cloud network you're interested in.
4. Click **Internet Gateways**.
5. Click the Actions icon (three dots) for the internet gateway, and then click **Terminate**.
6. Confirm when prompted.

This change disables direct internet connectivity for the VCN.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Internet Gateway](#)
- [DeleteInternetGateway](#)
- [Public IP Addresses](#)

### Instance has a public IP

**Issue:** An instance has a public IP address. This means the instance could be publicly addressable if other required components are present and configured correctly in the VCN.

### Basics:

- Carefully consider allowing internet access to any instances. For example, don't accidentally allow internet access to sensitive DB systems.
- For an instance to be publicly addressable:
  - The instance must have a public IP address and reside in a public subnet in the VCN (instances in private subnets cannot have public IP addresses).
  - The subnet's security list must be configured to allow traffic for all IPs (0.0.0.0/0) and all ports.
  - The VCN must have an internet gateway and be configured to route outbound traffic from the subnet to the internet gateway.
- An instance can have more than one public IP address. A given public IP is assigned to a private IP on a particular VNIC on the instance. An instance can have more than one VNIC, and each VNIC can have more than one private IP.

### To remove a public IP address from an instance:

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Instances**.
2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Click **Attached VNICs**.  
The primary VNIC and any secondary VNICs attached to the instance are displayed.
5. Click the VNIC you're interested in.  
The VNIC's primary private IP and any secondary private IPs are displayed.
6. For the private IP you're interested in, click the Actions icon (three dots), and then click **Edit**.
7. In the **Public IP Address** section, for **Public IP Type**, select the radio button for **No Public IP**.
8. Click **Update**.

The public IP is unassigned from the instance.

**More information:**

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Public IP Addresses](#)
- [DeletePublicIp](#)
- [Internet Gateway](#)

### Load balancer has no inbound rules or listeners

**Issue:** A load balancer's subnet security lists have no ingress rules, or a load balancer has no listener. In this case, the load balancer can't receive incoming traffic.

**Basics:**

- Load balancers provide automated traffic distribution from one entry point to multiple servers reachable from your virtual cloud network (VCN). Each load balancer exists in a subnet governed by security list rules. A load balancer receives incoming data traffic from one or more listeners.
- Security lists provide stateful and stateless firewall capability to control network access to your load balancer and backend servers.
  - If there are no ingress (inbound) rules in any of the subnet's security lists, no traffic is allowed to the instances in that subnet.
  - For defense in depth, configure ingress security list rules to state a specific known source and not an open source (0.0.0.0/0).
- A listener is a logical entity that checks for incoming traffic on the load balancer's IP address.
  - To handle TCP, HTTP, and HTTPS traffic, you must configure at least one listener per traffic type.
  - You can apply path route rules to a listener to route traffic to the correct backend set without using multiple listeners or load balancers. A *path route* is a string that

the listener matches against an incoming URI to determine the appropriate destination backend set.

- Ensure that your Oracle Cloud Infrastructure load balancers use inbound rules or listeners to allow access only from known resources.
- Exceptions can be configured in CASB to reduce alerts from exempted load balancers.

### To enable a listener to accept traffic:

The following procedure is for the Oracle Cloud Infrastructure Console.

To enable a listener to accept traffic, you must update your VCN's security list rules:

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.  
The list of VCNs in the current compartment appears.
2. Click the name of the VCN containing your load balancer, and then click **Security Groups** or **Security Lists**.  
A list of the security groups or lists in the cloud network appears.
3. Click the name of the NSG or security list that applies to your load balancer.
4. Add or edit the existing rules to allow access from the appropriate resources.  
An NSG's security rules appear on the **Network Security Group Details** page. From there you can add, edit, or remove rules.  
The **Security List Details** page provides access to separate tables in which you can add or edit **Ingress Rules** or **Egress Rules**.  
For details on rule configuration, see [Security Rules](#).

### To create a listener:

Usually, you create a listener as part of the load balancer creation workflow. To create a listener for an existing load balancer:

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.

2. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Listeners**, and then click **Create Listener**.
4. In the **Create Listener** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the listener. The name must be unique, and cannot be changed. Avoid entering confidential information.
  - **Hostname:** Optional. Select up to 16 [virtual hostnames](#) for this listener.



### Important

To apply a virtual hostname to a listener, the name must be part of the load balancer's configuration. If the load balancer has no associated hostnames, you can [create one](#) on the **Hostnames** page.

- **Protocol:** Required. Specify the protocol to use, either HTTP or TCP.
- **Port:** Required. Specify the port on which to listen for incoming traffic.
- **Use SSL:** Optional. Check this box to associate an SSL certificate bundle with the listener. The following settings are required to enable SSL handling. See [Managing SSL Certificates](#) for more information.
  - **Certificate Name:** Required. The friendly name of the SSL certificate bundle to use.
  - **Verify Peer Certificate:** Optional. Select this option to enable peer certificate verification.
  - **Verify Depth:** Optional. Specify the maximum depth for certificate chain verification.

- **Backend Set:** Required. Specify the default backend set to which the listener routes traffic.
- **Idle Timeout in Seconds:** Optional. Specify the maximum idle time in seconds. This setting applies to the time allowed between two successive receive or two successive send network input/output operations during the HTTP request-response phase.



### Tip

The maximum value is 7200 seconds. For more information, see [Connection Management](#).

- **Path Route Set:** Optional. Specify the name of the set of path-based routing rules that applies to this listener's traffic.



### Important

- To apply a [path route set](#) to a listener, the set must be part of the load balancer's configuration.
- To remove a path route set from an existing listener, choose **None** as the **Path Route Set** option. The path route set remains available for use by other listeners on this load balancer.

5. Click **Create**.

When you create a listener, you must also [update your VCN's security list rules](#) to allow traffic to that listener.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Security Lists](#)
- [Managing a Load Balancer](#)
- [Managing Load Balancer Listeners](#)
- [Managing Request Routing](#)

### Load balancer has no backend sets

**Issue:** A load balancer has no backend set. In this case, the load balancer has no place to distribute incoming data and no means to monitor backend server health.

#### **Basics:**

- A backend set is a logical entity defined by a load balancing policy, a health check policy, and a list of backend servers.
- The backend set determines the load balancer's traffic distribution policy, such as:
  - IP Hash
  - Least Connections
  - Weighted Round Robin
- You specify the test parameters to confirm the health of backend servers when you create a backend set.
- If you have an existing load balancer with no backend set, you can specify the backend servers that receive traffic from the load balancer after you create a backend set.
- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted load balancers.

#### **To create a backend set:**

The following procedure is for the Oracle Cloud Infrastructure Console.

Usually, you create a backend set as part of the load balancer creation workflow. To create a backend set for an existing load balancer:

1. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
2. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
3. In the **Resources** menu, click **Backend Sets** (if necessary), and then click **Create Backend Set**.
4. In the **Create Backend Set** dialog box, enter the following:
  - **Name:** Required. Specify a friendly name for the backend set. It must be unique within the load balancer, and it cannot be changed.  
Valid backend set names include only alphanumeric characters, dashes, and underscores. Backend set names cannot contain spaces. Avoid entering confidential information.
  - **Traffic Distribution Policy:** Required. Choose the load balancer policy for the backend set. The available options are:
    - **IP Hash**
    - **Least Connections**
    - **Weighted Round Robin**

For more information on these policies, see [How Load Balancing Policies Work](#).



### Tip

You cannot add a backend server marked as **Backup** to a backend set that uses the IP Hash policy.

- **Use SSL:** Optional. Check this box to associate an SSL certificate bundle with the backend set.

If there are no certificate bundles attached to the load balancer, this option is disabled.

- **Certificate Name:** Required. Select the certificate bundle to use. You can choose any certificate bundle that is attached to the current load balancer. See [Managing SSL Certificates](#) for more information.
  - **Verify Peer Certificate:** Optional. Select this option to enable peer certificate verification.
  - **Verify Depth:** Optional. Specify the maximum depth for certificate chain verification.
- **Session Persistence:** Optional. Specify how the load balancer manages session persistence.



### Important

See [Session Persistence](#) for important information on configuring these settings.

- **Disable Session Persistence:** Choose this option to disable cookie-based session persistence.
- **Enable Application Cookie Persistence:** Choose this option to enable persistent sessions from a single logical client when the response from a backend application server includes a `Set-cookie` header with the cookie name you specify.
  - **Cookie Name:** The cookie name used to enable session persistence. Specify `*` to match any cookie name. Avoid entering confidential information.

- **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
- **Enable Load Balancer Cookie Persistence:** Choose this option to enable persistent sessions based on a cookie inserted by the load balancer.
  - **Cookie Name:** Specify the name of the cookie used to enable session persistence. If blank, the default cookie name is `X-Oracle-BMC-LBS-Route`.

Ensure that any cookie names used at the backend application servers are different from the cookie name used at the load balancer. Avoid entering confidential information.
  - **Disable Fallback:** Check this box to disable fallback when the original server is unavailable.
  - **Domain Name:** Optional. Specify the domain in which the cookie is valid.

This attribute has no default value. If you do not specify a value, the load balancer does not insert the domain attribute into the `Set-cookie` header.
  - **Path:** Optional. Specify the path in which the cookie is valid. The default value is `/`.
  - **Expiration Period in Seconds:** Optional. Specify the amount of time the cookie remains valid. If blank, the cookie expires at the end of the client session.
  - **Attributes**
    - **Secure:** Specify whether the `Set-cookie` header should contain the `Secure` attribute. If selected, the client sends the cookie only using a secure protocol.

If you enable this setting, you cannot associate the corresponding backend set with an HTTP listener.

- **HTTP Only:** Specify whether the `Set-cookie` header should contain the `HttpOnly` attribute. If selected, the cookie is limited to HTTP requests. The client omits the cookie when providing access to cookies through non-HTTP APIs such as JavaScript channels.
- **Health Check:** Required. Specify the test parameters to confirm the health of backend servers.
  - **Protocol:** Required. Specify the protocol to use, either HTTP or TCP.



### Important

Configure your health check protocol to [match your application or service](#).

- **Port:** Optional. Specify the backend server port against which to run the health check.



### Tip

You can enter the value '0' to have the health check use the backend server's traffic port.

- **URL Path (URI):** (HTTP only) Required. Specify a URL endpoint against which to run the health check.
- **Interval in ms:** Optional. Specify how frequently to run the health check, in milliseconds. The default is 10000 (10 seconds).

- **Timeout in ms:** Optional. Specify the maximum time in milliseconds to wait for a reply to a health check. A health check is successful only if a reply returns within this timeout period. The default is 3000 (3 seconds).
- **Number of retries:** Optional. Specify the number of retries to attempt before a backend server is considered "unhealthy". This number also applies when recovering a server to the "healthy" state. The default is '3'.
- **Status Code:** (HTTP only) Optional. Specify the status code a healthy backend server must return.
- **Response Body Regex:** (HTTP only) Optional. Provide a regular expression for parsing the response body from the backend server.

5. Click **Create**.

After your backend set is provisioned, you must specify backend servers for the set. See [Managing Backend Servers](#) for more information.

### More information:

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Managing Backend Sets](#)
- [Editing Health Check Policies](#)
- [Managing a Load Balancer](#)

### Load balancer SSL certificate expires in *X* days

**Issue:** A load balancer's SSL certificate expires soon. When the certificate expires, data traffic can be interrupted and security compromised.

### Basics:

- To ensure continuous security and usability, SSL certificates must be rotated on a timely basis.

- You can configure an exception in Oracle CASB Cloud Service to reduce alerts from exempted load balancers.

### To rotate a load balancer's certificate bundle:

The following procedure is for the Oracle Cloud Infrastructure Console.

To ensure consistent service, you must update (rotate) expiring certificates:

1. Update your client or backend server to work with a new certificate bundle.

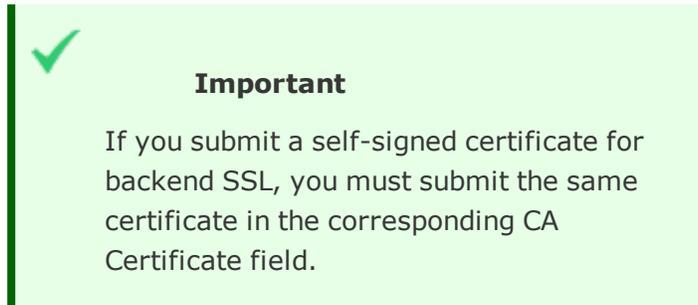


#### Note

The steps to update your client or backend server are unique to your system.

2. Upload the new SSL certificate bundle to the load balancer
  - a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
  - b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
  - c. Click the load balancer you want to configure.
  - d. In the **Resources** menu, click **Certificates**, and then click **Add Certificate**.
  - e. In the **Add Certificate** dialog box, enter the following:
    - **Certificate Name:** Required. Specify a friendly name for the certificate bundle. It must be unique within the load balancer, and it cannot be changed in the Console. (It can be changed using the API.) Avoid entering confidential information.
    - **Choose SSL Certificate File:** Required. Drag and drop the certificate file, in PEM format, into the **SSL Certificate** field.

Alternatively, you can choose the **Paste SSL Certificate** option to paste a certificate directly into this field.



- **Specify CA Certificate:** Optional. (Recommended for backend SSL termination configurations.) Select (check) this box if you want to provide a CA certificate.
  - **Choose CA Certificate File:** Drag and drop the CA certificate file, in PEM format, into the **CA Certificate** field.  
Alternatively, you can choose the **Paste CA Certificate** option to paste a certificate directly into this field.
- **Specify Private Key:** Optional. (Required for SSL termination.) Select (check) this box if you want to provide a private key for the certificate.
  - **Choose Private Key File:** Drag and drop the private key, in PEM format, into the **Private Key** field.  
Alternatively, you can choose the **Paste Private Key** option to paste a private key directly into this field.
  - **Enter Private Key Passphrase:** Optional. Specify the private key passphrase.

f. Click **Add Certificate**.

3. Edit listeners or backend sets (as needed) so they use the new certificate

### bundle

#### *EDITING A LISTENER:*

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Choose the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. In the **Resources** menu, click **Listeners**.
- d. For the listener you want to edit, click the Actions icon (three dots), and then click **Edit Listener**.
- e. In the **Certificate Name** drop-down list, choose the new certificate bundle.
- f. Click **Submit**.

#### *EDITING A BACKEND SET:*



#### **Warning**

Updating the backend set temporarily interrupts traffic and can drop active connections.

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. In the **Resources** menu, click **Backend Sets**, and then click the name of the backend set you want to edit.
- d. Click **Edit Backend Set**.

- e. In the **Edit Backend Set** dialog box, select (check) **Use SSL**.
- f. In the **Certificate Name** drop-down list, choose the new certificate bundle.
- g. Click **Save Changes**.

#### 4. (Optional) Remove the expiring SSL certificate bundle



##### **Important**

You cannot delete an SSL certificate bundle that is associated with a listener or backend set. Remove the bundle from any additional listeners or backend sets before deleting.

- a. Open the navigation menu. Under the **Core Infrastructure** group, go to **Networking** and click **Load Balancers**.
- b. Click the name of the **Compartment** that contains the load balancer you want to modify, and then click the load balancer's name.
- c. Click the load balancer you want to configure.
- d. In the **Resources** menu, click **Certificates**.
- e. For the certificate you want to delete, click the Actions icon (three dots), and then click **Delete**.
- f. Confirm when prompted.

##### **More information:**

- [Securing Networking: VCN, Load Balancers, and DNS](#)
- [Managing SSL Certificates](#)
- [Managing Load Balancer Listeners](#)

- [Managing Backend Sets](#)
- [Managing a Load Balancer](#)

### Object Storage

#### Public buckets detected

**Issue:** Public buckets were detected in your tenancy. Confirm that the creation of each public bucket is intentional and authorized. If the bucket is not sanctioned for public access, follow the procedure for changing the visibility of a bucket and make the bucket private.

#### Basics:

- Carefully assess the business requirement for public access to a bucket. When you enable anonymous access to a bucket, users can obtain object metadata, download bucket objects, and optionally list bucket contents.
- Changing the type of access is bi-directional. You can change a bucket's access from public to private or from private to public.
- Changing the type of access doesn't affect existing pre-authenticated requests. Existing pre-authenticated requests still work.

#### To change the visibility of a bucket (private or public):

The following procedure is for the Oracle Cloud Infrastructure Console.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.  
A list of the buckets in the compartment you're viewing is displayed. If you don't see the one you're looking for, verify that you're viewing the correct compartment (select from the list on the left side of the page).
2. Click the bucket name to see the bucket details.  
**Visibility:** shows the current bucket setting, which is **Private** by default.
3. Click **Edit Visibility**.
4. In the **Edit Visibility** dialog box, edit the visibility settings:

- **Visibility**
    - **Public**
    - **Private**
  - If you select **Public** to enable public access, decide whether or not you want to let users list the bucket contents. Click **Allow users to list objects from this bucket** to set the visibility of bucket object lists.
5. Click **Save Changes**.

### More information:

- [Securing Object Storage](#)
- [Managing Buckets](#)
- [Using Pre-Authenticated Requests](#)

## Oracle Cloud Security Response to Intel L1TF Vulnerabilities

Intel disclosed a new set of speculative execution side-channel processor vulnerabilities affecting their processors. For more information, see [Vulnerability Note VU#584653](#). These L1 Terminal Fault (L1TF) vulnerabilities affect a number of Intel processors, and they have received the following CVE identifiers:

- CVE-2018-3615, which impacts Intel Software Guard Extensions (SGX) and has a CVSS Base Score of 7.9.
- CVE-2018-3620, which impacts operating systems and System Management Mode (SMM) running on Intel processors and has a CVSS Base Score of 7.1.
- CVE-2018-3646, which impacts virtualization software and Virtual Machine Monitors (VMM) running on Intel processors and has a CVSS Base Score of 7.1.

See [Intel Processor L1TF vulnerabilities: CVE-2018-3615, CVE-2018-3620, CVE-2018-3646](#) for more information.

### Oracle Cloud Infrastructure

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

However, vulnerability CVE-2018-3620 could enable a rogue user mode process to read privileged kernel memory within the same virtual machine. As a result, if you manage your own operating systems (OS), you are advised to keep up with OS security patches to address this vulnerability.

The following sections contain the details of mitigations and actions.

#### **Oracle Cloud Infrastructure Compute**

For details and required actions related to the Compute service's VM and bare metal instances, see [Oracle Cloud Infrastructure Customer Advisory for L1TF Impact on the Compute Service](#).

#### **Oracle Cloud Infrastructure Database**

If you use Autonomous Data Warehouse and Autonomous Transaction Processing, you have no further action to take.

For details and required actions related to Oracle Cloud Infrastructure offerings for VM DB systems, bare metal DB systems, and Exadata DB systems, see [Oracle Cloud Infrastructure Customer Advisory for L1TF Impact on the Database Service](#).

#### **Platform Service and Kubernetes Services on Oracle Cloud Infrastructure**

Oracle has deployed technical mitigations designed to prevent malicious attacker's VM instance from accessing data from other VM instances on the same hypervisor.

However, vulnerability CVE-2018-3620 could enable a rogue user-mode process to read privileged kernel memory within the same virtual machine. As a result, Platform Service hosts managed by Oracle are being patched by Oracle. If you manage your own operating systems you're advised to keep up with the OS security patches to address this vulnerability.

### Other Oracle Cloud Infrastructure Services

Mitigations designed to protect all other Oracle Cloud Infrastructure services have been deployed. Oracle will notify and coordinate directly with customers for any additional required maintenance activities.

### Oracle Cloud Infrastructure Classic and Oracle Platform Service on Oracle Cloud Infrastructure Classic

For more information see [Oracle Cloud Infrastructure Classic](#).

Oracle is deploying technical mitigations designed for Infrastructure and Platform Services on Oracle Cloud Infrastructure Classic. Some customers may experience reboots or downtime associated while deploying these mitigations.

Vulnerability CVE-2018-3620 could enable a rogue user-mode process to read privileged kernel memory within the same virtual machine. As a result, Platform Service hosts managed by Oracle are being patched by Oracle. If you manage your own operating systems you're advised to keep up with the OS security patches to address this vulnerability.

### Oracle Cloud Infrastructure Customer Advisory for L1TF Impact on the Compute Service

Intel disclosed a new set of speculative execution side-channel processor vulnerabilities affecting their processors. For more information, see [Vulnerability Note VU#584653](#). These L1 Terminal Fault (L1TF) vulnerabilities affect a number of Intel processors, and they have received the following CVE identifiers:

- CVE-2018-3615, which impacts Intel Software Guard Extensions (SGX) and has a CVSS Base Score of 7.9.
- CVE-2018-3620, which impacts operating systems and System Management Mode (SMM) running on Intel processors and has a CVSS Base Score of 7.1.
- CVE-2018-3646, which impacts virtualization software and Virtual Machine Monitors (VMM) running on Intel processors and has a CVSS Base Score of 7.1.

See the [Oracle Cloud Security Response to Intel L1TF Vulnerabilities](#) for more information.

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

You should be aware that the vulnerability CVE-2018-3620 could enable a rogue user-mode process to read privileged kernel memory within the same operating system (OS). As a result, you are advised to keep up with OS security patches to address this vulnerability. See [Protecting your Compute Instance Against the L1TF Vulnerability](#) for instructions to patch the OS on the instances you manage.

### **Additional Guidance for Oracle Cloud Infrastructure Bare Metal Instances**

Bare metal instances in Oracle Cloud Infrastructure offer you full control of a physical server. Oracle Cloud Infrastructure's network virtualization is designed and configured to protect these instances from unauthorized access of other instances on the Oracle Cloud Infrastructure network, including other customer instances, both VM instances and other bare metal instances.

If you're running your own virtualization stack or hypervisors on bare metal instances, the L1TF vulnerability allows a virtual machine to access privileged information from the underlying hypervisor or other VMs on the same bare metal instance. You should review the [Intel recommendations](#) about vulnerabilities CVE-2018-3615, CVE-2018-3620, and CVE-2018-3646, and make changes to your configurations as you deem appropriate.

### **Protecting your Compute Instance Against the L1TF Vulnerability**

Intel disclosed a new set of speculative execution side-channel processor vulnerabilities affecting their processors, for more information, see [Vulnerability Note VU#584653](#). These L1 Terminal Fault (L1TF) vulnerabilities affect a number of Intel processors, and they have received the following CVE identifiers:

- CVE-2018-3615 which impacts Intel Software Guard Extensions (SGX) and has a CVSS Base Score of 7.9.
- CVE-2018-3620 which impacts operating systems and System Management Mode (SMM) running on Intel processors and has a CVSS Base Score of 7.1.

- CVE-2018-3646 which impacts virtualization software and Virtual Machine Monitors (VMM) running on Intel processors and has a CVSS Base Score of 7.1.

See the [Oracle Cloud Security Response to Intel L1TF Vulnerabilities](#) for more information.

### **RECOMMENDED ACTION**

Oracle recommends that you patch the operating systems for your existing bare metal and virtual machine (VM) instances, and verify that this includes the patch for the CVE-2018-3620 vulnerability. For VM instances, the Oracle Cloud Infrastructure team has implemented the necessary workarounds designed to mitigate the CVE-2018-3646 vulnerability. For bare metal instances using virtualization technology, you should also follow these instructions to ensure that they are mitigated against the CVE-2018-3646 vulnerability.

If you're running your own virtualization stack or hypervisors on bare metal instances, you should apply the patch for the CVE-2018-3646 vulnerability.

The information in the following sections detail the commands needed to update your running instances created from [Oracle-Provided Images](#).

The following Oracle-provided image releases have been updated with the recommended patches, so instances created from these images or new image releases include the recommended patches for the L1TF vulnerability.

### Oracle-provided images updated with recommended patches for the L1TF vulnerability

- [Oracle-Linux-7.5-2018.08.14-0](#)
- [Oracle-Linux-7.5-Gen2-GPU-2018.08.14-0](#)
- [Oracle-Linux-6.10-2018.08.14-0](#)
- [CentOS-7-2018.08.15-0](#)
- [CentOS-6.10-2018.08.15-0](#)
- [Canonical-Ubuntu-18.04-2018.08.15-0](#)
- [Canonical-Ubuntu-16.04-2018.08.15-0](#)

## CHAPTER 29 Security Guide and Announcements

---

- [Canonical-Ubuntu-16.04-Gen2-GPU-2018.08.15-0](#)
- [Canonical-Ubuntu-14.04-2018.08.15-0](#) (deprecated)
- [Windows-Server-2016-Standard-Edition-VM-Gen2-2018.08.16-0](#)
- [Windows-Server-2016-Datacenter-Edition-BM-Gen2-2018.08.16-0](#)
- [Windows-Server-2016-Datacenter-Edition-BM-Gen2-DenseIO-2018.08.16-0](#)
- [Windows-Server-2012-R2-Standard-Edition-VM-2018.08.14-0](#)
- [Windows-Server-2012-R2-Standard-Edition-VM-Gen2-2018.08.14-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-2018.08.15-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-Gen2-DenseIO-2018.08.15-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-Gen2-2018.08.15-0](#)
- [Windows-Server-2008-R2-Enterprise-Edition-VM-2018.08.15-0](#)
- [Windows-Server-2008-R2-Enterprise-Edition-VM-Gen2-2018.08.15-0](#)

For your running instances created from imported custom images, refer to the operating system (OS) vendor's guidance to patch the OS for the L1TF vulnerability.

### **PATCHING ORACLE LINUX INSTANCES**

For Oracle Linux, the patches for the CVE-2018-3620 and CVE-2018-3646 vulnerabilities are addressed by the same set of patches.

Bare metal instances must have the latest microcode updates from Intel. This step is not required for VM instances.

To install the latest microcode updates, run the following command:

```
sudo yum update microcode_ctl
```

The microcode RPM should be greater than or equal to `microcode_ctl-2.1-29.2.0.4.e17_5.x86_64.rpm`. This is the version of the microcode package that shipped for the Spectre v3a and Spectre v4 updates. No additional update is required. In addition to the microcode update, you should also patch your bare metal instances using the following set of instructions.

### To patch the OS for bare metal and VM instances with downtime

The `yum-plugin-security` package allows you to use `yum` to obtain a list of all of the errata that are available for your system, including security updates. You can also use Oracle Enterprise Manager 12c Cloud Control or management tools such as Katello, Pulp, Red Hat Satellite, Spacewalk, and SUSE Manager to extract and display information about errata.

1. To install the `yum-plugin-security` package, run the following command:

```
sudo yum install yum-plugin-security
```

2. Use the `--cve` option to display the errata that correspond to a specified CVE, and to install those required packages, by running the following commands:

```
sudo yum updateinfo list --cve CVE-2018-3620
sudo yum update --cve CVE-2018-3620
```

A system reboot will be required once the package is applied. By default, the boot manager will automatically enable the most recent kernel version. For more information on using `yum update`, visit [Installing and Using the Yum Security Plugin](#).

3. After the system reboots, ensure that the following file is populated.

```
cat /sys/devices/system/cpu/vulnerabilities/l1tf
```

### PATCHING WINDOWS INSTANCES

#### *PROTECTING NEW WINDOWS VM AND BARE METAL INSTANCES*

When you create a new VM or bare metal instance based on the latest Oracle-provided Windows images, the image includes the Microsoft-recommended patches to protect against the L1TF vulnerability. Windows bare metal instances also include the latest microcode updates from Intel.

There is no further action required from you to protect your new Windows-based VM or bare metal instances from the L1TF vulnerability. You should ensure that you keep the your instances updated with the latest patches as recommended by your OS vendor.

### *PROTECTING EXISTING WINDOWS VM AND BARE METAL INSTANCES*

#### To update the microcode for existing bare metal instances

Bare metal instances launched before the Oracle-provided Windows images were updated must have the latest microcode updates from Intel. You need to recycle your Windows bare metal instances in order to receive the latest Intel microcode update. This step is not required for VM instances.

1. Create a new custom image of your Windows bare metal instance, see [Creating Windows Custom Images](#) for more information.
2. Terminate your existing Windows bare metal instance.
3. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**. Find the custom image you want to use.
4. Click the Actions icon (three dots), and then click **Create Instance**.
5. Provide additional launch options as described in [Creating an Instance](#).

Once you have completed these steps, perform the steps in the next procedure to update the instance with the latest OS updates from Microsoft.

#### To patch the OS for bare metal and VM instances with downtime

Windows images include the Windows Update utility, which you can run to get the latest Windows updates from Microsoft. You have to configure the security list on the subnet on which the instance is running to allow instances to access Windows update servers. See [Windows OS Updates for Windows Images](#) and [Security Lists](#) for more information.

1. Verify that you have installed the latest Windows OS security update from Microsoft.
  - a. If automatic updates are turned on, the updates should be automatically delivered to the instance.
  - b. To manually check for the latest update, select **Start**.
  - c. In **Settings** select **Updates & security** and then select **Windows Update**.

- d. In **Windows Update**, click **Check for updates**.
- e. When you turn on automatic updates, this update will be downloaded and installed automatically. For more information about how to turn on automatic updates, see [Windows Update: FAQ](#).

For additional details see [Windows Server guidance to protect against L1 terminal fault](#).

### PATCHING UBUNTU OR CENTOS INSTANCES

When you create a new VM or bare metal instance based on the latest Oracle-provided Ubuntu or CentOS images, the image includes the recommended patches to protect against the L1TF vulnerability, see for more information [L1 Terminal Fault \(L1TF\)](#) and [L1TF - L1 Terminal Fault Attack - CVE-2018-3620 & CVE-2018-3646](#).

For existing VM or bare metal instances you should follow the guidance provided by the OS vendor for patching systems.

## Oracle Cloud Infrastructure Customer Advisory for L1TF Impact on the Database Service

Intel disclosed a new set of speculative execution side-channel processor vulnerabilities affecting their processors. For more information, see [Vulnerability Note VU#584653](#). These L1 Terminal Fault (L1TF) vulnerabilities affect a number of Intel processors, and they have received the following CVE identifiers:

- CVE-2018-3615, which impacts Intel Software Guard Extensions (SGX) and has a CVSS Base Score of 7.9.
- CVE-2018-3620, which impacts operating systems and System Management Mode (SMM) running on Intel processors and has a CVSS Base Score of 7.1.
- CVE-2018-3646, which impacts virtualization software and Virtual Machine Monitors (VMM) running on Intel processors and has a CVSS Base Score of 7.1.

See the [Oracle Cloud Security Response to Intel L1TF Vulnerabilities](#) for more information.

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

### **Autonomous Data Warehouse and Autonomous Transaction Processing**

Autonomous Data Warehouse provides fully managed databases optimized for running data warehouse workloads. Autonomous Transaction Processing provides fully managed databases optimized for running online transaction processing and mixed database workloads. Autonomous Data Warehouse and Autonomous Transaction Processing are not affected by the L1TF vulnerabilities, CVE-2018-3615, CVE-2018-3620, and CVE-2018-3646. No further action is required by customers.

### **Guidance for the DatabaseService on Bare Metal Instances**

The Database service on Oracle Cloud Infrastructure bare metal instances offer customers full control over their Oracle Database running on a physical server. Oracle Cloud Infrastructure's network virtualization is designed and configured to protect these instances from unauthorized access from other instances on the Oracle Cloud Infrastructure network, including other customer instances, both VM instances and other bare metal instances.

### **Actions for Customers with VM DB Systems, Bare Metal DB Systems, or Exadata DB Systems**

Vulnerability CVE-2018-3620 could enable a rogue user-mode process to read privileged kernel memory within the same operating system. As a result, you need to patch these systems once these patches are available. These patches will be available shortly and Oracle will update this page when the operating system (OS) patches are published. Oracle will update the Database base images with the latest patches for new instance launches.

Once the patches are available, use the following instructions to patch a running instance:

- For DB systems on bare metal instances, apply the OS patches following the instructions in [Updating a DB System](#).
- For DB systems on a VM instance, configured using the Oracle Cloud Infrastructure

Database service, apply the OS patches following the instructions in [Updating a DB System](#).

- For the DB systems on a VM instance configured using the Oracle Platform Service Manager, apply the OS patches following the instructions in [Applying Linux OS Security Patches by Using the dbaascli Utility](#).
- For Exadata DB systems, apply the OS patches following the instructions in [Updating an Exadata DB System](#).

## Oracle Cloud Security Response to Intel Microarchitectural Data Sampling (MDS) Vulnerabilities

Intel disclosed four new speculative execution side-channel processor vulnerabilities affecting Intel processors. These vulnerabilities have received the following CVE identifiers:

- CVE-2019-11091: Microarchitectural Data Sampling Uncacheable Memory (MDSUM)
- CVE-2018-12126: Microarchitectural Store Buffer Data Sampling (MSBDS)
- CVE-2018-12127: Microarchitectural Load Port Data Sampling (MLPDS)
- CVE-2018-12130: Microarchitectural Fill Buffer Data Sampling (MFBDS)

For more information, see <https://blogs.oracle.com/security/intelmds>.

### Oracle Cloud Infrastructure

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

However, if you manage your own operating systems (OS), you are advised to keep up with OS security patches to address this vulnerability.

The following sections contain the details of mitigations and actions.

### **Oracle Cloud Infrastructure Compute**

For details and required actions related to the Compute service's VM and bare metal instances, see [Oracle Cloud Infrastructure Customer Advisory for MDS Impact on the Compute Service](#).

### **Oracle Cloud Infrastructure Database**

If you use Autonomous Data Warehouse and Autonomous Transaction Processing, you have no further action to take.

For details and required actions related to offerings for VM DB systems, bare metal DB systems, and Exadata DB systems, see [Oracle Cloud Infrastructure Customer Advisory for MDS Impact on the Database Service](#).

### **Oracle Cloud Infrastructure Container Engine for Kubernetes**

To help secure your existing worker nodes for the Oracle Cloud Infrastructure Container Engine for Kubernetes Oracle recommends replacing your current node pools with new node pools. Please follow the instructions described in ['Upgrading' the version of Kubernetes running on worker nodes by creating a new node pool](#). All worker nodes created or upgraded after May 14th, 2019 are not impacted by this security issue.

### **Other Oracle Cloud Infrastructure Services**

Technical mitigations designed to protect all other Oracle Cloud Infrastructure services against the MDS processor vulnerabilities have been deployed. Oracle will notify customers if additional maintenance activities are required.

### **Oracle Cloud Infrastructure Classic and Oracle Platform Service on Oracle Cloud Infrastructure Classic**

For more information see [Oracle Cloud Infrastructure Classic](#).

In response to the MDS processor vulnerabilities, Oracle is performing mandatory maintenance for Infrastructure and Platform Services on Oracle Cloud Infrastructure Classic.

Platform Service hosts managed by Oracle are being patched by Oracle. If you manage your own operating systems, you are advised to keep up with the appropriate OS security patches to address these vulnerabilities.

### Oracle Cloud Infrastructure Customer Advisory for MDS Impact on the Compute Service

Intel disclosed four new speculative execution side-channel processor vulnerabilities affecting Intel processors. These vulnerabilities have received the following CVE identifiers:

- CVE-2019-11091: Microarchitectural Data Sampling Uncacheable Memory (MDSUM)
- CVE-2018-12126: Microarchitectural Store Buffer Data Sampling (MSBDS)
- CVE-2018-12127: Microarchitectural Load Port Data Sampling (MLPDS)
- CVE-2018-12130: Microarchitectural Fill Buffer Data Sampling (MFBDS)

For more information, see <https://blogs.oracle.com/security/intelmds>.

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

You are advised to keep up with OS security patches to address this vulnerability. See [Oracle Cloud Infrastructure Compute Content Impact](#) for instructions to patch the OS on the instances you manage.

#### **Additional Guidance for Oracle Cloud Infrastructure Bare Metal Instances**

Bare metal instances in Oracle Cloud Infrastructure offer customers full control of a physical server. Oracle Cloud Infrastructure's network virtualization is designed and configured to protect these instances from unauthorized access of other instances on the Oracle Cloud Infrastructure network, including other customer instances, both VM instances and other bare metal instances

However, for customers running their own virtualization stack on bare metal instances, the MDS vulnerabilities could allow a virtual machine to access privileged information from the underlying hypervisor or other VMs on the same bare metal instance. These customers should

review Intel's recommendations about these MDS vulnerabilities and make the recommended changes to their configurations, <https://www.intel.com/content/www/us/en/security-center/advisory/intel-sa-00233.html>.

### **Oracle Cloud Infrastructure Compute Content Impact**

Intel disclosed four new speculative execution side-channel processor vulnerabilities affecting Intel processors. These vulnerabilities have received the following CVE identifiers:

- CVE-2019-11091: Microarchitectural Data Sampling Uncacheable Memory (MDSUM)
- CVE-2018-12126: Microarchitectural Store Buffer Data Sampling (MSBDS)
- CVE-2018-12127: Microarchitectural Load Port Data Sampling (MLPDS)
- CVE-2018-12130: Microarchitectural Fill Buffer Data Sampling (MFBDS)

For more information, see <https://blogs.oracle.com/security/intelmds>.

### **RECOMMENDED ACTION**

Oracle recommends that customers patch the operating systems for their existing bare metal and virtual machine (VM) instances and verify that these OS updates include the patch for the MDS vulnerabilities. For VM instances, the Oracle Cloud Infrastructure team has implemented the necessary workarounds designed to mitigate for the MDS vulnerabilities. For bare metal instances using virtualization technology, you should also follow the following instructions

If you are running your own virtualization stack or hypervisors on bare metal instances, you should apply the appropriate patch required to address the MDS processor vulnerabilities.

The information in the following sections detail the commands needed to update your running instances created with [Oracle-Provided Images](#).

The following Oracle-provided image releases have been updated with the recommended patches, as a result instances created using these images or subsequent images include the recommended patches for the MDS vulnerabilities.

### Oracle-provided images updated with recommended patches for the MDS vulnerability

- [Oracle-Linux-6.10-2019.05.14-0](#)
- [Oracle-Linux-7.6-2019.05.14-0](#)
- [Oracle-Linux-7.6-Gen2-GPU-2019.05.14-0](#)
- [Windows-Server-2008-R2-Enterprise-Edition-VM-2019.05.14-0](#)
- [Windows-Server-2008-R2-Enterprise-Edition-VM-Gen2-2019.05.14-0](#)
- [Windows-Server-2012-R2-Standard-Edition-VM-2019.05.15-0](#)
- [Windows-Server-2012-R2-Standard-Edition-VM-Gen2-2019.05.14-0](#)
- [Windows-Server-2012-R2-Standard-Edition-VM-Gen2-E2-2019.05.15-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-Gen2-2019.05.14-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-Gen2-DenseIO-2019.05.15-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-Gen2-E2-2019.05.14-0](#)
- [Windows-Server-2012-R2-Datacenter-Edition-BM-2019.06.17-0](#)
- [Windows-Server-2016-Standard-Edition-VM-Gen2-2019.05.14-0](#)
- [Windows-Server-2016-Standard-Edition-VM-Gen2-E2-2019.05.14-0](#)
- [Windows-Server-2016-Datacenter-Edition-BM-Gen2-2019.05.14-0](#)
- [Windows-Server-2016-Datacenter-Edition-BM-Gen2-DenseIO-2019.05.14-0](#)
- [Windows-Server-2016-Datacenter-Edition-BM-Gen2-E2-2019.05.15-0](#)
- [CentOS-6.10-2019.05.15-0](#)
- [CentOS-7-2019.05.16-0](#)
- [Canonical-Ubuntu-14.04-2019.05.15-0](#) (deprecated)
- [Canonical-Ubuntu-16.04-2019.05.15-0](#)
- [Canonical-Ubuntu-16.04-Gen2-GPU-2019.05.15-0](#)
- [Canonical-Ubuntu-16.04-Minimal-2019.05.15-0](#)

## CHAPTER 29 Security Guide and Announcements

---

- [Canonical-Ubuntu-18.04-2019.05.15-0](#)
- [Canonical-Ubuntu-18.04-Minimal-2019.05.15-0](#)

Customers running instances created from imported third-party images should refer to the operating system (OS) vendor's guidance to patch the OS for the MDS vulnerability.

### PATCHING ORACLE LINUX INSTANCES

Oracle has released security patches for Oracle Linux 6, Oracle Linux 7, and Oracle VM Server for X86 products. In addition to the OS patches, customers should run the latest version of the microcode from Intel to mitigate these issues. For both bare metal and VM instances, please install the latest Ksplice via [uptrack-upgrade](#).



#### Note

See [Installing Ksplice Uptrack Within the Oracle Cloud Infrastructure](#) for how to install Ksplice.

For Oracle Linux, the patches for the MDS vulnerabilities are addressed by the same set of patches. For further information please see the following:

- <https://linux.oracle.com/cve/CVE-2018-12126.html>
- <https://linux.oracle.com/cve/CVE-2018-12130.html>
- <https://linux.oracle.com/cve/CVE-2018-12127.html>
- <https://linux.oracle.com/cve/CVE-2019-11091.html>

Bare metal instances must have the latest microcode updates from Intel. This step is not required for VM instances.

To install the latest microcode updates on bare metal instances, run the following command:

```
sudo yum update microcode_ctl
```

The required versions of microcode\_ctl rpms are:

- **Oracle Linux 7:** microcode\_ctl 2.1-47.0.4
- **Oracle Linux 6:** microcode\_ctl 1.17-1002

No additional update is required. In addition to the microcode update, you should also patch your bare metal instances using the following set of instructions.

### To patch the OS for bare metal and VM instances with downtime

The `yum-plugin-security` package allows you to use yum to obtain a list of all errata that are available for your system, including security updates. You can also use Oracle Enterprise Manager 12c Cloud Control or management tools such as Katello, Pulp, Red Hat Satellite, Spacewalk, and SUSE Manager to extract and display information about errata.

1. To install the `yum-plugin-security` package, run the following command:

```
sudo yum install yum-plugin-security
```

2. Use the `--cve` option to display the errata that correspond to a specified CVE, and to install those required packages, by running the following commands:

```
sudo yum updateinfo list --cve CVE-####-####
sudo yum update --cve CVE-####-####
```

Replace `####-####` in the above commands with the relevant CVE numbers.

3. A system reboot will be required once the package is applied. By default, the boot manager will automatically enable the most recent kernel version. For more information on using yum update, visit [Installing and Using the Yum Security Plugin](#).
4. After the system reboots, ensure that the following file is populated:

```
cat /sys/devices/system/cpu/vulnerabilities/mds
```

### PATCHING WINDOWS INSTANCES

#### *PROTECTING NEW WINDOWS VM AND BARE METAL INSTANCES*

When you create a new VM or bare metal instance based on the latest Oracle-provided Windows images, the image includes the Microsoft-recommended patches to protect against the MDS vulnerability. Windows bare metal instances also include the latest microcode

updates from Intel. To apply the MDS patch install the latest Windows Updates and reboot the instance. You should ensure that you keep your instances updated with the latest patches as recommended by your OS vendor.

### *PROTECTING EXISTING WINDOWS VM AND BARE METAL INSTANCES*

#### To update the microcode for existing bare metal instances

Bare metal instances launched before the Oracle-provided Windows images were updated must have the latest microcode updates from Intel. You need to recycle your Windows bare metal instances in order to receive the latest Intel microcode update. This step is not required for VM instances.

1. Create a new custom image of your Windows bare metal instance, see [Creating Windows Custom Images](#) for more information.
2. Terminate your existing Windows bare metal instance.
3. Open the navigation menu. Under **Core Infrastructure**, go to **Compute** and click **Custom Images**. Find the custom image you want to use.
4. Click the Actions icon (three dots), and then click **Create Instance**.
5. Provide additional launch options as described in [Creating an Instance](#).

Once you have completed these steps, perform the steps in the next procedure to update the instance with the latest OS updates from Microsoft

#### To patch the OS for bare metal and VM instances with downtime

Windows images include the Windows Update utility, which you can run to get the latest Windows updates from Microsoft. You have to configure the security list on the subnet on which the instance is running to allow instances to access Windows update servers. See Windows OS Updates for Windows Images and Security Lists for more information.

1. Verify that you have installed the latest Windows OS security update from Microsoft.
  - a. If automatic updates are turned on, the updates should be automatically delivered to the instance.
  - b. To manually check for the latest update, select **Start**.
  - c. In **Settings** select **Updates & security** and then select **Windows Update**.
  - d. In **Windows Update**, click **Check for updates**.
  - e. When you turn on automatic updates, this update will be downloaded and installed automatically. For more information about how to turn on automatic updates, see [Windows Update: FAQ](#).

For additional details see [Windows Server guidance to protect against speculative execution side-channel vulnerabilities](#).

### PATCHING UBUNTU OR CENTOS INSTANCES

The recommended patches to protect against the MDS vulnerabilities are included when you create a new VM or bare metal instance based on the latest Oracle-provided Ubuntu or CentOS images, see [Microarchitectural Data Sampling \(MDS\)](#) and [MDS - Microarchitectural Store Buffer Data - CVE-2018-12130, CVE-2018-12126, CVE-2018-12127, and CVE-2019-11091](#). For existing VM or bare metal instances you should follow the patching guidance provided by the original OS vendor.



#### Note

Any images published after May 14th, 2019 listed in the [image release notes](#) will include the MDS patches. If using earlier images already launched, follow patching instructions.

### Oracle Cloud Infrastructure Customer Advisory for MDS Impact on the Database Service

Intel disclosed 4 new speculative execution side-channel processor vulnerabilities affecting Intel processors. These vulnerabilities have received the following CVE identifiers:

- CVE-2019-11091: Microarchitectural Data Sampling Uncacheable Memory (MDSUM)
- CVE-2018-12126: Microarchitectural Store Buffer Data Sampling (MSBDS)
- CVE-2018-12127: Microarchitectural Load Port Data Sampling (MLPDS)
- CVE-2018-12130: Microarchitectural Fill Buffer Data Sampling (MFBDS)

For more information, see <https://blogs.oracle.com/security/intelmds>.

Oracle has deployed technical mitigations across Oracle Cloud Infrastructure systems designed to prevent a malicious attacker's virtual machine (VM) instance from accessing data from other VM instances.

#### **Autonomous Data Warehouse and Autonomous Transaction Processing**

Autonomous Data Warehouse provides fully managed databases optimized for running data warehouse workloads.

Autonomous Transaction Processing provides fully managed databases optimized for running online transaction processing and mixed database workloads.

Autonomous Data Warehouse and Autonomous Transaction Processing are not affected by MDS vulnerabilities. These services do not run on their own hypervisor and they do not allow for the execution of untrusted code in their services enclave. Customers can execute code within their own instances and each customer instance is isolated from that of another customer. No further customer action is currently required.

#### **Guidance for the DatabaseService on Bare Metal Instances**

The Database service on Oracle Cloud Infrastructure bare metal instances offer customers full control over their Oracle Database running on a physical server. Oracle Cloud Infrastructure's network virtualization is designed and configured to protect these instances from

unauthorized access from other instances on the Oracle Cloud Infrastructure network, including other customer instances, both VM instances and other bare metal instances. As a result, the Database service on bare metal instances are not affected by the MDS vulnerabilities.

### **Actions for Customers with VM DB Systems, Bare Metal DB Systems, or Exadata DB Systems**

Customers are advised to apply available patches at the earliest possible time. Use the following instructions to patch a running instance:

- For DB systems on bare metal instances, apply the OS patches following the instructions in [Updating a DB System](#).
- For DB systems on a VM instance, configured using the Oracle Cloud Infrastructure Database service, apply the OS patches following the instructions in [Updating a DB System](#).
- For the DB systems on a VM instance configured using the Oracle Platform Service Manager, apply the OS patches following the instructions in [Applying Linux OS Security Patches by Using the dbascli Utility](#).
- For Exadata DB systems, apply the OS patches following the instructions in [Updating an Exadata DB System](#).

# CHAPTER 30 Storage Gateway

This chapter explains how to use Storage Gateway to connect your on-premise applications with Oracle Cloud Infrastructure.

## Overview of Storage Gateway

Storage Gateway is a cloud storage gateway that lets you connect your on-premises applications with Oracle Cloud Infrastructure. Applications that can write data to an NFS target can also write data to the Oracle Cloud Infrastructure Object Storage, without requiring application modification to uptake the REST APIs.



### **Important**

Storage Gateway is the evolution of the Storage Software Appliance that was launched with Oracle Cloud Infrastructure Classic. Now that you're migrating to Oracle Cloud Infrastructure Object Storage, you'll use Storage Gateway, with its enhanced file-to-object transparency and improved scale and performance.

## Storage Gateway and Oracle Cloud Infrastructure Concepts

The following concepts are essential to working with Oracle Cloud Infrastructure Storage Gateway.

### **FILE SYSTEM**

A Storage Gateway file system on a local host maps its files and directories to objects with the same names in a corresponding Oracle Cloud Infrastructure Object Storage bucket.

### **FILE SYSTEM CACHE**

Storage Gateway's configurable file system cache enables asynchronous and optimized movement of data to the cloud. The file system cache serves as both a write buffer and a read cache for data storage and retrieval. The write buffer contains data that copied to the disk cache and queued for upload to Oracle Cloud Infrastructure. The read cache contains frequently retrieved data that's accessible locally for read operations.

Proper file system cache configuration is critical to Storage Gateway performance. See [Configuring the Cache for File Systems](#) for details.

### **METADATA**

The metadata associated with a Storage Gateway file is stored as custom metadata for the corresponding object in Oracle Cloud Infrastructure Object Storage. Examples of file metadata include object id, creation date, modification date, size, and permissions. Storage Gateway caches all metadata for the file system locally.

### **NFSv4**

NFS is an established and widely adopted distributed file system protocol for handling network storage. NFS lets client computers mount file systems on remote servers and access those remote file systems over the network as though they were local file systems. Storage Gateway performs the NFS to REST API translation needed to interact with Oracle Cloud Infrastructure Object Storage.

### **ORACLE CLOUD INFRASTRUCTURE**

Oracle Cloud Infrastructure is a set of complementary cloud services that lets you build and run a wide range of applications and services in a highly available hosted environment. Oracle Cloud Infrastructure offers high-performance compute capabilities (as physical hardware instances) and storage capacity in a flexible overlay virtual network that is securely accessible from your on-premises network.

### **TENANCY**

A tenancy is a secure and isolated partition within Oracle Cloud Infrastructure where you can create, organize, and administer your cloud resources.

### ORACLE CLOUD INFRASTRUCTURE OBJECT STORAGE AND ARCHIVE STORAGE

Oracle Cloud Infrastructure offers two distinct storage tiers for you to store your unstructured data. Use the [Object Storage Standard](#) tier for data to which you need fast, immediate, and frequent access. Use the [Archive Storage](#) service's **Archive** tier for data that you access infrequently, but which must be preserved for long periods of time. Both storage tiers use the same manageable resources (for example, [objects](#) and [buckets](#)). The difference is that when you upload a file to Archive Storage, the object is immediately archived. Before you can access an archived object, you must first restore the object to the Standard tier.



#### Note

In the Storage Gateway documentation, generic references to Object Storage encompass both the Standard and Archive storage tiers.

Both storage tiers are simple to use, perform well, and scale to an unlimited capacity.

### BUCKET

An Object Storage bucket is a logical container for storing objects. A file system created in Storage Gateway maps to a corresponding bucket by the same name in Object Storage. A bucket is associated with a single Oracle Cloud Infrastructure compartment. The compartment has policies that determine what actions a user can perform on the bucket the objects it contains.

### OBJECT

An individual file or directory written to a Storage Gateway file system on an NFS share creates an identically named object in the target Object Storage bucket. An object is composed of the object itself and metadata about the object.

### NAMESPACE

A logical entity that serves as a top-level container for all Oracle Cloud Infrastructure Object Storage buckets and objects. The namespace enables you to control bucket naming within your tenancy. Each tenancy has one unique and uneditable Object Storage [namespace](#) that is global, spanning all compartments and regions. Bucket names must be unique within your tenancy.

### COMPARTMENT

A collection of Oracle Cloud Infrastructure-related resources. Only users and groups with access permissions explicitly granted by an administrator can access the resources. Compartments help you organize resources and make it easier to control access to those resources. Object Storage automatically creates a root compartment when a compartment is provisioned. An administrator can then create more compartments within the root compartment and add access rules for those compartments. A bucket can exist in only one compartment.

## How Storage Gateway Works

Storage Gateway is installed in an Oracle Cloud Infrastructure compute instance or as a Linux Docker instance on one or more hosts in your on-premises data center. Applications store and retrieve objects from Oracle Cloud Infrastructure Object Storage through *file systems* that you create in Storage Gateway.

Storage Gateway exposes an NFS mount point that can be mounted to any host that supports an NFSv4 client. The Storage Gateway mount point maps to an Object Storage bucket.

There is file to object transparency between Storage Gateway and Object Storage:

- A Storage Gateway file system directory on a local host maps to a bucket with an identical name in Oracle Cloud Infrastructure Object Storage.
- Any file written to a Storage Gateway file system is written as an object with the same name in the associated Object Storage bucket. The system stores associated file attributes as object metadata.

- You can access Object Storage objects directly using native APIs, SDKs, third-party tools, the HDFS connector, and the Oracle Cloud Infrastructure CLI and Console. You use the **Refresh** operation in Storage Gateway to ingest any data that was added or modified directly in Object Storage.

Enterprise applications typically work with files in nested directories. Object Storage buckets, and the objects within those buckets, exist in a flat hierarchy. Storage Gateway flattens the directory hierarchy into nested object prefixes in Object Storage. See [Interacting With Object Storage](#) for details.

### Recommended Uses and Workloads

The following summarizes some of the ways that you can use Storage Gateway.

#### **DATA TRANSFER**

Use Storage Gateway to move data from your on-premises host to Oracle Cloud Infrastructure. Storage Gateway is not a replacement for general-purpose network attached storage (NAS), though it behaves similarly to NAS. Use Storage Gateway's integrated Cloud Sync feature to transfer and synchronize data to Oracle Cloud Infrastructure.

#### **CLOUD TIERING**

Use Storage Gateway to expand the capacity of on-premises storage solutions without capital expenditures. Configuring and connecting a Storage Gateway file system with a large cache to Oracle Cloud Infrastructure Object Storage provides unlimited scale to create a workflow in which files get automatically moved to the cloud and retrieved only on demand. Even though on-demand retrieval is slower than access to local storage, capital expenditures or changes to existing tools and software are not required.

#### **BACKUPS**

Use Storage Gateway to move files to Oracle Cloud Infrastructure Archive Storage as a cost-effective backup solution. You can move individual files and compressed or uncompressed ZIP or TAR archives. Storing secondary copies of data is an ideal use case for Storage Gateway.

### **ARCHIVAL**

Storage Gateway is ideal for archive use cases.

### **DISASTER RECOVERY**

Storage Gateway lets traditional applications move data to highly durable object storage. When you need to recover data, create a fresh instance of Storage Gateway to return the data from Object Storage.

## Uses and Workloads Not Supported

Storage Gateway does not support the following uses and workloads.

### **GENERAL-PURPOSE NETWORK STORAGE**

Storage Gateway isn't a general-purpose storage filer and must not be used as a replacement for traditional network storage appliances.

### **FILE SYNC AND SHARE**

Though Storage Gateway is an effective data mover, it's not a replacement for file sync and share services. Evaluate Oracle services like Oracle Document Cloud service if you need file sync and share functionality.

### **CONTENT COLLABORATION**

Storage Gateway does not support multiple Storage Gateway instances simultaneously reading from and writing to a single Object Storage bucket. Do not use Storage Gateway as a tool for distributed teams to collaborate on creating and managing content.

### **FREQUENTLY MODIFIED FILES**

Do not use Storage Gateway if you expect your data to be modified frequently. Each time a file is modified and closed, Storage Gateway creates an updated version and uploads it to Object Storage as a new object. Frequently modified data results in substantial inefficiency, in terms of both bandwidth consumption and capacity utilization.

### RENAMING LARGE DIRECTORY TREES

Renaming directories in the Storage Gateway works well for a small directory tree. However, renaming a parent directory with many children can be slow. The service updates the object ID of every corresponding child object in the object store to reflect the new path. If you do start a rename, ensure that the action finishes by monitoring the **Pending Uploads** field in the Storage Gateway user interface.

## Security Considerations

### ADMIN PASSWORD

Because Storage Gateway administrators can create, modify, and delete file systems, follow these password guidelines:

- Set a strong password.
- Ensure that the password is secure.
- Share passwords with others only on a need-to-know basis.

### DOCKER

Storage Gateway runs inside a Docker container for security and isolation. Follow these Docker-related guidelines and recommendations:

- Avoid or minimize Docker instance operations.
- Avoid logging in to the Docker container. If there is a genuine requirement to log in to the Docker container, use extreme caution to avoid service disruption. Do not change the Docker configuration or the Docker instance unless instructed to do so by Oracle support personal.
- Although the NFS protocol controls access to the file system from clients, Storage Gateway file systems are also locally mounted inside the Docker container. To prevent unauthorized access to file system data, ensure that a Docker container is accessible only by an administrator or an authorized user.

- Configure the Docker host to limit user access to the Storage Gateway Docker container.
- Files and directories in a Docker container are also visible in the Docker host - typically file systems and directories that are provisioned in the Docker host and mapped to the container. Set the appropriate ownership and modes to ensure that only an administrator or an authorized user can access these folders. We recommend the following:
  - A dedicated Storage Gateway host.
  - Limit who can access the Storage Gateway host.
  - Set firewall rules to limit access to the Docker host and Docker container.
  - Implement backup and retention policies for the files associated with Storage Gateway.

### ACCESS CONTROL

Default file system export options are too permissive. Set more restrictive export options so that only trusted NFS clients can access the file system data and metadata. Modify the advanced file system settings for **NFS Allowed Hosts** and **NFS Export Options** to restrict access to a file system. In addition to NFS protocol security, you can also set up and configure a firewall on the host to further control access to the file system. UID/GID/modes control access to files and directories. Set the appropriate ownership mode to protect sensitive data.

### OBJECT STORAGE

Files in a file system are uploaded to Oracle Cloud Infrastructure and stored as objects in an Object Storage bucket. Associated file attributes are stored as object metadata. Access control for Object Storage is different from access control for a traditional file system. Anyone with permission to read or modify any object in the bucket can read or modify all objects in the bucket. To protect sensitive data, set up Oracle Cloud Infrastructure IAM policies to limit who can access objects in the bucket.

Storage Gateway transfers data to Oracle Cloud Infrastructure using HTTPS, which encrypts data packets in flight between Storage Gateway and the cloud. Data written to Object Storage is always automatically encrypted in the cloud.

### Limits on Storage Gateway Resources

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

Other limits include:

- Ensure that the number of file systems per Storage Gateway doesn't exceed 10. For best performance, host each file system on a dedicated Storage Gateway.
- Ensure that the number of objects stored in a Storage Gateway file system doesn't exceed 100 million. For datasets that consist of more than 100 million objects, distribute the objects across multiple Storage Gateways.
- Ensure that you configure adequate local storage for file system cache. Storage Gateway warns you if you have configured less than the recommended 500 GB.
- The minimum amount of memory required for any Storage Gateway file system is 16 GB.
  - File systems with up to 50 million files require 32 GB of memory.
  - Large file systems with up to 100 million files require 64 GB of memory.
- The number of files in the cache is limited to 20,000 regardless of the specified cache size in bytes.
- To improve the efficiency of file ingestion, cloud upload operations, and to reduce the number of objects in the namespace, bin-pack or zip small files before writing them to Storage Gateway.

### Storage Gateway Release Notes

Release notes provide version-specific release information and important Storage Gateway issues that you need to be aware of:

<https://docs.cloud.oracle.com/iaas/releasenotes/services/storage-gateway/>

### Features of Storage Gateway

This topic highlights key features of Storage Gateway.

### POSIX-Compliant NFS Access to Oracle Cloud Infrastructure Object Storage

Using Storage Gateway, your applications can interact with Oracle Cloud Infrastructure Object Storage through standard NFSv4 protocols. You connect Storage Gateway file systems to Object Storage buckets. Storage Gateway stores files as objects in an Oracle Cloud Infrastructure Object Storage bucket and supports multipart uploads for large objects. Object Storage does not, however, support symbolic links, hard links, or special device files.

### Data Integrity with Checksum Verification

The built-in data integrity checks ensure that data is validated as it moves through the data path from Storage Gateway to Oracle Cloud Infrastructure Object Storage. Checksum verification helps ensure data integrity. Metadata integrity checks ensure that the metadata is in a consistent state. The checksum for each file can be read using a custom interface.

### Large File Support

The Oracle Cloud Infrastructure Object Storage service supports multipart uploads for faster, more efficient, and resilient uploads. Storage Gateway can use multipart upload for files larger than 128 MB. With multipart uploads, individual parts of an object can be uploaded in parallel to reduce the amount of time you spend uploading. Multipart uploads also minimize the impact of network failures by letting you retry a failed part upload instead of requiring you to retry an entire object upload. See [Using Multipart Uploads](#) for details.

Beginning with version 1.3, Storage Gateway provides partial update capabilities to:

- Reduce upload latency.
- Improve the use of available network bandwidth.
- Reduce the minimum required cache size.
- Enable ingestion of single files that are larger than the Storage Gateway cache size.

Storage Gateway's partial update capability leverages the Object Storage service multipart upload functionality to break a large file into smaller parts, and then upload the parts in parallel. After uploading, the service reconstructs them as a single object. With partial update, the service can upload only the modified file parts and reconstruct the object using the unchanged parts that already exist in Object Storage. The service does not need to overwrite the existing object by uploading the entire file.

Since the full file does not have to exist within the Storage Gateway cache at one time, partial update enables the service to ingest single files that are larger than the allocated cache size.

The default file part size is 1GB, with a maximum size of 10TB for the full file. Although the part size is configurable per file system, Oracle does not recommend that you use custom sizes. Custom part sizes do not improve upload performance, but can reduce the maximum supported full file size. The maximum file size is 10,000 times the file part size, with a hard cap of 10TB. If you customize the file part size, you do not need to restart the Storage Gateway, but you must reconnect to the affected file system.

All file parts in the cache are managed by the least recently used (LRU) cache management policy.

To use partial update for objects stored in the Archive Storage tier, you must first restore the object.

See [Upgrading Storage Gateway](#) for information about upgrading your Storage Gateway system to use partial update.

### Support for Data Archival

In addition to uploading to buckets in the Object Storage Standard tier, Storage Gateway supports uploading to and restoring objects from buckets in the Archive Storage tier.

When you create a file system, you specify the storage tier in which to create the corresponding Object Storage bucket.

- The default Standard Object Storage tier is used for storing data to which you need fast, immediate, and frequent access.

- The Archive Storage tier is used for storing data that is accessed infrequently and requires long retention periods.

While Archive Storage is more cost effective than Object Storage for preserving cold data, you must first *restore* the objects before you can access them. The restoration process can take up to four hours depending on the size of the object. See [Overview of Archive Storage](#) and [Restoring Files and Objects from Archive Storage](#) for details.

Storage Gateway supports Oracle Cloud Infrastructure Object Storage object lifecycle policies to manage the archiving and deletion of objects in a bucket according to a pre-defined schedule. Using object lifecycle policies, you can specify bucket creation in the Standard Object Storage tier, and then create a policy to schedule the subsequent movement of data to the Archive Storage tier. This lifecycle policy archival method is useful if you have on-premises applications that generate intermediary or temporary files and directories that are inappropriate for immediate archival. See [Using Object Lifecycle Management](#) for details.

### Automated Object Deletion

When you delete a Storage Gateway file from a file system, the corresponding object in Object Storage is automatically deleted.

### Quick Access to Select Files with Cache Pinning

Storage Gateway lets you *pin* files to the file system cache for quick access. You can pin files to the cache for file systems connected to either the Object Storage Standard or Archive tier.

When you write a file to your Storage Gateway file system, the file is initially stored in the file system cache, and then asynchronously uploaded to your Oracle Cloud Infrastructure bucket. After a file has been uploaded, the cache manager can remove the file from the file system cache. To meet the cache threshold specified for the file system, cache is reclaimed using the *Least Recently Used* (LRU) cache management policy. If you want specific files to be available in the cache for quick access, you can pin the files to the file system cache. Once pinned, files are not removed from the file system cache until you explicitly unpin them.

### Storage Gateway Health Check

The Storage Gateway performs automated "health checks" on the system to monitor the status of the following:

- Storage Gateway services and resources
- Local storage, file system cache, metadata storage, and log storage

### Integrated Cloud Transfer and Synchronization (Cloud Sync)

Storage Gateway provides an integrated cloud transfer and synchronization feature called Cloud Sync that lets you back up and transfer files on local storage to and from Oracle Cloud Infrastructure Object Storage buckets. This new feature replaces the independent, downloadable cloud sync utility that was available in the previous Storage Gateway version.

You can use the Storage Gateway management console or CLI to create, monitor, and manage Cloud Sync jobs similar to other enterprise NAS backup/replication offerings. Cloud Sync runs as part of the Storage Gateway software inside the Docker instance on the host.

### Getting Started With Storage Gateway

This topic provides recommendations for getting started with Storage Gateway.

### Recommended Reading

- If you have not done so already, read [Overview of Storage Gateway](#). That topic describes:
  - Key concepts for understanding both Storage Gateway and Object Storage.
  - Important security considerations.
  - Recommended uses and workloads.
  - Uses and workloads to avoid.

- Configuring the cache for file systems is key to Storage Gateway. Read [Configuring the Cache for File Systems](#) to understand the importance of the file system cache and the guidelines for configuring the cache when you add a file system.
- To understand the prerequisite tasks and requirements for interacting with Object Storage, read [Interacting With Object Storage](#) .

### Next Steps for Setting Up Storage Gateway

Key topics for setting up Storage Gateway include:

- [Installing Storage Gateway](#)
- [Logging In to the Storage Gateway Management Console](#)
- [Creating Your First File System](#)
- [Mounting File Systems on Clients](#)

### Next Steps for Using Storage Gateway

Key topics for using and managing Storage Gateway include:

- [Managing File Systems](#)
- [Managing Storage Gateway](#)
- [Monitoring Storage Gateway](#)

## Configuring the Cache for File Systems

Storage Gateway caches frequently retrieved data on the local host, minimizing the number of REST API calls to Oracle Cloud Infrastructure Object Storage and enabling faster data retrieval. You configure the cache for a file system when you create the file system. See [Creating Your First File System](#) and [Adding a File System](#).

### About File System Cache

The file system cache serves as both a read cache and a write buffer for data storage and retrieval. The read cache contains frequently retrieved data that's accessible locally for read operations. The write buffer contains data that has been copied to the disk cache and queued for upload to your Oracle Cloud Infrastructure tenancy.

When you retrieve data from Oracle Cloud Infrastructure, the data is stored in the Storage Gateway read cache. The read cache allows subsequent I/O operations to that data at local disk speed.

When the read cache is full or reaches the configured limit, Storage Gateway removes files from the cache based on a least recently used (LRU) algorithm. Files pending upload to your tenancy are not removed from cache. You can also preserve files that you do not want removed from cache.

For more information on how to preserve files in the read cache, see [Preserving Files in the File System Cache](#).

When an application transfers files through an NFS share, the files are written to the write buffer. The write buffer can contain many files that are queued and pending upload. If the host on which Storage Gateway is installed fails or Storage Gateway stops abruptly, the pending upload operations persist on the local disk. When Storage Gateway restarts, the pending upload operations resume and the data is uploaded to Oracle Cloud Infrastructure.

### Configuring Local Storage for File Systems and Cache

Storage Gateway uses local storage attached to the server (or virtual server) for hosting the file systems and cache. Files written to a file system in Storage Gateway are uploaded to the associated Object Storage bucket, with a portion of the file set maintained locally in the file system as a warm cache.

For optimal performance, reliability, and fault tolerance, follow these guidelines when configuring the local Storage Gateway storage:

- Allocate dedicated local storage for each Storage Gateway file system, and associated metadata and logs.
- Multiple disks (hard disk drives or solid-state drives) in a RAID10 set provide an optimal balance of performance, reliability, and fault tolerance. Alternatively, you can use RAID6.



### Important

Avoid RAID0 or single disk (no RAID) because of the potential for data loss upon disk failure.

- Provision sufficient space to accommodate the read cache **and** the write buffer (for ingesting new files) without ever becoming more than 80% full.  
In general, provision file system cache storage that is at least 1.5 times the size of the file set that you want to hold in the read cache. For example, assume that the entire file set requires 50 TB of space. You expect frequent access to 10% (5 TB) of that file set. Ensure that the file system cache storage has at least 7.5 TB of usable capacity. If the cache size reaches a near-full threshold, any data ingestion results in an out of space error in Storage Gateway.
- When you provision local storage at installation time, Oracle recommends that you configure the read cache to be equal to the file system cache size minus the desired write buffer size. If the file system cache is less than 300 GB, Storage Gateway generates a warning message.

## Determining File System Cache Size

The Storage Gateway file system cache serves as both a read cache and a write buffer. You can specify the maximum size of the read cache. The write buffer uses any remaining available space in the file system cache. You do not explicitly specify a size for the write buffer.



### Important

Oracle recommends that you configure the read cache to be equal to the file system cache size minus the desired write buffer size.

### Read Cache Size

The default maximum read cache size is 300 GB if the file system cache size is greater than 300 GB. Changing the default maximum read cache size is optional. The appropriate size depends on Storage Gateway workload. While the default setting for the read cache is appropriate for most workloads, consider increasing the size if Storage Gateway must retrieve a significant amount of data from the cloud.

Use the following guidelines to determine the appropriate setting for the read cache size:

- Do not set the read cache maximum equal to the size of the file system cache. Doing so allocates 100% of the space for the read cache and leaves no capacity for ingesting new files. If there is no available space for new file ingestion, Storage Gateway stops ingesting data and begins evicting files from the read cache to create space. Always preserve some space in the file system cache for ingestion.
- Set the read cache size to equal the amount of data that you anticipate to be accessed frequently, while leaving enough capacity for the write buffer.



### Note

If the total file system cache size is less than 300 GB, Storage Gateway automatically sets the read cache maximum size to 20% of the file system cache. The system does not honor custom read cache configuration settings for a file system cache less than 300 GB.

After you calculate the optimal file system cache size, you can configure the read cache when creating the file system or adjust it after monitoring the workload. See [Adding a File System](#) and [Changing the Properties of a File System](#).

### Write Buffer Size

Optimizing the space available for the write buffer is an important part of determining the appropriate file system cache size. The write buffer size increases when data is ingested in Storage Gateway and decreases after the data is uploaded to the cloud.



#### Important

- When the write buffer uses all available file system cache space, further data ingestion is blocked until a portion of the existing files are uploaded and evicted from the cache.
- Oracle recommends that you allow a minimum of 300 GB for the write buffer under any circumstances.

Use the following guidelines to determine the space needed for the write buffer:

- Identify the amount of data to be uploaded in Storage Gateway. If a large amount of data is uploaded, the Storage Gateway write buffer can reach its maximum size. Exceeding the write buffer leads to I/O failure as the file system cache has no space available. If you cannot regulate data ingestion, you can increase the file system cache space to avoid I/O failure. You can regulate I/O by pausing after a certain amount of data is ingested or by periodically allowing uploads to complete before ingesting more data. For example, you can use this approach for backups run as cron jobs when the file system cache space is less than the amount of data to be ingested.
- Calculate the amount of data that is ingested on any typical day or week in Storage Gateway. Also, calculate the amount of data that is uploaded over a time period, based

## CHAPTER 30 Storage Gateway

on the available bandwidth or historical data. Ensure that the difference between these calculations does not exceed the write buffer size.

- Some applications can handle I/O failure, and then resume writing data. In this case, consider setting the cache size to the amount of data that you'd like the application to tolerate before the cache space can be reclaimed.

As stated earlier, you want to avoid completely filling the file system cache. The write buffer grows by the difference between the ingest rate and the upload rate. The file system cache size must be larger than the *read cache plus the total number of bytes buffered* at any point during the job. If you have workloads that upload large amounts of data in parallel, you can use the following equation to determine the amount of space needed by the write buffer.

```
WB >= D * (1 - UR/IR) + E
```

**WB** = Recommended write buffer size

**D** = Total uploaded dataset size

**UR** = Upload rate (The upload rate is the lesser of the actual upload rate or the disk read speed.)

**IR** = Ingestion rate (disk write speed)

**E** = Extra margin (Oracle recommends at least 50 GB.)



### Note

This equation applies only if the upload rate is less than the ingest rate.

Run the following command to measure your system's ingestion rate:

```
sudo docker exec -it ocisg python /opt/oracle/gateway/python/packages/ocisg_helper/disk_speed_test.py
```

Example output:

```
Write speed:
2.1 GB copied in 9.4 seconds (228 MB/s)
```

## CHAPTER 30 Storage Gateway

---

```
Read speed:
2.1 GB copied in 6.8 seconds (315 MB/s)
```

If read and writes occur in parallel during the job, the read and write speeds are about 50% of the returned values.

Run the following command to measure your system's upload rate:

```
sudo docker exec ocisg cat /mnt/gateway/cache-phoenix/:::diag:oci-network-speed-test
```

Example output:

```
Average Upload Speed = 125 MB/s
```

### EXAMPLE 1

Daily dataset is 500 GB

Upload rate = 2 MB/s

Disk read = 600 MB/s

Ingestion rate (Disk write) = 600 MB/s

Apply the equation:

$$WB \geq 500 \text{ GB} * (1 - (2/600)) + 50 \text{ GB}$$
$$WB \geq 549 \text{ GB}$$

In this case, the upload rate is very slow compared to the ingestion rate, so the entire dataset needs to fit in the write buffer.

### EXAMPLE 2

Daily dataset is 1 TB

Upload rate = 300 MB/s

Disk read = 400 MB/s

Ingestion rate (Disk write) = 100 MB/s

In this case, we recommend the minimum write buffer size of 300 GB, since the upload rate is higher than the ingestion rate.

### EXAMPLE 3

Daily dataset is 5 TB

Upload rate = 250 MB/s

Disk read = 400 MB/s

Ingestion rate (Disk write) = 300 MB/s

Apply the equation:

$$WB \geq 5 \text{ TB} * (1 - (250/300)) + 50 \text{ GB}$$
$$WB \geq 0.88 \text{ TB}$$

WB is 880 GB.



#### Note

Storage Gateway begins throttling I/O when the free cache space falls below 15 GB. Determine if the application is able to handle the throttling or if you want to provision more cache space.

## Preserving Files in the File System Cache

When you write a file to your file system, the file is initially stored in the file system cache, and then uploaded to your Oracle Cloud Infrastructure tenancy. After a file has been uploaded, the cache manager can remove the file from the file system cache. To meet the cache threshold specified for the file system, cache is reclaimed using the *Least Recently Used* (LRU) cache management policy. If you want specific files to be available in the cache for quick access, you can *pin* the files to the file system cache. Once pinned, files are not removed from the file system cache until you explicitly unpin them. You can view the **Maximum Read Cache Size in GiB** for a selected file system in the management console under **Settings**.

You can pin files connected to both **Standard** and **Archive** storage tiers to file system cache. Files that you write to a file system are always uploaded to your tenancy, regardless of whether the files are pinned to the cache.

If the file that you want to pin to cache is not present in the cache, the file is automatically downloaded to the cache if the file system is connected to a **Standard** storage tier. If that file belongs to a file system connected to an **Archive** storage tier, you must first restore the file before the file can be downloaded to the cache. See [Restoring Files and Objects from Archive Storage](#) for details.



### Important

- By default, the cache pinning feature is enabled on all file systems.
- When selecting the files for cache pinning, consider the overall cache threshold and calculate the residual cache space that remains available for normal cache operations. For example, assume that your cache threshold is 1 TB and you estimate that the files you want to pin to cache occupy 300 GB. That leaves 700 GB of usable space in your cache after pinning the files.
- When you restore a file from the `Archive` storage tier, the file moves to the `Standard` storage tier. The file remains in `Standard` storage for 24 hours or the retention duration you specify. The continued availability of the file in the cache depends on the LRU operation. However, if you pin such a file to the cache, the restored file remains in the cache until you unpin the file.

### Enabling and Managing Cache Pinning

To perform cache pinning operations for a file system, run the following command from the NFS client on which the file system is mounted:

```
cat /path/to/mountpoint/<file_path>:::cache:cache_command[:argument]
```

The following table lists the cache pinning operations and the corresponding command and argument for each operation:

Operation	Cache Command	Argument
Enable cache pinning for a file system. By default, cache pinning is enabled for all file systems.	set-preserve-option	true
Get the cache pinning status for a file system.	get-preserve-option	No argument
Disable cache pinning for a file system.	set-preserve-option	false
List the files that are pinned to the cache.	list-preserve	No argument
Remove deleted files from the preserve list.	list-preserve-update	No argument
Add a file to the preserve list.	add-preserve	No argument
Remove a file from the preserve list.	remove-preserve	No argument
Clear the preserve list.	clear-preserve	No argument

### Example Commands

- To enable cache pinning for the `myFS` file system:

```
cat /mnt/gateway/myFS/:::cache:set-preserve-option:true
```

- To get the cache pinning status for `myFS`:

```
cat /mnt/gateway/myFS/:::cache:get-preserve-option
```

If cache pinning is enabled for the file system, the output of this command is `true`. Otherwise, the output is `false`.

- To disable cache pinning for the `myFS` file system:

```
cat /mnt/gateway/myFS/:::cache:set-preserve-option:false
```

- To add a file `myFile` of the `myFS` file system to the preserve list:

```
cat /mnt/gateway/myFS/myFile:::cache:add-preserve
```

- To find out which files are added to the preserve list of the `myFS` file system:

```
cat /mnt/gateway/myFS/:::cache:list-preserve
```

Sample output of the preceding command:

```
["/doNotDelete.txt", "/myFileMetadata", "/myFile"]
```

- To remove the file `myFile` from the preserve list

```
cat /mnt/gateway/myFS/myFile:::cache:remove-preserve
```

- To update the preserve list when the output of the `cache:list-preserve` command indicates that a pinned file has been removed from the file system:

```
cat /mnt/gateway/myFS/:::cache:list-preserve-update
```

Sample of the original preserve list:

```
["/doNotDelete.txt", "/myFileMetadata"]
```

Output of the `cache:list-preserve` command after the file `myFileMetadata` is removed from the cache:

```
["/doNotDelete.txt", "Status: 1 files appear to no longer exist. Please run list-preserve-update"]
```

Output of the `cache:list-preserve-update` command:

```
["/doNotDelete.txt"]
```

- To clear the preserve list for a file system:

```
cat /mnt/gateway/myFS/:::cache:clear-preserve
```

## Understanding Storage Gateway Performance

This topic covers the performance characteristics of Storage Gateway and the ways you can maximize its efficiency.

### Performance Characteristics

It is important to understand the basic performance characteristics of Storage Gateway:

- Because there is transactional overhead for each file, Storage Gateway generally exhibits better performance with large files than with small files. Storage Gateway can only upload data as fast as your connection and the storage host allows. Storage Gateway buffers the data in local disk storage while waiting to upload to Oracle Cloud Infrastructure. Once a file is uploaded, the file's local copy can be removed from the cache to free up space. If file system cache space falls below 10 GB, application I/O and file/directory creation can fail.
- Modifying a file involves uploading a whole new copy of that file. This behavior is not efficient for frequently modified large files.
- Storage Gateway does not support frequently modified files such as logs, databases, or virtual disks. Storage Gateway depends on the closing of a file to trigger an upload of that file. If a file never closes or is modified frequently, the upload event cannot successfully occur.
- Upload throughput to Object Storage (WAN) is typically slower than NFS client throughput (LAN). As a result, Storage Gateway can accumulate a large amount of data that is pending upload.
- You can attach Storage Gateway to an existing Object Storage bucket that contains data. The service is optimized for this type of initialization. Storage Gateway can

initialize about 700 thousand files per hour with a hard disk drive-based cache. It can initialize about 7 million files per hour with an NVMe SSD-based cache.

### Factors That Affect Performance

To get the maximum performance benefits from Storage Gateway, follow the practices documented at [Best Practices for Using Storage Gateway](#).

In addition to having sufficient memory and file system cache space, Oracle recommends that you use SSDs to help improve NFS ingestion rate.

Storage Gateway is tuned for maximum upload and download performance by default. No additional tuning is needed.

### Performance Testing

While measuring performance is complex and open to variability, we have observed the following in performance benchmark tests with 10-Gb/s link speed:

Workload (Upload/Download)	Configuration	Average Upload Throughput	Average Download Throughput
Single large file of 400 GB	CPU: 8 cores Memory: 32GB Disk Read: 340 MB/s Disk Write: 234 MB/s Link Speed: 10 Gb/s (1.25 GB/s)	239 MB/s	195 MB/s
Multiple files of 10-50 GB (Total data = 400GB)	CPU: 8 cores Memory: 32GB Disk Read: 340 MB/s Disk Write: 234 MB/s Link Speed: 10 Gb/s (1.25 GB/s)	260 MB/s	225 MB/s



### Note

Using FastConnect with Storage Gateway makes optimal use of full link speed. We have customers using FastConnect with 10-Gb/s link speed and have seen each gateway achieve 400–450 MB/s uploads to Oracle Cloud Infrastructure.

See [FastConnect Overview](#) for more information.

## Testing Network Bandwidth

Storage Gateway provides a diagnostic command that you can use to test the bandwidth in your environment and ensure that you get the expected upload and download speeds. The amount of data transferred depends on these factors:

Bandwidth \* Delay Product (bits) = total\_available\_bandwidth (bits/sec) x round\_trip\_time (sec)



### Note

Different buckets can have different upload and download speeds.

The round\_trip\_time can vary by region.

### To run the diagnostic command:

You need `root` permissions to run the `diag` command.

1. Using SSH, log in to the host on which you installed Storage Gateway.
2. Run the `diag` command, specifying the Storage Gateway file system name:

```
[root@ocisg-ashburn opc]# sudo docker exec ocisg cat /mnt/gateway/<file_system_name>/:::diag:oci-network-speed-test
```

The `diag` command responds with the average upload speed, for example:

```
Average Upload Speed = 217 MB/s
```

## Interacting With Object Storage

This topic helps you understand the Oracle Cloud Infrastructure Object Storage environment and how it interacts with a Storage Gateway.

### Creating the Required IAM Users, Groups, and Policies

An Oracle Cloud Infrastructure administrator must perform prerequisite tasks in preparation for data movement between Storage Gateway and Object Storage. If you are new to Oracle Cloud Infrastructure, we recommend that you read [Setting Up Your Tenancy](#).

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

Access to resources is provided to groups using policies and then inherited by the users that are assigned to those groups. For details on creating groups, see [Managing Groups](#).

For Storage Gateway, an administrator creates these groups with the following policies:

```
Allow group <group_name> to manage buckets in compartment <compartment_name>
```

```
Allow group <group_name> to manage objects in compartment <compartment_name>
```

### Content Consistency Between Storage Gateway and Object Storage

Changes to the files in Storage Gateway, including create, write, update, and delete, eventually are consistent with Object Storage. Uploads are asynchronous and buffered for performance, so Storage Gateway file changes might not yet be reflected in Object Storage.

You can access, modify, and upload objects directly to a bucket using Object Storage native APIs, SDKs, the CLI, the Console, or the HDFS connector. Objects modified in these ways do

not appear as files in Storage Gateway until you click **Refresh** in the Storage Gateway management console.

### Name Restrictions

Storage Gateway file and file system names must adhere to Object Storage bucket and object name restrictions and guidelines.

Use the following guidelines for naming file systems:

- Use from 1 to 256 UTF-8 characters.
- Valid characters are letters (upper or lower case), numbers, hyphens, underscores, and periods.



#### **Important**

Names cannot contain a slash (/) character because this character delimits Object Storage bucket and object names.

- Do not include confidential information.
- Make the name unique within a Storage Gateway instance.

Use the following guidelines for naming files:

- Use from 1 to 1024 characters.
- Valid characters are letters (upper or lower case), numbers, and characters other than linefeed, newline, and NULL.
- Use only Unicode characters for which the UTF-8 encoding does not exceed 1024 bytes. Clients are responsible for URL-encoding characters.
- Do not include confidential information.

- Make the name unique within the bucket. Do not use the name of an existing object within the bucket when naming an object unless you intend to overwrite the existing object with the contents of the new or renamed object.

### Custom Metadata

POSIX file and directory attributes are stored in custom metadata. These attributes include `uid`, `gid`, `mode`, `atime`, `ctime`, and `mtime`. If existing objects in Object Storage are missing the required custom metadata, Storage Gateway assigns the following default values:

- `uid=0`
- `gid=0`
- `mode=0644` for file and `0755` for directory

The custom metadata is not updated in Object Storage until a file operation triggers Storage Gateway to update the file in Object Storage. Timestamp metadata (`atime`, `ctime`, and `mtime`) are expressed in milliseconds. Access modes are expressed in octal and include file/directory bit.

The custom metadata names follow these guidelines:

- Only ASCII characters.
- A maximum of 128 bytes.

The custom metadata values follow these guidelines:

- Only UTF-8 characters.
- A maximum of 256 bytes.

### Understanding Directory and File Hierarchy Translations in Object Storage

Within an Object Storage namespace, buckets and objects exist in a flat hierarchy. Storage Gateway flattens the file system directory hierarchy into nested object prefixes in Object Storage.

## CHAPTER 30 Storage Gateway

---

For directories:

- A Storage Gateway file system called `myFS` that contains a directory called `myDir`, appears in Object Storage as:

```
n/<os_namespace>/b/myFS/o/myDir/
```

- A Storage Gateway file system called `myFS` that contains a `myDir` subdirectory called `mySubDir`, appears in Object Storage as:

```
n/<os_namespace>/b/myFS/o/myDir/mySubDir/
```

You can distinguish a Storage Gateway directory from a Storage Gateway file in the following ways:

- Directories have a trailing slash `/`.
- Directory size or length is **0** (zero).

For files:

- A Storage Gateway file system called `myFS` that contains a directory called `myDir` with a file called `file1`, appears in Object Storage as:

```
n/<os_namespace>/b/myFS/o/myDir/file1
```

- A Storage Gateway file system called `myFS` that contains a `myDir` subdirectory called `mySubDir` with a file called `file2`, appears in Object Storage as:

```
n/<os_namespace>/b/myFS/o/myDir/mySubDir/file2
```

You can distinguish a Storage Gateway file from a Storage Gateway directory in the following ways:

- Directories have a trailing `/` and files do not.
- File length can be **0** (zero) or non-zero, but directory length is always **0** (zero).

## Internal Storage Gateway Objects

Storage Gateway creates some special internal objects in Object Storage. These objects have a `/gateway` directory prefix. For example:

```
/n/<object_storage_namespace>/b/<bucket>/o//gateway
```



### Important

Do not modify or remove the objects in the special `/gateway` directory. These objects are critical for Storage Gateway operation.

## Installing Storage Gateway

This topic provides instructions for installing the Storage Gateway software.

### Prerequisites

These instructions assume that you are familiar with the administration and configuration commands of the operating system on your host machine. To install Storage Gateway, your host system must meet certain hardware and software requirements.

### Hardware Recommendations and Requirements

To run Storage Gateway, the host machine must meet the following requirements:

- Two dual-core CPUs or better. Oracle recommends 4-core CPUs.
- Minimum memory requirements:
  - 16 GB for required for any Storage Gateway file system.
  - 32 GB for file systems up to 50 million files.
  - 64 GB for file systems up to 100 million files.
- The recommended local storage disk size is 600 GB, which includes 500 GB for the file system cache, 80 GB for metadata storage, and 20 GB for log storage.



### Important

Provision local storage before installing Storage Gateway. For best performance, allocate dedicated local storage file systems for the Storage Gateway cache, the metadata, and the logs. The installation script prompts you for the paths to your Storage Gateway file system cache, metadata storage, and log storage locations. Follow the disk size recommendations provided by the installer.

Oracle recommends that you use the XFS file system for the file system cache, metadata, and logs. XFS is a 64-bit file system designed for parallel I/O. Parallel I/O allows a system to scale based on the number of I/O threads and file system bandwidth.

### Software Requirements

- Oracle Linux 7 with UEK Release 4 or later.



### Note

If you create an Oracle Cloud Infrastructure Compute instance to host Storage Gateway, the instance creation wizard provides an option to choose the operating system image.

- Docker 1.12.6 or newer. Docker is an open platform for building, shipping and running

distributed applications. For more information, see <https://www.docker.com/>.

- NFSv4.



### Note

The Storage Gateway installation software automatically installs Docker and the NFS protocol.

## Hosting Storage Gateway on an Oracle Cloud Infrastructure Compute Instance

To host Storage Gateway on an Oracle Cloud Infrastructure Compute instance, you need:

- An SSH key pair in PEM format.
  - To create a key pair, see [Creating a Key Pair](#).
  - If your public key is not in PEM format, use the following command to convert it:

```
ssh-keygen -f <key_name>.pub -e -m pem
```

- An Oracle Cloud Infrastructure user account with an API signing key (the public key from your SSH key pair).
  - If you need to create a user account, see [To create a user](#).
  - To upload an API signing key to an existing user account, see [To upload an API signing key](#).
- A VCN and related resources. For help creating a VCN, see [VCNs and Subnets](#).

The following configuration points apply to your VCN:

- Do not select the **Use DNS Hostnames in this VCN** check box unless you plan to use DNS hostnames for your Storage Gateway Compute instance.
- The security list must include a rule to allow SSL (443) ingress.
- After you install the Storage Gateway software your host machine, you must [add a security list rule](#) to allow communication with the management console port.

More information appears on this page after the Storage Gateway installation instructions.

- A Compute instance. See [Creating an Instance](#).  
The VM.Standard2.4 Compute shape meets the minimum required specifications for Storage Gateway. Large file systems might require an image with more resources.
- A Block Volume. See [Creating a Volume](#).
  - The recommended disk size is 600 GB.
  - Attach the volume to your Compute instance. See [Attaching a Volume](#).
  - If you specify [iSCSI](#) as the volume attachment type, you must also connect and mount the volume from the instance for the volume to be usable. For more information, see [Volume Attachment Types](#) and [Connecting to a Volume](#).

### Installing Storage Gateway

You can install Storage Gateway on an Oracle Cloud Infrastructure Compute instance or an on-premises host that meets the [hardware](#) and [software](#) requirements.

#### To install the Storage Gateway software

1. Connect to your Compute instance or on-premises host.  
For help connecting to an Oracle Cloud Infrastructure Compute instance, see [Connecting to an Instance](#).
2. If your host volume is new, you might need to format and mount the disk.



### Tip

This task describes the simplest way to create a functional file system to host a Storage Gateway. It uses one device and file system to host the cache, metadata, and log volumes. You specify the paths to those volumes later in this procedure. To optimize performance for your system, you can:

- Create a separate device and file system for each of the cache, metadata, and log volumes.
- Create a single device, but create logical volumes and file systems for the cache, metadata, and log volumes.

To format the disk and create a file system:

- a. Run *fdisk*:

```
sudo fdisk /dev/sdb
```

(Optional) Press **m** to view the *fdisk* options.

- b. Choose command **g - create a new empty GPT partition table**.
- c. Choose command **w - write table to disk and exit**.
- d. Run the following command to create an XFS (file system):

```
sudo mkfs -t ext3 /dev/sdb
```

You might see the following warning:

```
/dev/sdb is entire device, not just one partition!
Proceed anyway? (y,n)
```

Choose **y** to proceed.

To mount the formatted volume:

- a. Create a mount directory:

```
sudo mkdir /ocisg
```

- b. Mount the drive:

```
sudo mount /dev/sdb /ocisg
```

3. Download the [Storage Gateway 1.3 tar archive](#).
4. Use the SFTP tool of your choice to copy the tar archive into the `/tmp` folder of the host machine.
5. On the host machine, change directory to `/tmp` and extract the files from the tar archive:

```
cd /tmp
sudo tar xvzf ocisg-1.3.tar.gz
```

This command extracts the files from the tar archive into a subdirectory named `ocisg-1.3`.

6. Change directory to `ocisg-1.3` and run the installation script as `sudo` or `root` user:

```
cd ocisg-1.3
sudo ./ocisg-install.sh
```

### List of installation script options

Optionally, you can specify the following `ocisg-install.sh` script flags:

- **-a** Runs the installation in advanced configuration mode, which lets you specify ports and the Docker network mode.

In addition to prompting you for the paths to the metadata storage, cache storage, and log storage, advanced configuration mode also prompts you for:

- The Docker network mode (`host` or `bridge`).

Bridge mode is the default. It allows multiple instances of Storage Gateway to run on the same host.

Host mode improves network performance. If you plan to run only one instance of Storage Gateway on the host, Oracle recommends host mode. If you encounter issues with host mode, try bridge mode or contact [My Oracle Support](#).

- The host port to use for the management console.
- The host port to use for NFS access.
- The host port to use for the HTTP REST service.



### Tip

For each host port specification, you can designate a port or press **Enter** to let Storage Gateway dynamically allocate the port. You can use the `ocisg configure port` command to change the ports later. See [Managing Storage Gateway Using the CLI](#) for details.

- **-d** Installs Storage Gateway at the location you specify instead of the default location of `/opt/ocisg`. For example:

```
sudo ./ocisg-install.sh -d /opt/storagegateway
```

- **-h** Displays the installation script help information.
- **-p** Specifies that Storage Gateway is running behind a proxy server. You can specify multiple proxy arguments. For example:

```
./ocisg-install.sh -p http://myproxy.com:80 -p https://mysecureproxy.com:80
```

- **-q** Runs the installation in quiet mode.

If you supply the paths to the Storage Gateway cache, metadata, and log storage locations using **-m**<path\_to\_metadata\_storage>, **-c**<path\_to\_cache\_storage>, and **-l**<path\_to\_log\_storage>, you are not prompted for input. For example:

```
sudo ./ocisg-install.sh -q -m /ocisg/metadata -c /ocisg/cache -l /ocisg/log
```



### Note

Ignore the `devicemapper` warning message if it appears during the installation.

The script guides you through the Storage Gateway installation. Depending on your host machine configuration, some steps can require your input:

- a. Docker does not appear to be installed. Do you want to install docker engine with yum? [y/N]

Press **y**, and then press **Enter**.

The installation script automatically installs Docker and configures the storage driver for use with Storage Gateway.



### Important

If Docker is already installed on your system, the installation script does not automatically configure the storage driver and returns a warning message:

```
Checking that docker is installed and using the
correct version
Found docker version Docker version 18.03.1-ol,
build 0d51d18
The storage appliance requires to set devicemapper
as the docker storage driver.
Please follow the setup link below to enable
devicemapper and rerun the install.
```

Manually verify and update the Docker storage driver to be `devicemapper` as required. See [Verifying and Updating the Storage Driver in Docker](#).

- b. NFS server does not appear to be enabled. Do you want to enable NFS?  
[y/N]

Press **y**, and then press **Enter**.

- c. When prompted, press **Enter** accept the default installation location.
- d. When prompted, specify the paths to your Storage Gateway cache, metadata, and log storage locations.

The following examples represent paths for a simple system. Your setup might include paths to separate devices and file systems for each location.

- i. Enter a path for the file system cache. For example:

```
/ocisg/sg/cache
```

- ii. Enter the path for metadata storage. For example:

```
/ocisg/sg/metadata
```

- iii. Enter the path for log storage. For example:

```
/ocisg/sg/log
```

If you receive warnings about cache, metadata, and log storage existing on the same volume, enter **y** to proceed with the installation.

After a successful installation, the script provides the following information:

- The URL to log in to the Storage Gateway management console.
- The NFS port number.
- An example command for mounting your Storage Gateway file systems.

If you installed Storage Gateway on a Compute instance, see [Security List Requirements for Compute Instance Installations](#).

### Security List Requirements for Compute Instance Installations

If you installed Storage Gateway on an Oracle Cloud Infrastructure Compute instance, that instance must be able to receive HTTPS connections from other hosts and allow communication with the Storage Gateway management console. To open the necessary port, add an ingress rule to the security list governing the instance's host subnet. To learn about VCN security control, see [Security Lists](#).



### Important

This installation task assumes that your existing security list already allows traffic to port 443, as described in the [Hosting Storage Gateway on an Oracle Cloud Infrastructure Compute Instance](#) section of this page. If port 443 is not open, you must add a security list rule to open it.

### To add a security list rule

1. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **Virtual Cloud Networks**.
2. Click the name of the cloud network (VCN) that hosts your Compute instance.
3. Click **Security Lists**.
4. Click the name of the security list that governs the subnet hosting your Compute instance.
5. Click **Edit All Rules**.
6. Add an ingress rule:
  - a. Leave **Stateless** unmarked.
  - b. Select **CIDR** for the **Source Type**.
  - c. Enter **0.0.0.0/0** to indicate all IP addresses.
  - d. Select **TCP** as the **IP Protocol**.
  - e. Enter **All** in the **Source Port Range** field.
  - f. Specify your Storage Gateway management console port in the **Destination Port Range** field. For example:

32769

## CHAPTER 30 Storage Gateway

---

If you do not know the management console port for your Storage Gateway installation, run the following command on the host machine:

```
sudo ocisg info
```

The management console port appears at the end of the management console URL:

```
Management Console: https://exampleCompute:32769
```

### 7. Click **Save Security List Rules**.

You can now connect to the compute instance using the public IP address (https: <public\_IP\_address>). See [Getting the Instance Public IP Address](#) for details.

## Verifying and Updating the Storage Driver in Docker

To verify the storage driver in Docker:

### 1. Start docker:

```
sudo systemctl start docker
```

### 2. Verify the information in docker:

```
sudo docker info
```

### 3. Look for Storage Driver in the output. For example:

```
Containers: 0
 Running: 0
 Paused: 0
 Stopped: 0
Images: 0
Server Version: 18.03.1-ol
Storage Driver: overlay2
 Backing Filesystem: xfs
 Supports d_type: true
 Native Overlay Diff: false
Logging Driver: json-file
Cgroup Driver: cgroupfs
Plugins:
```

## CHAPTER 30 Storage Gateway

```
Volume: local
Network: bridge host macvlan null overlay
Log: awslogs fluentd gcplogs gelf journald json-file logentries splunk syslog
Swarm: inactive
Runtimes: runc
Default Runtime: runc
Init Binary: docker-init
containerd version: 773c489c9c1b21a6d78b5c538cd395416ec50f88
runc version: 4fc53a81fb7c994640722ac585fa9ca548971871
init version: 949e6fa
Security Options:
 seccomp
 Profile: default
 selinux
Kernel Version: 4.1.12-124.15.4.el7uek.x86_64
Operating System: Oracle Linux Server 7.5
OSType: linux
Architecture: x86_64
CPUs: 4
Total Memory: 13.45GiB
Name: ocisg-mahesh
ID: OJ2H:QUSK:BWQZ:25L6:VI5V:CXGX:WFXT:NNNP:RK60:OS4P:4ABE:JWMV
Docker Root Dir: /var/lib/docker
Debug Mode (client): false
Debug Mode (server): false
Registry: https://index.docker.io/v1/
Labels:
Experimental: false
Insecure Registries:
 127.0.0.0/8
Live Restore Enabled: false

Registries: docker.io (secure)
```



### Note

You can ignore the `devicemapper` warning message, if it appears.

If Storage Driver is *not* `devicemapper`, do the following:

- a. Stop docker:

```
sudo systemctl stop docker
```

- b. Look for `/etc/docker/daemon.json` in the host.  
If the file `daemon.json` does not exist, create it.

- c. In the `daemon.json` file, set the `storage-driver` variable to `devicemapper`:

```
{
 "storage-driver": "devicemapper"
}
```

- d. Restart docker:

```
sudo systemctl start docker
```

- e. Verify the information in docker:

```
sudo docker info
```

Look for `Storage Driver` in the output and verify that the storage driver is `devicemapper`.

### Next Step

[Logging In to the Storage Gateway Management Console](#)

## Logging In to the Storage Gateway Management Console

Use the Storage Gateway management console to create, manage, and monitor file systems.

Storage Gateway provides the URL to access the management console after a successful installation. When you access the management console for the first time, a wizard prompts you to create the administrator credentials and your first file system.



### Note

Storage Gateway uses a self-signed certificate for the HTTPS connection. Your browser might warn that the SSL certificate couldn't be verified. If you entered the correct public IP address of the Storage Gateway instance, you can safely ignore this warning. The steps to ignore this warning and go to the management console vary depending on the browser you use.

### To log in to the management console:

1. Enter one of the following URLs in a supported web browser:
  - If you installed the software on an on-premises host, enter the URL provided at the end of the Storage Gateway installation script:

```
https://<storagegateway_hostname>:<port_number>
```

For example:

```
https://myStorageGatewayHost:3775
```



### Note

If you cannot access Storage Gateway using the hostname, contact your network administrator. Depending on your network configuration, your Storage Gateway hostname might need to be added to DNS or you might need to use an IP address rather than the hostname.

- If you installed the software in an Oracle Cloud Infrastructure compute instance, enter the URL as follows:

```
https://<instance_public_IP_address>:<port_number>
```

For example:

```
https://192.168.14.5:3775
```



### Note

See [Getting the Instance Public IP Address](#) for details.

The console log-in page appears. The page prompts you to set and confirm a password for the Storage Gateway `admin` user.

2. Enter a password that meets the following requirements:
  - From 8 to 32 characters.
  - At least one special character, one numerical character, one uppercase character, and one lowercase character.

## Next Step

[Creating Your First File System](#)

## Creating Your First File System

This topic guides you through creating your first Storage Gateway file system.

Think of a file system as a namespace containing a dataset that's accessible through Storage Gateway. A Storage Gateway file system in this context represents a mapping between a directory on your on-premises host and a bucket in Oracle Cloud Infrastructure Object Storage. When you create a Storage Gateway file system, you define the connection credentials that Storage Gateway uses to connect to your Oracle Cloud Infrastructure tenancy.

When you log in to the management console for the first time, a wizard prompts you to create the administrator credentials and your first file system.

### To create your first file system

1. Log in to the management console.
2. Click **File Systems** on the upper-right area of the management console.
3. Click **Create File System**.
4. Enter the required information in **Create a File System**:
  - **File System Name:** A unique, friendly name for the file system. Avoid entering confidential information. Use the following guidelines when naming a file system:
    - Use from 1 to 256 characters.
    - Valid characters are letters (upper or lower case), numbers, hyphens, underscores, and periods.



#### **Important**

The name **cannot** contain the following:

- A slash (/) character because this character delimits bucket and object names in Oracle Cloud Infrastructure Object Storage
- The strings "pub" or "priv"

If an Object Storage bucket by this file system name doesn't exist in your tenancy, the bucket is created.

If a corresponding Object Storage bucket by this file system name exists in your tenancy and there is data in the bucket, Storage Gateway works asynchronously to sync the bucket and file system contents.

- **Select the Object Storage tier in which you want to store your data**



### Important

Once set, you cannot change the storage tier in which a bucket resides.

You can use the Oracle Cloud Infrastructure Object Storage object lifecycle policies feature to manage the archiving and deletion of objects in a bucket according to a predefined schedule. See [Using Object Lifecycle Management](#) for details.

- **Standard:** The Standard tier is the primary default Object Storage tier for storing data that requires frequent and fast access. See [Overview of Object Storage](#) for more information.
- **Archive:** The Archive tier is a special tier for storing data that is accessed infrequently and requires long retention periods. See [Overview of Archive Storage](#) for more information. Access to data in the Archive tier is not immediate since you must restore archived data before it's accessible (see [Restoring Files and Objects from Archive Storage](#)).

- **Object Storage endpoint:** Required. The Object Storage API endpoint for your service instance. To find the object storage API endpoint for your Oracle Cloud Infrastructure Object Storage tenancy, see the [API documentation for Oracle Cloud Infrastructure Object Storage](#).



### Important

The following information is required to connect your Storage Gateway file systems to Oracle Cloud Infrastructure. See [Required Keys and OCIDs](#) for detailed information on how to generate the required keys and where to obtain these OCIDs.

- **Compartment OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage compartment.
- **Tenant OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage tenancy.
- **User OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage user.
- **Public Key's Finger Print:** Required. Your Oracle Cloud Infrastructure Object Storage public key fingerprint.
- **Private Key:** Required. Your Oracle Cloud Infrastructure Object Storage private key.

- **Private Key Passphrase:** Required if a passphrase was specified during key creation. Your Oracle Cloud Infrastructure Object Storage private key passphrase.



### Note

Your private key and passphrase are securely stored in the Storage Gateway docker. The Storage Gateway installation generates a pair of public and private keys. The system uses the public key to encrypt sensitive data.

5. Click **Save**.

The values you entered must match your Oracle Cloud Infrastructure credentials. If you get an error message, verify your entries, update any incorrect values, and click **Save** again.

6. Click **Show Advanced File System Configuration**.

Enter the required configuration information or click **Use Default** to accept the default values:

- **NFS Allowed Hosts:** A comma-separated list of hosts allowed to connect to the NFS export. You can also specify **\*** to allow all hosts to connect.

For example: `2001:db8:9:e54::/64, 192.0.2.0/24`

- **NFS Export Options:** The NFS export options.

Example: `rw, sync, insecure, no_subtree_check, no_root_squash`



### Important

Do not specify the `fsid` option.

- **Concurrent Uploads:** The number of concurrent uploads to Oracle Cloud Infrastructure.

This field indicates the maximum number of files that can be concurrently uploaded in Storage Gateway. If the value is 15, the concurrent file uploads can be between 1-15.

The allowed range is from 1 to 30.

- **Sync Policy:** The metadata operations are flushed to the disk based on the sync policy, but do not affect on-disk consistency. Currently, only **Posix Standard** is supported. Only the synchronous transactions (like `fsync`, `ODSYNC`, and `OSYNC`) are committed to the disk. All other transactions are handled asynchronously.

- **Cloud Read-ahead:** The number of blocks to be downloaded and used to *read ahead* when reading files for improved performance.

Default value: 50

- **Maximum Read Cache Size in GiB:** The maximum read cache.

When the read cache is full or reaches the configured limit, Storage Gateway removes files from the cache based on a least recently used (LRU) algorithm. Files pending upload to your tenancy are not removed from cache. You can also preserve files that you do not want removed from cache.



### Note

The number of files in cache is limited to 20,000, regardless of the specified cache size in bytes.

See [Configuring the Cache for File Systems](#) for details.

The default value depends on how you provisioned local storage before installing Storage Gateway. The recommended local storage disk size is 600 GB (500 GB for file system cache, 80 GB for metadata, 20 GB for log). If you followed the documented recommendations, the default value for the read cache is approximately 300 GB.

7. Click **Save**.

The file system is created and appears in the **File Systems** listing.

### Next Steps

Connect the file system to a directory on the Storage Gateway host. For more information, see [Connecting a File System](#).

You can also do the following in the management console:

- Set up the NFS export. This directory acts as a mount point. For more information, see [Mounting File Systems on Clients](#).
- Add more file systems. For more information, see [Adding a File System](#).
- View the details of a file system. For more information, see [Viewing the Details of a File System](#).

## Managing File Systems

A Storage Gateway file system connects a directory on a local host to an Object Storage bucket in Oracle Cloud Infrastructure. This topic describes how to manage Storage Gateway file systems.

### Adding a File System

You can add file systems in Storage Gateway and connect each file system to an Object Storage bucket in your tenancy.

#### To add a file system

1. Log in to the management console.
2. Click **File Systems** on the upper-right area of the management console.
3. Click **Create File System**.
4. Enter the required information in **Create a File System**:

- **File System Name:** A unique, friendly name for the file system. Avoid entering confidential information. Use the following guidelines when naming a file system:
  - Use from 1 to 256 characters.
  - Valid characters are letters (upper or lower case), numbers, hyphens, underscores, and periods.



### Important

The name **cannot** contain the following:

- A slash (/) character because this character delimits bucket and object names in Oracle Cloud Infrastructure Object Storage
- The strings "pub" or "priv"

If an Object Storage bucket by this file system name doesn't exist in your tenancy, the bucket is created.

If a corresponding Object Storage bucket by this file system name exists in your tenancy and there is data in the bucket, Storage Gateway works asynchronously to sync the bucket and file system contents.

- **Select the Object Storage tier in which you want to store your data**



### Important

Once set, you cannot change the storage tier in which a bucket resides.

You can use the Oracle Cloud Infrastructure Object Storage object lifecycle policies feature to manage the archiving and deletion of objects in a bucket according to a predefined schedule. See [Using Object Lifecycle Management](#) for details.

- **Standard:** The Standard tier is the primary default Object Storage tier for storing data that requires frequent and fast access. See [Overview of Object Storage](#) for more information.
- **Archive:** The Archive tier is a special tier for storing data that is accessed infrequently and requires long retention periods. See [Overview of Archive Storage](#) for more information. Access to data in the Archive tier is not immediate since you must restore archived data before it's accessible (see [Restoring Files and Objects from Archive Storage](#)).

- **Object Storage endpoint:** Required. The Object Storage API endpoint for your service instance. To find the object storage API endpoint for your Oracle Cloud Infrastructure Object Storage tenancy, see the [API documentation for Oracle Cloud Infrastructure Object Storage](#).



### Important

The following information is required to connect your Storage Gateway file systems to Oracle Cloud Infrastructure. See [Required Keys and OCIDs](#) for detailed information on how to generate the required keys and where to obtain these OCIDs.

- **Compartment OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage compartment.
- **Tenant OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage tenancy.
- **User OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage user.
- **Public Key's Finger Print:** Required. Your Oracle Cloud Infrastructure Object Storage public key fingerprint.
- **Private Key:** Required. Your Oracle Cloud Infrastructure Object Storage private key.

- **Private Key Passphrase:** Required if a passphrase was specified during key creation. Your Oracle Cloud Infrastructure Object Storage private key passphrase.



### Note

Your private key and passphrase are securely stored in the Storage Gateway docker. The Storage Gateway installation generates a pair of public and private keys. The system uses the public key to encrypt sensitive data.

5. Click **Save**.

The values you entered must match your Oracle Cloud Infrastructure credentials. If you get an error message, verify your entries, update any incorrect values, and click **Save** again.

6. Click **Show Advanced File System Configuration**.

Enter the required configuration information or click **Use Default** to accept the default values:

- **NFS Allowed Hosts:** A comma-separated list of hosts allowed to connect to the NFS export. You can also specify **\*** to allow all hosts to connect.

For example: `2001:db8:9:e54::/64, 192.0.2.0/24`

- **NFS Export Options:** The NFS export options.

Example: `rw, sync, insecure, no_subtree_check, no_root_squash`



### Important

Do not specify the `fsid` option.

- **Concurrent Uploads:** The number of concurrent uploads to Oracle Cloud Infrastructure.

This field indicates the maximum number of files that can be concurrently uploaded in Storage Gateway. If the value is 15, the concurrent file uploads can be between 1-15.

The allowed range is from 1 to 30.

- **Sync Policy:** The metadata operations are flushed to the disk based on the sync policy, but do not affect on-disk consistency. Currently, only **Posix Standard** is supported. Only the synchronous transactions (like `fsync`, `ODSYNC`, and `OSYNC`) are committed to the disk. All other transactions are handled asynchronously.

- **Cloud Read-ahead:** The number of blocks to be downloaded and used to *read ahead* when reading files for improved performance.

Default value: 50

- **Maximum Read Cache Size in GiB:** The maximum read cache.

When the read cache is full or reaches the configured limit, Storage Gateway removes files from the cache based on a least recently used (LRU) algorithm. Files pending upload to your tenancy are not removed from cache. You can also preserve files that you do not want removed from cache.



### Note

The number of files in cache is limited to 20,000, regardless of the specified cache size in bytes.

See [Configuring the Cache for File Systems](#) for details.

The default value depends on how you provisioned local storage before installing Storage Gateway. The recommended local storage disk size is 600 GB (500 GB for file system cache, 80 GB for metadata, 20 GB for log). If you followed the documented recommendations, the default value for the read cache is approximately 300 GB.

7. Click **Save**.

The file system is created and appears in the **File Systems** listing.

### Connecting a File System

After you create a file system, you must connect the file system to an Oracle Cloud Infrastructure Object Storage bucket before you can store and retrieve data through the file system.



#### Warning

If network connectivity with Oracle Cloud Infrastructure is lost, your file system is disconnected.

### To connect a file system

1. Log in to the Storage Gateway management console.
2. On the **Dashboard** tab, identify the file system that you want to connect to your Object Storage bucket.

Click **Connect**.

3. If a bucket with the same name as the file system exists in Object Storage, the file system is connected to that bucket. Any existing data cached in the Storage Gateway file system is deleted to ensure consistency with the data stored in the bucket. If a bucket by that name doesn't exist, the bucket is created and the file system is connected to the bucket. If the compartment OCID was specified during file system creation, then the bucket is created in that compartment. Otherwise, the bucket is created in the `root` compartment by default.



### Important

You can mount a read/write file system on only one Storage Gateway at a time.

If the file system that you're importing is connected to another Storage Gateway, the **File System: Claim Ownership** window appears. You can claim ownership and confirm that the other Storage Gateway can be disconnected. Enter the following information, and then click **Claim Ownership**:

- Public key finger print
- Private key
- Private key passphrase

Claiming ownership ensures that you don't inadvertently connect a new file system to a bucket that's already connected to another Storage Gateway file system.

## Mounting File Systems on Clients

Each Storage Gateway file system maps a directory to an Oracle Cloud Infrastructure Object Storage bucket. To establish the connection between Storage Gateway and an NFS client, you must mount the Storage Gateway file system on the NFS client.

Any Linux/UNIX NFS client certified to work with NFSv4 server running on Oracle Linux 7.x is compatible with Storage Gateway.



### Note

Storage Gateway does not currently support NFS clients running on Windows or Mac OS.

### To mount a Storage Gateway file system

1. Log in to the Storage Gateway host.
2. Start Storage Gateway:

```
sudo ocisg up
```

3. Find the NFS port number:

```
sudo ocisg info
```

Note the NFS port number from the output. For example:

```
Management Console: https://myStorageGatewayHost.example.com:32775
```

If you have already configured a OCISG File System via the Management Console, you can access the NFS share using the following port.

```
NFS Port: 32774
```

```
Example: mount -t nfs -o vers=4,port=32774 myStorageGatewayHost.example.com:<OCISG File System name> /local_mount_point
```

In the sample output:

- myStorageGatewayHost.example.com is the Storage Gateway host name.
  - 32775 is the management console port number.
  - 32774 is the NFS port number.
4. Log in to the NFS client from which you want to access your service instance through Storage Gateway.
  5. Create a directory on the NFS client.

### 6. Mount the file system on the directory that you created on the NFS client:

```
sudo mount -t nfs -o vers=4,port=<NFS_port_number> <storage_gateway_host_name>:/<ocisg_file
system_name> /<local_mount_point>
```

In this command:

- Replace *<NFS\_port\_number>* with the NFS port number you noted earlier.
- Replace *<storage\_gateway\_host\_name>* with the server name or IP address of the server on which Storage Gateway is installed.
- Replace *<ocisg\_file system\_name>* with the file system name that you want to mount.
- Replace *<local\_mount\_point>* with the path to the directory you created on the NFS client.

For example:

```
sudo mount -t nfs -o vers=4,port=32774 myStorageGatewayHost.example.com:/myFirstFS /home/xyz/abc
```

In this example,

- 32774 is the NFS port.
- myStorageGatewayHost.example.com is the Storage Gateway host name.
- myFirstFS is the file system name.
- /home/xyz/abc is the path to the directory `abc` on the NFS client.

The Storage Gateway file system is now mounted on the NFS client directory. You can now access the Storage Gateway file system from the NFS client.

For more information, see [Using Storage Gateway File Management Operations](#).

## Viewing the Details of a File System

You can view the configuration details of a file system and monitor the upload activity through the management console of Storage Gateway.

To view the details of a file system

1. Log in to the management console.
2. Click the name of the file system.
  - The **Details** tab displays the Oracle Cloud Infrastructure service type, storage tier, and the identity domain associated with your tenancy. If the file system is connected, you can see mount point connection information to help you with mounting that file system. For example:  
**NFS Client Mount Command:** `mount -t nfs -o vers=4,port=<nfs_mount_port> 129.213.122.84:/perftest01 /<local_mount_point>`
  - The **Settings** tab displays the following details:
    - Details of the tenancy and scope specified for the file system.
    - File system properties.
    - NFS and cache settings for the file system.

You can edit these settings. If you make changes, remember to click **Save**.

If your file system is connected, you can also see:

- The **Activity** tab, which shows ongoing and pending file upload activity.  
If you contact Oracle Support Services about any issue with the file system, you might need to provide the file system log to help diagnose the issue. To view or download the file system log, click **View Streaming Logs** near the lower-right corner of the **Details** tab.
- The **Completed Uploads** tab, which shows the last 100 files that were uploaded to Oracle Cloud Infrastructure Object Storage during the current browser session.



### Note

The file list doesn't persist across browser sessions. If you refresh the page or open the **Completed Uploads** tab in another browser after the files are uploaded, the list will be empty.

- You can also disconnect the file system. See [Disconnecting a File System](#).

## Changing the Properties of a File System

You can change the properties of a file system using the Storage Gateway management console.

### To change the properties of a file system

1. Log in to the management console.
2. In the **Dashboard**, click the name of the file system that you want to edit.
3. In the **Settings** tab, edit the file system properties and advanced settings, such as the cache limits.
4. Click **Save**.
5. For the changes to take effect, disconnect and reconnect the file system.

## Refreshing a File System

The auto-refresh feature triggers file system refreshes based on a time interval you specify. The system schedules the next refresh after any in-progress refresh completes. That means the elapsed time between the beginning of any two successive refreshes is equal to the specified auto-refresh interval plus the time required to run a file system refresh.

## CHAPTER 30 Storage Gateway

---

Use the following command to configure the auto-refresh feature:

```
ocisg set <file_system_name> dataset.refreshInterval=<interval_in_minutes>
```

The configuration command works on created and connected file systems. The configuration does not take effect until the file system is disconnected and reconnected or the Storage Gateway application restarts. To apply the changes, run:

```
ocisg down
ocisg up
```

Attribute caching can cause NFS clients to be unaware of files, corresponding to new objects in the bucket, that are created in a Storage Gateway file system during a refresh. You can use the `noac` mount option to turn off attribute caching. Turning off attribute caching can affect system performance.

When you run a refresh, the system reads attributes and fetches information about all objects in the corresponding bucket. Use a larger refresh interval for buckets with many objects.



### Note

After you refresh a file system, or create one for a bucket that already contains objects, Oracle recommends that you check for any files that might have been missed due to network connectivity issues.

To check for missing files, run the following command:

```
zgrep -ni "failed to get the object for" <path_to_gateway_logs>/<file_system_name>.*
```

For example, if the path to the gateway logging directory is `/ocisg/log` and the file system name is `my-fs-1`, the command is:

```
zgrep -ni "failed to get the object for" /ocisg/log/my-fs-1.*
```

Files listed in the output of this command were not successfully registered with the gateway. If any file names appear in the list, refresh the file system again.

## Disconnecting a File System

When a file system is disconnected, no one can access or modify that file system.

We recommend disconnecting file systems that are not in use. Disconnecting a file system frees up the resources associated with that file system, making those resources available to file systems that are active (connected).

### To disconnect a file system

1. Log in to the management console.
2. In the **Dashboard**, click the name of the file system that you want to disconnect.
3. Click **Disconnect**.

When you disconnect a file system, the bucket to which the file system was previously connected and the contents of that bucket remain intact.

4. For the changes to take effect, disconnect and reconnect the file system.

You can resume storing and retrieving data by connecting the file system again. You can delete the disconnected file system when you no longer need it. For more information, see [Deleting a File System](#).

### Importing an Existing File System

Before you import an existing file system from another Storage Gateway, ensure that any pending file uploads to Oracle Cloud Infrastructure Object Storage are complete.

#### To import an existing file system

1. Log in to the management console.
2. Click the **Create File System** navigation link.
3. Click **Create File System** in the navigation pane on the left.  
The **Create a File System** page appears.
4. Enter the required information in **Create a File System**.  
For the file system name, enter the name of the existing file system that you want to import to this Storage Gateway.
5. Click **Save**.
6. Select the options that you'd like to enable in the file system.
7. Click **Show Advanced** and enter the required information.
8. Click **Save**.  
The file system is created and appears on the **Dashboard** tab.
9. Click **Connect** for the file system that you want to import.

- a. If the file system that you're importing is connected to another Storage Gateway, the **File System: Claim Ownership** window appears so you can claim ownership and confirm that the other Storage Gateway can be disconnected. Enter the following information and click **Claim Ownership**:
  - Public key finger print
  - Private key
  - Private key passphrase
- b. If you connect to a file system that previously belonged to a different gateway, you must restart the new owning gateway:

```
ocisg down
ocisg up
```

10. Mount the file system to a directory on the Storage Gateway host and set up the NFS export. For example:

```
sudo mount -t nfs -o vers=4,port=<NFS_port_number> <storage_gateway_host_name>:<ocisg_file_system_name> /<local_mount_point>
```

In this command:

- Replace *<NFS\_port\_number>* with the NFS port number you noted earlier.
- Replace *<storage\_gateway\_host\_name>* with the server name or IP address of the server on which Storage Gateway is installed.
- Replace *<ocisg\_file system\_name>* with the file system name that you want to mount.
- Replace *<local\_mount\_point>* with the path to the directory you created on the NFS client.

### Deleting a File System

You can delete a file system from Storage Gateway when you no longer need it.

To delete a file system:

1. Log in to the management console.
2. On the **Dashboard**, identify the file system that you want to delete.



### Important

When you disconnect a file system, the bucket to which the file system was previously connected and the contents of that bucket remain intact.

Deleting a file system does not automatically delete the objects in the bucket. If you want to remove objects from the Object Storage bucket, set the **Delete Old File Versions** property for the file system and delete all the files before disconnecting the file system.

3. Ensure that the file system is disconnected. If it's still connected, click **Disconnect**.
4. After the file system is disconnected, click its name.  
The details of the file system appear.
5. Click **Delete**.  
The file system is deleted from Storage Gateway.

## Managing Storage Gateway

This topic describes some basic Storage Gateway management tasks.

### Managing Storage Gateway Using the CLI

You can use the `oci_sg` command line interface (CLI) to manage Storage Gateway. To use the CLI, open an `ssh` connection and log in to the host on which you installed Storage Gateway.



### Note

You can use the `ocisg` command line interface (CLI) to create, manage, and monitor Storage Gateway Cloud Sync jobs. See [Using Storage Gateway Cloud Sync](#) for details.

The CLI supports the following Storage Gateway management tasks:

- To start Storage Gateway:

```
sudo ocisg up
```

- To stop Storage Gateway:

```
sudo ocisg down
```



### Note

If the server with a Storage Gateway system fails, you can reinstall and start a new one. All the configuration and system data is automatically downloaded and applied. The pending upload and download activities resume when the new Storage Gateway system runs.

If a disk cache is unrecoverable on the Storage Gateway server, data might be lost since the file might not have been transferred to the bucket in your tenancy. To ensure efficient data protection, see [Best Practices for Using Storage Gateway](#).

- To view details about Storage Gateway and how to access the management console:

```
sudo ocisg info
```

## CHAPTER 30 Storage Gateway

---

- To find the version of Storage Gateway:

```
sudo ocisg version
```

- To configure Storage Gateway to use a proxy server for connections to Oracle Cloud Infrastructure Object Storage:

```
sudo ocisg configure proxy <proxy_server_URL>
```



### Note

After configuring the proxy server, you must stop and restart Storage Gateway.

By default, no proxy server is specified.

- To remove previously configured proxy server details in Storage Gateway:

```
sudo ocisg configure proxy [remove]
```

- To configure Storage Gateway to use SSL when communicating with the management console and REST APIs:

```
sudo ocisg configure ssl true
```

SSL is enabled by default.



### Note

After configuring Storage Gateway to use SSL, you must stop and restart Storage Gateway.

To disable SSL:

```
sudo ocisg configure ssl false
```

- To specify ports for the Storage Gateway services:

```
sudo ocisg configure port <service> <port_number>
```

## CHAPTER 30 Storage Gateway

---

- **<service>**: Specify `admin`, `nfs`, or `rest`.
- **<port\_number>**: Ensure that the port number you specify is not already in use on the Storage Gateway host.

By default, the port number is assigned dynamically for the Storage Gateway services when you start Storage Gateway.



### Note

For the port assignment to take effect, you must stop and start Storage Gateway.

- To remove the static port assignment for a service:

```
sudo ocisg configure port <service> remove
```

- To allocate memory for the Storage Gateway host:

```
sudo ocisg configure memory <memory_in_GB>
```

- To remove the memory allocation:

```
sudo ocisg configure memory remove
```

By default, Storage Gateway uses 4 GB of the available memory on the host server. You can delete the memory information by using the `remove` parameter.



### Note

After configuring memory for Storage Gateway, you must stop and restart Storage Gateway.

- To specify the docker network mode:

```
sudo ocisg configure network mode
```

The mode can be either *host* or *bridge*.

The default mode is `bridge`. In this mode, you can run multiple instances of Storage Gateway on your host.

In the `host` mode, you can run only a single instance of Storage Gateway. Network performance is better in `host` mode.



### Note

After specifying the docker network mode, you must stop and restart Storage Gateway.

- To change the Storage Gateway `admin` password:

```
sudo ocisg do password:reset
```

Set a new password:

```
sudo ocisg password:set <new_password>
```

Enter a password that meets the following requirements:

- Uses from 8 to 32 characters.
  - Includes at least one special character, one numerical character, one uppercase character, and one lowercase character.
- To view help for the available commands:

```
sudo ocisg help
```

## Using Storage Gateway File Management Operations

This topic describes how to use the Storage Gateway file management operations.

Exercise caution when using the REST API, Java library, or any other client to retrieve, create, update, or delete objects directly in a bucket that's mapped to a file system in Storage Gateway. Until you **Refresh** the Storage Gateway file system, Storage Gateway is not aware of the changes and data is inconsistent between Storage Gateway and Object Storage.

### Uploading Files to Buckets

Before you connect the file system to the Oracle Cloud Infrastructure Object Storage bucket, make a note of the Oracle Cloud Infrastructure Object Storage tenancy details such as namespace, tenant OCID, and compartment OCID.

Copy the files to the mounted directory on the Storage Gateway or the NFS client host. Storage Gateway writes the files to the disk cache. The system queues and asynchronously uploads the files to an Object Storage bucket. Corresponding objects are created in the storage tier you specified during file system creation, either **Standard** or **Archive**. See [Creating Your First File System](#) or [Managing File Systems](#) for details.



#### Note

Storage Gateway automatically performs multipart upload for files larger than 128 MB. See [Using Multipart Uploads](#) for details.

You can view files uploaded to your tenancy during the current browser session. See the **Completed Uploads** tab in [Viewing the Details of a File System](#).

### Reading Files

When you write a file to a Storage Gateway file system, the system stores the file in the local disk cache. You can read the file directly from the mounted directory. Storage Gateway asynchronously copies the file to the corresponding Object Storage bucket in your tenancy. To retrieve the data from the bucket using Storage Gateway, read the files from the mounted directory.

If space is available, Storage Gateway automatically places the files in the read cache. If the file is in the read cache, you can retrieve the file immediately. If the file is not available in the read cache and it is stored in the `Archive` tier, you must restore the object. For more information, see [Restoring Files and Objects from Archive Storage](#).



### Note

You cannot read or write to a file that is stored in the **Archive** tier and does not exist in the read cache. This action returns an `Input/Output error`.

Storage Gateway checks data integrity using checksum verification on uploads. The system might not be able to perform data integrity validation on a partial read, since checksum verification works only on a whole file or object.

To read the upload checksum for a file in a file system, run the following command from the NFS client on which the file system is mounted:

```
sudo docker exec ocisg bash -c "cd /mnt/gateway/${filesystem} && cat ${filepath}:::meta:csn"
```

## Restoring Files and Objects from Archive Storage

You can initiate a file restore from the Storage Gateway command line. You can also initiate an object restore from Archive Storage in Oracle Cloud Infrastructure. You can read the corresponding file using Storage Gateway after the object has been restored to Object Storage.



### Note

Storage Gateway supports Oracle Cloud Infrastructure Object Storage object lifecycle policies to manage the archiving and deletion of objects in a bucket according to a pre-defined schedule. Using object lifecycle policies, you can specify bucket creation in the Standard Object Storage tier, and then create a policy to schedule the subsequent movement of data to the Archive Storage tier. This lifecycle policy archival method is useful if you have on-premises applications that generate intermediary or temporary files and directories that are inappropriate for immediate archival. See [Using Object Lifecycle Management](#) for details.

### Restoring Archived Files Using the Storage Gateway Command Line

#### To restore one or more archived files

Open a command prompt on the Storage Gateway host and run the `ocisg archive restore` command. Specify the full path to a directory or to a file.

```
ocisg archive restore <file_system_name> <full/path/to/directory/or/file> [<#_of_hours>]
```

By default, the file remains restored for 24 hours after restoration. However, you can optionally specify [*<#\_of\_hours>*] with an integer value of from 1 to 240 hours.

For example, to restore all the files in a directory:

```
ocisg archive restore myFS myDir/mySubDir 240
```

For example, to restore a single file:

```
ocisg archive restore myFS myDir/mySubDir/file2 240
```

### To check the archive status of one or more files

You can get the archive status for all files in a file system, for all files in a directory, or for an individual file.

Open a command prompt on the Storage Gateway host and run the `ocisg archive restore-status` command.

```
ocisg archive restore-status <file_system_name> [<full/path/to/directory/or/file>]
```

The status can be one of the following:

- Archived
- In progress
- Restored

For example, to check the archive status for all files in a file system:

```
ocisg archive restore-status myFS
```

To check the archive status for all files in directory:

```
ocisg archive restore-status myFS myDir/mySubDir
```

To check the archive status for an individual file:

```
ocisg archive restore-status myFS myDir/mySubDir/file2
```

### To check the restoration job status for a file system

You can get the status for all restoration jobs that have been initiated for a file system.

Open a command prompt on the Storage Gateway host and run the `ocisg archive restore-jobs` command.

```
ocisg archive restore-jobs <file_system_name>
```

For example:

```
ocisg archive restore-jobs myFS
```

### Restoring Archived Files Using Oracle Cloud Infrastructure



#### Important

If you use Oracle Cloud Infrastructure to restore archived objects, use the **Refresh** operation in Storage Gateway to display the data that was added or modified directly in Object Storage.

To restore an archived object using the Oracle Cloud Infrastructure Console



#### Tip

You need `OBJECT_RESTORE` permissions to restore Archive Storage objects.

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment your bucket is in.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. To restore a single object, click the Actions icon (three dots) to the right of the object you want to restore, and then click **Restore**. To restore multiple objects, select the check boxes to the left of each object you want to restore, then click **Restore**.
6. Optionally, specify the **Time Available for Download in Hours**.  
By default, you have 24 hours to download an object after restoration. However you can alternatively specify a download time of from 1 to 240 hours. You can find out how much download time is remaining by looking at **Available for Download** in object **Details**

or by looking at the Actions icon (three dots) menu to the right of **Download**. Refresh the browser to obtain up-to-date remaining download time information.

After the allotted download time expires, the object returns to Archive Storage.

7. Click **Restore Objects**.

Error messages are generated if there is a problem with restoring the selected objects.

You can optionally click **Retry failed restore option**.

### To check the status of an object restoration using the Oracle Cloud Infrastructure Console

1. Open the navigation menu. Under **Core Infrastructure**, click **Object Storage**.
2. Choose the compartment your bucket is in.  
A list of buckets is displayed.
3. Click the bucket name that contains your object.
4. Click **Objects** under **Resources**.  
A list of objects in the bucket is displayed.
5. Click the Actions icon (three dots) to the right of the object you want to check the restoration or download status of, then click **Details**.
6. Check the **Status**.  
**Status** displays one of the following:
  - Archived
  - Restoring
  - Restored

To restore an archived object using the Oracle Cloud Infrastructure CLI



### Tip

You need OBJECT\_RESTORE permissions to restore Archive Storage objects.

Open a command prompt and run `oci os object restore` to restore an object from Archive Storage:

```
oci os object restore -ns <object_storage_namespace> -bn <archive_bucket_name> --name <archived_object_name> [--hours <#_of_hours>]
```

By default, you have 24 hours to download an object after restoration. However, you can optionally specify `--hours` with an integer value of download time of from 1 to 240 hours.

To check the status of an object restoration using the Oracle Cloud Infrastructure CLI

Open a command prompt and run `oci os object restore-status` to check the status of restoring an object from Archive Storage:

```
oci os object restore-status -ns <object_storage_namespace> -bn <archive_bucket_name> --name <archived_object_name>
```

## Deleting Files

Remove the files that you no longer need from the NFS client by deleting them from the directory on which the file system is mounted.

## Monitoring Storage Gateway

This topic describes how to monitor Storage Gateway file system upload activity, system health, and storage usage. The topic also describes how to view system notifications and how to receive notifications in email.



### Important

Monitoring file system activity in the Storage Gateway management console consumes system resources. Monitoring file system activity is not recommended when Storage Gateway is under high load.

## Monitoring Upload Activity

When you upload a file to a file system, you can view the status of the upload activity. The **Activity** tab shows the ongoing and pending upload activity in a file system.

### To monitor upload activity

1. Log in to the management console.
2. Click **File Systems** in the upper-right corner of the management console.
3. On the **File Systems** page, select a file system.
4. Click **Activity**.

You can see the upload progress of the file in the **Uploading** pane.

## Monitoring System Health Status

You can monitor the overall system health status using **System Status** on the right side of the management console.

Storage Gateway performs an automated "health check" on the system to monitor the status of:

- Storage Gateway resources and services.
- Local storage, file system cache, metadata storage, and log storage.

**System Status** shows the result of Storage Gateway health analysis, either **Healthy** or **Unhealthy**.

**System Status** also provides information on:

- Any system warnings or errors.
- **Local I/O Mode**, which depends on local disk usage:
  - **Normal**  
The available disk space is greater than 10 GB in Storage Gateway. You can upload files in Storage Gateway and upload them to your Oracle Cloud Infrastructure tenancy.
  - **Rejecting I/O**  
The available disk space is less than 10 GB in Storage Gateway. Storage Gateway runs in protection mode and does not allow any writes to its local disk. All read operations work normally. All Storage Gateway metadata operations fail except for deletions and truncation.  
  
To return to **Normal** mode, you must wait until all ongoing upload activities complete and the files are removed from the read cache.
- **Throughput**  
The approximate upload throughput to Object Storage. If there is no recent activity, **Throughput** shows **Idle**.
- **Available Read Cache**  
The amount of read cache available. For optimal performance, reliability, and fault tolerance, follow the guidelines for configuring cache storage . See [Configuring Local Storage for File Systems and Cache](#) for details.
- **Pending Uploads**  
The number of files or directories, for all file systems, awaiting upload to Object Storage. If **Pending Uploads** is 0 (zero), all files and directories have been uploaded.

## Monitoring Storage Usage

You can track the storage usage and availability.

### To monitor storage usage

1. Log in to the management console.
2. Click **System** in the upper-right corner of the management console.
3. Click **System Stats**.

The system data appears in three panes:

- **Local Storage**
- **Local I/O**
- **Local Resources**

#### **Local Storage**

This pane provides a graphical representation of the amount of storage being used and the available free storage on the Storage Gateway host. You can see:

- Available local storage.
- Storage used for pending uploads and preserved cache files.
- Storage used for metadata.
- Storage used for logging.
- Storage used for other applications.

#### **Local I/O**

This pane displays the local I/O mode of Storage Gateway based on the local disk space usage on the Storage Gateway host.

#### **Local Resources**

This pane shows the overall memory usage and availability for Storage Gateway from the following fields:

- **Available Cores:** The number of CPUs being used by Storage Gateway.
- **Maximum Memory Available to Storage Gateway:** The total RAM available for Storage Gateway.
- **Memory Used by Storage Gateway:** The amount of memory being used by the file systems in Storage Gateway.
- **Free Memory:** The amount of free RAM available in Storage Gateway host.

### Viewing System Notifications

The **System Notifications** tab shows system notifications and helps you track overall system performance.

#### To view system notifications

1. Log in to the management console.
2. Click **System** in the upper-right corner of the management console.
3. Click **System Notifications**.

You can view a list of warnings or critical system notifications.

### Configuring Email Notification

You can configure Storage Gateway to notify you by email about system health and Cloud Sync job completion.

#### To configure email notification

1. Log in to the management console.
2. Click **System** in the upper-right corner of the management console.
3. Click **System Notifications**.

4. If email notifications are not yet configured, click **Click here to configure**.
5. Enter the required information for the following fields:
  - SMTP server.
  - Email addresses to receive notifications.
6. Click **Show Advanced Options** and enter the required information in the advanced configuration fields:
  - SMTP port.
  - SMTP User name.
  - SMTP Password.
  - Sender's Email Address.  
The default value is: `noreply@oracle.com`
7. Click **Save**.
8. Click **Test Email Notification** to verify that the specified email address receives a system notification email.

## Using Storage Gateway Cloud Sync

Use Cloud Sync to move on-premises datasets from a local file system to Storage Gateway, where the data is then moved asynchronously to Oracle Cloud Infrastructure Object Storage. You can also use Cloud Sync to synchronize Storage Gateway file system changes back to the local file system.

Cloud Sync generates the following for each sync job:

- Sync status (Created, Running, Completed, Failed, or Canceled).
- Number of files and directories to be copied from the source to the target.
- File and directory upload progress .
- Number of files and directories uploaded to the target.
- Time the job started, the time the job ended, and the run duration.

- The number of files skipped. (Cloud Sync skips non-regular files, such as symlinks, in the source directory.)
- A list of skipped files.

You can use the Storage Gateway [management console](#) or the [command line interface \(CLI\)](#) to create, manage, and monitor Cloud Sync jobs.

### About Cloud Sync

Cloud Sync is an enhanced wrapper around the Linux `rsync` command and relies on `rsync` to detect new and changed files using size and modification time. Cloud Sync also relies on `rsync` to verify the files once the data transfer is complete using checksums of the files. Cloud Sync calls the Storage Gateway diagnostic commands to provide detailed data movement progress and status between your local server, Storage Gateway, and Oracle Cloud Infrastructure.

### Understanding Failure Behavior

Because the file system is mounted locally, the NFS client running on the host handles any issues resulting from file system availability or connectivity.

Cloud Sync does the following:

- Reports and logs any failures to list or copy specific files (for example, resulting from permission issues).
- Monitors and reports on the pending and completed uploads to Object Storage.

Storage Gateway handles any connectivity and access issues to Object Storage and performs retry operations as needed.

### Prerequisites for Cloud Sync

- Create the Storage Gateway file system that serves as either the source or target destination for the sync operation. A file system on the Storage Gateway host maps to a bucket with an identical name in Oracle Cloud Infrastructure Object Storage. See [Creating Your First File System](#) or [Adding a File System](#) for details.
- Obtain the proper credentials to mount the file system share from the local server.
- The local server source must:
  - Have names services set up correctly so that UIDs and GIDs are preserved and are not remapped to `nobody`.
  - Be exported with root permissions to read and traverse the entire source tree.

### Using the Management Console

You can use the Storage Gateway management console to create, manage, and monitor Cloud Sync jobs.

#### To create a Cloud Sync job that syncs all files and directories at the specified source location

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.  
By default, a list of all the Cloud Sync jobs that have already been created is displayed.
3. Click **Create Cloud Sync Job**.
4. In **Create Cloud Sync Job** page, specify the attributes of the job:
  - **Job Name:** Required. A unique, user-friendly name for the job. Avoid entering confidential information.
  - **Source Path:** Path to the Cloud Sync source directory containing the files to sync.

- If you are syncing a local file system to Oracle Cloud Infrastructure, specify the source path as:

```
/cloudsync/mounts/<user_mount>[/<path_to_directory>]
```

- If you are syncing Oracle Cloud Infrastructure to a local file system, specify the source path as:

```
<storage_gateway_file_system>/<path_to_directory>
```

- **Target Path:** Path to the Cloud Sync target directory for the synced files.
  - If you are syncing a local file system to Oracle Cloud Infrastructure, specify the target path as:

```
<storage_gateway_file_system>/<path_to_directory>
```
  - If you are syncing Oracle Cloud Infrastructure to a local file system, specify the target path as:

```
/cloudsync/mounts/<user_mount>[/<path_to_directory>]
```
- **Enable Auto-Deletion:** Enable this option if you want Cloud Sync to automatically delete files from the target when:
  - Files are deleted from the source.
  - Source files have been renamed.

By default, when a file is deleted on the source, Cloud Sync does not automatically delete the file on the target unless you enable **Enable Auto-Deletion**. Also, when a source file is renamed, the file with the old name is deleted and a file with the new name is created. By default, Cloud Sync does not delete the file with the old name on the target (retaining both a file with the old name and a file with the new name) unless you choose **Enable Auto-Deletion**. The names of all deleted files are stored in a job-specific log directory.

### 5. Click **Create Cloud Sync Job**.

A Cloud Sync job is created and displayed in the list of jobs.

### To list Cloud Sync jobs

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.  
By default, a list of all the Cloud Sync jobs is displayed.
3. Optionally, you can filter the job listing by status (Created, Running, Completed, Failed, or Canceled) and type of Cloud Sync job (upload or download) by clicking one of the **Quick Filters**.

### To run a Cloud Sync job

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.
3. In the list of jobs, find the Cloud Sync job that you want to run.
4. Click **Run** to the right of the job name.  
Cloud Sync runs the job. The management console displays the status of the job just below the job name.

### To get the status of a Cloud Sync job

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.
3. In the list of jobs, find the Cloud Sync job for which you want status.  
The management console displays the status of the job (Created, Running, Completed, Failed, or Canceled) just below the job name.

### To cancel a Cloud Sync job

You can only cancel a job if it is running.

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.

3. In the list of jobs, find the Cloud Sync job that you want to cancel.



### Tip

Use **Quick Filters** to get a list of **Running** jobs.

4. Click **Cancel** to the right of the job name.

### To delete a Cloud Sync job

You cannot delete a running job.



### Tip

Cancel a running job before you try to delete it.

1. Log in to the management console.
2. Click **Cloud Sync** in the upper-right corner of the management console.
3. In the list of jobs, find the created, canceled, or failed Cloud Sync job that you want to delete.
4. Click **Delete** to the right of the job name.

### Using the CLI

You can use the `ociisg` command line interface (CLI) to create, manage, and monitor Cloud Sync jobs. Using `ssh`, log in to the host on which you installed Storage Gateway to use the CLI.

To create a Cloud Sync job that syncs all files and directories at the specified

## source location

**Warning**

Avoid entering confidential information in the job name.

Open a command prompt and run `ocisg cloudsync create` to create a job:

```
sudo ocisg cloudsync create [--schedule=<schedule>] [--auto-delete] [--parallel=<number>] [--files-from=<file>] [--exclude-from=<file>] --verify-contents <job_name> <source_path> <target_path>
```

`--schedule` is an option to automate the Cloud Sync job so it runs according to the specified schedule. Set the schedule value using the format used for *cron* jobs. For example, `--schedule="*/5 * * * *"` runs the job every five minutes.

`--auto-delete` is an option to direct Cloud Sync to automatically delete files from the target when files are deleted from the source, and old names of files that have been renamed. By default, Cloud Sync does not automatically delete the files on the target unless you specify this option. The names of all deleted files are stored in a job-specific log directory.

`--parallel=<number>` is an option to specify the number of processes for data synchronization. By default, the number of processes is set to one. With a single process using a hard disk drive, you can expect a sync rate of 60-70 MB/s. If your system has higher disk throughput, you can use the number of processes that is proportional to the available bandwidth. Oracle recommends from two to five processes for optimal performance. The maximum number of processes allowed is 10.

`--files-from=<file>` is an option to specify a set of files that you want to sync to the target. If you do not set this option, the service syncs all files. The `<file>` should be a file under the `/cloudsync/` directory. For example, `--files-from="/cloudsync/files.list"`.

`--exclude-from=<file>` is an option to specify a set of files that you want to exclude from the Cloud Sync job. The `<file>` should be a file under the `/cloudsync/` directory. For example, `--exclude-from="/cloudsync/exclude.list"`.

## CHAPTER 30 Storage Gateway

---

--verify-contents is an option to enable verification of the destination contents against the source. If you do not explicitly enable verification, new files added to Cloud Sync sources after the job has started are not reported as errors when they do not appear in the destination.

The Cloud Sync job is created and displayed in the list of jobs.

### To list Cloud Sync jobs

Open a command prompt and run `ocisg cloudsync list` to list jobs:

```
sudo ocisg cloudsync list [-s <status>] [<job_name_or_type>]
```

Optionally, you can filter the list of jobs by specifying job status (Created, Running, Completed, Failed, or Canceled). You can also list a single job by specifying the name of that job.

For example:

```
sudo ocisg cloudsync list sync_to_os1
```

### To run a Cloud Sync job

Open a command prompt and run `ocisg cloudsync run` to run a job:

```
sudo ocisg cloudsync run <job_name>
```

For example:

```
sudo ocisg cloudsync run sync_to_os1
```

### To get the status of a Cloud Sync job

Open a command prompt and run `ocisg cloudsync status` to get the status of a job:

```
sudo ocisg cloudsync status <job_name>
```

For example:

## CHAPTER 30 Storage Gateway

---

```
sudo ocisg cloudsync status sync_to_os1
```

### To cancel a Cloud Sync job

You can cancel a job only if the job is in progress.

Open a command prompt and run `ocisg cloudsync cancel` to cancel a job:

```
sudo ocisg cloudsync cancel <job_name>
```

For example:

```
sudo ocisg cloudsync cancel sync_to_os1
```

### To delete a Cloud Sync job

Open a command prompt and run `ocisg cloudsync delete` to delete a job:

```
sudo ocisg cloudsync delete <job_name>
```

For example:

```
sudo ocisg cloudsync delete sync_to_os1
```

## Best Practices for Using Storage Gateway

Apply the recommendations found in the following topics to optimize the manageability, performance, reliability, and security of your Storage Gateway.

- [Security Considerations](#)
- [Understanding Storage Gateway Performance](#)
- [Configuring Local Storage for File Systems and Cache](#)
- [Determining File System Cache Size](#)
- [Recommended Uses and Workloads](#)
- [Uses and Workloads Not Supported](#)

- [Renaming Directory Trees](#)
- [Limits on Storage Gateway Resources](#)

# Troubleshooting Storage Gateway

This topic covers some common Storage Gateway issues and how to address them.

## I installed docker and NFS on my host, but I can't install Storage Gateway

1. Add the docker group to the existing groups in your host:

```
sudo groupadd docker
```

2. Add your user id to the docker group:

```
usermod -a -G docker <username>
```

3. Shut down your host:

```
shutdown -r now
```

4. Log in to your host and run the Storage Gateway installation script:

```
sudo ./ocisg-install.sh
```

## I can't access the management console

1. Run the `ocisg info` command:

```
sudo ocisg info
```

If Storage Gateway is not running, start Storage Gateway:

```
sudo ocisg up
```

2. Make a note of the management console port number from the output:

```
Management Console: https://myStorageGatewayHost.example.com:32775
```

If you have already configured a OCISG File System via the Management Console, you can access the NFS share using the following port.

## CHAPTER 30 Storage Gateway

---

```
NFS Port: 32774
```

```
Example: mount -t nfs -o vers=4,port=32774 myStorageGatewayHost.example.com:<OCISG File System name> /local_mount_point
```

In the example, `myStorageGatewayHost.example.com` is the Storage Gateway host name and `32775` is the management console port number.

3. Ensure that Storage Gateway is running `docker` on the Storage Gateway host.
4. Check that the management console port number in the output from `ocisg info` matches the port you're using to access the management console.
5. Ensure that you are using `https` if you have enabled SSL. SSL is enabled by default.

### I am unable to mount a file system

1. Run the `ocisg info` command:

```
sudo ocisg info
```

If Storage Gateway is not running, start Storage Gateway:

```
sudo ocisg up
```

2. Make a note of the management console port number and NFS port number from the output:

```
Management Console: https://myStorageGatewayHost.example.com:32775
```

If you have already configured a OCISG File System via the Management Console, you can access the NFS share using the following port.

```
NFS Port: 32774
```

```
Example: mount -t nfs -o vers=4,port=32774 myStorageGatewayHost.example.com:<OCISG File System name> /local_mount_point
```

In the sample output:

## CHAPTER 30 Storage Gateway

---

- myStorageGatewayHost.example.com is the Storage Gateway host name.
  - 32775 is the management console port number.
  - 32774 is the NFS port number.
4. Log in to the NFS client from which you want to access your service instance through Storage Gateway.
  5. Create a directory on the NFS client.
  6. Mount the file system on the directory that you created on the NFS client:

```
sudo mount -t nfs -o vers=4,port=<NFS_port_number> <storage_gateway_host_name>:/<ocisg_file
system_name>/<local_mount_point>
```

In this command:

- Replace *<NFS\_port\_number>* with the NFS port number.
- Replace *<storage\_gateway\_host\_name>* with the server name or IP address of the server on which Storage Gateway is installed.
- Replace *<ocisg\_file system\_name>* with the name of the file system you want to mount.
- Replace *<local\_mount\_point>* with the path to the directory you created on the NFS client.

For example:

```
sudo mount -t nfs -o vers=4,port=32774 myStorageGatewayHost.example.com:/myFirstFS /home/xyz/abc
```

In this example,

- 32774 is the NFS port number.
  - myStorageGatewayHost.example.com is the Storage Gateway host name.
  - myFirstFS is the file system name.
  - /home/xyz/abc is the path to the directory abc on the NFS client.
7. Ensure that Storage Gateway is running `docker` on the Storage Gateway host.
  8. Ensure that the NFS protocol is running:

## CHAPTER 30 Storage Gateway

---

```
sudo systemctl enable nfs-server
```

9. Check that the NFS port number in the output from `ocisg info` matches the port you're using to connect to with your NFS client.

### I cannot delete a bucket after canceling a Cloud Sync job

If you cancel an active or stalled Cloud Sync job and disconnect the file system, you might not be able to delete the associated Object Storage bucket. If file uploads were in progress when you canceled the job, the Object Storage service might expect a commit that never completed. In this case, the service does not allow bucket deletion and returns the error "multipart upload pending". You can use the CLI to resolve the issue.

1. List the bucket's pending multipart uploads:

```
oci os multipart list -bn <bucket_name>
```

Be sure to note the relevant object names and upload IDs.

2. Delete all pending uploads:

```
oci os multipart abort -bn <bucket_name> --object-name <object_name> --upload-id <upload_id>
```

3. Delete the bucket:

```
oci os bucket delete -ns <object_storage_namespace> --name <bucket_name>
```

### Additional NFS Troubleshooting

The Storage Gateway installation software installs the NFS, if needed, and automatically configures it. After the installation, the NFS is configured and a file system created. You can then mount the filesystem from a remote client. Sometimes this mount can fail.

To troubleshoot a mount failure:

1. Ensure that the NFS port is included in the Storage Gateway's subnet security list and that it is available there. If the port does not appear in the subnet security list, add it and retry the mount.

2. Run `rpcinfo -p`. The command should return:

```
100003 4 tcp 2049 nfs
```

This result means that NFS is ready, available, and the mount succeeds.

3. If `nfs` does not appear in the response to the `rpcinfo -p` command, enable and restart both `rpcbind` and NFS:

```
sudo systemctl enable rpcbind
sudo systemctl enable nfs
sudo systemctl start rpcbind
sudo systemctl start nfs
```

4. Run the `rpcinfo -p` command again to verify that NFS is now available.
  - a. If NFS still is not available, reboot the Storage Gateway.
  - b. Run the `rpcinfo -p` command again to confirm.
5. If you remain unable to mount the file system, contact [My Oracle Support](#).

## Contacting Oracle Support

If you need technical support or help with Storage Gateway, you can go to [My Oracle Support](#) and create a service request. See [Creating a Service Request](#) for information.

## Upgrading Storage Gateway

Storage Gateway notifies you when there is a new version available for you to download and install:

- A pop-up notification appears in the management console after you log in.
- A small banner notification appears at the top of the **Dashboard**.
- If you have configured email notifications, the system sends an email notification. See [Configuring Email Notification](#) for details.



### Important

If you are upgrading from Storage Gateway 1.0, underlying database and schema changes require you to recreate your Storage Gateway file systems. See [Recreating Your File Systems](#) for details.

### Before You Begin

- Plan for downtime appropriately since the upgrade takes some time to complete. The downtime varies depending on the system resources, the number of file systems, and the number of files.
- Wait for any pending or ongoing write operations from the NFS client instances to complete, then unmount all file systems.
- Wait for pending uploads to Oracle Cloud Infrastructure Object Storage to complete. On the **Dashboard** under **System Status**, ensure that the **Pending Uploads** field shows 0.
- Disconnect all of the file systems.
- Ensure that there is no ongoing activity in the **Activity** tab for each file system in the management console.



### Important

To enable partial update capabilities in Storage Gateway, there must be no pending uploads for any associated file systems before the upgrade. The upgrade process purges the existing local cache across all Storage Gateway file systems.

### Upgrading Storage Gateway

#### To upgrade to Storage Gateway 1.3

1. Log in to your Storage Gateway host.
2. [Download the Storage Gateway 1.3 tar archive](#).
3. Extract the files from the downloaded `ocisg-1.3.tar.gz` file:

```
tar xvzf ocisg-1.3.tar.gz
```

This command extracts the file's contents to a subdirectory named `ocisg-1.3`.

4. Change directory to `ocisg-1.3`:

```
cd ocisg-1.3
```

5. Run the installation script as `sudo` or `root` user:

```
sudo ./ocisg-install.sh
```

If you encounter any interruption during the upgrade, such as lost connectivity, rerun the installation script to resume the upgrade.

If you are upgrading from Storage Gateway 1.0, you must [recreate the file systems](#) that were created in the 1.0 version of the Storage Gateway software. Connect the file systems in the management console and claim ownership if there's a bucket ownership prompt. Storage Gateway version 1.3 rebuilds the local metadata for existing buckets in Object Storage. The more objects there are in the buckets, the more time it takes to rebuild the metadata.

### Recreating Your File Systems

When you created file systems in Storage Gateway 1.0, corresponding Oracle Cloud Infrastructure Object Storage buckets were created. Recreate those file systems in Storage Gateway so that you can connect to the same buckets and automatically see the files that have already been uploaded to Object Storage.

When you recreate Storage Gateway file systems, data that you already uploaded to Oracle Cloud Infrastructure Object Storage is automatically included in the newly created file system.

### To recreate your file systems

1. Log in to the management console.
2. Click **File Systems** on the upper-right area of the management console.
3. Click **Create File System**.
4. Enter the required information in **Create a File System**:
  - **File System Name:** A unique, friendly name for the file system. Avoid entering confidential information. Use the following guidelines when naming a file system:
    - Use from 1 to 256 characters.
    - Valid characters are letters (upper or lower case), numbers, hyphens, underscores, and periods.



#### Important

The name **cannot** contain the following:

- A slash (/) character because this character delimits bucket and object names in Oracle Cloud Infrastructure Object Storage
- The strings "pub" or "priv"

If an Object Storage bucket by this file system name doesn't exist in your tenancy, the bucket is created.

If a corresponding Object Storage bucket by this file system name exists in your tenancy and there is data in the bucket, Storage Gateway works asynchronously to sync the bucket and file system contents.

- **Select the Object Storage tier in which you want to store your data**



**Important**

Once set, you cannot change the storage tier in which a bucket resides.

You can use the Oracle Cloud Infrastructure Object Storage object lifecycle policies feature to manage the archiving and deletion of objects in a bucket according to a predefined schedule.

See [Using Object Lifecycle Management](#) for details.

- **Standard:** The Standard tier is the primary default Object Storage tier for storing data that requires frequent and fast access. See [Overview of Object Storage](#) for more information.
- **Archive:** The Archive tier is a special tier for storing data that is accessed infrequently and requires long retention periods. See [Overview of Archive Storage](#) for more information. Access to data in the Archive tier is not immediate since you must restore archived data before it's accessible (see [Restoring Files and Objects from Archive Storage](#)).

- **Object Storage endpoint:** Required. The Object Storage API endpoint for your service instance. To find the object storage API endpoint for your Oracle Cloud Infrastructure Object Storage tenancy, see the [API documentation for Oracle Cloud Infrastructure Object Storage](#).



### Important

The following information is required to connect your Storage Gateway file systems to Oracle Cloud Infrastructure. See [Required Keys and OCIDs](#) for detailed information on how to generate the required keys and where to obtain these OCIDs.

- **Compartment OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage compartment.
- **Tenant OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage tenancy.
- **User OCID:** Required. Unique identifier of your Oracle Cloud Infrastructure Object Storage user.
- **Public Key's Finger Print:** Required. Your Oracle Cloud Infrastructure Object Storage public key fingerprint.
- **Private Key:** Required. Your Oracle Cloud Infrastructure Object Storage private key.

- **Private Key Passphrase:** Required if a passphrase was specified during key creation. Your Oracle Cloud Infrastructure Object Storage private key passphrase.



### Note

Your private key and passphrase are securely stored in the Storage Gateway docker. The Storage Gateway installation generates a pair of public and private keys. The system uses the public key to encrypt sensitive data.

5. Click **Save**.

The values you entered must match your Oracle Cloud Infrastructure credentials. If you get an error message, verify your entries, update any incorrect values, and click **Save** again.

6. Click **Show Advanced File System Configuration**.

Enter the required configuration information or click **Use Default** to accept the default values:

- **NFS Allowed Hosts:** A comma-separated list of hosts allowed to connect to the NFS export. You can also specify **\*** to allow all hosts to connect.

For example: `2001:db8:9:e54::/64, 192.0.2.0/24`

- **NFS Export Options:** The NFS export options.

Example: `rw, sync, insecure, no_subtree_check, no_root_squash`



### Important

Do not specify the `fsid` option.

- **Concurrent Uploads:** The number of concurrent uploads to Oracle Cloud Infrastructure.

This field indicates the maximum number of files that can be concurrently uploaded in Storage Gateway. If the value is 15, the concurrent file uploads can be between 1-15.

The allowed range is from 1 to 30.

- **Sync Policy:** The metadata operations are flushed to the disk based on the sync policy, but do not affect on-disk consistency. Currently, only **Posix Standard** is supported. Only the synchronous transactions (like `fsync`, `ODSYNC`, and `OSYNC`) are committed to the disk. All other transactions are handled asynchronously.

- **Cloud Read-ahead:** The number of blocks to be downloaded and used to *read ahead* when reading files for improved performance.

Default value: 50

- **Maximum Read Cache Size in GiB:** The maximum read cache.

When the read cache is full or reaches the configured limit, Storage Gateway removes files from the cache based on a least recently used (LRU) algorithm. Files pending upload to your tenancy are not removed from cache. You can also preserve files that you do not want removed from cache.



### Note

The number of files in cache is limited to 20,000, regardless of the specified cache size in bytes.

See [Configuring the Cache for File Systems](#) for details.

The default value depends on how you provisioned local storage before installing Storage Gateway. The recommended local storage disk size is 600 GB (500 GB for file system cache, 80 GB for metadata, 20 GB for log). If you followed the documented recommendations, the default value for the read cache is approximately 300 GB.

7. Click **Save**.

The file system is created and appears in the **File Systems** listing.

# Uninstalling Storage Gateway

This topic describes how to uninstall Storage Gateway.

## Uninstalling

To uninstall Storage Gateway

1. Log in to the on-premises host or compute instance from which you want to uninstall Storage Gateway.
2. Stop Storage Gateway:

```
sudo ocisg down
```

3. If the `ocisg_data` container exists in `docker ps -a` output, remove it:

```
sudo docker rm -v ocisg_data
```

4. Delete the image in `docker`:

```
sudo docker rmi $(sudo docker images | grep ocisg | awk '{print $3}')
```

5. Delete all the files in `/usr/bin/` that begin with `ocisg`:

```
sudo rm /usr/bin/ocisg*
```

6. View the contents of the file `gateway_config`:

```
cat /etc/gateway_config
```

### Sample output:

```
$ cat /etc/gateway_config
DATASTORAGE=/ocisg/cache
MDSTORAGE=/ocisg/metadata
LOGSTORAGE=/ocisg/log
PROXY=
```

## CHAPTER 30 Storage Gateway

---

```
USE_SSL=
MEMORY=
NETWORK=bridge
HTTP_FRAMEWORK=
ADMINPORT=443
NFSPORT=32769
RESTPORT=32768
```

5. Delete the `DATASTORAGE` directory, for example:

```
sudo rm -rf /ocisg/cache
```

6. Delete the `MDSTORAGE` directory , for example:

```
sudo rm -rf /ocisg/metadata
```

7. Delete the `LOGSTORAGE` directory , for example:

```
sudo rm -rf /ocisg/log
```

8. Delete the `gateway_config` file:

```
sudo rm /etc/gateway_config
```

9. Delete the Storage Gateway installation directory `ocisg`:

```
sudo rm -rf /opt/ocisg
```

## Getting Help with Storage Gateway

This topic provides information about getting help with Oracle Cloud Infrastructure Storage Gateway.

### Contacting Oracle Support

If you need technical support or help with Storage Gateway, you can go to [My Oracle Support](#) and create a service request. See [Creating a Service Request](#) for information.

### Downloading the Support Bundle

If you contact Oracle Support about any issue with Storage Gateway, you might need to provide a support bundle to help the Oracle Support technicians diagnose the issue.

1. Log in to the management console.
2. Click **System** in the upper-right corner of the management console.
3. Click **Help**.
4. Click **Download Support Bundle** in **System Logs**.  
You can download and save the support bundle.

### Contents of the Support Bundle

The support bundle contains the following information:

- All of the logs needed for diagnostics.
- Local storage usage information.
- Basic system information such as memory size, Docker version, and the Storage Gateway version.
- A list of file systems.
- Cloud Sync job details.

# CHAPTER 31 Streaming Service Overview

The Oracle Cloud Infrastructure Streaming service provides a fully managed, scalable, and durable storage solution for ingesting continuous, high-volume streams of data that you can consume and process in real time. Streaming can be used for messaging, ingesting high-volume data such as application logs, operational telemetry, web click-stream data, or other use cases in which data is produced and processed continually and sequentially in a publish-subscribe messaging model.



## Note

Streaming is not available in Oracle Cloud Infrastructure Government Cloudrealms.

## Streaming Usage Scenarios

Here are some of the many possible uses for Streaming:

- **Metric and log ingestion:** Use the Streaming service as an alternative for traditional file-scraping approaches to help make critical operational data more quickly available for indexing, analysis, and visualization.
- **Messaging:** Use Streaming to decouple components of large systems. Streaming provides a pull/buffer-based communication model with sufficient capacity to flatten load spikes and the ability to feed multiple consumers with the same data independently. Key-scoped ordering and guaranteed durability provide reliable primitives to implement various messaging patterns, while high throughput potential allows for such a system to scale well.
- **Web/Mobile activity data ingestion:** Use Streaming for capturing activity from websites or mobile apps (such as page views, searches, or other actions users may take). This information can be used for real-time monitoring and analytics, as well as in data warehousing systems for offline processing and reporting.

- **Infrastructure and apps event processing:** Use Streaming as a unified entry point for cloud components to report their life cycle events for audit, accounting, and related activities.

## Streaming Concepts

The following concepts are essential to working with Streaming.

### **STREAM**

A partitioned, append-only log of messages.

### **PARTITION**

A partition is a section of a stream. Partitions allow you to distribute a stream by splitting messages across multiple nodes. Each partition can be placed on a separate machine to allow for multiple consumers to read from a stream in parallel.

### **CURSOR**

A pointer to a location in a stream. This location could be a pointer to a specific offset or time in a partition, or to a groups' current location.

### **MESSAGE**

A Base64-encoded record that is published to a stream.

### **PRODUCER**

An entity that publishes messages to a stream.

### **CONSUMER**

An entity that reads messages from one or more streams.

### **CONSUMER GROUP**

A consumer group is a set of instances which coordinates messages from all of the partitions in a stream. Instances in a consumer group maintain group membership through interaction; lack of interaction for a period of time results in a timeout, removing the instance from the group.

**TOPIC/KEY**

An identifier used to group related messages.

**OFFSET**

The location of a message within a partition. You can use the offset to restart reading from a stream.

**PERMISSIONS**

You can use IAM to set permissions on the following operations: list, get, update, create, and delete streams.

**STREAM ARCHIVE**

You can archive a stream to a bucket in Object Storage.

## How Streaming Works

The Streaming service provides a robust, scalable mechanism that you can use to produce and consume high volumes of data between application components.

Here's how Streaming works: a *producer* publishes *messages* to a *stream*, which is an append-only log. These messages are distributed among the partitions using the message's key.

Streams are divided into a number of *partitions* for scalability. Partitions allow you to distribute a stream by splitting messages across multiple nodes (or brokers). Each partition can be placed on a separate machine to allow multiple consumers to read a stream in parallel. Multiple consumers can read from any partition regardless of where the partition is hosted.

A *consumer* can read messages from one or more streams. Each message within a stream is marked with an offset value, so a consumer can pick up where it left off if it is interrupted.

You can use the Streaming service by:

- Creating a stream using the Console or API.
- Using a producer to publish data to the stream.
- Building consumers to read and process messages from a stream using the [GetMessages](#) API .
- Archiving the stream to an Object Storage bucket.

### Stream Archiving

You can archive a stream to an OCI Object Storage bucket. You can set up a stream archiver to start archiving from the latest position in the stream, or from the oldest position in the stream.

For information on setting up a stream archiver, see [Managing Streams](#).

### Limits on Streaming Resources

The Streaming service has the following limitations:

- Message retention of up to a maximum 7 days
- Throughput is limited to 1MB per second per partition
- Each partition can handle up to 1MB maximum message size
- Each partition can handle 5 [GetMessages](#) API calls per second
- Each partition can support up a maximum total data write rate of 1MB per second
- Each enterprise tenancy has a limit of 5 partitions (non-enterprise partitions have a limit of 0, but you can request more)

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

## Resource Identifiers

Most types of Oracle Cloud Infrastructure resources have a unique, Oracle-assigned identifier called an Oracle Cloud ID (OCID). For information about the OCID format and other ways to identify your resources, see [Resource Identifiers](#).

# Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

For general information about using the API, see [REST APIs](#).

## Using Streaming

To get started with Streaming, see the following topics:

- For instructions on how to manage streams, see [Managing Streams](#).
- For information about publishing messages to a stream, see [Publishing Messages](#).
- For information on how to consume messages, see [Consuming Messages](#).
- For SDK information, see [Oracle Cloud Infrastructure SDKs](#).

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

For common policies used to authorize Streaming users, see [Common Policies](#).

For in-depth information on granting users permissions for the Streaming service, see [Details for the Streaming Service](#) in the IAM policy reference.

## Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

## Moving Resources to a Different Compartment

You can move streams from one compartment to another. For more information, see [Managing Compartments](#).

*Last edited: 11/26/2019 11:10 AM*

## Managing Streams

This topic describes how to work with streams.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

## CHAPTER 31 Streaming Service Overview

---

For administrators: The policy in [Let streaming users manage streams](#) lets the specified group do everything with streaming and related Streaming service resources.

To set up and use a stream archiver, you must have read access to the stream and write access to the Object Storage. For example:

```
allow service stream-processing to use stream-pull in tenancy
allow service stream-processing to manage objects in tenancy
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for the Streaming service, see [Details for the Streaming Service](#) in the IAM policy reference.

### Using the Console

#### To create a stream

1. Click the **Create Stream** button at the top of the topic list.
2. In the **Create Stream** dialog box, configure your topic:
  - a. **Compartment** : Required. Select a compartment from the drop-down list.
  - b. **Stream Name**: Required. Specify a friendly name for the topic. It does not have to be unique, but it cannot be changed. Avoid entering confidential information.
3. In the **Stream Settings** panel:
  - a. **Retention (in Hours)**: Enter the number of hours (from 24 to 168) to retain messages. The default value is 24.
  - b. **Number of Partitions**: Enter the number of partitions for the stream. The maximum number is based on the limits for your tenancy.
    - i. You can optionally configure the stream settings based on estimated capacity. To do this, click on the **Configure stream settings based on estimated capacity** link to reveal the **Capacity Estimates** panel.
      - i. Fill out the capacity estimates and then click **the Estimate Stream Size** button. This will automatically populate the **Number of**

**partitions** field.

4. **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Submit**.

### To delete a stream

1. Click the check box next to the topics you want to delete and then click **Delete Topic**.
2. Confirm when prompted.

### To produce a test message

1. Click a stream to display the stream details page.
2. Click the **Produce Test Message** button.
3. On the **Test Stream** dialog, enter the text-only message to produce in the **Data** text box.
4. Click the **Produce** button.

### To show recent messages on a stream

1. Click a stream to display the stream details page.
2. In the **Recent Messages** panel, click the **Refresh** button.

### To archive a stream

1. Click a stream to display the stream details page.

2. In the **Stream Archiver** panel, click the **Configure and create archiver** button to display the **Create Archiver** dialog.
3. Choose whether you want to create a new bucket or use an existing bucket for your stream archiver.
4. Type the name of a new bucket or select the name of an existing bucket.
5. Select where you want to start archiving the stream:
  - a. Select **Latest** to start from the most recent position in the stream
  - b. Select **Trim Horizon** to start archiving from the oldest position in the stream
6. You can optionally expand the **Object Storage Settings** panel to configure advanced options for your stream archiver:
  - a. Set the maximum object size (in MB) in the **Object Rollover Size** field.
  - b. Set the maximum timeout before a new object is created in the **Object Rollover Time Threshold** field.
7. Click the **Create** button to display the details page for your new Stream archiver. It will take a few moments for the streaming service to provision the resource for the archiver.
8. Once the resources are provisioned for the Stream archiver, you can press the **Start** button to begin archive the stream.

### To move a stream to a different compartment

1. Click a stream to display the stream details page.
2. Find the stream you want to move in the list, click the the Actions icon (three dots), and then click **Move Resource**.
3. Choose the destination compartment from the list.
4. Click **Move Resource**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage streams:

- [CreateStream](#)
- [ListStreams](#)
- [GetStream](#)
- [UpdateStream](#)
- [DeleteStream](#)

### Publishing Messages

This topic covers how to emit messages to a stream.

#### Overview

To publish messages:

- Create a stream using [CreateStream](#).
- Publish a message to a specified stream using [PutMessages](#).

#### Publishing Messages

Once a stream is created and active, you can publish messages.

A message is composed of a key and a value. Both the key and the value are byte arrays. If you don't explicitly supply a key, one is generated automatically.

The message is published to a partition in the stream. If there is more than one partition, the partition where the message is published is calculated using the message's key. If the key is null, the partition is calculated using a random 16-byte value.

For messages with a null key, do not expect messages with same value to go on the same partition, since the partitioning scheme may change. Passing a null key will put the message

in a random partition. If you want to ensure that messages with the same value go to the same partition, you should use the same key for those messages.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let streaming users manage streams](#) lets the specified group do everything with streaming and related Streaming service resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Streaming Service](#) in the IAM policy reference.

### Using the Console

#### To publish a message

1. Click on the stream that you want to publish on.
2. Click the **Produce Test Message** button.
3. Type the message in the **Data** text box.
4. Click **Produce**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to produce messages:

- [PutMessages](#)

## Consuming Messages

This topic covers how to consume messages from a stream.

### Overview

Consuming messages requires you to:

- Create a cursor
- Use the cursor to read messages

### Using Cursors

A cursor is a pointer to a location in a stream. This location could be a pointer to a specific offset or time in a partition, or to a groups' current location.

To consume messages, use the [CreateCursor](#) or [CreateGroupCursor](#) API to create a cursor on a partition to indicate where to start consuming messages.

There are five supported cursor types:

- `TRIM_HORIZON` - Start consuming from the oldest available message in the stream. Create a cursor at the `TRIM_HORIZON` to consume all messages in a stream.
- `AT_OFFSET` - Start consuming at a specified offset. The offset must be greater than or equal to the offset of the oldest message and less than or equal to the latest published offset.
- `AFTER_OFFSET` - Start consuming after the given offset. This cursor has the same restrictions as the `AT_OFFSET` cursor.
- `AT_TIME` - Start consuming from a given time. The timestamp of the returned message will be on or after the supplied time.
- `LATEST` - Start consuming messages that were published after you created the cursor.

### Consuming Messages

Once you've created a cursor, you can start to consume messages using [GetMessages](#). Each call to [GetMessages](#) returns the cursor to use in the next [GetMessages](#) call. The returned cursor will never be null and expires in 5 minutes. As long as you keep consuming, you should never have to re-create a cursor.

### Consuming Messages using a Consumer Group

Consumers can be configured to consume messages as part of a group. Stream partitions are distributed among members of a group, such that messages from any single partition are only sent to a single consumer.

Partition assignments are rebalanced as consumers join or leave the group. Group consumption is accomplished using the same cursor mechanism as with single consumers, but using a different kind of cursor.

To create a consumer, create a group cursor, providing a group, instance name, and cursor type. Groups are created on the first request to create a cursor; their retention period is the same as their assigned stream:

```
CreateGroupCursorRequest groupRequest = CreateGroupCursorRequest.builder()
 .streamId(streamId)
 .createGroupCursorDetails(CreateGroupCursorDetails.builder()
 .groupName(groupName)
 .instanceName(instanceName)
 .type(CreateGroupCursorDetails.Type.TrimHorizon)
 .commitOnGet(true)
 .build())
 .build();

CreateGroupCursorResponse groupCursorResponse = streamClient.createGroupCursor(groupRequest);
String groupCursor = groupCursorResponse.getCursor().getValue();
// this groupCursor can be used in the same message loop a described above; subsequent getMessages calls
return an updated groupCursor.
```

Once you've created a cursor, you can start to consume messages using [GetMessages](#). Each call to [GetMessages](#) returns the cursor to use in the next [GetMessages](#) call. The returned cursor will never be null and expires in 5 minutes. As long as you keep consuming, you should never have to re-create a cursor.

### Required IAM Policy

To use Oracle Cloud Infrastructure, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

For administrators: The policy in [Let streaming users manage streams](#) lets the specified group do everything with streaming and related Streaming service resources.

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for databases, see [Details for the Streaming Service](#) in the IAM policy reference.

### Using the Console

You cannot use the console to consume messages.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to consume messages:

- [CreateCursor](#)
- [GetMessages](#)

### Using the Streaming SDK

This topic covers how to use the Streaming SDK.

### Getting Started

For information on installing and configuring the Oracle Cloud Infrastructure SDKs, see [Developer Tools](#).

### Architecture Overview

Streaming contains the following key building blocks:

- **Stream:** A partitioned, append-only log of messages.
- **Partition:** A section of a stream. Partitions allow you to distribute a stream by splitting messages across multiple nodes. Each partition can be placed on a separate machine to allow for multiple consumers to read from a topic in parallel.
- **Producer:** An entity that publishes messages to a stream.
- **Consumer:** An entity that reads messages from a stream.
- **Consumer group:** A group of consumers that can read independently read messages from separate partitions of a stream.

### Streaming Clients

The Streaming SDK is encapsulated in two clients: the `StreamAdminClient` and the `StreamClient`.

The `StreamAdminClient` incorporates the control plane operations of the streaming service. You can use it to create, delete, update, modify, and list streams.

To instantiate the `StreamAdminClient` object:

```
StreamAdminClient adminClient = new StreamAdminClient([authProvider]);
adminClient.setEndpoint("https://streaming.r2.oracleiaas.com"); // You cannot use the setRegion method
```

The `StreamClient` is used to publish and consume messages.

To instantiate a `StreamClient` object:

## CHAPTER 31 Streaming Service Overview

---

```
// First you have to get the stream you want to consume/publish.
// You can either make a CreateStream, GetStream, or ListStream call. They all return a
"messagesEndpoint" as part of a Stream object.
// That endpoint NEEDS to be used when creating the StreamClient object.
GetStreamRequest getStreamRequest = GetStreamRequest.builder().streamId(streamId).build();
Stream stream = adminClient.getStream(getStreamRequest).getStream();

StreamClient streamClient = new StreamClient([authProvider]);
streamClient.setEndpoint(stream.getMessagesEndpoint());
```

### Creating a Stream

To create a stream, use the `createStream` method of `StreamAdminClient`. Creating a stream is an asynchronous operation. You can check on the completion of the create operation by checking that the `lifecycleStateDetails` property of your new stream is either `Active` or `Failed`.

The following is an example showing how to create a stream:

```
// No error handling
CreateStreamDetails createStreamDetails = CreateStreamDetails.builder()
 .partitions(5) // number of partitions you want in your stream

 .name("myStream") // the name of the stream - only used in the console
 .compartmentId(tenancy) // the compartment id you want your stream to live in
 .build();

// You can also add tags to the createStreamDetails object.
CreateStreamRequest createStreamRequest =
 CreateStreamRequest.builder()
 .createStreamDetails(createStreamDetails)
 .build();

Stream stream = adminClient.createStream(createStreamRequest).getStream();

while (stream.getLifecycleState() != Stream.LifecycleState.Active && stream.getLifecycleState() !=
Stream.LifecycleState.Failed) {

 GetStreamRequest getStreamRequest = GetStreamRequest.builder().streamId(stream.getId()).build();
 stream = adminClient.getStream(getStreamRequest).getStream();
}

// Handle stream Failure
```

### Deleting a Stream

To delete a stream, use the `deleteStream` method API of the `StreamAdminClient`. Deleting a stream is an asynchronous operation; the stream state changes to `Deleted` once the delete operation is finished. During the deletion process, the stream can't be used for consuming or producing messages.

The following example shows how to use the `deleteStream` method to delete a stream:

```
// No error handling
DeleteStreamRequest deleteStreamRequest =
 DeleteStreamRequest.builder()
 .streamId(stream.getId())
 .build();
adminClient.deleteStream(deleteStreamRequest);
```

### Listing Streams

Use the `listStreams` method to return a list of streams for a given compartment.

You can filter the returned list by OCID, life cycle state, and name.

The results can be sorted in ascending or descending order by name or creation time.

The results are passed back in a paginated list. A token is passed back with each page of results; pass this token back to the `getOpcNextPage` method to retrieve the next page of results. A null token returned from `getOpcNextPage` indicates that no more results are available.

For example:

```
// No error handling
ListStreamsRequest listStreamsRequest =
 ListStreamsRequest.builder()
 .compartmentId(tenancy)
 .build();
// You can filter by OCID (exact match only) [builder].id(streamId) -> This will return 0..1 item
// You can filter by name (exact match only) [builder].name(name) -> This will return 0..n items
// You can order the result per TimeCreated or Name [builder].sortBy(SortBy.[TimeCreated|Name])
```

## CHAPTER 31 Streaming Service Overview

---

```
// You can change the ordering [builder].sortOrder(SortOrder.[Asc|Desc])
// You can filter by lifecycleState [builder].lifecycleState(lifecycleState)

String page;
do {
 ListStreamsResponse listStreamsResponse = adminClient.listStreams(listStreamsRequest);
 List<StreamSummary> streams = listStreamsResponse.getItems();
 // Do something with the streams
 page = listStreamsResponse.getOpcNextPage();
} while (page != null);
```

### Retrieving Stream Details

To get details about a stream, use the `getStream` method and then examine the properties of the stream. For example:

```
// No error handling
GetStreamRequest getStreamRequest =
 GetStreamRequest.builder()
 .streamId(streamId)
 .build();

Stream stream = adminClient.getStream(getStreamRequest).getStream();
```

### Publishing Messages

Once a stream is created and active, you can publish messages using the `streamClient.putMessages` method.

A message is composed of a key (which can be null) and a value. Both key and value are byte arrays.

The message is published to a partition in the stream. If there is more than one partition, the partition where the message is published is calculated using the message's key. If the key is null, the partition is calculated using a subset of the value. For messages with a null key, do not expect messages with same value to go on the same partition since the partitioning scheme may change. Sending a null key will put the message in a random partition. If you want to ensure that messages with the same value go to the same partition, you should use the same key for those messages.

## CHAPTER 31 Streaming Service Overview

---

The following code shows how to publish a message:

```
// No error handling
List<PutMessagesDetailsEntry> messages = new ArrayList<>();

for (int i = 0; i < 40; i++) {
 byte[] key = "myKey".getBytes(Charsets.UTF_8); // In that case, all messages will go on the same
partition since the key is the same.
 byte[] value = UUID.randomUUID().toString().getBytes(Charsets.UTF_8);
 messages.add(new PutMessagesDetailsEntry(key, value));
}

PutMessagesDetails putMessagesDetails =
 PutMessagesDetails.builder()
 .messages(messages)
 .build();

PutMessagesRequest putMessagesRequest =
 PutMessagesRequest.builder()
 .putMessagesDetails(putMessagesDetails)
 .build();

PutMessagesResult putMessagesResult = streamClient.putMessages(putMessagesRequest).getPutMessagesResult
();
// It's not because the call didn't fail that the messages were successfully published!
int failures = putMessagesResult.getFailures();
// If failures is > 0, it means we have a partial-success call.
List<PutMessagesResultEntry> entries = putMessagesResult.getEntries();
// entries is a list of the same size as the list of messages you sent.
// It is guaranteed that the order of the messages is the same as when you sent them.
// Each entry contains either "offset/partition/timestamp" if the message was successfully published
// or "error/errorMessage" if it failed.
if (failures != 0) {
 entries.forEach(entry -> {
 if (StringUtil.isEmpty(entry.getError())) {
 // That particular message failed to get published.
 // It could be a throttle error and in that case error would be "429" and errorMessage would
contain a meaningful message.
 // Or it could be an internal error on our side and error would be "500".

 // Possible solution would be to republish only failed messages.
 }
 });
}
```

```
});
}
```

### Consuming Messages

Consuming messages requires the use of a cursor, which is a pointer to an offset into a partition.

There are five supported cursor types:

- **TRIM\_HORIZON** - Start consuming from the oldest available message in the stream. Create a cursor at the TRIM\_HORIZON to consume all messages in a stream.
- **AT\_OFFSET** - Start consuming at a specified offset. The offset must be greater than or equal to the offset of the oldest message and less than or equal to the latest published offset.
- **AFTER\_OFFSET** - Start consuming after the given offset. This cursor has the same restrictions as the AT\_OFFSET cursor.
- **AT\_TIME** - Start consuming from a given time. The timestamp of the returned message will be on or after the supplied time.
- **LATEST** - Start consuming messages that were published after you created the cursor.

To create a TRIM\_HORIZON cursor, which starts consuming starting from the oldest available message:

```
// No error handling
CreateCursorDetails createCursorDetails =
 CreateCursorDetails.builder()
 .type(Type.TrimHorizon)
 .partition("0")
 .build();

// If using AT_OFFSET or AFTER_OFFSET you need to specify the offset [builder].offset(offset)
// If using AT_TIME you need to specify the time [builder].time(new Date(xxx))

CreateCursorRequest createCursorRequest =
 CreateCursorRequest.builder()
 .createCursorDetails(createCursorDetails)
 .build();
```

## CHAPTER 31 Streaming Service Overview

---

```
String cursor = streamClient.createCursor(createCursorRequest).getCursor().getValue();
// Cursor will then be used to get messages from the stream.
```

Once you've created a cursor, you can start to consume messages using the `GetMessages` method. Each call to [GetMessages](#) returns the cursor to use in the next [GetMessages](#) call. The returned cursor is not null and expires in 5 minutes. As long as you keep consuming, you should never have to re-create a cursor.

Here's an example of using a cursor to retrieve messages:

```
// No error handling (there is a high chance of getting a throttling error using a tight loop)
while (true) { // or your own exit condition
 GetMessagesRequest getMessagesRequest =
 GetMessagesRequest.builder()
 .cursor(cursor)
 .build();

 GetMessagesResponse getMessagesResponse = streamClient.getMessages(getMessagesRequest);

 // This could be empty, but we will always return an updated cursor
 getMessagesResponse.getItems().forEach(message -> {
 // Process the message
 });

 cursor = getMessagesResponse.getOpcNextCursor(); Consuming Messages
}
```

### Using Consumer Groups

Consumers can be configured to consume messages as part of a group. Stream partitions are distributed among members of a group, such that messages from any single partition are only sent to a single consumer.

Partition assignments are rebalanced as consumers join or leave the group. Group consumption is accomplished using the same cursor mechanism as with single consumers, but using a different kind of cursor.

### How Consumer Groups Work

A consumer group is a set of instances which coordinate to consume messages from all of the partitions in a stream. Instances maintain membership to a group through interaction; lack of interaction for a period of time results in a timeout, removing the instance from the group. Partitions are reserved for specific instances in a group; reservations are rebalanced in response to specific group events, such as an instance joining the group, or instance time-out.

While an instance maintains membership to a group, calls to get messages will return messages from only partitions reserved for that instance. Partition reservation attempts to ensure that messages are only processed by a single instance.

The set of instances in a group is expected to change over time; transience can occur for many reasons, commonly because of server failure, scaling patterns, or operational management.

Persisting consumer processing state beyond the lifetime of an instance, allows instances to come and go, picking up where appropriate in context of the group.

An instance commits offsets of processed messages to ensure that they are not processed again by other instances in the same group.

### Consuming Messages with a Consumer Group

To create a consumer group, create a group cursor, providing a group, instance name, and cursor type. Groups are created on the first request to create a cursor; their retention period is the same as their assigned stream:

```
CreateGroupCursorRequest groupRequest = CreateGroupCursorRequest.builder()
 .streamId(streamId)
 .createGroupCursorDetails(CreateGroupCursorDetails.builder()
 .groupName(groupName)
 .instanceName(instanceName)
 .type(CreateGroupCursorDetails.Type.TrimHorizon)
 .commitOnGet(true)
 .build())
 .build();

CreateGroupCursorResponse groupCursorResponse = streamClient.createGroupCursor(groupRequest);
String groupCursor = groupCursorResponse.getCursor().getValue();
// this groupCursor can be used in the same message loop a described above; subsequent getMessages calls
return an updated groupCursor.
```

Once you've created a cursor, you can start to consume messages using [GetMessages](#). Each call to [GetMessages](#) returns the cursor to use in the next [GetMessages](#) call. The returned cursor is never null and expires in 5 minutes. As long as you keep consuming, you should never have to re-create a cursor.

## Streaming Metrics

This topic describes the metrics emitted by the Streaming service using the metric namespace `oci_oss`.

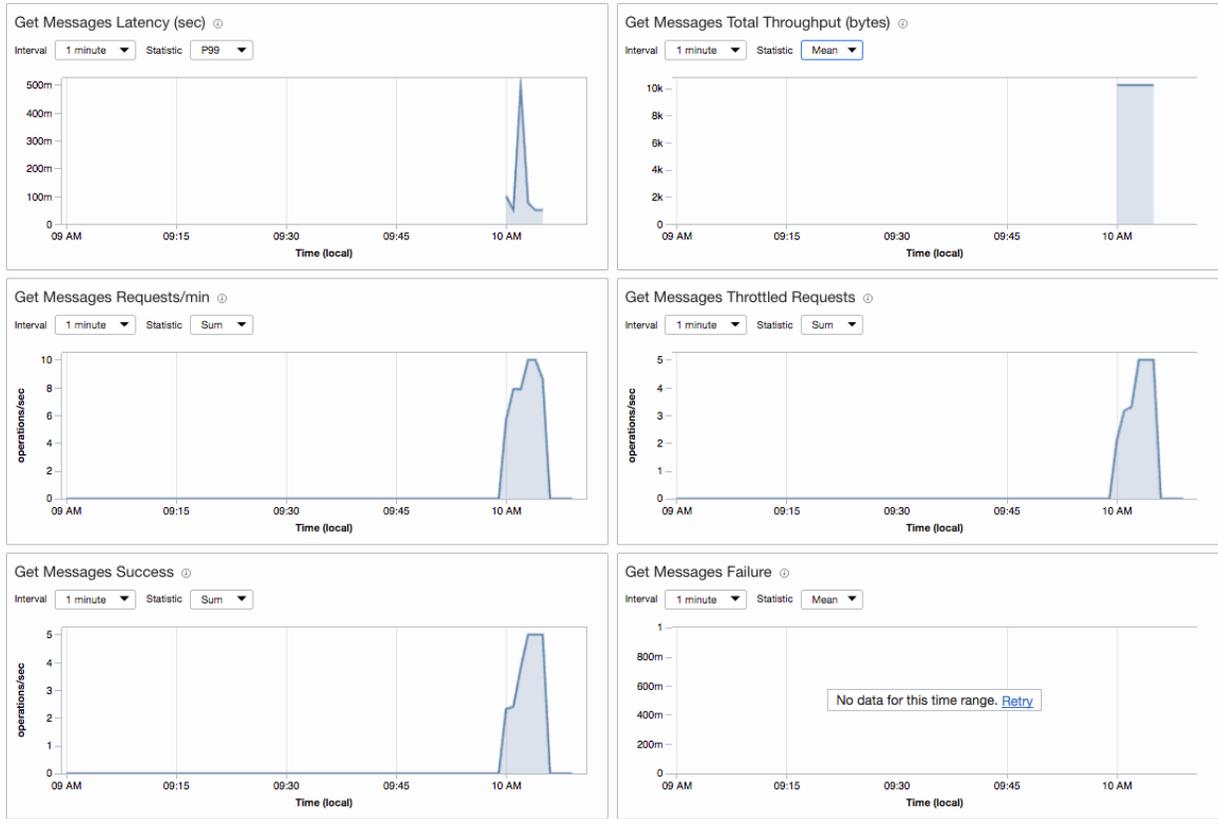
### Overview of Streaming Metrics

The Streaming service provides metrics showing how the service is performing. These metrics are automatically available.

You can use these metrics to:

- Understand the produce/consume latency for a real-time application.
- Calculate and validate the price of service usage.
- Monitor changes in throughput over time.
- Check the time that the last message was consumed.

## CHAPTER 31 Streaming Service Overview



To view a default set of metrics charts in the Console, navigate to the Service Metrics page and then select the `oci_oss` metric namespace.

### Terminology

The following terms are used when discussing Streaming service metrics:

- **Namespace:** A namespace is a container for Streaming metrics. The namespace identifies the application or service sending the metrics. The namespace for the Streaming is `oci/oss`.
- **Metrics:** Metrics are the fundamental concept in telemetry and monitoring. Metrics define a time-series set of datapoints. Each metric is uniquely defined by namespace,

metric name, compartment identifier, and a set of one or more dimensions, and a unit of measure. Each datapoint has a time stamp, a value, and a count associated with it.

- *Dimensions*: A dimension is a key-value pair that defines the characteristics associated with the metric; for example, `resourceId(streamOcid.`
- *Statistics* : Statistics are metric data aggregations over specified periods of time. Aggregations are done using the namespace, metric name, dimensions, and the data point unit of measure within the time period specified.
- *Alarms*: Alarms are used to automate operations monitoring and performance. An alarm keeps track of changes that occur over a specific period of time and performs one or more defined actions, based on the rules defined for the metric.

### Available Metrics

The following tables describe the available Streaming metrics.

#### **Producers**

## CHAPTER 31 Streaming Service Overview

Metric	Metric Display Name	Unit	Description	Dimensions
PutMessagesLatency.Time	Put Messages Latency	time (ms)	Time taken for put messages operation measured over time range	streamOcid,regionId
PutMessagesThroughput.Bytes	Put Messages Total Throughput	bytes	Bytes pushed to the stream measured over time	
PutMessagesThroughput.Count	Put Messages Records/second	count	Count of messages pushed to stream measured over time	
PutMessagesThrottling.Count	Put Messages Throttled Records/second	count	Number of put messages throttled either due to volume or requests measured over time	

## CHAPTER 31 Streaming Service Overview

Metric	Metric Display Name	Unit	Description	Dimensions
PutMessagesSuccess.Count	Put Messages Success	count	Successful Requests for put messages per stream measured over time	
PutMessagesFault.Count	Put Messages Failure	count	Total Failed putMessage Requests per stream measured over time	
PutMessagesRecords.Count	Put Messages Requests	count	Number of records published to a stream measured over time	
PutMessages.Bytes	Put Messages Bytes	bytes	Bytes pushed to a stream over time	
PutMessages.Count	Put Messages Count	count	Number of messages pushed over time	

## CHAPTER 31 Streaming Service Overview

---

### Consumers

## CHAPTER 31 Streaming Service Overview

Metric	Metric Display Name	Unit	Description	Dimensions
GetMessagesLatency.Time	Get Messages Latency	time (ms)	Time taken for get messages operation measured over time range	streamOcid,regionId
GetMessagesThroughput.Bytes	Get Messages Total Throughput	bytes	Bytes retrieved from stream measured over time	
GetMessagesThroughput.Count	Get Messages Requests/second	count	Count of messages read from stream measured over time	
GetMessagesThrottling.Count	Get Messages Throttled Requests	count	Number of get messages throttled either due to volume or requests measured over time	

## CHAPTER 31 Streaming Service Overview

Metric	Metric Display Name	Unit	Description	Dimensions
GetMessagesSuccess.Count	Get Messages Success	count	Successful Requests for get messages per stream measured over time	
GetMessagesFault.Count	Get Messages Failure	count	Total Failed getMessage Requests per stream measured over time	
GetMessages.Bytes	Get Messages Bytes	bytes	Bytes retrieved from a stream over time	
GetMessages.Count	Get Messages Count	count	Number of messages read over time	

### Consumer Groups

## CHAPTER 31 Streaming Service Overview

---

<b>Metric</b>	<b>Metric Display Name</b>	<b>Unit</b>	<b>Description</b>	<b>Dimensions</b>
<code>totalActiveConsumers.Count</code>	Total Active Consumer Groups	count	Number of active consumers.	streamOcid,regionId
<code>totalActiveConsumersGroups.Count</code>	Total Active Consumers	count	Number of active consumer groups.	

# CHAPTER 32 Tagging

This chapter explains how to use tags to add metadata to your resources.

## Tagging Overview

Oracle Cloud Infrastructure Tagging enables you to add metadata to resources. This allows you to define keys and values and associate them with resources. You can then use the tags to help you organize and list resources based on your business needs.

## How Tagging Works

The Tagging service provides two ways for you to approach adding tags to resources. Each approach offers a different type of tag for you to work with:

- Defined tags - tag administrators manage resource metadata.
- Free-form tags - unmanaged metadata applied to resources by users.

One approach involves a tag administrator creating and managing all the tags that users will apply to resources. You use IAM policy to limit who can *create* tags to a few select tag administrators, while granting all others in the tenancy only the ability to *apply* tags. The benefit to this approach is that you can create and manage the keys and values used to tag resources. This avoids typos that weaken automation based on tags and provides better reporting based on tags.

The other approach is to allow users to add tags to resources. Each tag is edited or applied at the resource by you or a user creating or modifying a resource.

You can use both types of tags throughout your tenancy. Most of the Tagging features require defined-tags. "Tag" is used generically to refer to defined tags. To create metadata that you can trust to manage resources and collect data, use defined tags. With defined tags, the following scenarios become possible:

- You can create default tags that are applied to all resources in compartments. See [Managing Tag Defaults](#).
- Specify that users must apply tags to resources to successfully create resources in compartments.
- If you make a typo using defined tags, you can correct it by editing or even deleting. When you delete a defined tag, Oracle removes the key (and any value) for that tag from all resources. See [Deleting Tag Key Definitions and Namespace](#).
- Associate a list of predefined values for a defined tag. See [Using Predefined Values](#).
- Use system variables to automatically generate values for defined tags or tag defaults. See [Using Tag Variables](#).
- Track costs based on defined tags. See [Using Cost-Tracking Tags](#).

## Tagging Concepts

Here's a list of the basic tagging concepts:

### **TAG NAMESPACE**

You can think of a tag *namespace* as a container for your tag keys. It consists of a name, and zero or more tag key definitions. Tag namespaces are not case sensitive, and must be unique across the tenancy. The namespace is also a natural grouping to which administrators can apply policy. One policy on the tag namespace applies to all the tag definitions contained in it.

### **TAG KEY**

The name you use to refer to the tag. Tag keys are case insensitive (for example, "mytagkey" duplicates "MyTagKey"). Tag keys for defined tags must be created in a namespace. A tag key must be unique within a namespace.

### **TAG VALUE TYPE**

The tag value type specifies the data type allowed for the value. Currently two data types are supported: string and a list of strings.

### **KEY DEFINITION**

A key definition defines the schema of a tag and includes a namespace, tag key, and tag value type.

### **TAG VALUE**

The tag value is the value the user applying the tag adds to the tag key. Tag values support two data types: strings and lists of strings. You can define a list of values for the user to select from when you define the tag key, or you can allow the user to enter any value when the tag is applied to the resource. If you select a string tag value when you create the key, the user can leave the value blank when they apply the key.

In the example:

```
Operations.CostCenter="42"
```

Operations is the namespace, CostCenter is the tag key, and 42 is the tag value.

### **TAG (OR DEFINED TAG)**

A tag is the instance of a key definition that is applied to a resource. It consists of a namespace, a key, and a value. "Tag" is used generically to refer to defined tags.

### **FREE-FORM TAG**

A basic metadata association that consists of a key and a value only. Free-form tags have limited functionality. See [Understanding Free-form Tags](#).

### **COST TRACKING**

Cost tracking is a feature supported for defined tags. Tags enabled for cost tracking (also called cost-tracking tags) help you to manage costs in your tenancy. You can use cost-tracking tags to track projected costs and set budgets and receive alerts when your costs reach a particular threshold. See [Using Cost-Tracking Tags](#).

### **TAG DEFAULT**

Tag defaults let you specify tags to be applied automatically to all resources, in a specific compartment, at the time of creation, regardless of the permissions of the user who creates the resource. See [Managing Tag Defaults](#).

### **RETIRE**

You can *retire* a tag key definition or a tag namespace. Retired tag namespaces and key definitions can no longer be applied to resources. However, retired tags are not removed from the resources to which they have already been applied. You can still specify retired tags when searching, filtering, reporting, and so on.

### **REACTIVATE**

You can *reactivate* a tag namespace or tag key definition that has been retired to reinstate its usage in your tenancy.

### **TAG VARIABLE**

You can use a variable to set the value of a tag. When you add or update a tag on a resource, the variable resolves to the data it represents. See [Using Tag Variables](#).

### **PREDEFINED VALUES**

You can use a variable to set the value of a tag. When you add or update a tag on a resource, the variable resolves to the data it represents. See [Using Predefined Values](#).

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

For administrators: Use the following topics to find example of IAM policy for Tagging:

## CHAPTER 32 Tagging

---

- [Required Permissions for Working with Defined Tags](#)
- [Required Permissions for Working with Tag Defaults](#)
- [Required Permissions for Working with Free-form Tags](#)

### Region Availability

Tagging is currently available in all regions.

### Ways to Access Oracle Cloud Infrastructure

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface) or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide. For a list of available SDKs, see [Software Development Kits and Command Line Interface](#).

To access the [Console](#), you must use a supported browser. Oracle Cloud Infrastructure supports the latest desktop versions of Google Chrome, Microsoft Edge, Internet Explorer 11, Safari, Firefox, and Firefox ESR. Note that private browsing mode is not supported for Firefox, Internet Explorer, or Edge. Mobile browsers are not supported.

### Limits on Tags

See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase.

- Tag namespaces: 100 tag namespaces per tenancy
- Tags per Tag namespace: 100 tags per tag namespace
- Tags per resource: 10 free-form tags and 64 defined tags
- Tags enabled for cost-tracking: 10 per tenancy (includes both active and retired tags)
- Total tag data size: 5 K (JSON). The total tag data size includes all tag data for a single resource (all applied tags and tag values). Sizing is per UTF-8.

Resource	Supported Characters	Max Length
Tag namespace	Printable ASCII, excluding periods (.) and spaces	100 characters
Tag key name (free-form and defined)	Printable ASCII, excluding periods (.) and spaces	100 characters
Tag value (free-form and defined)	Unicode characters	256 characters

## Managing Tags and Tag Namespaces

Oracle Cloud Infrastructure supports two kinds of tags: free-form tags and defined tags.



### Tip

Watch a video to introduce you to the concepts and features of tagging: [Introduction to Tagging](#).

## Required IAM Policy

If you're in the Administrators group, then you have the required access for managing tag namespaces and tags. For more policy samples specific to working with tags and tag namespaces, see [Required Permissions for Working with Defined Tags](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

## Overview of Tags and Tag Namespaces

Defined tags provide more features and control than free-form tags. Before you create a defined tag key, you first set up a tag *namespace* for it. You can think of the tag namespace as a container for a set of tag keys. When you create the tag key definition, you must choose the type of value (which also determines how the user applying the tag adds the value):

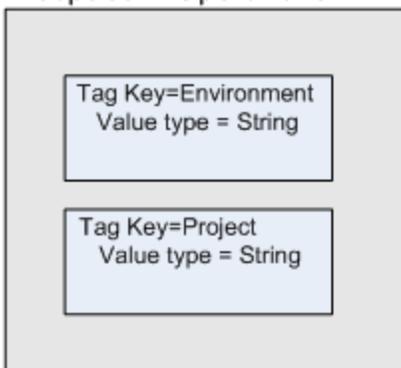
- You can leave it empty so that a user can fill in the value
- You can create a list of values so that the user must choose from those values

To apply a defined tag to a resource, a user first selects the tag namespace, then the tag key within the namespace, and then they can assign the value. If the tag key contains a blank value, the user can type in a value or leave it blank. If the tag key contains a list, the user must select a value from the list.

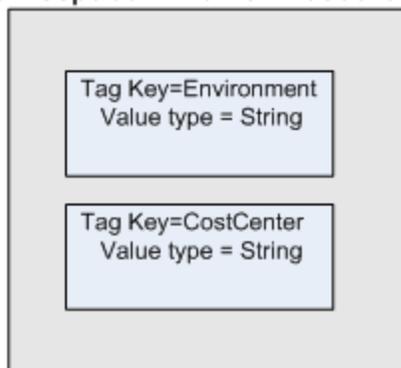
Defined tags support policy to allow you to control who can apply your defined tags. The tag namespace is the entity to which you can apply policy. Administrators can control which groups of users are allowed to use each namespace.

The following diagrams illustrate defined tags. Two tag namespaces are set up: Operations and HumanResources. The tag keys are defined in the namespaces. Within each namespace, the tag keys must be unique, but a tag key name can be repeated across namespaces. In the example, both namespaces include a key named "Environment."

Namespace = Operations



Namespace = HumanResources



## CHAPTER 32 Tagging

---

The first instance is tagged with two tags from the Operations tag namespace, indicating that it belongs to the Operations production environment and the Operations project "Alpha". The second instance is tagged with tags from both the HumanResources tag namespace and the Operations tag namespace, indicating that it belongs to the HumanResources "Production" environment, the HumanResources cost center "42", and also the Operations project "Beta".



### Working with Defined Tags

You must set up the tag namespace and tag keys in your tenancy before users can apply a defined tag to a resource. As part of the set up process, you must also grant permissions to the user groups that will need to use the namespace.

Features of defined tags include:

- Consist of a tag namespace, a key, and a value.
- The tag namespace and tag key definition must be set up in your tenancy before users can apply a defined tag to a resource.
- You can create the tag key with either a tag value type of string (for the user to add a value or leave blank) or a list of values (from which the user must choose).
- When applying a defined tag, users select from the list of tag keys.
- Users can apply a defined tag during resource creation or to an existing resource.

## CHAPTER 32 Tagging

---

- Defined tag keys are case insensitive.
- Defined tag values are case sensitive. For example, "alpha" and "Alpha" are distinct values.
- You can use tag variables.
- You can create a list of predefined variables to associate with a tag key. Users that apply the tag to a resource must select a value from the list you create.

### Required Permissions for Working with Defined Tags

To apply, update, or remove defined tags for a resource, a user must be granted permissions on the resource *and* permissions to use the tag namespace.

Users must be granted `use` access on the *tag namespace* to apply, update, or remove a defined tag for a resource.

Some example policies for tag namespaces:

To allow a group to simply view the tag namespaces in the tenancy (or in a compartment) requires `inspect` access:

```
Allow group GroupA to inspect tag-namespaces in tenancy
```



#### Important

To apply tags to a resource when using the Console, a user must have permissions to `inspect tag-namespaces in tenancy`. If the user does not have this permission, the list of tag namespaces cannot be presented to the user in the dialog menu.

To allow a group to read the tag definitions contained in tag namespaces requires `read` access:

```
Allow group GroupA to read tag-namespaces in tenancy
```

## CHAPTER 32 Tagging

---

To allow a group to apply, update, or remove a defined tag for a resource requires the `use` access on the tag namespace:

```
Allow group GroupA to use tag-namespaces in tenancy
```

To allow usage of a specific tag namespace or namespaces, use a `where` clause with the `target.tag-namespace.name` variable. For example:

```
Allow group GroupA to use tag-namespaces in tenancy where target.tag-namespace.name='Operations'
```

or to specify multiple tag namespaces:

```
Allow group GroupA to use tag-namespaces in tenancy where any {target.tag-namespace.name='Operations', target.tag-namespace.name='HumanResources'}
```

To manage tag namespaces and the tag definitions in them, requires `manage` access:

```
Allow group GroupA to manage tag-namespaces in tenancy
```

In addition to the permissions to work with the tag namespace, to apply or remove defined tags on a resource you must have the update permission for the resource. For many resources, the update permission is granted with the `use` verb. For example, users who can [use instances](#) in `CompartmentA`, can also apply, update, or remove defined tags for instances in `CompartmentA`.

Some resources don't include the update permission with the `use` verb. To allow a group to apply, update, or remove defined tags for these resources without granting the full permissions of `manage`, you can add a policy statement to grant only the `<RESOURCE>_UPDATE` permission from the `manage` verb. For example, to allow a group `NetworkUsers` to work with defined tags with VCNs in `CompartmentA`, you could write a policy like:

```
Allow group NetworkUsers to use vcns in compartment CompartmentA
```

```
Allow group NetworkUsers to manage vcns in compartment CompartmentA where request.permission='VCN_UPDATE'
```

The `inspect` permission for a resource grants permissions to view defined tags for that resource. For example, users who can `inspect` instances can also view any defined tags applied to the instance.

For information about resource permissions, see [Policy Reference](#).

### Example Scenario

Your company has an Operations department. Within the Operations department are several cost centers. You want to be able to tag resources that belong to the Operations department with the appropriate cost center.

1. Create a tag namespace definition called Operations.
2. In the Operations namespace, create a tag key definition called CostCenter.

Alice already belongs to the group [InstanceLaunchers](#). Alice can manage instances in CompartmentA. You want Alice and other members of the InstanceLaunchers group to be able to apply the Operations.CostCenter tag to instances in CompartmentA.

To grant the InstanceLaunchers group access to the Operations tag namespace (and only the Operations tag namespace), add the following statements to the [InstanceLaunchers policy](#):

```
Allow group InstanceLaunchers to use tag-namespaces in compartment CompartmentA where target.tag-namespace.name='Operations'
```

Alice can now apply the Operations.CostCenter tag to resources in CompartmentA.

### Retiring Key Definitions and Namespace Definitions

You can retire a tag key definition or a tag namespace definition.

When you retire a tag key definition, you can no longer apply it to resources. However, the tag is not removed from the resources that it was applied to. The tag still exists as metadata on those resources and you can still call the retired tag in operations (such as listing, sorting, or reporting).



### Important

Retiring a tag stops cost tracking for the tag, but if you do not disable the cost-tracking option on the tag key definition, the retired tag continues to count against your maximum of 10 cost-tracking tags for your tenancy. Ensure that you disable cost tracking before you retire the tag key definition. To disable cost-tracking after a tag is retired, you must reactivate the tag key definition to update it. You can't update tag key definitions that are in the retired state.

When you retire a tag namespace, all the tag keys in the tag namespace are retired. As described earlier, this means that all tags in the tag namespace can no longer be applied to resources, though existing tags are not removed. No new keys can be created in the retired tag namespace.

## Reactivating Tag Key Definitions and Namespace Definitions

You can reactivate retired tag key definitions and tag namespace definitions.

When you reactivate a tag key, it is again available for you to apply to resources.

When you reactivate a tag namespace, you can once again create tag key definitions in the namespace. However, if you want to use any of the tag key definitions that were retired with the namespace, you must explicitly reactivate each tag key definition.

## Moving Tag Namespaces to a Different Compartment

You can move a tag namespace to a different compartment. The tag namespace can be active or retired when you move it. When you move the tag namespace, all its tag key definitions are moved along with it.

This functionality is useful if you need to reorganize your compartment hierarchy, or if you need to delete a compartment that contains a retired tag namespace. Remember that you can't delete a compartment that still contains resources. A retired tag namespace, even though it is retired, is still an existing resource. Moving the retired tag namespace to a different compartment can enable you to delete its original containing compartment.

To move a tag namespace, you must be allowed to `manage tag-namespaces` in both compartments.

See the procedure [To move a tag namespace to a different compartment](#).

### Deleting Tag Key Definitions and Namespace

You can delete a tag key definition or a tag namespace.

When you delete a tag key definition, you begin a process that removes the tag from all resources in your tenancy. These things happen immediately:

- If the tag was a cost-tracking tag, it no longer counts against your 10 cost-tracking tags limit, whether you first disabled it or not.
- If the tag was used with dynamic groups, none of the rules that contain the tag will be evaluated against the tag.

The delete action is asynchronous and initiates a work request. Once you start the delete operation, the state of the tag changes to deleting and tag removal from resources begins. This process can take up to 48 hours depending on the number of resources that were tagged as well as the regions in which those resources reside. When all tags have been removed, the state changes to deleted. You cannot restore a deleted tag. Once the deleted tag changes its state to Deleted, you can use the same tag name again.

To delete a tag key definition, you must first retire it.

To delete a tag namespace, you must first delete all its tag key definitions.



### Tip

To delete a tag namespace that contains tag key definitions, first retire the tag namespace. When you retire a tag namespace, all the tag keys in the tag namespace are retired.

## Using the Console



### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

## Managing Tag Namespaces

### To create a tag namespace

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click **Create Namespace Definition**.
3. Enter the following:
  - **Create in Compartment:** The compartment in which you want to create the namespace definition.

- **Namespace Definition Name:** A unique name for this set of tags. The name must be unique within your tenancy. Tag namespace is case insensitive. You cannot change this value later.
  - **Description:** A friendly description. You can change this value later if you want to.
4. Click **Create Namespace Definition**.

### To update a tag namespace's description

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your tenancy is displayed.
2. Click the tag namespace you want to update.  
The namespace's details are displayed. The description is displayed under the namespace's name.
3. Click the pencil next to the description.
4. Edit the description and save it.

### To retire a tag namespace

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace you want to retire.  
The namespace's details are displayed.
3. Click **Retire Namespace**.
4. Confirm when prompted.

### To reactivate a tag namespace

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace you want to reactivate.  
The namespace's details are displayed.
3. Click **Reactivate Namespace**.
4. Confirm when prompted.

### To move a tag namespace to a different compartment

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace you want to move.  
The namespace's details are displayed.
3. Click **Move Tag Namespace**.
4. Select the **Target Compartment** that you want to move the tag namespace to.
5. Click **Move Tag Namespace**.

### To delete a tag namespace

To delete a tag namespace, you must first [delete](#) all of its tag key definitions.

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the empty tag namespace you want to delete.

The namespace's details are displayed.

3. Click **Delete Tag Namespace**.
4. Confirm when prompted.

## Managing Key Definitions

### To create a key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.

2. Click the tag namespace you want to add the tag key definition to.  
A list of the tag key definitions that belong to the namespace is displayed.

3. Click **Create Tag Key Definition**.

4. Enter the following:

- **Tag Key:** Enter the key. The key can be up to 100 characters in length. Tag keys are case insensitive and must be unique within the tag namespace.
- **Description:** Enter a friendly description.
- **Cost-tracking:** Select the check box to enable this tag for cost tracking. You have a limit of 10 cost-tracking tags in your tenancy. For more information, see [Using Cost-Tracking Tags](#).

5. Under **Tag Value Type**, choose one of the following:

- **Static Value:** Specifies that the user applying the tag can specify any value for this key.
- **A List of Values:** Specifies that the user must apply a value from a list you create. When you select this option, the **Values** box appears. Type the values from which the user can select. Separate multiple values with new lines. You must have at least one value. You can't have blank lines or duplicate values.

6. Click **Create Tag Key Definition**.

### To update a tag key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to update.  
A list of the tag key definitions is displayed.
3. Click the tag key definition you want to update.  
The key definition's details are displayed.
4. Click **Edit Tag Key Definition**.  
The edit dialog appears.  
You can change the description, the tag value type, and enable or disable cost tracking.  
If you chose a list of values, the **Values** box appears, and you must add at least one value. You can't have blank lines or duplicate values in the **Values** box.  
You have a limit of 10 cost-tracking tags in your tenancy. For more information, see [Using Cost-Tracking Tags](#).
5. Make your changes and save it.

### To retire a tag key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to retire.  
A list of the tag key definitions is displayed.
3. Click the tag key definition you want to retire.

The tag key definition's details are displayed. If this tag is a cost-tracking tag, disable the cost-tracking flag. If you don't disable cost-tracking, this tag will still count against your tenancy maximum of 10 cost-tracking tags, even after it is retired.

4. Click **Retire Tag Key Definition**.
5. Confirm when prompted.

### To reactivate a tag key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to reactivate.  
A list of the tag key definitions is displayed.
3. Click the tag key definition you want to reactivate.  
The tag key definition's details are displayed.
4. Click **Reactivate Tag Key Definition**.
5. Confirm when prompted.

### To delete a tag key definition

To delete a tag key definition, you must first [retire it](#).

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to delete.  
A list of the tag key definitions is displayed.
3. Click the retired tag key definition you want to delete.  
The tag key definition's details are displayed.

4. Click **Delete Tag Key Definition**.
5. Confirm when prompted.  
After you click OK, the state changes to Deleting. Track the progress of the operation using the work request.

### To track the progress of the delete operation

- a. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
- b. Click the tag namespace that includes the tag key definition you deleted.  
A list of the tag key definitions is displayed.
- c. Click **Work Requests**.  
The work requests for deleted tag definitions are displayed.

### To monitor a work request for a deleted tag

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you deleted.  
A list of the tag key definitions is displayed.
3. Click **Work Requests**.  
The work requests for deleted tag definitions are displayed.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

## CHAPTER 32 Tagging

---

Use these API operations to manage tag namespaces:

- [GetTagNamespace](#)
- [ListTagNamespaces](#)
- [CreateTagNamespace](#)
- [UpdateTagNamespace](#) - use to retire or reactivate a tag namespace
- [DeleteTagNamespace](#)
- [ChangeTagNamespaceCompartment](#)

Use these API operations to manage tag key definitions:

- [GetTag](#) - gets the tag key definition
- [ListTags](#) - lists tag key definitions
- [ListCostTrackingTags](#) - lists the tags that have been enabled for cost-tracking (can be performed in the root compartment only)
- [CreateTag](#) - creates a tag key definition
- [UpdateTag](#) - updates the tag key definition (use this operation to retire or reactivate a tag key)
- [DeleteTag](#) - deletes the tag key definition

Use these API operations to manage work requests spawned by the DeleteTag operation:

- [ListTaggingWorkRequests](#)
- [ListTaggingWorkRequestErrors](#)
- [ListTaggingWorkRequestLogs](#)
- [GetTaggingWorkRequest](#)

## Understanding Free-form Tags

Oracle Cloud Infrastructure supports two kinds of tags: free-form tags and defined tags. This topic describes free-form tags.

## CHAPTER 32 Tagging

---

Because free-form tags are limited in functionality, Oracle recommends that you only use them when you are first getting started with tagging, to try out the tagging feature in your system. For more information about the features and limitations of free-form tags, see [Working with Free-form Tags](#).

### Required IAM Policy

If you're in the Administrators group, then you have the required access for free-form tags. For more policy samples specific to working with free-form tags, see [Required Permissions for Working with Free-form Tags](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

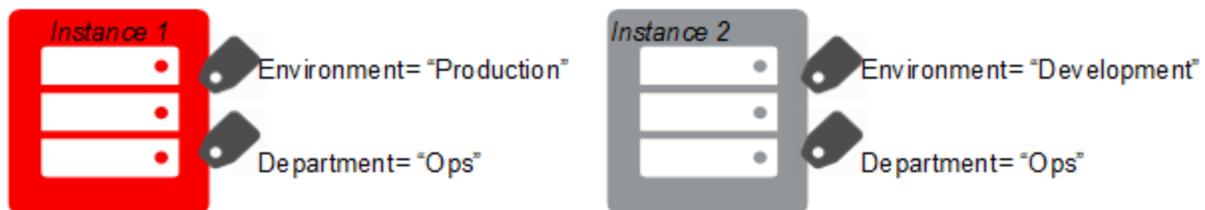
### Overview of Free-form Tags

Free-form tags consist simply of a key and a value, for example:

Environment: Production

where "Environment" is the key and "Production" is the value.

You can apply multiple free-form tags to a single resource (up to the [limit](#)).



### Working with Free-form Tags

Free-form tags consist simply of a key-value pair. Free-form tags have limited features. To experience the full feature set of tagging, use [defined tags](#).

Features of free-form tags include:

## CHAPTER 32 Tagging

---

- Consist of a key and a value. Free-form tags do not belong to a namespace.
- You can apply free-form tags during resource creation or to an existing resource.
- Free-form tag keys are case sensitive. For example, "Project" and "project" are distinct keys.
- Free-form tag values are case sensitive. For example, "alpha" and "Alpha" are distinct values.

Limitations of free-form tags include:

- When applying a free-form tag, you can't see a list of existing free-form tags, so you don't know what tags and values have already been used.
- You can't see a list of existing free-form tags in your tenancy.
- You can't use free-form tags to control access to resources (that is, you can't include free-form tags in IAM policies).
- You can't use tag variables in free-form tags.
- You can't use predefined values in free-form tags.

### REQUIRED PERMISSIONS FOR WORKING WITH FREE-FORM TAGS

To apply, update, or remove free-form tags for a resource, you must have the update permission on the resource. For many resources, the update permission is granted with the `use` verb. For example, users who can [use instances](#) in `CompartmentA`, can also apply, update, or remove free-form tags for instances in `CompartmentA`.

Some resources don't include the update permission with the `use` verb. To allow a group to apply, update, or remove free-form tags for these resources without granting the full permissions of `manage`, you can add a policy statement to grant only the `<RESOURCE>_UPDATE` permission from the `manage` verb. For example, to allow a group `NetworkUsers` to work with free-form tags with VCNs in `CompartmentA`, you could write a policy like:

```
Allow group NetworkUsers to use vcn in compartment CompartmentA
```

```
Allow group NetworkUsers to manage vcn in compartment CompartmentA where request.permission='VCN_UPDATE'
```

The `inspect` verb for a resource grants permissions to view free-form tags for that resource. So users who can `inspect` instances in `CompartmentA` can also view any free-form tags applied to the instance.

For information about resource permissions, see [Policy Reference](#).

## Using Cost-Tracking Tags

You can use cost-tracking tags to help manage costs in your tenancy. Use cost-tracking tags to do any of the following:

- Filter projected costs
- Set budgets

You can only use cost-tracking tag with defined tags. You cannot specify free-form tags as cost-tracking tags.

## Required IAM Policy

Cost tracking is a feature of defined tags. To allow users to work with cost tracking, use the same IAM policy for working with managing tag namespaces and tags. For more information, see [Required Permissions for Working with Defined Tags](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

## Working with Cost-Tracking Tags

Suppose you have a defined tag key definition called `Finance.CostCenter`. You enable this tag definition for cost tracking. You apply the tag with a value of `"W1"` (`Finance.CostCenter="W1"`) to some resources, and you apply the tag with a value of `"C2"` (`Finance.CostCenter="C2"`) to some other resources. These tags enable the following scenarios:

- When you view your Cost Analysis, you can filter the usage information to show you only the costs generated by the resources tagged with "Finance.CostCenter=W1" . You can then filter your usage summary by "Finance.CostCenter=C2".
- Create a budget for resources tagged "Finance.CostCenter=W1" and another budget for resources tagged "Finance.CostCenter=C2". If spending surpasses a certain amount or is forecast to exceed a particular threshold, you can set up alerts that notify you.

Use cost-tracking tags to compare and track resource usage based on tags, or to set up budgets based on resources grouped by tags, rather than by compartments.

### Viewing Usage by Cost-Tracking Tags

To view usage by filtering projected costs or creating a budget

- Cost Analysis provides easy-to-use visualization tools to help you track and optimize your spending. For more information, see [Checking Your Balance and Usage](#).
- Budgets can be used to set thresholds for your spending. You can set alerts on your budget to let you know when you might exceed your budget, and you can view all of your budgets and spending from one single place Console. See [Budgets Overview](#) for more information.

### Limits on Cost-Tracking Tags

- You can have a maximum of 10 tags enabled for cost-tracking in your tenancy at a time. This limitation includes both active and retired tags.

### Using the Console

You can enable cost-tracking when you create a tag key definition, or you can update an existing tag key definition to enable cost tracking.

### To create a key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace you want to add the tag key definition to.  
A list of the tag key definitions that belong to the namespace is displayed.
3. Click **Create Tag Key Definition**.
4. Enter the following:
  - **Tag Key:** Enter the key. The key can be up to 100 characters in length. Tag keys are case insensitive and must be unique within the tag namespace.
  - **Description:** Enter a friendly description.
  - **Cost-tracking:** Select the check box to enable this tag for cost tracking. You have a limit of 10 cost-tracking tags in your tenancy.
5. Under **Tag Value Type**, choose one of the following:
  - **Static Value:** Specifies that the user applying the tag can specify any value for this key.
  - **A List of Values:** Specifies that the user must apply a value from a list you create. When you select this option, the **Values** box appears. Type the values from which the user can select. Separate multiple values with new lines. You must have at least one value. You can't have blank lines or duplicate values.
6. Click **Create Tag Key Definition**.

### To update a tag key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to update.

A list of the tag key definitions is displayed.

3. Click the tag key definition you want to update.

The key definition's details are displayed.

4. Click **Edit Tag Key Definition**.

The edit dialog appears.

You can change the description, the tag value type, and enable or disable cost tracking.

If you chose a list of values, the **Values** box appears, and you must add at least one value. You can't have blank lines or duplicate values in the **Values** box.

You have a limit of 10 cost-tracking tags in your tenancy. .

5. Make your changes and save it.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [ListCostTrackingTags](#) - lists the tags that have been enabled for cost-tracking (can be performed in the root compartment only)

### Using Predefined Values

You can create a list of values and associate that list with a tag key definition. When users apply the tag to a resource, they must select a value from the list. Use lists of predefined values to impose limits on the values that users can apply to tags.

You can use predefined values with defined tags and default tags. You cannot create lists of predefined values for free-form tags.

### Required IAM Policy

Predefined values are a feature of defined tags. To allow users to work with predefined values, use the same IAM policy for working with managing tag namespaces and tags. For more information, see [Required Permissions for Working with Defined Tags](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

### Working with Predefined Values

You can update existing tags to use predefined values.

Every predefined list you create must contain at least one value. Your lists can't contain duplicate values or blank entries. With predefined values, users applying tags can't set the value of a tag to `null`. For more information, see [Using the Console](#).

### Predefined Values and Default Tags

You can use predefined values and default tags to ensure that users creating resources apply tag values you trust.

Here's how it works:

1. You define a list of predefined values for a tag key.
2. You create a default tag that uses the key with the list of predefined values *and* requires that users who create resources in the compartment add the value to the tag.
3. Oracle prompts all users creating resources in the compartment to enter a tag value. Since the tag key contains a predefined list you created, the value the user applies is a value you trust.

These features help to ensure that new resources contain the values you expect. For more information, see [Managing Tag Defaults](#).

### Using the Console



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To create a key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace you want to add the tag key definition to.  
A list of the tag key definitions that belong to the namespace is displayed.
3. Click **Create Tag Key Definition**.
4. Enter the following:
  - **Tag Key:** Enter the key. The key can be up to 100 characters in length. Tag keys are case insensitive and must be unique within the tag namespace.
  - **Description:** Enter a friendly description.
  - **Cost-tracking:** Select the check box to enable this tag for cost tracking. You have a limit of 10 cost-tracking tags in your tenancy. For more information, see [Using Cost-Tracking Tags](#).
5. Under **Tag Value Type**, choose one of the following:
  - **Static Value:** Specifies that the user applying the tag can specify any value for this key.

- **A List of Values:** Specifies that the user must apply a value from a list you create. When you select this option, the **Values** box appears. Type the values from which the user can select. Separate multiple values with new lines. You must have at least one value. You can't have blank lines or duplicate values.
6. Click **Create Tag Key Definition**.

### To update a tag key definition

1. Open the navigation menu. Under **Governance and Administration**, go to **Governance** and click **Tag Namespaces**.  
A list of the tag namespaces in your current compartment is displayed.
2. Click the tag namespace that includes the tag key definition you want to update.  
A list of the tag key definitions is displayed.
3. Click the tag key definition you want to update.  
The key definition's details are displayed.
4. Click **Edit Tag Key Definition**.  
The edit dialog appears.  
You can change the description, the tag value type, and enable or disable cost tracking. If you chose a list of values, the **Values** box appears, and you must add at least one value. You can't have blank lines or duplicate values in the **Values** box.  
You have a limit of 10 cost-tracking tags in your tenancy. For more information, see [Using Cost-Tracking Tags](#).
5. Make your changes and save it.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [CreateTag](#) - creates a tag key definition
- [UpdateTag](#) - updates the tag key definition (use this operation to retire or reactivate a tag key)

## Using Tag Variables

You can use a variable to set the value of a defined tag. When you add the tag to a resource, the variable resolves to the data it represents. You can use tag variables in defined tags and default tags. You cannot use tag variables in free-form tags.

### Required IAM Policy

Tag variables are a feature of defined tags. To allow users to work with tag variables, use the same IAM policy for working with managing tag namespaces and tags. For more information, see [Required Permissions for Working with Defined Tags](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for groups or other IAM components, see [Details for IAM](#).

### Working with Tag Variables

Consider the following example:

```
Operations.CostCenter="{iam.principal.name} at {oci.datetime}"
```

Operations is the namespace, CostCenter is the tag key, and the tag value contains two tag variables `{iam.principal.name}` and `{oci.datetime}`. When you add this tag to a resource, the variable resolves to your user name (the name of the principal that applied the tag) and a time date stamp for when you added the tag.

```
user_name at 2019-06-18T18:00:57.604Z
```

The variable is replaced with data at the time you apply the tag. If you later edit the tag, the variable is gone and only the data remains. You can edit the tag value in all the ways you would edit any other tag value.

To create a tag variable, you must use a specific format.

## CHAPTER 32 Tagging

---

```
${<variable>}
```

Type a dollar sign followed by open and close curly brackets. The tag variable goes between the curly brackets. You can use tag variables with other tag variables and with string values.

### Supported Tag Variables

The following tag variables are supported.

Variable	Description
<code>\${iam.principal.name}</code>	The name of the principal that tagged the resource.
<code>\${iam.principal.type}</code>	The type of principal that tagged the resource.
<code>\${oci.datetime}</code>	The date and time that the tag was created.

### Managing Tag Defaults

This topic describes how to you can specify tags to be automatically applied to resources at creation time.

#### Required IAM Policy

If you're in the Administrators group, then you have the required access for managing tag defaults and tag namespaces. For specific policy information for this feature, see [Required Permissions for Working with Tag Defaults](#).

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). If you want to dig deeper into writing policies for tagging or other IAM components, see [Details for IAM](#).

#### Overview of Tag Defaults

Tag defaults let you specify tags to be applied automatically to all resources, at the time of creation, in a specific compartment. This feature allows you to ensure that appropriate tags

are applied at resource creation without requiring the user who is creating the resource to have access to the tag namespaces. Consider the following example:

Alice is a finance administrator and has access to the restricted tag namespace Finance. Alice can set up a tag default to apply the Finance.CostCenter tag to all resources with a value of W1234. Eli can create resources but does not have access to the Finance tag namespace. When Eli creates a resource, the Finance.CostCenter tag is automatically applied with a value of W1234. Eli cannot change this tag, and Alice is confident that it will always be applied correctly and not changed by the users who create or edit resources.

Tag defaults allow tenancy administrators to create secure permissions boundaries between users concerned with governance and users who need to create and administer resources.

### Where to Manage Tag Defaults

Tag defaults are defined for a specific compartment, and in the Console you manage them on the Compartment Details page.

## CHAPTER 32 Tagging

ORACLE Cloud

Identity » Compartments » Compartment Details » Tag Defaults

### CompartmentA

Compartment for department A

Rename Compartment Edit Description Add Tag(s) Delete

Compartment Information Tags

Parent Compartment: (root)  
OCID: ...ossqa Show Copy  
Created: Mon, 05 Nov 2018 18:17:23 GMT

### Tag Defaults

Default tags will be applied to any resource created in this compartment, even if they are not shown during create resource. Default tags may be overwritten on resource creation if the user has sufficient permissions.

Create Tag Default

Tag Namespace	Tag Key	Default Value	Tag Key Status	Cost Tracking	
CostCenter	Operations	A1	Active	No	:

Showing 1 Item(s)

Resources

- Child Compartments (1)
- Work Requests (0)
- Tag Defaults (1)

### Required Permissions for Working with Tag Defaults

To create or edit a tag default for a compartment, you must be granted, at a minimum, the following combination of permissions:

- `manage tag-defaults` access on the compartment where you are adding the tag default
- `use tag-namespaces` access on the compartment where the tag namespace resides
- `inspect tag-namespaces` access on the tenancy

For the full mapping of permissions to API operations, see [Details for IAM](#).

For example, assume you have created a set of tag namespaces in CompartmentA. To give a group named TagAdmins access to add tag defaults to CompartmentA, write a policy with the following statements:

## CHAPTER 32 Tagging

---

```
Allow group TagAdmins to manage tag-defaults in compartment CompartmentA
```

```
Allow group TagAdmins to use tag-namespaces in compartment CompartmentA
```

```
Allow group TagAdmins to inspect tag-namespaces in tenancy
```

Now assume you want to allow TagAdmins to also create tag defaults in CompartmentA using tag namespaces that reside in CompartmentZ. Add a statement to allow TagAdmins to use tag namespaces in CompartmentZ:

```
Allow group TagAdmins to use tag-namespaces in compartment CompartmentZ
```

Now when TagAdmins create tag defaults in CompartmentC, they can use tag namespaces that reside in either CompartmentA or CompartmentZ.



### Important

See this [known issue](#) for information about the policy statements that are required to use tag defaults with Compute instance pools and cluster networks.

## Working with Tag Defaults

Tag defaults can only be set up for defined tags. Free-form tags are not supported for tag defaults.

You can define up to 5 tag defaults per compartment.

After a tag default is created, the default tag is applied to any new resources created in the compartment. Previously existing resources in the compartment are not tagged retroactively. Similarly, if you change the default value of the tag default, existing occurrences are not updated. And, if you delete the tag default from the compartment, existing occurrences of the tag are not removed from resources.

### Required Tag Values

For tag defaults, you must include a tag value, but you have a choice about how the value is applied.

- Default value
- User-applied value

If you use a default value, then you must create it. This value is applied to all resources.

If you specify that a user-applied value is required, then the user creating the resource must enter the value for the tag at resource creation time. Users cannot create resources without entering a value for the tag.



### Tip

You can use predefined values and user-applied values to ensure that users only apply a value that you trust. See [Using Predefined Values](#).

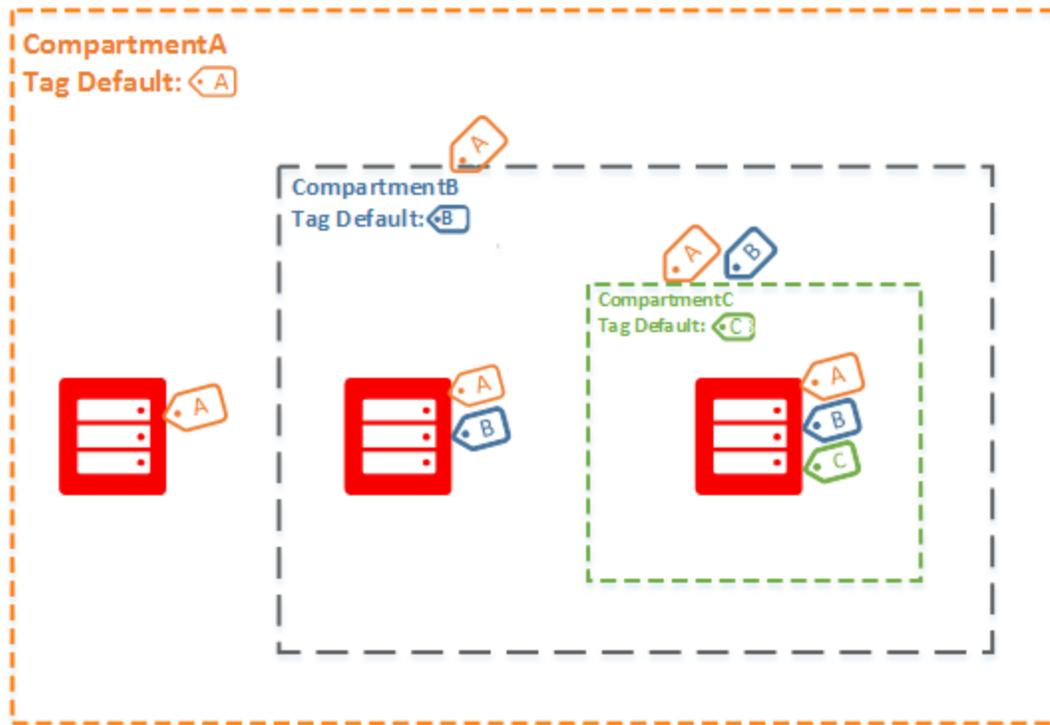
### Tag Inheritance

The default tag is applied to all resources that get created in the compartment, including child compartments and the resources created in the child compartments.

Example:

- In CompartmentA, you create a tag default, TagA.  
Resources (and compartments) created in CompartmentA are automatically tagged with TagA.
  - In the subcompartment, CompartmentB, you create tag default, TagB.  
Resources and compartments created in CompartmentB are automatically tagged with TagA and TagB.
    - In the sub-subcompartment, CompartmentC, you create tag default TagC.  
Resources created in CompartmentC are automatically tagged with TagA, TagB, and TagC.

This example is illustrated in the following graphic:



### Overriding Tag Defaults

Tag defaults can be overridden at the time of resource creation by users who have the appropriate permissions to both create the resource and use the tag namespace.

Example: CompartmentA has a tag default defined to apply `CostCenter.Operations="42"`. Pradeep belongs to a group that grants him permissions to create instances in CompartmentA and also to use tag namespaces in CompartmentA. He creates an instance in CompartmentA, and in the Create Instance dialog, he applies the tag `CostCenter.Operations="50"`. Because he has the appropriate permissions, when the instance is created, the tag default is overridden, and the instance is tagged with `CostCenter.Operations="50"`.

After a resource is created and tagged, users with the appropriate permissions to both update the resource and use the tag namespace can modify the default tags that were applied at resource creation.

### Tag Defaults and Retired Tags

Retired tags can't be applied to new resources. Therefore, if the tag namespace or tag key specified in a tag default is retired, resources can no longer be created in the compartment, because the retired tag cannot be applied. You must delete the tag default that specifies the retired tag to continue to create resources in the compartment.

### Tag Defaults and Tag Variables

You can use tag variables in tag defaults. Tag variables dynamically resolve at resource creation time. For example, you enter a tag variable for principal name as the tag default in a particular compartment.

```
Operations.CostCenter="{iam.principal.name}"
```

Davis and Garcia each create buckets in that compartment. The buckets that Davis creates include default tags that contain his name as the value, while the buckets that Garcia creates have his name.

```
Operations.CostCenter="Davis"
```

```
Operations.CostCenter="Garcia"
```

Meanwhile, the tag default still contains the original variables. See [Using Tag Variables](#).

### Limits on Tag Defaults

There is a limit of 5 tag defaults that can be defined per compartment.

See [Limits on Tags](#) for more limits on tags.

### Using the Console



#### Warning

Avoid entering confidential information when assigning descriptions, tags, or friendly names to your cloud resources through the Oracle Cloud Infrastructure Console, API, or CLI.

### To create a tag default

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments you have access to is displayed.
2. In the list, find the name of the compartment you want to add a tag default to and click its name.
3. On the compartment details page, click **Tag Defaults**.  
The list of existing tag defaults is displayed.
4. Click **Create Tag Default**.
5. Enter the following (all fields are required):
  - **Tag Namespace:** Select the tag namespace for the tag default.
  - **Tag Key:** Select the tag key.
6. Specify the type of value you want this tag to have:
  - **Default Value:** Enter the value you want this tag to have.
  - **User-applied Value:** Users that create resources must enter the value as resources are created.
7. Click **Create Tag Default**.

### To update the default value of a tag default

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments you have access to is displayed.
2. In the list, click the name of the compartment that has the tag default you want to update.
3. On the compartment details page, click **Tag Defaults**.  
The list of existing tag defaults is displayed.
4. Find the tag default you want to update. Go to the the Actions icon (three dots) and click **Edit**.
5. Specify the type of value you want this tag to have:
6.
  - **Default Value:** Enter the value you want this tag to have.
  - **User-applied Value:** Users that create resources must enter the value as resources are created.
7. Click **Save Changes**.

### To delete a tag default

1. Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Compartments**.  
A list of the compartments you have access to is displayed.
2. In the list, click the name of the compartment that has the tag default you want to delete.
3. On the compartment details page, click **Tag Defaults**.  
The list of existing tag defaults is displayed.
4. Find the tag default you want to delete. Go to the the Actions icon (three dots) and click **Delete**.
5. Confirm when prompted.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use these API operations to manage tag defaults:

- [GetTagDefault](#)
- [ListTagDefaults](#)
- [CreateTagDefault](#)
- [UpdateTagDefault](#)
- [DeleteTagDefault](#)

# CHAPTER 33 Web Application Firewall

This chapter explains how to make your endpoints more secure by monitoring and filtering out potentially malicious traffic.

## Overview of the Web Application Firewall Service

Oracle Cloud Infrastructure Web Application Firewall (WAF) is a cloud-based, Payment Card Industry (PCI) compliant, global security service that protects applications from malicious and unwanted internet traffic. WAF can protect any internet facing endpoint, providing consistent rule enforcement across a customer's applications.

WAF provides you with the ability to create and manage rules for internet threats including Cross-Site Scripting (XSS), SQL Injection and other OWASP-defined vulnerabilities. Unwanted bots can be mitigated while tactically allowed desirable bots to enter. Access rules can limit based on geography or the signature of the request.

The global Security Operations Center (SOC) will continually monitor the internet threat landscape acting as an extension of your IT infrastructure.

## Web Application Firewall Service Components

### **WEB APPLICATION FIREWALL POLICY**

WAF policies encompass the overall configuration of your WAF service, including origin management, protection rule settings, and bot detection features.

### **ORIGIN**

Your web application's origin host server. An origin must be defined in your WAF policy in order to set up protection rules or other features.

### **PROTECTION RULES**

Protection rules can be configured to either allow, block, or log network requests when they meet the specified criteria of a protection rule. The WAF will observe traffic to your web

application over time and suggest new rules to apply. To view a list of available WAF rules, see [Supported Protection Rules](#).

### **BOT MANAGEMENT**

The WAF service includes several features that allow you to detect and either block or allow identified bot traffic to your web applications. Bot management features include: JavaScript Challenge, CAPTCHA Challenge, and GoodBot whitelists. For more information, see [Bot Management](#).

## Ways to Access the WAF Service

You can access Oracle Cloud Infrastructure using the Console (a browser-based interface), [command line interface \(CLI\)](#), or the [REST API](#). Instructions for the Console and API are included in topics throughout this guide.

To access the Console, you must use a supported browser. You can use the Console link at the top of this page to go to the sign-in page. Enter your tenancy, user name, and your password.

## Authentication and Authorization

Each service in Oracle Cloud Infrastructure integrates with IAM for authentication and authorization, for all interfaces (the Console, SDK or CLI, and REST API).

An administrator in your organization needs to set up groups, compartments, and policies that control which users can access which services, which resources, and the type of access. For example, the policies control who can create new users, create and manage the cloud network, launch instances, create buckets, download objects, etc. For more information, see [Getting Started with Policies](#). For specific details about writing policies for each of the different services, see [Policy Reference](#).

If you're a regular user (not an administrator) who needs to use the Oracle Cloud Infrastructure resources that your company owns, contact your administrator to set up a user ID for you. The administrator can confirm which compartment or compartments you should be using.

### Note About The API

The WAF service is powered by the Oracle Cloud Infrastructure [Web Application Acceleration and Security \(WAAS\) API](#). All WAF related calls must be made using the WAAS API. To create a WAF configuration using the API, you must first create a [WAAS policy](#) with a defined origin and domain using the API. For the purposes of access control, you must provide the OCID of the compartment where you want the service to reside. For information about access control and compartments, see [Overview of the IAM Service](#).

### WAF Service Capabilities and Limits

The WAF service is limited to 50 policies per tenant and 100 access rules per policy. See [Service Limits](#) for a list of applicable limits and instructions for requesting a limit increase. To set compartment-specific limits on a resource or resource family, administrators can use [compartment quotas](#).

The WAF service allows a total run time of 20 minutes for upload and download processes through the WAF.

### Required IAM Service Policy

To use Oracle Cloud Infrastructure, you must be given access in a policy for waas-policy. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in.

#### Policy examples:

- To allow a specific user group to manage policies in the WAF:

```
Allow group <GroupName> to manage waas-policy in compartment <CompartmentName>
```

```
Allow group <GroupName> to read waas-work-request in compartment <CompartmentName>
```

- To allow a specific user group to manage certificates in the WAF:

```
Allow group <GroupName> to manage waas-certificate in compartment <CompartmentName>
```

## CHAPTER 33 Web Application Firewall

---

- To allow a specific user group view policies in the WAF

```
Allow group <GroupName> to read waas-policy in tenancy <TenancyName>
```

If you're new to policies, see [Getting Started with Policies](#) and [Common Policies](#). For more details about policies for WAF, see [Details for the WAF Service](#).

### Moving WAF Policies to a Different Compartment

You can move WAF policies from one compartment to another. After you move a WAF policy to the new compartment, inherent policies apply immediately and affect access to the WAF policies through the Console, SDK or CLI. For more information, see [Managing Compartments](#).

### Monitoring Resources

You can monitor the health, capacity, and performance of your Oracle Cloud Infrastructure resources by using metrics, alarms, and notifications. For more information, see [Monitoring Overview](#) and [Notifications Overview](#).

For information about available WAF service metrics and how to view them, see [WAF Metrics](#).

### Creating Automation with Events

You can create automation based on state changes for your Oracle Cloud Infrastructure resources by using event types, rules, and actions. For more information, see [Overview of Events](#).

### Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs. You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see [Resource Tags](#).

# Getting Started with WAF

If you're new to Oracle Cloud Infrastructure WAF, this topic gives guidance on how to proceed.

## Before You Begin

To begin using the WAF service, you must have the following available if you plan to run your site on HTTPS/443:

- Public certificate for the fully qualified domain name (FQDN) of the application.
- Corresponding private key for the site.
- IP address of the LBaaS or other public facing endpoint of application.
- Ability to update DNS records for the domain.

## Securing Your WAF

To secure your WAF, you must configure your servers to accept traffic from the WAF servers. Configure your origin's ingress rules to only accept connections from the following CIDR ranges:

- 192.157.18.0/23
- 205.147.88.0/21
- 192.69.118.0/23
- 198.181.48.0/21
- 199.195.6.0/23

## Create a Policy to Route Traffic Through the WAF

To begin, create a policy to route traffic through the WAF without rules enabled. Creating a policy without rules enabled ensures that there are no regressions by having a reverse proxy in front of the application.

### To create a policy

1. Select the region and compartment where the policy should be maintained (there is no constraint around the WAF co-existing with Load Balancing or other application resources in Oracle Cloud Infrastructure).
2. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
3. Click **Create WAF Policy**.
4. In the **Create WAF Policy** dialog box, enter the following:
  - **Policy Name:** A unique name for the policy. Avoid entering confidential information.
  - **Domains:**
    - **Primary Domain:** The fully qualified domain name (FQDN) of the application where the policy will be applied.
    - **Additional Domains:** (Optional) Subdomains where the policy will be applied.
  - **WAF Origin:** The host or IP address of the public internet facing application that is being protected by the application.
    - **Origin Name:** A unique name for the origin.
    - **URI** - Enter the public facing endpoint (IPv4 or FQDN) of the application.
    - **HTTPS Port:** The port used for secure HTTP connection. The default port is 443.
    - **HTTP Port:** The HTTP port the origin listens on. The default port is 80.
    - **Header(s):** (Optional)
      - **Header Name:** The name displayed in the HTTP request header and the header value that can be added and passed to the origin server with all requests.
      - **Header Value:** Specifies the data requested by the header.

- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
5. Click **Create WAF Policy**. The WAF Policy overview appears. Expect the policy to become active within 15 minutes of creation.

### Update DNS to Enable WAF

In this step, you update the CNAME for your zone to route requests from internet clients to WAF. Use the following instructions to make this DNS change in the Console. If your DNS setup resides with another provider, refer to their documentation for instructions.

#### To update the CNAME for your zone

1. In the Policy Information tab of the WAF Policy overview, select the CNAME Target.
2. Copy the CNAME target to your clipboard.
3. Open the navigation menu. Under **Core Infrastructure**, go to **Networking** and click **DNS Zone Management**.
4. Click the Zone Name of the primary domain where you want to update the record. Zone details and a list of records appear.
5. Select the check box for the CNAME record and select **Edit** from the **Actions** drop-down menu.
6. In the **Edit Record** dialog box, update the Target field with the CNAME Target from your clipboard.
7. Click **Submit**.
8. Click **Publish Changes**.
9. In the confirmation dialog box, click **Publish Changes**.

### Upload Your Certificate and Key

This step assumes that your site runs on HTTPS/443.

#### To upload your certificate and key

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of your WAF Policy. The WAF Policy overview appears.
3. Click **Settings**.
4. Click **Edit**.
5. In the **Edit Settings** dialog box, enter the following:
  - **WAF Origin:** Select the name and IP address of the origin.
  - **Enable HTTPS Support:** Select this option so that communications between the browser and web app are encrypted. Enter the following information:
    - **SSL Certificate:** Drag and drop, select, or paste a valid SSL certificate in PEM format. You must also include intermediate certificates (the website certificate must be first). The following is an example:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate_Base64_encoded_certificate>
-----END CERTIFICATE-----
```

- **Private Key:** Drag and drop, select, or paste a valid private key in PEM format in this field. The private key cannot be protected by a passphrase. The following is an example:

```
-----BEGIN PRIVATE KEY-----
<Base64_encoded_private_key>
-----END PRIVATE KEY-----
```

- **Self Signed Certificate:** Enable this field when using a self-signed certificate to show an SSL warning in the browser.
  - **HTTP to HTTPS Redirect:** When enabled, all HTTP traffic is automatically redirected to HTTPS.
6. Click **Save**. Updates to your WAF policy appear in the list to be published in Unpublished Changes.
  7. In the WAF Policy overview, under **Unpublished Changes**, click **View**.
  8. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
  9. Click **Publish All**.
  10. In the Publish Changes dialog box, click **Publish All**.

### Test Your Application

In this step, you ensure that requests are being routed to the WAF and that your application continues to function normally with a reverse proxy in the topology.

#### To test your application

1. Open a browser.
2. Enter the FQDN of the website protected by WAF.
3. Test the functionality of the application.
4. Inspect HTTP Response Headers to see if traffic is flowing through WAF. Some HTTP Response Headers to look for are:
  - X-Cdn: Served-By-Zenedge
  - Server: ZENEDGE
5. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
6. Click the name of your WAF Policy. The WAF Policy overview appears.

7. Under **Logs**, click **View**. Logs for the WAF policy appear.



### Note

You may experience a one-minute delay on logs aggregated and available through the console.

8. Copy the IP address and User-Agent value of your requests to your clipboard. You can use this information when you [enable WAF to passively detect for access rules](#).

## Enable WAF to Passively Detect Rules

In this step, you enable WAF to detect protection rules without blocking requests. Enabling WAF to passively detect rules helps you visualize the traffic that may pose a threat to your site and help you tune the WAF to exclude false positives.

### To enable WAF to detect protection rules

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure rule settings for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Use the [Protection Rules table](#) to locate the rules you want to detect.
5. Enter the Rule IDs you located from the table into the **Rule ID** filter. For this example, enter **941140** (Cross-Site Scripting) in the **Rule ID** filter.
6. Select **Detect** from the **Actions** drop-down menu for the protection rules you filtered.

### To enable WAF to detect access rules

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure access rules for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Click **Add Access Rule**.
5. In the Add Access Rule dialog box, enter the following:
  - a. **Name:** DetectRequestsFromMySpecificBrowser
  - b. **Rule Condition:** Select **IP Address is** from the drop-down and enter the IP address you copied to your clipboard while [testing your application](#) in the **IP Address** field.
  - c. Click **+Additional Condition**.
  - d. **Rule Condition:** Select **User Agent is** from the drop-down and enter the agent value you copied to your clipboard while [testing your application](#) in the **User Agent Header** field.
  - e. **Rule Action:** Select **Detect Only**.



#### Note

Both the IP Address and the User Agent in the preceding example must match for the rule to be triggered. If a different User Agent is used to test your application, the request will not be detected.

6. Click **Add Access Rule**.
7. Click **Unpublished Changes**.

8. Click **Publish All**.

### Test the Rules

When the policy is active, you can test that your rules are detected by WAF.

### To initiate requests

1. Use the same browser you used when you [tested your application](#) to do the following:
  - a. Request the FQDN of your application.
  - b. Request the FQDN of your application with the following query parameter appended: `?id=<script>alert("TEST");</script>`.
2. Use a different browser on the same machine and repeat the preceding requests. All requests should go through the application.

### To verify that WAF is detecting requests

To verify that WAF is detecting requests identified as a risk:

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view logs for. The WAF Policy overview appears.
3. Click **Logs**. Logs for the WAF policy appear.
4. Select the **Detect** check box from the **Actions** filter.
5. Verify that there are two entries for the protection rule triggered by the Cross-Site Scripting request and one entry for detecting the User Agent and IP Address.

### View Recommendations

#### To view protection rule recommendations

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view protection rule recommendations for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Click the **Recommendations** tab. This list is generated based on the traffic the WAF detects flowing through the WAF. If nothing appears in this list, keep testing the FQDN of your application and check back later.
5. Select the protection rules with a **Detect** recommended action and then click **Accept Recommendations**.



### Tip

You can use the **Recommended Action** filter to locate a recommendation by **Detect**.

## Enable WAF to Actively Block Requests

After you verify that requests are being detected, you can start blocking the undesired traffic.

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure rule settings for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Enter rule ID **941140** in the **Rule ID** filter.
5. (Optional) To search for rules with a Detect action, select the **Detect** check box from the **Rule Action** filter.
6. Select **Block** from the **Actions** drop-down menu for rule ID **941140** and any other protection rules you filtered.

7. Under WAF Policy, click **Unpublished Changes**.
8. Click **Publish All**.
9. In the Publish Changes dialog box, click **Publish All**.
10. [Test the rules](#) again by initiating requests. You should get 403 Forbidden errors when testing with the JavaScript on the URL.

## Managing WAF Policies

The Oracle Cloud Infrastructure WAF service enables you to create a WAF policy and origin.

### Order of Processing

The order in which rules and handlers are processed is:

1. IP Whitelists/Blacklists/Good Bot Whitelists
2. Access Rules
3. JavaScript Challenge
4. Device Fingerprinting Challenge (available in the API)
5. Human Interaction Challenge (available in the API)
6. Captcha Challenge
7. Protection Rules
8. Rate Limiting (available in the API)

### Using the Console

#### Create and Manage WAF Policies

##### To create a WAF policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Security**

and click **WAF Policies**.

2. Click **Create WAF Policy**.

3. In the **Create WAF Policy** dialog box, enter the following:

- **Policy Name:** A unique name for the policy. Avoid entering confidential information.
- **Domains:**
  - **Primary Domain:** The fully qualified domain name (FQDN) of the application where the policy will be applied.
  - **Additional Domains:** (Optional) Subdomains where the policy will be applied.
- **WAF Origin:** The host or IP address of the public internet facing application that is being protected by the application.
  - **Origin Name:** A unique name for the origin. Avoid entering confidential information.
  - **URI:** The IPv4 address or fully qualified domain name (FQDN) of the origin. The URI can be a full URI, not just a host/IP.
  - **HTTPS Port:** The port used for secure HTTP connection. The default port is 443.
  - **HTTP Port:** The HTTP port the origin listens on. The default port is 80.
  - **Header(s):** (Optional)
    - **Header Name:** The name displayed in the HTTP request header and the header value that can be added and passed to the origin server with all requests.
    - **Header Value:** Specifies the data requested by the header.
- **Tags:** Optionally, you can apply tags. If you have permissions to create a resource, you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.



### Note

You can add multiple origins to your WAF policy and load balance them accordingly using the `origins` and `originGroups` field of the [UpdateWaasPolicy](#) operation in the WAAS API.

4. Click **Create WAF Policy**. The WAF Policy overview appears. You can access Origin Management, Access Control, WAF, Bot Management, Alerts, and any unpublished changes. While the policy is being created, no changes can be made until the process has completed. Expect the policy to become active within 15 minutes of creation. A CNAME target is generated for each policy. The CNAME target is a hyphenated version of your FQDN within the Oracle Cloud Infrastructure domain (for example, myapp-mydomain-com.oraclecloud.net).
5. In your DNS zone, update the CNAME record entry with the value of the CNAME target that is generated. This enables traffic to be routed through the WAF before the application. This value is presented soon after you publish your policy the first time on the main page of the policy.

### To update a WAF policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to update. The WAF Policy overview appears.



### Tip

You can use the Date Created sort filter to sort policies by the date they were created in ascending or descending order.

3. Click **Edit**.
4. In the Edit WAF Policy dialog box, make the needed changes and then click **Save**.

### To delete a WAF policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Select the check box for the policy you want to delete.



### Tip

You can use the Date Created sort filter to sort policies by the date they were created in ascending or descending order.

3. Click **Delete**.
4. In the confirmation dialog box, click **Delete**.  
The status of the policy changes from **Active** to **Deleting**. Deleted policies are maintained for a short time before they are unavailable in the Console.

### To publish changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes. Pending

changes do not persist across browser sessions. Once you publish changes, it cannot be edited until changes propagate to the edge nodes.

1. In the WAF Policy overview, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
3. Click **Publish All**.
4. In the Publish Changes dialog box, click **Publish All**.

### To manage tags for a WAF policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view. The WAF Policy overview appears.
3. Click the **Tags** tab to view or edit existing tags. Or click **Apply tag(s)** to add new ones.

For more information, see [Resource Tags](#).

### To move a WAF policy to a different compartment

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. In the **Scope** section, select a compartment.
3. Find the WAF policy in the list, click the the Actions icon (three dots), and then click **Move Resource to a Different Compartment**.
4. Choose the destination compartment from the list.
5. Click **Move Resource**.

### Using the CLI

Open a command prompt and run the following command to get the details of a WAAS policy:

## CHAPTER 33 Web Application Firewall

---

```
oci waas waas-policy get --waas-policy-id <policy_ocid>
```

This can be useful in retrieving the necessary information when opening a ticket with Oracle Cloud Infrastructure support. For more information about how to access and use the CLI, see [Command Line Interface \(CLI\)](#).

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [CreateWaasPolicy](#)
- [GetWaasPolicy](#)
- [UpdateWaasPolicy](#)
- [DeleteWaasPolicy](#)
- [ChangeWaasPolicyCompartment](#)

### Origin Management

An origin is an endpoint (typically an IP address) of the application protected by the WAF. An origin can be an Oracle Cloud Infrastructure load balancer public IP address. A load balancer IP address can be used for high availability to an origin. Multiple origins can be defined, but only a single origin can be active for a WAF.

You can set HTTP headers for outbound traffic from the WAF to the origin server. These name value pairs are then available to the application.

### Securing Your WAF

To secure your WAF, you must configure your servers to accept traffic from the WAF servers. Configure your origin's ingress rules to only accept connections from the following CIDR ranges:

- 192.157.18.0/23
- 205.147.88.0/21
- 192.69.118.0/23
- 198.181.48.0/21
- 199.195.6.0/23

### Using the Console

#### To add an origin to your WAF policy

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to add an origin to. The WAF Policy overview appears.
3. Click **Origin Management**.
4. Click **Add Origin**.
5. In the **Add Origin** dialog box, enter the following:
  - **Origin Name**: A unique name for the origin. Avoid entering confidential information.
  - **URI**: The IPv4 address or fully qualified domain name (FQDN) of the origin. The URI can be a full URI, not just a host/IP.
  - **Set as WAF origin?**: Enable this field to designate this origin as the WAF origin.
  - **HTTPS PORT**: The port used for secure HTTP connection. The default is 443.
  - **HTTP PORT**: The HTTP port the origin listens on. The default is 80.
  - **Header(s)**: (Optional) The name displayed in the HTTP request header and the header value that can be added and passed to the origin server with all requests. Additional headers can be added here.
6. Click **Add Origin**. The origin is added to the Origin Management list. You can now

configure WAF rules.



### Note

You can add multiple origins to your WAF policy and load balance them accordingly using the `origins` and `originGroups` field of the [UpdateWaasPolicy](#) operation in the WAAS API.

### To edit an origin

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to edit the origin for. The WAF Policy overview appears.
3. Click **Origin Management**.
4. In the Origin Management view, select the check box for the origin you want to edit.
5. Click **Edit**.
6. In the Edit Origin dialog box, make the needed changes.
7. Click **Save Origin**.

### To delete an origin

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to delete an origin from. The WAF Policy overview appears.
3. Click **Origin Management**.

4. In the Origin Management view, select the check box for the origin you want to delete.
5. Click **Delete**.
6. In the confirmation dialog box, click **Delete**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [CreateWaasPolicy](#)
- [GetWaasPolicy](#)
- [UpdateWaasPolicy](#) To remove an origin from the policy using the API, use the `UpdateWaasPolicy` method and leave the origin field empty upon update.

Each origin has a unique name (key). The name of the origin to be used by the WAF must be referenced in the `wafConfig` portion of the settings. For example, if you have the following origins in your configuration:

```
{
 "compartmentId": "ocidl.compartment.oc1..aaaaatsdfssdfsdsgxz",
 "lifecycleState": "ACTIVE",
 "displayName": "myWAFprotectedApp",
 "origins": {
 "primaryorigin": {
 "httpPort": 80,
 "httpsPort": 443,
 "uri": "67.205.161.231",
 "customHeaders": []
 },
 "secondaryorigin": {
 "httpPort": 80,
 "httpsPort": 443,
 "uri": "54.175.154.7",
```

```
 "customHeaders": [
 {
 "name": "OriginHeader",
 "value": "true"
 },
 {
 "name": "OriginHeader2",
 "value": "true"
 }
]
 }
}
```

Then within the `wafConfig`, the origin in use would be referenced by name:

```
"wafConfig": {
 "deviceFingerprintChallenge": {"isEnabled": false},
 "origin": "secondaryorigin",
 "whitelists": [],
}
```

In this example, the WAF is actively using `secondaryorigin`.

## Bot Management

Bot Management enables you to mitigate undesired bot traffic from your site using CAPTCHA and JavaScript detection tools, while enabling known published bot providers to bypass these controls.

Non-human traffic makes up most of the traffic to sites. Bot Manager is designed to detect and block, or otherwise direct, non-human traffic that may interfere with site operations. The Bot Manager features mitigate bots that conduct content and price scraping, vulnerability scanning, comment spam, brute force attacks, and application-layer DDoS attacks. You can also whitelist good bots.



### Warning

When you enable Bot Management, you incur a higher rate on requests to the WAF.

### JavaScript Challenge

JavaScript Challenge validates that the client can accept JavaScript with a binary decision. JavaScript Challenge is generally the first level of bot mitigation, but not sufficient with more advanced bot tools, which require more advanced challenges. Additional functionality, like detecting Network Address Translation (NAT) traffic, can mitigate the risk of blocking legitimate user traffic from users behind a shared IP address.

The **Action Threshold** parameter defines the number of requests that fail the challenge before the action is taken. The requests that fail under this threshold are not logged. For example, if you set the JavaScript challenge action to **Block** and the **Action Threshold** to 10, and a client that doesn't accept JavaScript makes 11 requests within the **Action Expire Time**, the first 10 requests will be allowed through to origin (assuming there are no other rules) and logs will show one **Block** entry action taken for the JavaScript Challenge.

### CAPTCHA Challenge

If a specific URL should be accessed only by a human, you can control it with CAPTCHA protection. You can customize the comments for the CAPTCHA Challenge for each URL. Bots are kept from accessing protected web application functionality using CAPTCHA images designed to be out of reach of computer vision and OCR technologies.

### Good Bot Whitelist

Good Bots provides the list of bots managed by known providers, such as Baidu or Google. You can allow the access from a specific good bot, or block the bot if they serve no business purpose. Allowed good bots from this section are whitelisted.

Whitelisted bots are flagged with a Bypass action in the WAF policy Logs. You can select the **Bypass** check box from the **Action** filter in Logs to search for the traffic allowed from these rules. Logged good bot events are categorized as a Threat Intelligence Leads log type, however, they are not a threat when the action taken is to Bypass.

The list of good bots on this menu are managed and continuously updated. Additional good bots can be added as a new access control rule in Access Control.

### Using the Console

#### To configure JavaScript Challenge settings

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure JavaScript Challenge settings for. The WAF Policy overview appears.
3. Click **Bot Management**.
4. Click **Enable JavaScript Challenge**.
5. In the JavaScript Challenge dialog box, select the **Enable JavaScript Challenge** check box.
6. In the JS Challenge Action section, choose one of the following methods:
  - **Detect Only:** Select this option if you want to be alerted for every matched request.
    - **Set Header for Failed Request:** (Optional) Select this check box to specify an additional HTTP header to requests that fail the challenge.
  - **Block:** Select this option to block requests by returning a response code, error page, or CAPTCHA.
    - **Block Action:** Select the action that will be taken when a matching request is blocked.
    - **Block Response Code:** Select a status code to return in response to blocked requests.
7. Enter the following information:
  - **Action Threshold:** Specify the number of failed requests before taking action.
  - **Action Expire Time:** Enter the number of seconds between challenges to the same IP address.
8. Click **Save**.

The JavaScript Challenge is added to the list of changes to be published.

### To edit JavaScript Challenge settings

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to edit JavaScript Challenge settings for. The WAF Policy overview appears.
3. Click **Bot Management**.
4. Click **Edit JavaScript Challenge**.
5. In the **Edit JavaScript Challenge** dialog box, make the needed changes.
6. Click **Save**.

### To add a CAPTCHA Challenge

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to edit CAPTCHA challenge settings for. The WAF Policy overview appears.
3. Click **Bot Management**.
4. Click the **CAPTCHA Challenge** tab.
5. Click **Add CAPTCHA Challenge**.
6. In the Add CAPTCHA Challenge dialog box, enter the following information:
  - **CAPTCHA Title**: Enter the text for the CAPTCHA page title.
  - **CAPTCHA URL Path**: Enter the URL path challenged by CAPTCHA.
  - **Session Duration**: Enter the number of seconds after which the CAPTCHA challenge cannot be resubmitted to the same user.

- **CAPTCHA Header:** Enter the text that will appear before the CAPTCHA image (for example, "I am not a robot").
  - **Footer Text:** Enter the text that will be shown after the CAPTCHA input box and before the submit button.
  - **Incorrect CAPTCHA Text:** Enter the text that will appear when incorrect text is entered (for example, "The CAPTCHA was incorrect. Please try again.>").
  - **Submit button:** Enter the text for the Submit button (for example, "Yes, I am human.>").
7. Click **Preview CAPTCHA** to preview the CAPTCHA challenge in a new tab.
  8. Click **Add**.

### To edit a CAPTCHA Challenge

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to edit CAPTCHA Challenge settings for. The WAF Policy overview appears.
3. Click **Bot Management**.
4. Click the **CAPTCHA Challenge** tab.
5. Select the check box for the CAPTCHA you want to edit.
6. Select **Edit** from the **Actions** drop-down menu.
7. Update the CAPTCHA Challenge and then click **Save**.

### To delete a CAPTCHA Challenge

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to delete CAPTCHA Challenge settings for.

The WAF Policy overview appears.

3. Click **Bot Management**.
4. Click the **CAPTCHA Challenge** tab.
5. Select the check box for the CAPTCHA Challenge you want to delete.
6. Select **Delete** from the **Actions** drop-down menu.
7. In the Confirm dialog box, click **Delete**.

The deleted CAPTCHA Challenge is added to the list of changes to be published.

### To manage the Good Bots Whitelist

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure Bot Management for. The WAF Policy overview appears.
3. Click **Bot Management**.
4. Click the **Good Bots Whitelist** tab.
5. Select each bot you want to designate as a good bot.

The designated good bots are added to the list of changes to be published.

### To publish changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes. Pending changes do not persist across browser sessions. Once you publish changes, it cannot be edited until changes propagate to the edge nodes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.

3. Click **Publish All**.
4. In the Publish Changes dialog box, click **Publish All**.

### To discard changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
3. Select the check box for the change you want to discard.
4. Click **Discard**.
5. In the Discard Change dialog box, click **Discard**.

### Using the CLI

You can use the CLI to enable rate limiting, device fingerprinting, and human interaction challenges.

#### To enable rate limiting

Open a command prompt and run the following command to enable rate limiting:

```
oci waas address-rate-limiting update-waf --is-enabled true --allowed-rate-per-address 1 --max-delayed-count-per-address 2 --waas-policy-id <policy_ocid>
```

This default rate limit setting will allow one request per second before starting to delay. It will delay for two requests until the traffic falls within the threshold boundaries. It will use the default error response code of 503.

#### To enable device fingerprinting to detect

Open a command prompt and run the following command to enable device fingerprinting to

## CHAPTER 33 Web Application Firewall

---

detect:

```
oci waas device-fingerprint-challenge update --is-enabled true --action DETECT --failure-threshold 2 --
action-expiration-in-seconds 240 --failure-threshold-expiration-in-seconds 600 --max-address-count 2 --
max-address-count-expiration-in-seconds 255 --waas-policy-id <policy_ocid>
```

To enable the human interaction challenge to detect

Open a command prompt and run the following command to enable the human interaction challenge to detect:

```
oci waas human-interaction-challenge update --is-enabled true --waas-policy-id <policy_ocid>
```

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [UpdateJsChallenge](#)
- [GetJsChallenge](#)
- [UpdateCaptchas](#)
- [GetCaptchas](#)
- [GetDeviceFingerprintChallenge](#)
- [UpdateDeviceFingerprintChallenge](#)
- [GetHumanInteractionChallenge](#)
- [UpdateHumanInteractionChallenge](#)
- [GetWafAddressRateLimiting](#)
- [UpdateWafAddressRateLimiting](#)

# WAF Protection Rules

Protection rules match web traffic to rule conditions and determine the action to be taken when the conditions are met. Protection Rule Settings allow you to define the parameters for enforcement any time a protection rule is matched. Recommendations aid in the optimization of your WAF security profile. The Security Operations team proactively monitors all events to provide recommendations about the action of a specific ruleset. See [Supported Protection Rules](#) for additional information.

## Using the Console

### To apply an action to a protection rule

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure rule settings for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Click the **Rules** tab.
5. Select the protection rule you want to apply an action to.



#### Tip

You can use the **Rule ID** or **Rule Action** filters to locate a protection rule.

6. Select one of the following actions from the **Actions** drop-down menu:
  - **Detect**: Matching requests generate an alert and the request is proxied.
  - **Block**: Matching requests are blocked.

- **Off:** The rule is disabled.

The protection rule action is added to the list to be published.

### To edit rule settings

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure rule settings for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Click the **Rule Settings** tab.
5. Click **Edit Rule Settings**.
6. In the Edit Rule Settings dialog box, enter the following:
  - **Block Action:** The action taken on malicious requests blocked by WAF.
  - **Block Response Code:** Provides information indicating why the request was blocked.
  - **Max Number of Arguments:** The maximum number of arguments allowed in the request. The recommended setting is 255.
  - **Max Length of Argument:** The maximum argument length allowed in the request. The recommended setting is 400.
  - **Max Total Argument Length:** The maximum argument length for all arguments in the request. The recommended setting is 64000.
  - **Recommendations Period:** The period in days to analyze for recommended actions.
  - **Allowed HTTP Methods:** The list of allowed HTTP protocol methods.
7. Click **Save**.

The accepted protection rules are added to the list to be published.

### To accept recommendations

Recommendations will begin appearing once sufficient traffic has gone through the WAF to profile the right security posture.

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to configure rule settings for. The WAF Policy overview appears.
3. Click **Protection Rules**.
4. Click the **Recommendations** tab.
5. Select the protection rules you want to accept.



#### Tip

You can use the **Recommended Action** filter to locate a recommendation by **Detect** or **Block**.

6. Click **Accept Recommendations**.

The accepted protection rules are added to the list to be published.

### To publish changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes. Pending changes do not persist across browser sessions. Once you publish changes, it cannot be edited until changes propagate to the edge nodes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.

3. Click **Publish All**.
4. In the Publish Changes dialog box, click **Publish All**.

### To discard changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
3. Select the check box for the change you want to discard.
4. Click **Discard**.
5. In the Discard Change dialog box, click **Discard**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [GetProtectionRule](#)
- [ListProtectionRules](#)
- [UpdateProtectionRules](#)
- [GetProtectionSettings](#)
- [UpdateProtectionSettings](#)
- [ListRecommendations](#)
- [AcceptRecommendations](#)

### Listing and Accepting Protection Rule Recommendations

Use the following operations to get the list of recommended rules:

- [ListRecommendations](#)

```
{
 "name": "SQL authentication bypass attempts",
 "action": "OFF",
 "description": "Detects basic SQL authentication bypass attempts.",
 "exclusions": [],
 "key": "981244",
 "tags": "SQL Injections, Recommended"
},

{
 "modSecurityRuleIds": [
 "950001",
 "959070",
 "959071",
 "959072",
 "950908",
 "959073"
],
 "name": "Common SQL Injections",
 "action": "OFF",
 "description": "detects common SQL injection attacks",
 "exclusions": [],
 "key": "950001",
 "tags": "SQL Injections, WASCTC, OWASP, A1, PCI, Recommended"
},
```

Using the key values from the output of the GET call above, you can accept one or more of the recommendations using the following operation passing an array of the keys:

## CHAPTER 33 Web Application Firewall

- [AcceptRecommendations](#)

Body:

```
[
 "981244",
 "950001"
]
```

### Protection Rule Specific Settings

Several protection rule settings are settings for specific protection rules.

Setting	Rule ID	Rule Name
Allowed HTTP Methods	911100	Restrict HTTP Request Methods
Max Total Argument Length	960341	Total Arguments Limits
Max Number of Arguments	960335	Number of Arguments Limits
Max Length of Argument	960208	Values Limits

The term "Arguments" refers to either query parameters or body parameters in a PUT/POST request. For instance, if the Max Number of Arguments is 2 and RuleID 960335 is set to BLOCK, any of the following requests would be blocked:

```
GET /myapp/path?query=one&query=two&query=three
```

```
POST /myapp/path with Body {"arg1":"one","arg2":"two","arg3":"three"}
```

```
POST /myapp/path?query=one&query=two with Body {"arg1":"one"}
```

**Max Length of Argument** is the length of either a name or the value of the argument. **Total Argument Length** refers to the sum of the name and value length.

## Exclusions

Sometimes a protection rule can trigger a false positive. You can configure an exception if the request(s) generating the false positive have a particular argument or cookie that can be used to identify that request be excluded from the action normally taken on the rule. Exclusions have to be created through the API. The following exclusion parameters can be used:

Name	Value
REQUEST_COOKIES	Cookie Value
REQUEST_COOKIES_NAMES	Cookie Name (value is irrelevant)
ARGS	Argument (Query Parameter or POST/PUT data)
ARGS_NAMES	Query Parameter Name (value is irrelevant)

### Example

In this example, a block is applied to WAF Rule 911100 (Restrict HTTP Request Methods) with an exception to allow requests with an argument that contains "passthrough".

```
PUT / waasPolicies /<policy_ocid>/wafConfig/protectionRules
```

With the body:

```
[
 {
 "key": "911100",
 "action": "BLOCK",
 "exclusions": [
 {
 "target": "REQUEST_COOKIES",
 "exclusions": ["yourcompany.com", "Wed, 21 Oct 2015 07:28:00 GMT", "12345", "219ffwef9w0f"]
 },
 {
 "target": "REQUEST_COOKIES_NAMES",
```

## CHAPTER 33 Web Application Firewall

```
 "exclusions":["OAMAuthnCookie", "JSESSIONID", "HCM-PSJSESSIONID"]
 },
 {
 "target":"ARGS",
 "exclusions":["passthrough"]
 }
]
}
```

This will return a 202 Accepted HTTP status, which means the policy will enter an UPDATING state until changes are provisioned to the edge nodes.

### Supported Protection Rules

The Oracle Cloud Infrastructure WAF service supports many protection rule types. The following list provides a brief explanation of the purpose of each protection rule type.

#### Protection Rules

Rule ID/Key	Name	Description
90001	Filter Profanity	Detects profanity used in request headers and body.
90004	Executable File Upload Attempt	Detects attempts to upload executable files through input forms.
90006	Credit Card Leakage in Request: GSA SmartPay	Detects GSA SmartPay credit card numbers in user input.
90007	Credit Card Leakage in Request: MasterCard	Detects MasterCard credit card numbers in user input.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
90008	Credit Card Leakage in Request: Visa	Detects Visa credit card numbers in user input.
90009	Credit Card Leakage in Request: American Express	Detects American Express credit card numbers in user input.
90010	Credit Card Leakage in Request: Diners Club	Detects Diners Club credit card numbers in user input.
90011	Credit Card Leakage in Request: enRoute	Detects enRoute credit card numbers in user input.
90012	Credit Card Leakage in Request: Discover	Detects Discover credit card numbers in user input.
90013	Credit Card Leakage in Request: JCB	Detects JCB credit card numbers in user input.
120123	Joomla! Core CVE-2015-8562 Remote Code Execution Vulnerability Prevention	Detects Joomla! Core CVE-2015-8562 Remote Code Execution Vulnerability payload.
900032	HTTP Parameter Pollution (HPP) Detection	Detects requests that have multiple arguments with the same name indicative of an HPP attack.
911100	Restrict HTTP Request Methods	Allows only request methods specified by the configurable "Allowed http methods" parameter.
920100	Invalid HTTP Request Line	Detects an invalid HTTP request line.
920280	Missing/Empty Host Header	Detects a missing/empty host header.
920350	Invalid HTTP Request Line	Detects invalid HTTP request lines.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
941100	Cross-Site Scripting (XSS) Attempt: Libinjection - XSS Detection	Detects XSS Libinjection attempt.
941101	Cross-Site Scripting (XSS) Attempt: SS Attack Detected via libinjection	Detects an SS attack via libinjection.
941110	Cross-Site Scripting (XSS) Attempt: XSS Filters - Category 1	Detects script tag-based XSS vectors, for example, <code>&lt;script&gt; alert(1)&lt;/script&gt;</code> .
941120	Cross-Site Scripting (XSS) Attempt: XSS Filters - Category 2	Detects XSS vectors making use of event handlers like <code>onerror</code> , <code>onload</code> etc., for example, <code>&lt;body onload="alert(1)"&gt;</code> .
941130	Cross-Site Scripting (XSS) Attempt: XSS Filters - Category 3	Detects XSS vectors making use of attribute vectors.
941140	Cross-Site Scripting (XSS) Attempt: XSS Filters - Category 4	Detects XSS vectors making use of javascript URI and tags, for example, <code>&lt;p style="background:url(javascript:alert(1))"&gt;</code> .
941150	Cross-Site Scripting (XSS) Attempt: XSS Filters - Category 5	Detects HTML attributes - <code>src</code> , <code>style</code> , and <code>href</code> .
941160	Cross-Site Scripting (XSS) Attempt: NoScript XSS Filters	Detects NoScript InjectionChecker: HTML Injection.

## CHAPTER 33 Web Application Firewall

<b>Rule ID/Key</b>	<b>Name</b>	<b>Description</b>
941170	Cross-Site Scripting (XSS) Attempt: NoScript XSS Filters	Detects NoScript InjectionChecker: Attributes injection.
941180	Cross-Site Scripting (XSS) Attempt: Blacklist Keywords from Node-Validator	Detects Blacklist Keywords from Node-Validator.
941190	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941200	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941210	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941220	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941230	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941240	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
941250	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941260	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941270	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941280	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941300	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941310	Cross-Site Scripting (XSS) Attempt: US-ASCII encoding bypass listed on XSS filter evasion	Cross-Site Scripting (XSS) Attempt: US-ASCII encoding bypass listed on XSS filter evasion
941320	Cross-Site Scripting (XSS) Attempt: HTML Tag Handler	Cross-Site Scripting (XSS) Attempt: HTML Tag Handler
941330	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.

## CHAPTER 33 Web Application Firewall

---

<b>Rule ID/Key</b>	<b>Name</b>	<b>Description</b>
941340	Cross-Site Scripting (XSS) Attempt: XSS Filters from Internet Explorer	Detects XSS Filters from IE.
941350	Cross-Site Scripting (XSS) Attempt: UTF-7 encoding XSS filter evasion for IE	Cross-Site Scripting (XSS) Attempt: UTF-7 encoding XSS filter evasion for IE.
950002	Common System Command Access Attempt	Detects access attempts to common system commands, such as map, telnet, ftp, rcms, cmd.
950005	Common System Files Access Attempt	Detects access attempts to common system files, such as access, passwd, groupm global.asa, httpd.conf, boot.ini, /etc.
950006	Injection for Common System Commands	Detects injections for common system commands such as telnet, map, blocalgroup, ftp, rcmd, echo, cmd, chmod, passwd, mail.
950007	Blind SQL Injection	Detects common blind SQL injection attacks.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
950009	Session Fixation	Detects Session Fixation, an attack technique that forces a user's session ID to an explicit value. Depending on the functionality of the target website, several techniques can be utilized to "fix" the session ID value. These techniques range from Cross-site Scripting exploits to peppering the website with previously made HTTP requests. After a user's session ID has been fixed, the attacker will wait for that user to log in. Once the user does so, the attacker uses the predefined session ID value to assume the same online identity.
950010	LDAP Injection	Detects common LDAP data constructions injections.
950011	SSI Injection	Detects common Server-Side-Include format data injections.
950012	HTTP Request Smuggling	Detects specially crafted requests that under certain circumstances could be seen by the attacked entities as two different sets of requests. This allows certain requests to be smuggled through to a second entity without the first one realizing it.
950018	UPDF XSS Injection	Detects submitted links that contains the # fragment in a query_string.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
950019	Email Injection	Detects mail command injections targeting mail servers and webmail applications that construct IMAP/SMTP statements from user-supplied input that is not properly sanitized.
950103	Path/Directory Traversal	Detects path traversal attempts, also known as directory traversal or "../" attacks.
950107	URL Encodings Validation	Detects URL encoding inconsistencies, encoding abuse, and invalid formatting.
950116	Unicode Encoding/Decoding Validation	blocks full-width Unicode encoding as decoding evasions could be possible.
950117	URL Contains an IP Address	Detects a common RFI attack, when a URL contains an IP address.
950118	PHP Include() Function	Detects a common RFI php include() function attacks.
950119	Data Ends with Question Mark (s) (?)	Detects a common RFI attack, when data ends with question mark(s) (?).
950120	Host Doesn't Match Localhost	Detects a common RFI attack, when host doesn't match localhost.
950801	UTF Encoding Validation	Detects UTF encoding inconsistencies and invalid formatting.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
950907	OS Command Injection	Detects OS command injection in an application to elevate privileges, execute arbitrary commands, compromise the underlying operating system and install malicious toolkits such as those to participate in botnet attacks.
950910	HTTP Response Splitting	Detects Carriage Return + Linefeed characters in the response header that could cause attacked entities to interpret it as two separate responses instead of one.
958000	Addimport XSS Attack	Detects usage of addimport in request, cookies, or arguments.
958001	document Cookie XSS Attack	Detects usage of document.cookie in request, cookies, or arguments.
958002	execscript XSS Attack	Detects usage of execscript in request, cookies, or arguments.
958003	fromcharcode XSS Attack	Detects usage of fromcharcode in request, cookies, or arguments.
958004	innerHTML XSS Attack	Detects usage of innerhtml in request, cookies, or arguments.
958005	cdata XSS Attack	Detects usage of cdata in request, cookies, or arguments.
958006	body background XSS Attack	Detects usage of <body background in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958007	onload XSS Attack	Detects usage of onload in request, cookies, or arguments.
958008	input type image XSS Attack	Detects usage of <input type image in request, cookies, or arguments.
958009	import XSS Attack	Detects usage of import in request, cookies, or arguments.
958010	activexobject XSS Attack	Detects usage of activexobject in request, cookies, or arguments.
958011	background-image: XSS Attack	Detects usage of background-image: in request, cookies, or arguments.
958012	copyparentfolder XSS Attack	Detects usage of copyparentfolder in request, cookies, or arguments.
958013	createtextrange XSS Attack	Detects usage of createtextrange in request, cookies, or arguments.
958016	getparentfolder XSS Attack	Detects usage of getparentfolder in request, cookies, or arguments.
958017	getspecialfolder XSS Attack	Detects usage of getspecialfolder in request, cookies, or arguments.
958018	href javascript: XSS Attack	Detects usage of href javascript: in request, cookies, or arguments.
958019	href schell XSS Attack	Detects usage of href schell in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958020	href vbscript: XSS Attack	Detects usage of href vbscript: in request, cookies, or arguments.
958022	livescript: XSS Attack	Detects usage of livescript: in request, cookies, or arguments.
958023	lowsrc javascript: XSS Attack	Detects usage of lowsrc javascript: in request, cookies, or arguments.
958024	lowsrc shell XSS Attack	Detects usage of lowsrc shell in request, cookies, or arguments.
958025	lowsrc vbscript XSS Attack	Detects usage of lowsrc vbscript in request, cookies, or arguments.
958026	mocha: XSS Attack	Detects usage of mocha: in request, cookies, or arguments.
958027	onabort XSS Attack	Detects usage of onabort in request, cookies, or arguments.
958028	settimeout XSS Attack	Detects usage of settimeout in request, cookies, or arguments.
958030	src http: XSS Attack	Detects usage of src http: in request, cookies, or arguments.
958031	javascript: XSS Attack	Detects usage of javascript: in request, cookies, or arguments.
958032	src and shell XSS Attack	Detects usage of src and shell in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958033	vbscript: XSS Attack	Detects usage of vbscript: in request, cookies, or arguments.
958034	style bexpression XSS Attack	Detects usage of style bexpression in request, cookies, or arguments.
958036	type application x-javascript XSS Attack	Detects usage of type application x-javascript in request, cookies, or arguments.
958037	type application x-vbscript XSS Attack	Detects usage of type application x-vbscript in request, cookies, or arguments.
958038	type text ecmaascript XSS Attack	Detects usage of type text ecmaascript in request, cookies, or arguments.
958039	type text javascript XSS Attack	Detects usage of type text javascript in request, cookies, or arguments.
958040	type text jscript XSS Attack	Detects usage of type text jscript in request, cookies, or arguments.
958041	type text vbscript XSS Attack	Detects usage of type text vbscript in request, cookies, or arguments.
958045	url javascript: XSS Attack	Detects usage of url javascript: in request, cookies, or arguments.
958046	url shell XSS Attack	Detects usage of <url shell in request, cookies, or arguments.
958047	url vbscript: XSS Attack	Detects usage of url vbscript: in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958049	?meta XSS Attack	Detects usage of ?meta in request, cookies, or arguments.
958051	?script XSS Attack	Detects usage of < ?script in request, cookies, or arguments.
958052	alert XSS Attack	Detects usage of alert in request, cookies, or arguments.
958054	lowsrc and http: XSS Attack	Detects usage of lowsrc and http: in request, cookies, or arguments.
958056	iframe src XSS Attack	Detects usage of iframe src in request, cookies, or arguments.
958057	?iframe XSS Attack	Detects usage of ?iframe in request, cookies, or arguments.
958059	asfunction: XSS Attack	Detects usage of asfunction: in request, cookies, or arguments.
958291	Range Header Validation	Detects range header inconsistencies and invalid formatting.
958295	Connection Header Validation	Detects connection header inconsistencies and invalid formatting.
958404	onerror XSS Attack	Detects usage of onerror in request, cookies, or arguments.
958405	onblur XSS Attack	Detects usage of onblur in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958406	onchange XSS Attack	Detects usage of onchange in request, cookies, or arguments.
958407	onclick XSS Attack	Detects usage of onclick in request, cookies, or arguments.
958408	ondragdrop XSS Attack	Detects usage of ondragdrop in request, cookies, or arguments.
958409	onfocus XSS Attack	Detects usage of onfocus in request, cookies, or arguments.
958410	onkeydown XSS Attack	Detects usage of onkeydown in request, cookies, or arguments.
958411	onkeypress XSS Attack	Detects usage of onkeypress in request, cookies, or arguments.
958412	onkeyup XSS Attack	Detects usage of onkeyup in request, cookies, or arguments.
958413	onload XSS Attack	Detects usage of onload in request, cookies, or arguments.
958414	onmousedown XSS Attack	Detects usage of onmousedown in request, cookies, or arguments.
958415	onmousemove XSS Attack	Detects usage of onmousemove in request, cookies, or arguments.
958416	onmouseout XSS Attack	Detects usage of onmouseout in request, cookies, or arguments.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
958417	bonmouseover XSS Attack	Detects usage of bonmouseover in request, cookies, or arguments.
958418	onmouseup XSS Attack	Detects usage of onmouseup in request, cookies, or arguments.
958419	onmove XSS Attack	Detects usage of onmove in request, cookies, or arguments.
958420	onresize XSS Attack	Detects usage of onresize in request, cookies, or arguments.
958421	onselect XSS Attack	Detects usage of onselect in request, cookies, or arguments.
958422	onsubmit XSS Attack	Detects usage of onsubmit in request, cookies, or arguments.
958423	onunload XSS Attack	Detects usage of onunload in request, cookies, or arguments.
959151	php Code Injection	Detects a common injections attack, when request contain any php code e.g. "<\?>".
960000	File Name Validation	Detects multipart/form-data file name evasion attempts.
960007	Missing Host Header	Detects missing request host header.
960009	Missing User-Agent Header	Detects missing request user-agent header.
960011	GET/HEAD Requests Validation	Detects if GET/HEAD requests contain request body since it is not a common practice.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
960012	Content-Length Header Validation	Detects if content-length header is provided with every POST request.
960013	Require Content-Length to be provided with every HTTP/1.1 POST request that has no Transfer-Encoding header	Detects HTTP/1.1 requests that do not comply with HTTP 1.1 spec by having no content-length header when transfer-encoding is also absent.
960014	URI Validation	Ensures that URI and canonical server name are matching.
960015	Missing Accept Header	Detects missing request accept header.
960016	Content-Length Header Validation	Detects if content-length HTTP header is not numeric.
960017	Host Header Is IP Address	Detects if host header is a numeric IP address as it could be indicative of automated client access.
960020	Pragma Header Validation	Ensures that pragma, cache-control headers and HTTP protocol version supplied by the client are matching.
960022	Expect Header Validation	Ensures that expect header and HTTP protocol version supplied by the client are matching.
960024	Repetitive Non-Word Chars	Attempts to identify when 4 or more non-word characters are repeated in sequence.
960208	Values Limits	Detects HTTP requests with value length exceeding the configurable "Max length of argument" parameter.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
960209	Arguments Limits	Detects HTTP requests with argument name length exceeding the 100 symbols.
960335	Number of Arguments Limits	Detects HTTP requests with number of arguments exceeding the configurable "Max amount of arguments" value.
960341	Total Arguments Limits	Detects HTTP requests with total length of all arguments exceeding the configurable "Max total argument length" parameter.
960901	Character Set Validation	Ensures that only a specific character set(s) is used.
960902	Content-Encoding Header Validation	Ensures that identity is not specified in content-encoding header.
960904	Missing Content-Type Header	Detects missing content-type header or if combination of content-length and content-type headers is invalid.
960911	Request Line Format Validation Against the HTTP RFC	Uses rule negation against the regex for positive security. The regex specifies the proper construction of URI request lines such as: "http:" "://" host [ ":" port ] [ abs_path [ "?" query ] ]. It also outlines proper construction for CONNECT, OPTIONS, and GET requests.
960912	Malformed request bodies	Checks for request body parsing errors.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
960914	Strict Multipart Parsing Checks	Uses strict checks for what is accepted in the multipart/form-data request body. If the rule proves to be too strict for your environment consider changing it to Off.
960915	Multipart Unmatched Boundary Check	Checks for signs of evasions during file upload requests.
970901	5XX Status Code Information Leakage	Detects if an application generates 500-level status code. For example, 500 Internal Server Error, 501 Not Implemented...505 HTTP Version Not Supported.
973300	Common Direct HTML Injection	Detects tags that are the most common direct HTML injection points.
973306	Embedded JavaScript in Style Attribute	Detects embedded JavaScript in style attribute.
973307	Embedded Scripts Within JavaScript Fragments	Detects common JavaScript fragments like fromcharcode, alert, eval that can be used for attacks.
973309	CSS Fragments Attacks	Detects common CSS fragments attacks like <div style="background-image: url (javascript:...)"> or <img style="x:expression (document.write(1))">.
973310	Embedded Scripts Within Alert Fragments	Detects attacks like alert('xss'), alert("xss"), alert(/xss/).

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
973311	String.fromCharCode (88,83,83) attacks	Detects String.fromCharCode(88,83,83) attacks.
973312	";!--"<XSS>=&{()} Attacks	Detects ";!--"<XSS>=&{()} attacks.
973313	&{alert('xss')} Attacks	Detects &{alert('xss')} attacks.
973314	Doctype Entity Inject	Detects Doctype Entity inject attacks.
973331	Internet Explorer XSS Filters	Detects common IE XSS attacks.
973336	Embedding Scripts Within Scripts	Detects script tag-based XSS vectors. For example, <script> alert(1)</script>.
973337	Embedded Scripts Within Event Handlers	Detects event handler based XSS vectors. For example, <body onload="alert(1)">.
973338	Embedded Scripts Within URI Schemes	Detects "data", "javascript", "src" or other URI schemes/attributes based XSS vectors. For example, <p style="background:url (javascript:alert(1))">.
981136	Generic XSS Attacks	Detects common XSS attacks embedded within non-script elements. For example, jscript onsubmit cyparentfolder document javascript meta onchange onmove onkeydown onkeyup activexobject onerror onmouseup ecmascript bexpression onmouseover vbscript.
981172	SQL Character Anomaly Scoring	Attempts to gauge when there is an excessive use of meta-characters within a single parameter payload.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
981227	Request URI Validation	Detects invalid URI in request.
981242	-°lassic SQL Injection Probing	Detects classic SQL injection probings.
981244	SQL Authentication Bypass Attempts	Detects basic SQL authentication bypass attempts.
981245	SQL Authentication Bypass Attempts	Detects basic SQL authentication bypass attempts.
981246	SQL Authentication Bypass Attempts	Detects basic SQL authentication bypass attempts.
981272	SQL Injection Using sleep() or benchmark()	Detects blind SQL injection tests using sleep() or benchmark() functions.
981300	SQL Keyword Anomaly Scoring	Detects common SQL keywords anomalies.
981318	String Termination/Statement Ending	Identifies common initial SQLi probing requests where attackers insert/append quote characters to the existing normal payload to see how the app/db responds.
1000000	Shellshock Exploit Attempt	Detects the ability to unintentionally execute commands in Bash (CVE-2014-6271).
2017100	Apache Struts 2 Multipart Parser CVE-2017-5638 Remote Code Execution Vulnerability Prevention	Detects Apache Jakarta CVE-2017-5638 Remote Code Execution Vulnerability Payload.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
2018100	CVE-2018-6389 WordPress Parameter Resource Consumption Remote DoS	Detects WordPress parameter resource consumption remote DoS on jquery-ui-core .
2100019	/_layouts/scriptresx.ashx sections Parameter XSS	Detects Microsoft SharePoint /_layouts/scriptresx.ashx sections parameter XSS attacks.
2100023	/owssrv.dll List Parameter XSS	Detects Microsoft SharePoint /owssrv.dll List Parameter XSS attacks.
2100026	_layouts/Chart/WebUI/WizardList.aspx skey Parameter XSS	Detects Microsoft SharePoint _layouts/Chart/WebUI/WizardList.aspx skey Parameter XSS attacks.
2100027	_layouts/themeweb.aspx XSS	Detects Microsoft SharePoint _layouts/themeweb.aspx ctI00\$PlaceholderMain\$ctl82 \$customizeThemeSection\$accent6 Parameter XSS attacks.
2100028	_layouts/inplview.aspx ListViewPageUrl Parameter XSS	Detects Microsoft SharePoint _layouts/inplview.aspx ListViewPageUrl Parameter XSS attacks.
2100032	owssrv.dll View Parameter XSS	Detects Microsoft SharePoint owssrv.dll View Parameter XSS attacks.
2100033	NewForm.aspx TextField_spSave Parameter XSS	Detects Microsoft SharePoint NewForm.aspx TextField_spSave Parameter XSS attacks.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
2100034	/Lists/Calendar/calendar.aspx CalendarDate Parameter XSS	Detects Microsoft SharePoint /Lists/Calendar/calendar.aspx CalendarDate Parameter XSS attacks.
2100035	_layouts/Picker.aspx XSS	Detects Microsoft SharePoint _layouts/Picker.aspx ctI00\$PlaceholderDialogBodySection\$ctI04\$hiddenSpanData Parameter XSS attacks.
2100048	_layouts/help.aspx cid0 Parameter XSS	Detects Microsoft SharePoint _layouts/help.aspx cid0 Parameter XSS attacks.
2100062	_layouts/ScriptResx.ashx name Parameter LFI	Detects Microsoft SharePoint _layouts/ScriptResx.ashx name Parameter LFI attacks.
2100063	_layouts/OSSSearchResults.aspx k Parameter XSS	Detects Microsoft SharePoint _layouts/OSSSearchResults.aspx k Parameter XSS attacks.
2100069	wiki pages multiple Parameter XSS	Detects Microsoft SharePoint wiki pages multiple Parameter XSS (CVE-2013-3180) attacks.
2100070	/Lists/Links/AllItems.aspx XSS	Detects Microsoft SharePoint /Lists/Links/AllItems.aspx ctI00\$m\$g_2085a732_4692_4d3e_99d2_4d90ea5108d2\$ctI00\$ctI05\$ctI00\$ctI00\$ctI04\$ctI00\$ctI00\$UrlFieldUrl Parameter XSS attacks.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
2100082	Drupal - pre-auth SQL Injection Vulnerability	Detects Drupal pre-auth SQL injection vulnerability. A malicious user can inject arbitrary SQL queries and thereby control the complete Drupal site. This leads to a code execution as well. Drupal 7.32 fixed this bug.
2100083	Gerber WebPDM XSS Vulnerability	Detects cross-site scripting vulnerability in Gerber WebPDM Product Data Management System.
2100084	Gerber WebPDM SQL Injection Vulnerability	Detects SQL Injection Vulnerability in Gerber WebPDM Product Data Management System.
2100085	High X-SharePointHealthScore	Detects Microsoft SharePoint High X-SharePointHealthScore - potential DoS attack/availability risk.
2100086	Response Header Found	Detects Microsoft SharePoint SharePointError Response Header Found.
2100087	x-virus-infected Response Header Found	Detects x-virus-infected Response Header Found.
2100088	Rights Management (IRM) Error Response Header Found	Detects Microsoft SharePoint Information Rights Management (IRM) Error Response Header Found.
2100089	/_layouts/mobile/editform.aspx XSS	Detects Microsoft SharePoint /_layouts/mobile/editform.aspx XSS attacks.
2200924	IRC Botnet Attacks	Detects common IRC Botnet attack commands.



## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
2250127	e107 Plugin my_gallery Exploit	Detects e107 Plugin my_gallery Exploit "http://" . \$site. "e107_plugins/my_gallery/image.php?file=../../e107_config.php".
2250128	Opencart Remote File Upload Vulnerability	Detects Opencart remote file upload vulnerability.
2250129	Zen Cart Local File Disclosure Vulnerability	Detects Zen Cart local file disclosure vulnerability.
20182056	CVE-2003-1567 CVE-2004-2320 CVE-2010-0360 TRACE & CONNECT Attempts	Detects TRACE method attempts.
201821375	CVE-2012-0209, Remote Execution Backdoor Attempt Against Horde	Detects remote execution backdoor attempt against Horde.
201821438	CVE-2012-1723, CVE-2012-1889, CVE-2012-4681, Blackhole Exploit Kit JavaScript Carat String Splitting with Hostile Applet	Detects Blackhole exploit kit JavaScript carat string splitting with hostile applet.
201822063	CVE-2012-1823, CVE-2012-2311, CVE-2012-2335, CVE-2012-2336, PHP-CGI Remote File Include Attempt	Detects PHP-CGI remote file include attempts.
201826834	CVE-2012-4681, CVE-2012-5076, CVE-2013-2423, Sweet Orange Exploit Kit Landing Page in.php base64 uri	Detects Sweet Orange exploit kit landing page in.php base64 uri attacks.

## CHAPTER 33 Web Application Firewall

Rule ID/Key	Name	Description
201826947	CVE-2013-2423, DotkaChef/Rmayana/DotCache Exploit Kit Inbound Java Exploit Download	Detects DotkaChef/Rmayana/DotCache exploit kit inbound java exploit download attacks.
201826948	CVE-2013-1493, DotkaChef/Rmayana/DotCache Exploit Kit Inbound Java Exploit Download	Detects DotkaChef/Rmayana/DotCache exploit kit inbound java exploit download attacks.
201827040	CVE-2013-0422, CVE-2013-2423, Styx Exploit Kit Plugin Detection Connection Jorg	Detects Styx exploit kit plugin detection connection jorg attacks.
201841409	CVE-2017-3823, CVE-2017-6753, Cisco WebEx Explicit Use of Web Plugin	Detects Cisco WebEx explicit use of web plug-in.
201843811	CVE-2017-9812, Kaspersky Linux File Server WMC Directory Traversal Attempt	Detects Kaspersky Linux file server WMC directory traversal attempts.

### Custom Protection Rules

The WAF service allows you to define and apply custom protection rules from open source firewall modules to your WAF configurations, such as ModSecurity modules. This topic describes how to format, create, and implement custom protection rules in your WAF policies using the [WAAS API](#). For a list of protection rules already available in the service, see [Supported Protection Rules](#).



### Note

Custom protection rules can only be created using the [WAAS API](#).

### Custom Protection Rule Syntax

All custom protection rules are expressed in ModSecurity Rule Language. For more information about ModSecurity syntax, see [Making Rules: The Basic Syntax](#).

Additionally, each rule must include two placeholder variables that are updated by the WAF service upon publication of the rule.

**id: {{id\_1}}** - This field is updated with a unique rule ID generated by the WAF service which identifies a `SecRule`. More than one `SecRule` can be defined in the `template` field of a [CreateCustomProtectionRule](#) call. The value of the first `SecRule` must be `id: {{id_1}}` and the `id` field of each subsequent `SecRule` should increase by one, as shown in the example.

**ctl:ruleEngine={{mode}}** - The action to be taken when the criteria of the `SecRule` are met, either `OFF`, `DETECT` or `BLOCK`. This field is updated with the corresponding value of the `action` field of the `CustomProtectionRuleSetting` object when using the [UpdateWafConfig](#) operation.

### Example of a custom protection rule format:

```
SecRule REQUEST_COOKIES "regex matching SQL injection - part 1/2" \
 "phase:2, \
 msg:'Detects chained SQL injection attempts 1/2.', \
 id: {{id_1}}, \
 ctl:ruleEngine={{mode}}, \
 deny" \
SecRule REQUEST_COOKIES "regex matching SQL injection - part 2/2" \
 "phase:2, \
 msg:'Detects chained SQL injection attempts 2/2.', \
 id: {{id_2}}, \
 ctl:ruleEngine={{mode}}, \
 deny"
```

### Actions

The WAF service can take an action on an HTTP request when the criteria of a custom protection rule are met.

- **DETECT** - Logs the request when the criteria of the custom protection are met.
- **BLOCK** - Blocks the request when the criteria of the custom protection rule are met.
- **OFF** - The custom protection rule is inactive and will take no action.

### Create a Custom Protection Rule

Custom protection rules can be created and added to a compartment using the [CreateProtectionRule](#) call in the WAAS API. Using ModSecurity Rule Language formatting, populate the `template` field with the criteria of the rule.

#### EXAMPLE:

```
{
 "compartmentId":
"ocidl.compartment.region1..aaaaaaaa5n4ocdo64z66dzfftevo6ntalkpojfiwgmxohi4tlakqosvdrmq",
 "description": "The description text for the rule being created",
 "displayName": "Custom Protection Rule Name",
 "template": "SecRule REQUEST_URI / \"phase:2, t:none, capture, msg:'Custom (XSS) Attack. Matched Data:
 %{TX.0} found within %{MATCHED_VAR_NAME}: %{MATCHED_VAR}', id:{{id_1}}, ctl:ruleEngine={{mode}},
 tag:'Custom', severity:'2'\""
}
```

### Adding Custom Protection Rules to a WAF Configuration

Custom protection rules can be added to a WAF configuration using the [UpdateWafConfig](#) call in the WAAS API.

Add the OCID and the desired action to take to the `CustomProtectionRuleSetting` object of the [UpdateWafConfig](#) schema.

#### EXAMPLE:

```
[
 {
 "action": "BLOCK",
```

```
"id":
"ocidl.waascustomprotectionrule.oc1..aaaaaaaaalxd4jrws4rbbnddzlnotu3giuzo53kopbj747mbvarttr7vyy7ja"
},
{
 "action": "DETECT",
 "id":
"ocidl.waascustomprotectionrule.oc1..aaaaaaaaamx5r72ntmmhwgeaspzpdqcwsgprpuvwsa7xoshnyo3xhtpwcobeq"
}
]
```

To view a list of available custom protection rules in a compartment and their corresponding OCIDs, use the [ListCustomProtectionRules](#) call in the WAAS API.

## Access Control

As a WAF administrator you can define explicit actions for requests that meet various conditions. Conditions use various operations and regular expressions. A rule action can be set to log and allow, detect, or block requests.

The available conditions are shown in the following table:

## CHAPTER 33 Web Application Firewall

Criteria Type	Criteria
URL	<p>Users shall be able to define one or more criteria based on:</p> <ul style="list-style-type: none"><li>• URL is</li><li>• URL is not</li><li>• URL starts with</li><li>• URL ends with</li><li>• URL contains</li><li>• URL regex</li></ul> <p>The URL regex matching uses Perl-compatible regular expressions.</p>
IP Address	<p>Users shall be able to define one or more criteria based on:</p> <ul style="list-style-type: none"><li>• Client IP Address is</li><li>• Client IP Address is not</li></ul> <p>These values can be a valid IPv4 address, subset, or CIDR notation for a range. IPv6 is not yet supported.</p>
Country/Region	<p>Users shall be able to define one or more criteria based on:</p> <ul style="list-style-type: none"><li>• Country/Region is</li><li>• Country/Region is not</li></ul> <p>For the API, use a 2-letter country code.</p>

Criteria Type	Criteria
User Agent	<p>User Agent is a value that identifies the browser client.</p> <ul style="list-style-type: none"> <li>• User Agent is</li> <li>• User Agent is not</li> </ul>
HTTP Header	<p>HTTP Request headers can be evaluated as criteria:</p> <ul style="list-style-type: none"> <li>• HTTP Header contains</li> </ul> <p>The HTTP Header contains value should be entered with colon-delimited &lt;name&gt;:&lt;value&gt; .</p>
HTTP Method	<p>HTTP Methods can be evaluated as criteria:</p> <ul style="list-style-type: none"> <li>• HTTP method is</li> <li>• HTTP method is not</li> </ul> <p>Available methods include GET, POST, PUT, DELETE, HEAD, CONNECT, OPTIONS, TRACE, and PATCH.</p>

You can use the IP Whitelist tab to manage whitelists containing trusted IP addresses that bypass all rules and challenges.

## Using the Console



### Note

The WAF uses a first-match algorithm so that once an Access Rule criteria matches, it will stop evaluating future rules. The order of rules matters. Use the API to reorder rules.

### To add an access rule

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view IP Address Whitelists for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **Access Rules** tab.
5. Click **Add Access Rule**.
6. In the Add Access Rule dialog box, enter the following:
  - **Name:** A unique name for the access rule. Avoid entering confidential information.
  - **Rule Condition:** Select the condition that must be met before the rule is matched and specify the details of the condition. Additional conditions can be added in this section.
  - **Rule Action:** Determines the response to a request when the rule is matched. Select one of the following options:
    - **Log and Allow:** A log will be created for all matched requests and no further action will be taken.
    - **Detect Only:** A detection alert will be created for all matched requests and no further action will be taken.
    - **Block:** All matched requests will be blocked and a browser page for the selected response code will be returned.
      - **Block Action:** Select the action that will be taken when a matching request is blocked.
      - **Block Response Code:** Select a response code that will be returned when the request has been blocked. The response code provides

information indicating why the request was blocked. The default response code is 403 "Forbidden".

7. Click **Add Access Rule**. The access rule is added to the access rule list.

### To edit an access rule

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view access rules for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **Access Rules** tab.
5. Select the check box for the access rule you want to update and then select **Edit** from the **Actions** drop-down menu.
6. In the Edit Access Rule dialog box, make the necessary updates and then click **Save**.

### To delete an access rule

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view access rules for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **Access Rules** tab.
5. Select the check box for the access rule you want to delete and then select **Delete** from the **Actions** drop-down menu.

### To add an IP address whitelist

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view IP Address Whitelists for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **IP Whitelist** tab.
5. Click **Add IP Address Whitelist**.
6. In the Add IP Address Whitelist dialog box, enter the following:
  - **Whitelist Name:** A name for the IP addresses used in the list.
  - **IP Addresses:** Select an IP address or enter an IP address and select it to add it. This field supports CIDR notation.
7. Click **Add IP Address Whitelist**.  
The IP Address Whitelist is added to the list of changes to be published.

### To edit an IP Address Whitelist

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view IP Address Whitelists for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **IP Whitelist** tab.
5. Select the check box for the IP Address Whitelist name you want to edit.
6. Click **Edit**.
7. In the Edit IP Address Whitelist dialog box, make the needed changes.
8. Click **Save**.

The IP Address Whitelist change is added to the list of changes to be published.

### To delete an IP Address Whitelist

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view alerts for. The WAF Policy overview appears.
3. Click **Access Control**.
4. Select the **IP Whitelist** tab.
5. Select the check box for the IP Address Whitelist name you want to delete.
6. Click **Delete**.

The deleted IP Address Whitelist is added to the list of changes to be published.

### To publish changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes. Pending changes do not persist across browser sessions. Once you publish changes, it cannot be edited until changes propagate to the edge nodes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
3. Click **Publish All**.
4. In the Publish Changes dialog box, click **Publish All**.

### To discard changes

Updates to your WAF policy appear in the list to be published in Unpublished Changes.

1. Under WAF Policy, click **Unpublished Changes**.
2. In the Unpublished Changes list, click the drop-down arrow beside an unpublished change to review the change.
3. Select the check box for the change you want to discard.
4. Click **Discard**.
5. In the Discard Change dialog box, click **Discard**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the following operations to get an array of all access rules in the policy:

- [UpdateAccessRules](#)
- [ListAccessRules](#)
- [UpdateWhitelists](#)
- [ListWhitelists](#)

To create an access rule:

```
PUT /waasPolicies/{waasPolicyId}/wafConfig/accessRules
```

```
[
 {
 "name": "DetectRequestsToHealthCheck",
 "criteria": [
 {
 "condition": "URL_IS",
 "value": "/health/check"
 }
],
 "action": "DETECT",
 }
]
```

### Address Lists

Use the following API operations to create and manage address lists that can be applied to access rules:

- [CreateAddressList](#)
- [ListAddressLists](#)
- [GetAddressList](#)
- [UpdateAddressList](#)
- [DeleteAddressList](#)

#### Example

To create an address list:

```
POST /addressLists
{
 "addresses": [
 "198.51.100.0",
 "198.51.255.45",
 "198.51.145.55"
],
 "compartmentId": "ocidl.compartment.region1...",
 "displayName": "example IP addresses"
}
```

### Caching Rules

Caching rules allow you to selectively cache requested content on Oracle Cloud Infrastructure's edge servers, such as web pages or certain file types.



### Note

Caching rules are currently only configurable using the [WAAS API](#).

Use the following API operations to create and manage caching rules that can be applied to your WAF configurations:

- [ListCachingRules](#)
- [UpdateCachingRules](#)
- [PurgeCache](#)

## Available Cache Rules Criteria

The criteria of the caching rule determines if the requested content should be cached.

- **URL\_IS** - Matches if the concatenation of requested URL path and query is identical to the contents of the `value` field. For example, if this rule is set to cache the content of `www.example.com/products`, only HTTP requests for `www.example.com/products` will cache.
- **URL\_STARTS\_WITH** - Matches if the concatenation of requested URL path and query starts with the contents of the `value` field. For example, if this rule is set to cache content from `www.example.com/products`, all HTTP requests requesting URLs starting with `www.example.com/products` will be cached and subsequent requests will receive content from the cache, including requests for `www.example.com/products/new-product` and `www.example.com/products/old-product`.
- **URL\_PART\_ENDS\_WITH** - Matches if the concatenation of requested URL path and query ends with the contents of the `value` field. For example, if the rule is set to cache content from URLs that end with `/product.jpg`, HTTP requests for the URLs `www.example.com/products/new-product/product-banner.jpg` and `www.example.com/products/old-product/product-banner.jpg` will be cached and subsequent requests will receive content from the cache.

- **URL\_PART\_CONTAINS** - Matches if the concatenation of requested URL path and query contains the contents of the `value` field. If the rule is set to cache content from URLs that contain `/product-banner`, HTTP requests for the URLs `www.example.com/products/new-product/product-banner/blue.jpg` and `www.example.com/products/new-product/product-banner/red.jpg` will be cached and subsequent requests will receive content from the cache.

### Available Cache Rule Actions

A caching rule can be set to take one of two available actions when receiving a request:

- **CACHE** - Requests matching the criteria of the rule will be cached and subsequent requests will receive content from the cache.
- **BYPASS\_CACHE** - Requests matching the criteria of the rule will bypass the cache and be directed to the origin.

### Cache Duration

Content can be cached for a specified period of time on Oracle Cloud Infrastructure's edge servers or cached locally by the client. The duration is set in the `cachingDuration` and `clientCachingDuration` fields, in ISO 8601 extended format.

### Example of a Caching Rule

```
[
 {
 "action": "CACHE",
 "cachingDuration": "PT20M",
 "clientCachingDuration": "PT20M",
 "criteria": [
 {
 "condition": "URL_IS",
 "value": "/path/to-cache"
 }
],
 "isClientCachingEnabled": true,
 }
]
```

```
"name": "Caching Rule 1"
},
{
 "action": "BYPASS_CACHE",
 "criteria": [
 {
 "condition": "URL_PART_ENDS_WITH",
 "value": "urp-part-not-to-cache"
 }
],
 "isClientCachingEnabled": false,
 "name": "Do not cache"
}
]
```

### Best Practices

The order the caching rules is specified in are important. The rules are processed in the order they are specified in and the first matching rule will be used when processing a request. It is best to add rules that bypass cache to the top of the order and caching rules below any bypass rules.

### Purge Caches

Caches can be purged using the [PurgeCache](#) operation. Caches can either be selectively purged by specifying the URL path of a resource or all caches can be purged for the WAF by not specifying any resources to pass to the API.

#### EXAMPLES

Purge the cache for specified resources:

```
{
 "resources": [
 "/path/to-purge",
 "/multiple-paths"
]
}
```

## Threat Intelligence

WAF has several sources of known IP address threats that are updated daily. The IP address threats are displayed in the following table:

Source	Description
ABUSE  <sup>ch</sup>	Blacklist of "bad" SSL certificates identified by abuse.ch to be associated with malware or botnet activities.
Bambenek Consulting	Active and non-sinkholed Command & Control (C&C) IP addresses.
BlockList.de	Includes IP addresses for hosting phishing sites and other kinds of fraud activities such as ad-click or gaming fraud.
BruteForceBlocker Project	Feed of known IP addresses from blocked SSH brute force attacks.
Proofpoint ET Labs	IP addresses involved in suspicious and malicious activity.
Feodo IP Blocklist	IP addresses used as C&C communication channel by the Feodo Trojan.
Palevo	IP addresses which are being used as botnet C&C for the Palevo crimeware.
Webroot BotNets	Botnet C&C channels and infected zombie machine controlled by Bot master.
Webroot Denial of Service	Includes DOS, DDOS, anomalous sync flood, and anomalous traffic detection.
Webroot Mobile Threats	IP addresses of malicious and unwanted mobile applications. This category leverages data from the Webroot mobile threat research tea.

## CHAPTER 33 Web Application Firewall

Source	Description
Webroot Phishing	IP addresses hosting phishing sites and other kinds of illicit activities such as ad-click or gaming fraud.
Webroot Proxy	IP addresses providing proxy and def services.
Webroot Reputation	IP addresses currently known to be infected with malware. This category also includes IP addresses with an average low Webroot Reputation Index score.
Webroot Scanners	Includes all reconnaissance such as probes, host scan, domain scan and password brute force attacks.
Webroot Spam Sources	Includes tunneling spam messages through proxy, anomalous SMTP activities, and forum spam activities.
Webroot Tor Proxy	Includes IP addresses acting as exit nodes for the Tor Network. Exit nodes are the last point along the proxy chain and make a direct connection to the originator's intended destination.
Webroot Web Attacks	Includes known IP addresses involved in cross site scripting, iFrame injection, SQL injection, cross domain injection, or domain password brute force attacks.
Webroot Windows Exploits	Includes active IP addresses offering or distributing malware, shell code, rootkits, worms or viruses.
ZueS	IP blocklist including known C&C servers/hosts.

### Using the CLI

You can use the CLI to enable threat intelligence sources to block.

Open a command prompt and run the following command to list the keys for all of the threat intelligence:

```
oci waas threat-feed list --waas-policy-id <policy_ocid>
```

## CHAPTER 33 Web Application Firewall

Then parse the keys to block and add them to the JSON:

```
oci waas threat-feed update --threat-feeds '[{"key":"<key_id>","action":"BLOCK"}]' --waas-policy-id <policy_oid>
```

For example:

```
oci waas threat-feed update --threat-feeds '[{"key":"0998d237-bce8-4612-82c8-alca126c0492","action":"BLOCK"}]' --waas-policy-id ocid1.waaspolicy.oc1..aaaaaaaapfa5zrwnn575kru7mrlzkkmcdevp7w551d3phjxtg14s2phuepjq
```

### Using the API

Enabling Threat Intelligence can only be performed by using the API at this time.

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

To return a set of keys for the threat intelligence:

- [ListThreatFeeds](#)



#### Note

Do not use the keys in the example below, as keys are unique across each policy.

```
{
 "8d3f7f1b-673f-4e3a-ba49-08226f385df3": "OFF",
 "0ff7b308-6afe-4b83-91e0-e3ca04afed6e": "OFF",
 "ea5d7c67-1326-43c9-ac31-1df034b9c063": "OFF",
 "87b420ca-5fbb-4ad4-aeba-1b02a9e60b30": "OFF",
 "2168fc70-2d05-466a-9db5-c13c0e32177d": "OFF",
 "7d080a4a-58ce-4370-a02c-f600b3a84e7b": "OFF",
 "a36c7c50-e99e-4b84-9140-5653fc68ce8d": "OFF",
 "5de7bbc1-313f-4995-9810-f6f77cfd30c9": "OFF",
 "fd2152cc-14f5-4471-a58b-d94cc8a61444": "OFF",
 "cfacd3d3-65d9-4368-93e0-62c906e7a748": "OFF",
```

## CHAPTER 33 Web Application Firewall

```
"6eb86368-01ea-4e94-ac1b-49bf0e551443": "OFF",
"aabb45d9-0d75-481d-9568-58ecad217e1e": "OFF",
"3805ecc2-1d6d-428b-a03e-2a0fe77fd46f": "OFF",
"c3452861-4910-4f3a-9872-22cf92d424eb": "OFF",
"4cf31deb-11af-460e-a46a-ecc1946a6688": "OFF",
"eff34d63-6235-4081-976d-acd39248bdc3": "OFF",
"1d1c94d9-038b-45eb-acd4-fb422e281f4c": "OFF",
"687b5ff4-b1b6-4d12-8dba-3ea90b4536a1": "OFF",
"65cf274d-991b-41f8-adda-6fe60ba2704f": "OFF"
}
```

To set all threats to DETECT:

- [UpdateThreatFeeds](#)

With body:

```
[
{"action":"DETECT","key":"8d3f7f1b-673f-4e3a-ba49-08226f385df3"},
{"action":"DETECT","key":"0ff7b308-6afe-4b83-91e0-e3ca04afed6e"},
{"action":"DETECT","key":"ea5d7c67-1326-43c9-ac31-1df034b9c063"},
{"action":"DETECT","key":"87b420ca-5fbb-4ad4-aeba-1b02a9e60b30"},
{"action":"DETECT","key":"2168fc70-2d05-466a-9db5-c13c0e32177d"},
{"action":"DETECT","key":"7d080a4a-58ce-4370-a02c-f600b3a84e7b"},
{"action":"DETECT","key":"a36c7c50-e99e-4b84-9140-5653fc68ce8d"},
{"action":"DETECT","key":"5de7bbcl-313f-4995-9810-f6f77cfd30c9"},
{"action":"DETECT","key":"fd2152cc-14f5-4471-a58b-d94cc8a61444"},
{"action":"DETECT","key":"cfacd3d3-65d9-4368-93e0-62c906e7a748"},
{"action":"DETECT","key":"6eb86368-01ea-4e94-ac1b-49bf0e551443"},
{"action":"DETECT","key":"aabb45d9-0d75-481d-9568-58ecad217e1e"},
{"action":"DETECT","key":"3805ecc2-1d6d-428b-a03e-2a0fe77fd46f"},
{"action":"DETECT","key":"d9cfc537-dd50-427d-830e-a612f535c11f"},
{"action":"DETECT","key":"c3452861-4910-4f3a-9872-22cf92d424eb"},
{"action":"DETECT","key":"4cf31deb-11af-460e-a46a-ecc1946a6688"},
{"action":"DETECT","key":"eff34d63-6235-4081-976d-acd39248bdc3"},
{"action":"DETECT","key":"1d1c94d9-038b-45eb-acd4-fb422e281f4c"},
{"action":"DETECT","key":"687b5ff4-b1b6-4d12-8dba-3ea90b4536a1"},
{"action":"DETECT","key":"65cf274d-991b-41f8-adda-6fe60ba2704f"}
]
```

This will return a 202 Accepted HTTP status, which means the policy will enter an UPDATING state until changes are provisioned to the edge nodes.

### Settings

To use SSL with your WAF policy, you must add a certificate bundle. The certificate bundle you upload includes the public certificate and the corresponding private key. Self-signed certificates can be used for the internal communication within Oracle Cloud Infrastructure.

### Working with SSL Certificates

Oracle Cloud Infrastructure accepts third-party and self-signed certificates in PEM format only. The following is an example PEM encoded certificate:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
```

### Obtaining Third-Party SSL Certificates

You can purchase an SSL certificate from a trusted Certificate Authority such as [Symantec](#), [Thawte](#), [RapidSSL](#), or [GeoTrust](#). The certificate issuer will provide an SSL certificate that includes a certificate, intermediate certificate, and private key. Use this information, including the intermediate certificate, when adding an SSL certificate to Oracle Cloud Infrastructure.

### Converting to PEM format

If you receive your certificates and keys in formats other than PEM, you must convert them before you can upload them to the system. You can use [OpenSSL](#) to convert certificates and keys to PEM format.

### Uploading Certificate Chains

If you have multiple certificates that form a single certification chain, you must include all relevant certificates in one file before you upload them to the system. The following example of a certificate chain file includes four certificates:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
```

## CHAPTER 33 Web Application Firewall

---

```
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
```

### Submitting Private Keys

If your private key submission returns an error, the most common reasons are your private key is malformed or the system does not recognize the encryption method used for your key.

#### PRIVATE KEY CONSISTENCY

If you receive an error related to the private key, you can use OpenSSL to check its consistency:

```
openssl rsa -check -in <private_key>.pem
```

This command verifies that the key is intact, the passphrase is correct, and the file contains a valid RSA private key.

#### DECRYPTING A PRIVATE KEY

If the system does not recognize the encryption technology used for your private key, decrypt the key. Upload the unencrypted version of the key with your certificate bundle. You can use OpenSSL to decrypt a private key:

```
openssl rsa -in <private_key>.pem -out <decrypted_private_key>.pem
```

### Using the Console

#### To edit WAF settings

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view settings for. The WAF Policy overview appears.
3. Click **Settings**.
4. Click **Edit**.
5. In the **Edit Settings** dialog box, enter the following:
  - **WAF Origin:** Select the name and IP address of the origin.
  - **Enable HTTPS Support:** When enabled, all communications between the browser and web app are encrypted. Enter the following information:
    - **SSL Certificate:** Drag and drop, select, or paste a valid SSL certificate in PEM format. You must also include intermediate certificates (the website certificate must be first). The following is an example:

```
-----BEGIN CERTIFICATE-----
<Base64_encoded_certificate>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate_Base64_encoded_certificate>
-----END CERTIFICATE-----
```

- **Private Key:** Drag and drop, select, or paste a valid private key in PEM format in this field. The private key cannot be protected by a passphrase. The following is an example:

```
-----BEGIN PRIVATE KEY-----
<Base64_encoded_private_key>
-----END PRIVATE KEY-----
```

- **Self Signed Certificate:** Enable this field when using a self-signed certificate to show an SSL warning in the browser.

- **HTTP to HTTPS Redirect:** When enabled, all HTTP traffic is automatically redirected to HTTPS.
6. Click **Save**.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

- [CreateCertificate](#)
- [GetCertificate](#)
- [DeleteCertificate](#)

### Logs

Logs displays log activity and the details of each logged event within a specified time frame. Logs enable you to understand what rules and countermeasures are triggered by requests and are used as a basis to move request handling into block mode. Logs can come from Access Control, Protection Rules, or Bot events.



#### Note

If you have concerns over General Data Protection Regulation (GDPR) requirements, Logs can be disabled for the WAF service. You can use My Oracle Support to file a service request to disable Logs.

### Using the Console

#### To view Logs

1. Open the navigation menu. Under **Governance and Administration**, go to **Security** and click **WAF Policies**.
2. Click the name of the WAF Policy you want to view logs for. The WAF Policy overview appears.
3. Click **Logs**. Logs for the WAF policy appear.
4. To help find a log, you can use the following filter options:
  - To view alerting activity data for a specific time range, enter a **Start Date, Start Time, End Date, or End Time**.
  - To view logs for a URL, enter a **Request URL**.
  - To view logs for a client IP address, select an address from the **Client IP Address** drop-down menu.
  - To view logs for a country, select a country from the **Country Name** drop-down list.
  - To find a type of action, select an **Action** check box.
  - To find a type of log, select a **Log Type** check box from the following options:
    - Access Rules
    - CAPTCHA Challenge
    - JavaScript Challenge
    - Protection Rules
    - Human Interaction Challenge
    - Device Fingerprinting Challenge
    - Threat Intelligence Feeds
    - Address Rate Limiting
    - Access

5. Click the plus sign next to the Alert Type you want to view.

### Using the API

For information about using the API and signing requests, see [REST APIs](#) and [Security Credentials](#). For information about SDKs, see [Software Development Kits and Command Line Interface](#).

Use the [ListWafLogs](#) API operation to display log activity. You can use filters to help find a log.

### Example

You can filter logs by the following logType:

- ACCESS\_RULES
- CAPTCHA\_CHALLENGE
- JAVASCRIPT\_CHALLENGE
- PROTECTION\_RULES
- HUMAN\_INTERACTION\_CHALLENGE
- DEVICE\_FINGERPRINT\_CHALLENGE
- THREAT\_INTELLIGENCE\_FEEDS
- ADDRESS\_RATE\_LIMITING
- ACCESS

Logs can be filtered by logType by making the following request:

```
GET /20181116/waasPolicies/<unique_ID>
/wafLogs?logType=
<logType>
&timeObservedGreaterThanOrEqualTo=<timestamp>&timeObservedLessThan=<timestamp>&compartmentId=<unique_ID>
```

For example:

## CHAPTER 33 Web Application Firewall

---

```
GET /20181116/waasPolicies/ocid1.waaspolicy.oc1..wafLogs?logType=PROTECTION_
RULES&timeObservedGreaterThanOrEqualTo=2019-10-24T13:00:00+00:00&timeObservedLessThan=2019-10-
24T13:47:00+00:00&compartmentId=ocid1.compartment.oc1..
```

The following response output for the filtered logs is returned:

```
[
 {
 "action": "BLOCK",
 "clientAddress": "192.0.2.0",
 "countryCode": "US",
 "countryName": "United States",
 "domain": "example.com",
 "httpHeaders": {
 "Accept": "*/*",
 "Host": "example.com",
 "Referer": "",
 "Request-Id": "2019-10-24T13:46:25Z|fa68cab479|192.0.2.0|uwDPcqR0Qt",
 "User-Agent": "curl/7.54.0",
 "X-Client-Ip": "192.0.2.0",
 "X-Country-Code": "US",
 "X-Forwarded-For": "192.0.2.0, 192.0.2.0"
 },
 "httpMethod": "GET",
 "httpVersion": "HTTP/1.1",
```

## CHAPTER 33 Web Application Firewall

```
"incidentKey": "2019-10-24T13:46:25Z|fa68cab479|192.0.2.0|uwDPcqR0Qt",

"logType": "PROTECTION_RULES",

"protectionRuleDetections": {

 "950002": {

 "Message": "System Command Access. Matched Data: cmd.exe found within ARGS:abc:
cmd.exe",

 "Message details": "Access denied with code 403 (phase 2). Pattern match \"\\\\\\\\b(?:?:n
(?:map|et|c)|w(?:guest|sh)|telnet|rcmd|ftp)\\\\\\\\.exe\\\\\\\\b|cmd(?:
(?:32)?\\\\\\\\.exe\\\\\\\\b|\\\\\\\\b\\\\\\\\W*?\\\\\\\\/c))\" at ARGS:abc."

 }

},

"requestUrl": "/?abc=cmd.exe",

"timestamp": "Thu, 24 Oct 2019 13:46:25 GMT",

"userAgent": "curl/7.54.0"

},

{

 "action": "BLOCK",

 "clientAddress": "192.0.2.0",

 "countryCode": "US",

 "countryName": "United States",

 "domain": "example.com",

 "httpHeaders": {
```

## CHAPTER 33 Web Application Firewall

---

```
"Accept": "*/*",

"Host": "example.com",

"Referer": "",

"Request-Id": "2019-10-24T13:46:25Z|43bd96b710|192.0.2.0|E04WECJbcY",

"User-Agent": "curl/7.54.0",

"X-Client-Ip": "192.0.2.0",

"X-Country-Code": "US",

"X-Forwarded-For": "192.0.2.0, 192.0.2.0"

},

"httpMethod": "GET",

"httpVersion": "HTTP/1.1",

"incidentKey": "2019-10-24T13:46:25Z|43bd96b710|192.0.2.0|E04WECJbcY",

"logType": "PROTECTION_RULES",

"protectionRuleDetections": {

 "950002": {

 "Message": "System Command Access. Matched Data: cmd.exe found within ARGS:abc:

cmd.exe",

 "Message details": "Access denied with code 403 (phase 2). Pattern match \"\\\\\\\\\\\\b(?:?:n

(?:map|et|c)|w(?:guest|sh)|telnet|rcmd|ftp)\\\\\\\\.exe\\\\\\\\b|cmd(?::

(?:32)?\\\\\\\\.exe\\\\\\\\b|\\\\\\\\b\\\\\\\\W*?\\\\\\\\/c))\" at ARGS:abc."

 }

},

"requestUrl": "/?abc=cmd.exe",
```

## CHAPTER 33 Web Application Firewall

---

```
 "timestamp": "Thu, 24 Oct 2019 13:46:25 GMT",

 "userAgent": "curl/7.54.0"

 }

]
```

### Example

Logs can be filtered by `clientAddress` and time range by making the following request:

```
GET /20181116/waasPolicies/<unique_ID>/wafLogs?clientAddress=<IP
address>
&timeObservedGreaterThanOrEqualTo=<timestamp>&timeObservedLessThan=<timestamp>&compartmentId=<unique_ID>
```

For example:

```
GET
/20181116/waasPolicies/ocid1.waaspolicy.oc1../wafLogs?clientAddress=192.0.2.0&timeObservedGreaterThanOrEqualTo=2019-10-24T13:26:47+00:00&timeObservedLessThan=2019-10-24T13:26:56+00:00&compartmentId=ocid1.compartment.oc1..
```

The following response output for the filtered logs is returned:

```
[

 {

 "clientAddress": "192.0.2.0",

 "countryName": "Unknown",

 "domain": "example.com",

 "fingerprint": "-",

 "httpHeaders": {

 "Accept": "*/*",

 "Host": "example.com",
```

## CHAPTER 33 Web Application Firewall

---

```
 "Referer": "",
 "User-Agent": "curl/7.54.0",
 "X-Client-Ip": "192.0.2.01",
 "X-Country-Code": "AU",
 "X-Forwarded-For": "192.0.2.0, 192.0.2.0"
 },
 "httpMethod": "GET",
 "httpVersion": "1.1",
 "incidentKey": "2019-10-24T13:26:55Z|43bd96b710|192.0.2.0|ytQbBpuerK",
 "logType": "ACCESS",
 "originAddress": "130.35.212.39:80",
 "originResponseTime": "0.2500",
 "requestUrl": "/",
 "responseCode": 200,
 "responseSize": 4978,
 "timestamp": "Thu, 24 Oct 2019 13:26:55 GMT",
 "userAgent": "curl/7.54.0"
},
{
 "clientAddress": "192.0.2.0",
 "countryName": "Unknown",
```

## CHAPTER 33 Web Application Firewall

---

```
"domain": "example.com",
"fingerprint": "-",
"httpHeaders": {
 "Accept": "*/*",
 "Host": "example.com",
 "Referer": "",
 "User-Agent": "curl/7.54.0",
 "X-Client-Ip": "192.0.2.0",
 "X-Country-Code": "AU",
 "X-Forwarded-For": "192.0.2.0, 192.0.2.0"
},
"httpMethod": "GET",
"httpVersion": "1.1",
"incidentKey": "2019-10-24T13:26:53Z|4d7583f67c|192.0.2.0|KR8qhtyJnG",
"logType": "ACCESS",
"originAddress": "198.51.100.0:24",
"originResponseTime": "0.5070",
"requestUrl": "/",
"responseCode": 200,
"responseSize": 4978,
"timestamp": "Thu, 24 Oct 2019 13:26:54 GMT",
```

## CHAPTER 33 Web Application Firewall

---

```
 "userAgent": "curl/7.54.0"
 }
]
```

# CHAPTER 34 Developer Tools

This chapter includes general information about using the Oracle Cloud Infrastructure REST API and developer tools.

## Software Development Kits and Command Line Interface

Oracle Cloud Infrastructure provides a number of Software Development Kits (SDKs) and a Command Line Interface (CLI) to facilitate development of custom solutions.

- **Software Development Kits (SDKs)**  
Build and deploy apps that integrate with Oracle Cloud Infrastructure services. Each SDK provides the tools you need to develop an app, including code samples and documentation to create, test, and troubleshoot. In addition, if you want to contribute to the development of the SDKs, they are all open source and available on GitHub.
  - [SDK for Java](#)
  - [Python SDK](#)
  - [Ruby SDK](#)
  - [Go SDK](#)
- **[Command Line Interface \(CLI\)](#)**  
The CLI provides the same core capabilities as the Oracle Cloud Infrastructure Console, plus additional commands that can extend the Console's functionality. Convenient for developers or anyone who prefers the command line to a GUI.

### SDK for Java

The SDK for Java enables you to write code to manage Oracle Cloud Infrastructure resources.

This SDK and sample is dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

**Download:** [GitHub](#) or [Maven](#).

### Requirements

To use the SDK for Java, you must have the following:

- An Oracle Cloud Infrastructure account.
- A user created in that account, in a group with a policy that grants the desired permissions. This can be a user for yourself, or another person/system that needs to call the API. For an example of how to set up a new user, group, compartment, and policy, see [Adding Users](#). For a list of typical policies you may want to use, see [Common Policies](#).
- A key pair used for signing API requests, with the public key uploaded to Oracle. Only the user calling the API should be in possession of the private key. For more information, see [Configuring Credentials](#).
- Java 8
- A TTL value of 60. For more information, see [Java Virtual Machine TTL for DNS Name Lookups](#).

### Services Supported

- Analytics Cloud
- Announcements
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Compute Work Requests
- Content and Experience
- Core Services (Networking, Compute, Block Volume)
- Data Transfer
- Database
- Digital Assistant

- DNS
- Email Delivery
- Events
- File Storage
- Functions
- Health Checks
- IAM
- Integration
- Key Management
- Limits
- Load Balancing
- Monitoring
- Notifications
- Object Storage
- Quotas
- Resource Manager
- Search
- Streaming
- Web Application Acceleration and Security

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the SDK for Java is released, subscribe to the [Atom feed](#).

### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) and [oci-java-sdk](#) tags in your post
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Getting Started

This topic describes how to install and configure the Oracle Cloud Infrastructure SDK for Java.

#### DOWNLOADING THE SDK FROM GITHUB

You can download the SDK for Java as a zip archive from [GitHub](#). It contains the SDK, all of its dependencies, documentation, and examples. For best compatibility and to avoid issues, use the version of the dependencies included in the archive. Some notable issues are:

- Bouncy Castle: The SDK bundles 1.60, but if you need FIPS compliance, you must download and use a FIPS-certified version. The SDK supports bc-fips 1.0.1 and bcpkix-fips 1.0.1. You can download them at: <https://www.bouncycastle.org/fips-java/>
- Jersey Core and Client: The SDK bundles 2.24.1, which is required to support large object uploads to Object Storage. Older versions will not support uploads greater than ~2.1 GB.
- Jax-RS API: The SDK bundles 2.0.1 of the spec. Older versions will cause issues.



#### Note

The SDK for Java is bundled with Jersey, but you can also use your own JAX-RS implementation. For details, see [Using Your Own JAX-RS Implementation](#)

#### DOWNLOADING THE SDK FROM MAVEN OR JCENTER

The SDK for Java is available on [Maven Central](#) and [JCenter](#).

To use the Oracle Cloud Infrastructure SDK for Java in your project, import the `oci-java-sdk-bom`, followed by your project dependencies. For example:

```
<dependencyManagement>
 <dependencies>
 <dependency>
 <groupId>com.oracle.oci.sdk</groupId>
 <artifactId>oci-java-sdk-bom</artifactId>
 <!-- replace the version below with your required version -->
 <version>1.5.2</version>
 <type>pom</type>
 <scope>import</scope>
 </dependency>
 </dependencies>
</dependencyManagement>
<dependencies>
<dependency>
 <groupId>com.oracle.oci.sdk</groupId>
 <artifactId>oci-java-sdk-audit</artifactId>
</dependency>
<dependency>
 <groupId>com.oracle.oci.sdk</groupId>
 <artifactId>oci-java-sdk-core</artifactId>
</dependency>
<dependency>
 <groupId>com.oracle.oci.sdk</groupId>
 <artifactId>oci-java-sdk-database</artifactId>
</dependency>
<!-- more dependencies if needed -->
```

### CONFIGURING THE SDK

The SDK services need two types of configuration: credentials and client-side HTTP settings.

#### *CONFIGURING CREDENTIALS*

First, you need to set up your credentials and config file. For instructions, see [SDK and CLI Configuration File](#).

Next you need to set up the client to use the credentials. The credentials are abstracted through an `AuthenticationDetailsProvider` interface. Clients can implement this however you choose. We have included a simple POJO/builder class to help with this task (`SimpleAuthenticationDetailsProvider`).

- You can load a config with or without a profile:

```
ConfigFile config
 = ConfigFileReader.parse("~/oci/config");
ConfigFile configWithProfile
 = ConfigFileReader.parse("~/oci/config", "DEFAULT");
```

- The private key supplier can be created with the file path directly, or using the config file:

```
Supplier<InputStream> privateKeySupplier
 = new SimplePrivateKeySupplier("~/oci/oci_api_key.pem");
Supplier<InputStream> privateKeySupplierFromConfigEntry
 = new SimplePrivateKeySupplier(config.get("key_file"));
```

- To create an auth provider using the builder:

```
AuthenticationDetailsProvider provider
 = SimpleAuthenticationDetailsProvider.builder()
 .tenantId("myTenantId")
 .userId("myUserId")
 .fingerprint("myFingerprint")
 .privateKeySupplier(privateKeySupplier)
 .build();
```

- To create an auth provider using the builder with a config file:

```
AuthenticationDetailsProvider provider
 = SimpleAuthenticationDetailsProvider.builder()
 .tenantId(config.get("tenancy"))
 .userId(config.get("user"))
 .fingerprint(config.get("fingerprint"))
 .privateKeySupplier(privateKeySupplier)
 .build();
```

- Finally, if you use standard config file keys and the standard config file location, you can simplify this further by using `ConfigFileAuthenticationDetailsProvider`:

## CHAPTER 34 Developer Tools

---

```
AuthenticationDetailsProvider provider
 = new ConfigFileAuthenticationDetailsProvider("ADMIN_USER");
```

### CONFIGURING CLIENT-SIDE OPTIONS

Create a client-side configuration through the `ClientConfiguration` class. If you do not provide your own configuration, the SDK for Java uses a default configuration. To provide your own configuration, use the following:

```
ClientConfiguration clientConfig
 = ClientConfiguration.builder()
 .connectionTimeoutMillis(3000)
 .readTimeoutMillis(60000)
 .build();
```

After you have both a credential configuration and the optional client configuration, you can start creating service instances.

### CONFIGURING CUSTOM OPTIONS

In the config file, you can insert custom key-value pairs that you define, and then reference them as necessary. For example, you could specify a frequently used compartment ID in the config file like so (highlighted in red italics):

```
[DEFAULT]
user=ocidl.user.oc1..aaaaaaaat5nvwcn5j6aqzjcm5eqbb6qt2jvvpkanghtgdaqedqw3rynjq
fingerprint=20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34
key_file=~/.oci/oci_api_key.pem
tenancy=ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2bcm5dt2j6rx32uzr4h25vqstifsfdsq
custom_compartment_
id=ocidl.compartment.oc1..aaaaaaaayzfqeibduyox6iib3o1cmdar3ugly4fmameq4h7lcdlihrvur7xq
```

Then you can retrieve the value like so:

```
ConfigFile config
 = ConfigFileReader.parse("~/oci/config");

String compartmentId = config.get("custom_compartment_id");
```

### Configuration

This topic provides details on compatibility, advanced configurations, and add-ons for the Oracle Cloud Infrastructure SDK for Java.

#### SECURITY MANAGER PERMISSIONS

If your application needs to run inside the [Java Security Manager](#), you must grant additional permissions by updating a policy file, or by specifying an additional or a different policy file at runtime.

The SDK requires the following permissions:

- Required by Jersey:

```
permission java.lang.RuntimePermission "getClassLoader";
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";
permission java.lang.RuntimePermission "accessDeclaredMembers";
permission java.util.PropertyPermission "*", "read,write";
permission java.lang.RuntimePermission "setFactory";
```

- Required by the SDK to overwrite reserved headers:

```
permission java.util.PropertyPermission "sun.net.http.allowRestrictedHeaders", "write";
```

- Required by the SDK to open socket connections:

```
permission java.net.SocketPermission "*", "connect";
```

To include another policy file, in addition to Java Runtime Environment's default policy file, launch the Java Virtual Machine with:

```
java -Djava.security.manager -Djava.security.policy=</path/to/other_policy>
```

To replace the default policy file, launch the Java Virtual Machine with:

```
java -Djava.security.manager -Djava.security.policy==</path/to/other_policy>
```



### Note

Use a single equals sign (=) when supplying an additional policy file. Use a double equals sign (==) only if you wish to replace the default policy file.

### JAVA VIRTUAL MACHINE TTL FOR DNS NAME LOOKUPS

The Java Virtual Machine (JVM) caches DNS responses from lookups for a set amount of time, called *time-to-live* (TTL). This ensures faster response time in code that requires frequent name resolution.

The JVM uses the [networkaddress.cache.ttl](#) property to specify the caching policy for DNS name lookups. The value is an integer that represents the number of seconds to cache the successful lookup. The default value for many JVMs, `-1`, indicates that the lookup should be cached forever.

Because resources in Oracle Cloud Infrastructure use DNS names that can change, we recommend that you change the the TTL value to 60 seconds. This ensures that the new IP address for the resource is returned on next DNS query. You can change this value globally or specifically for your application:

- To set TTL globally for all applications using the JVM, add the following in the `$JAVA_HOME/jre/lib/security/java.security` file:

```
networkaddress.cache.ttl=60
```

- To set TTL only for your application, set the following in your application's initialization code:

```
java.security.Security.setProperty("networkaddress.cache.ttl" , "60");
```

### APACHE CONNECTOR ADD-ON

The `oci-java-sdk-addons-apache` is an optional add-on to the SDK for Java that allows for configuring a client connection pool and an HTTP proxy. The add-on leverages the Jersey `ApacheConnectorProvider` instead of the SDK's default `HttpURLConnectionProvider` when making service calls. The add-on can be found in the `bmc-addons` directory of the SDK.

## CHAPTER 34 Developer Tools

---

For details on installation and configuration, see the [Readme](#) for the add-on.

### USING YOUR OWN JAX-RS IMPLEMENTATION

The SDK for Java is bundled with Jersey, but you can also use your own JAX-RS implementation. For an example of how to configure your own implementation, see [RETEasy Client Configurator Add-On](#) and accompanying code samples.

#### *RESTEASY CLIENT CONFIGURATOR ADD-ON*

The `oci-java-sdk-addons-resteasy-client-configurator` is provided to demonstrate how to configure an alternate JAX-RS implementation. The add-on can be found in the `bmc-addons` directory of the SDK.

For details on installation and configuration, see the [Readme](#) for the add-on.

For code samples that demonstrate how to configure the client, see:

- [ResteasyClientExample.java](#)
- [ResteasyClientWithObjectStorageExample.java](#)
- [InstancePrincipalsAuthenticationDetailsProviderWithResteasyClientExample.java](#)

### USING SLF4J FOR LOGGING

Logging in the SDK is done through [SLF4J](#). SLF4J is a logging abstraction that allows the use of a user-supplied logging library (e.g., log4j). For more information, see the [SLF4J manual](#).

The following is an example that enables basic logging to standard out. More advanced logging options can be configured by using the log4j binding.

1. Download the SLF4J Simple binding jar: [SLF4J Simple Binding](#)
2. Add the jar to your classpath (e.g., add it to the `/third-party/lib` directory of the SDK download)
3. Add the following VM arg to enable debug level logging (by default, info level is used): -  
`Dorg.slf4j.simpleLogger.defaultLogLevel=debug`

### Concepts

This topic explains some of the key concepts for using the Oracle Cloud Infrastructure SDK for

Java.

### SYNCHRONOUS CALLS

To make synchronous calls, create an instance of the synchronous client. The general pattern for synchronous clients is that for a service named Example, there will be an interface named ExampleService, and the synchronous client implementation will be called ExampleServiceClient. Here's an example of creating an Object Storage client:

```
AuthenticationDetailsProvider provider = ...;
ObjectStorage clientWithDefaultClientConfig = new ObjectStorageClient(provider);
clientWithDefaultClientConfig.setRegion(Region.US_ASHBURN_1);

ClientConfiguration clientConfig = ...;
ObjectStorage clientWithExplicitClientConfig = new ObjectStorageClient(provider, clientConfig);
clientWithExplicitClientConfig.setRegion(Region.US_ASHBURN_1);
```

Synchronous calls will block until the response is available. All SDK APIs return a response object (regardless of whether or not the API sends any content back). The response object typically contains at least a request ID that you can use when contacting Oracle support for help on a particular request.

```
ObjectStorage client = ...;
GetBucketResponse response = client.getBucket(
 GetBucketRequest.builder().namespaceName("myNamespace").bucketName("myBucket").build());
String requestId = response.getOpcRequestId();
Bucket bucket = response.getBucket();
System.out.println(requestId);
System.out.println(bucket.getName());
```

### ASYNCHRONOUS CALLS

To make asynchronous calls, create an instance of the asynchronous client. The general pattern for asynchronous clients is that for a service named Example, there will be an interface named ExampleServiceAsync, and the asynchronous client implementation will be called ExampleServiceAsyncClient. Here's an example of creating an Object Storage client:

```
AuthenticationDetailsProvider provider = ...;
ObjectStorageAsync clientWithDefaultClientConfig = new ObjectStorageAsyncClient(provider);
clientWithDefaultClientConfig.setRegion(Region.US_ASHBURN_1);
```

## CHAPTER 34 Developer Tools

---

```
ClientConfiguration clientConfig = ...;
ObjectStorageAsync clientWithExplicitClientConfig = new ObjectStorageAsyncClient(provider,
clientConfig);
clientWithExplicitClientConfig.setRegion(Region.US_ASHBURN_1);
```

Asynchronous calls return immediately. You need to provide an `AsyncHandler` that will be invoked after the call completes either successfully or unsuccessfully:

```
ObjectStorageAsync client = ...;

AsyncHandler<GetBucketRequest, GetBucketResponse> handler = new AsyncHandler<GetBucketRequest,
GetBucketResponse>() {
 @Override
 public void onSuccess(GetBucketRequest request, GetBucketResponse response) {
 String requestId = response.getOpcRequestId();
 Bucket bucket = response.getBucket();
 System.out.println(requestId);
 System.out.println(bucket.getName());
 }

 @Override
 public void onError(GetBucketRequest request, Throwable error) {
 error.printStackTrace();
 }
};

Future<GetBucketResponse> future = client.getBucket(
 GetBucketRequest.builder().namespaceName("myNamespace").bucketName("myBucket").build(),
 handler);
```

### POLLING WITH WAITERS

The SDK offers waiters that allow your code to wait until a specific resource reaches a desired state. A waiter can be invoked in both a blocking or a non-blocking (with asynchronous callback) manner, and will wait until either the desired state is reached or a timeout is exceeded. Waiters abstract the polling logic you would otherwise have to write into an easy-to-use single method call.

Waiters are obtained through the service client (`client.getWaiters()`). Both a `Get<Resource>Request` and the desired lifecycle state are passed in to the `waiters.for<Resource>` method. For example:

## CHAPTER 34 Developer Tools

```
public static Instance waitForInstanceProvisioningToComplete(ComputeClient computeClient, String
instanceId) throws Exception {

 ComputeWaiters waiters = computeClient.getWaiters();
 GetInstanceResponse response = waiters.forInstance(
 GetInstanceRequest.builder().instanceId(instanceId).build(),
 Instance.LifecycleState.Running)
 .execute();

 return response.getInstance();
}
```

Each `waiters.for<Resource>` method has two versions:

- One version uses the default polling values. For example:

```
waiters.forInstance(GetInstanceRequest, LifecycleState)
```

- The other version gives you full control over how long to wait and how much time between polling attempts. For example:

```
waiters.forInstance(GetInstanceRequest, LifecycleState, TerminationStrategy, DelayStrategy)
```

### THREADING MODEL

A client becomes thread-safe when it is initialized. After setting its endpoint, you can safely use a client in multiple threads and concurrently call methods on it.

You can reuse a client on multiple requests, both across concurrent threads or within a single thread. Unless the environment's resources are constrained, you should only close the client immediately before it goes out of scope.



#### Note

This guarantee applies only to the default JAX-RS implementation, Jersey. When using an alternate implementation, you must manage thread safety yourself. For more information, see [Using Your Own JAX-RS Implementation](#)

### UPLOADING LARGE OBJECTS

The Object Storage service supports multipart uploads to make large object uploads easier by splitting the large object into parts. The SDK for Java supports raw multipart upload operations for advanced use cases, as well as a higher level upload class that uses the multipart upload APIs. [Managing Multipart Uploads](#) provides links to the APIs used for multipart upload operations. Higher level multipart uploads are implemented using the [UploadManager](#), which will: split a large object into parts for you, upload the parts in parallel, and then recombine and commit the parts as a single object in storage.

The [UploadObject](#) example shows how to use the UploadManager to automatically split an object into parts for upload to simplify interaction with the Object Storage service.

### RAW REQUESTS

Raw requests are useful, and in some cases necessary. Typical use cases are: when using your own HTTP client, making a OCI-authenticated request to an alternate endpoint, and making a request to a OCI API that is not currently supported in the SDK. The SDK for Java exposes the `DefaultRequestSigner` class that you can use to create a `RequestSigner` instance for non-standard requests.

The [Raw Request](#) example on GitHub shows how to:

- create an authentication provider and request signer
- integrate with an HTTP client (Jersey in this example) to authenticate requests

### SETTING THE ENDPOINTS

Service endpoints can be set in three ways.

- Call `setEndpoint()` on the service instance. This lets you specify a full host name (for example, `https://www.example.com`).
- Call `setRegion()` on the service instance. This selects the appropriate host name for the service for the given region. However, if the service is not supported in the region you set, the Java SDK returns an error.
- Pass the region in the configuration file. For more information, see [SDK and CLI Configuration File](#).

Note that a service instance cannot be used to communicate with different regions. If you need to make requests to different regions, create multiple service instances.

### FORWARD COMPATIBILITY AND ENUMS

If you have conditional logic based on an enum, be sure that your code handles the `UnknownEnumValue` case to ensure forward compatibility. Some response fields are of type enum, but in the future, individual services may return values not covered by existing enums for that field. To address this possibility, every response field of type enum has an additional value named `UnknownEnumValue`. If a service returns a value that is not recognized by your version of the SDK, then the response field will be set to this value.

### NEW REGION SUPPORT

If you are using a version of the SDK released prior to the announcement of a new region, you can use a workaround to reach it.

A *region* is a localized geographic area. For more information on regions and how to identify them, see [Regions and Availability Domains](#).

A *realm* is a set of regions that share entities. You can identify your realm by looking at the domain name at the end of the network address. For example, the realm for `xyz.abc.123.oraclecloud.com` is `oraclecloud.com`.

You must first call `Region.register` to register the new region, and then you can set the region by either using the configuration file or by calling the `setRegion` method.



#### Note

Once a region is registered, the federation endpoint is no longer required while using instance principals. For an example, see <https://github.com/oracle/oci-java-sdk/blob/master/bmc-examples/src/main/java/NewRegionAndRealmSupportWithoutSDKUpdate.java>.

## CHAPTER 34 Developer Tools

---

### ORACLECLOUD.COM REALM

For regions in the `oraclecloud.com` realm, you can pass new region names just as you would pass ones that are already defined in the [Region](#) enum for your SDK version.



#### Note

For the following code samples, be sure to supply the appropriate endpoints for your region.

If you are using version 1.2.34 or later of the SDK for Java, you can pass the new region name as a string using one of the following methods:

- To set the region on a previously created client:

```
client.setRegion("ca-toronto-1");
```

- To set a region when building a new client:

```
Identity identityClient = IdentityClient.builder()
 .region("ca-toronto-1")
 .build(provider);
```

- You can also pass the region in the configuration file. For more information, see [SDK and CLI Configuration File](#).

### OTHER REALMS

For regions in realms other than `oraclecloud.com`, you can use the following workarounds to reach new regions with earlier versions of the SDK.

To specify the endpoint:

```
AuthenticationDetailsProvider provider =
 new ConfigFileAuthenticationDetailsProvider(configurationFilePath, profile);

IdentityClient client = IdentityClient.builder()
 .endpoint("https://identity.ca-toronto-1.oraclecloud.com")
 .build(provider);
```

If you are authenticating via instance principals, you can set the endpoint and

federationEndpoint via the following process:

```
InstancePrincipalsAuthenticationDetailsProvider provider =
InstancePrincipalsAuthenticationDetailsProvider.builder()
 .federationEndpoint("https://auth.ca-toronto-1.oraclecloud.com/v1/x509")
 .build();

IdentityClient identityClient = IdentityClient.builder()
 .endpoint("https://identity.ca-toronto-1.oraclecloud.com")
 .build(provider);
```

### PAGINATED RESPONSES

Some APIs return paginated result sets, so you must check for additional items and if necessary, fetch the next page. You can do so manually or you can use an iterator.

#### *MANUALLY FETCHING PAGES*

The Response objects contain a method to fetch the next page token. If the token is null, there are no more items. If it is not null, you can make an additional request, by setting the token on the Request object, to get the next page of responses.



#### **Note**

Some APIs may return a token even if there are no additional results. Be sure to also check whether any items were returned and stop if there are none.

This example shows how to handle page tokens returned by the Object Storage API:

```
ObjectStorage client = ...;

ListBucketsRequest.Builder builder =
 ListBucketsRequest.builder().namespaceName(namespace);
String nextPageToken = null;
do {
 builder.page(nextPageToken);
 ListBucketsResponse listResponse = client.listBuckets(builder.build());
 List<Bucket> buckets = listResponse.getItems();
 // handle buckets
```

## CHAPTER 34 Developer Tools

---

```
nextPageToken = listResponse.getOpcNextPage();
} while (nextPageToken != null);
```

### USING AN ITERATOR

Instead of manually working with page tokens, you can use an iterator. Each service client exposes a `getPaginators()` method that returns a `Paginator` object. This object contains methods to return objects of type [Iterable](#). We support two approaches to using **iterable**:

- **Response Iterator:** You can iterate over the `Response` objects that are returned by the list operation. These are referred to as *ResponseIterators*, and their methods are suffixed with "ResponseIterator," for example, *listUsersResponseIterator*.

```
Iterable<ListUsersResponse> responseIterator = identityClient.getPaginators
().listUsersResponseIterator(request);
for (ListUsersResponse response : responseIterator) {
 for (User user : response.getItems()) {
 System.out.println(user);
 }
}
```

- **Record Iterator:** You can iterate over the resources/records that are listed. These are referred to as *RecordIterator*, and their methods are suffixed with "RecordIterator," for example, *listUsersRecordIterator*.

```
Iterable<User> recordIterator = identityClient.getPaginators().listUsersRecordIterator(request)
for (User user : recordIterator) {
 System.out.println(user);
}
```

### EXCEPTION HANDLING

When handling an exception, you can get more information about the HTTP request that caused it, such as the status code or timeout. You can also get the request ID when handling a `BmcException` by using the `getOpcRequestId` method.

This example shows a try-catch block that handles a `BmcException` and prints the request ID.

```
ObjectStorage client = ...;
try {
 GetBucketResponse response = client.getBucket(
 GetBucketRequest.builder().namespaceName("myNamespace").bucketName("myBucket").build());
}
```

## CHAPTER 34 Developer Tools

---

```
String requestId = response.getOpcRequestId();
System.out.println(requestId);
} catch (BmcException e) {
 String requestId = e.getOpcRequestId();
 System.out.println(requestId);
 e.printStackTrace();
}
```

Exceptions in the SDK for Java are runtime exceptions (unchecked), so they do not show up in method signatures. All APIs can throw a `BmcException`.

### Examples

Examples of SDK usage can be found on [GitHub](#), including:

- [Example: Synchronous Object Storage](#)
- [Example: Asynchronous Object Storage](#)
- [Example: Create an instance](#)
- [Example: Get an instance's public IP address](#)

The examples are also in the downloadable .zip file for the SDK. Examples for older versions of the SDK are in the downloadable .zip for the specific version, available [on GitHub](#).

If you'd like to see another example not already covered, file a [GitHub issue](#).

### RUNNING EXAMPLES

1. Download the SDK to a directory named `oci`. See [GitHub](#) for the download.
2. Unzip the SDK into the `oci` directory. For example: `tar -xf oci-java-sdk-dist.zip`
3. Create your configuration file in your home directory (`~/.oci/config`). See [Configuring the SDK](#).
4. Use `javac` to compile one of the previous example classes from the `examples` directory, ex:

```
javac -cp lib/oci-java-sdk-full-<version>.jar:third-party/lib/*
examples/ObjectStorageSyncExample.java
```

## CHAPTER 34 Developer Tools

---

5. You should now have a class file in the `examples` directory. Run the example:

```
java -cp examples:lib/oci-java-sdk-full-<version>.jar:third-party/lib/* ObjectStorageSyncExample
```

### THIRD-PARTY DEPENDENCIES AND SHADING

The SDK requires a number of third-party dependencies, which are available in the `third-party/lib` directory. To use the SDK library `lib/oci-java-sdk-full-<version>.jar`, all of the third-party dependencies in `third-party/lib` have to be on the class path.

The SDK also includes a second version of the SDK library, `shaded/lib/oci-java-sdk-full-shaded-<version>.jar`, which contains most of the third-party dependencies already. Only a few more third-party libraries in `shaded/third-party/lib` have to be on the class path when you use this version of the SDK library.

These two versions of the SDK library are functionally the same, however the second version, `shaded/lib/oci-java-sdk-full-shaded-<version>.jar` can simplify dealing with different versions of third-party dependencies. This is because all the dependencies that are included in `shaded/lib/oci-java-sdk-full-shaded-<version>.jar` were shaded, which means they will not interfere with other versions of themselves you may want to include along with this SDK.

You can use either `lib/oci-java-sdk-full-<version>.jar` or `shaded/lib/oci-java-sdk-full-shaded-<version>.jar`, but not both. When using `lib/oci-java-sdk-full-<version>.jar`, use all the third-party libraries in `third-party/lib`. When using `shaded/lib/oci-java-sdk-full-shaded-<version>.jar`, use all the third-party libraries in `shaded/third-party/lib`.

To use the shaded version of the SDK, replace the `javac` commands in steps 4 and 5 with the following:

- **Step 4:**

```
javac -cp shaded/lib/oci-java-sdk-full-shaded-<version>.jar:shaded/third-party/lib/*
examples/ObjectStorageSyncExample.java
```

- **Step 5:**

```
java -cp examples:shaded/lib/oci-java-sdk-full-shaded-<version>.jar:shaded/third-party/lib/*
ObjectStorageSyncExample
```

### Troubleshooting

This section contains troubleshooting information for the Oracle Cloud Infrastructure SDK for Java.

#### **OBJECTSTORAGE CLIENT DOES NOT CLOSE CONNECTIONS WHEN CLIENT IS CLOSED.**

Too many file descriptors are opened up, and it takes too long to close existing ones. An exception may look like this:

```
Caused by: java.io.FileNotFoundException: classes/casptest.pem (Too many open files)
at java.io.FileInputStream.open0(Native Method)
at java.io.FileInputStream.open(FileInputStream.java:195)
at java.io.FileInputStream.<init>(FileInputStream.java:138)
```

Use one of the following workarounds to fix this issue.

- Make this call before creating a client: `System.setProperty("http.keepAlive", "false");`
- Use this command line argument when running Java: `-Dhttp.keepAlive=false`

#### **SERIALIZATION ERRORS WHEN MAKING REQUESTS OR HANDLING RESPONSES**

If you encounter an `UnrecognizedPropertyException` error when handling a response to a call against the SDK for Java, this indicates that the version of the Jackson library in use does not support a feature that was injected at runtime from another dependency in your application's class path. This happens even if the `FAIL_ON_UNKNOWN_PROPERTIES` deserialization property is set to `false` for the configured `ObjectMapper`.

#### **Solution:**

Determine which version of Jackson libraries are referenced in your application's class path and, if necessary, upgrade to version 2.9.5. For a complete list of Jackson libraries that the SDK for Java depends on, please refer to the [pom.xml file](#) that is hosted on [GitHub](#).



### Note

If you customize a client when instantiated in your application, ensure that you reference the preconfigured `ObjectMapper` from the `RestClientFactory` using the `RestClientFactory#getObjectMapper()` method.

An alternative solution is to use the shaded version of the SDK for Java jar file, which includes a bundled version of the Jackson libraries.

### ENCRYPTION KEY SIZE ERRORS

By default, the SDK for Java can only handle keys of 128 bit or lower key length. Users get "Invalid Key Exception" and "Illegal key size" errors when they use longer keys, such as AES256.

Use one of the following workarounds to fix this issue.

- Use a 128 bit key, such as AES128.
- Install the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction from the following location:  
<http://www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html>

### TROUBLESHOOTING SERVICE ERRORS

Any operation resulting in a service error will cause an exception of type `com.oracle.bmc.model.BmcException` to be thrown by the SDK. For information about common service errors returned by OCI, see [API Errors](#).

## Python SDK

The Python SDK enables you to write code to manage Oracle Cloud Infrastructure resources.

## CHAPTER 34 Developer Tools

---

This SDK and sample is dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

**Download:** The Python SDK is available on [GitHub](#) or the [Python Package Index \(PyPi\)](#).

**Documentation:** Available on [readthedocs.io](#).

### Services Supported

- Analytics Cloud
- Announcements
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Compute Work Requests
- Content and Experience
- Core Services (Networking, Compute, Block Volume)
- Data Transfer
- Database
- Digital Assistant
- DNS
- Email Delivery
- Events
- File Storage
- Functions
- Health Checks
- IAM
- Integration Cloud

## CHAPTER 34 Developer Tools

---

- Key Management
- Limits
- Load Balancing
- Monitoring
- Notifications
- Object Storage
- Quotas
- Resource Manager
- Search
- Streaming
- Web Application Acceleration and Security

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the Python SDK is released, subscribe to the [Atom feed](#).

#### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) and [oci-python-sdk](#) tags in your post
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Ruby SDK

The Ruby SDK enables you to write code to manage Oracle Cloud Infrastructure resources.

This SDK and sample are dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

**Download:** The Ruby SDK is available on [GitHub](#) or [RubyGems](#).

**Documentation:** [Ruby SDK documentation](#).

#### Services Supported

- Analytics Cloud
- Announcements
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Content and Experience
- Compute Work Requests
- Core Services (Networking, Compute, Block Volume)
- Data Transfer
- Database
- Digital Assistant
- DNS
- Email Delivery
- Events
- File Storage
- Functions
- Health Checks

## CHAPTER 34 Developer Tools

---

- IAM
- Integration
- Key Management
- Limits
- Load Balancing
- Monitoring
- Notifications
- Object Storage
- Quotas
- Resource Manager
- Search
- Streaming
- Web Application Acceleration and Security

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the Ruby SDK is released, subscribe to the [Atom feed](#).

#### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) and [oci-ruby-sdk](#) tags in your post

- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Go SDK

The Go SDK enables you to write code to manage Oracle Cloud Infrastructure resources.

This SDK and sample is dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

**Download:** Download the SDK from [GitHub](#).

**Documentation:** Available on [godoc](#).

### Services Supported

- Analytics Cloud
- Announcements
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Compute Work Requests
- Content and Experience
- Core Services (Networking, Compute, Block Volume)
- Data Transfer
- Database
- Digital Assistant
- DNS
- Email Delivery

- Events
- File Storage
- Functions
- Health Checks
- IAM
- Integration Cloud
- Key Management
- Limits
- Load Balancing
- Monitoring
- Notifications
- Object Storage
- Quotas
- Resource Manager
- Search
- Streaming
- Web Application Acceleration and Security

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the Go SDK is released, subscribe to the [Atom feed](#).

### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) and [oci-go-sdk](#) tags in your post
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

## Command Line Interface (CLI)

The CLI is a small footprint tool that you can use on its own or with the Console to complete Oracle Cloud Infrastructure tasks. The CLI provides the same core functionality as the Console, plus additional commands. Some of these, such as the ability to run scripts, extend the Console's functionality.

This CLI and sample is dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

The CLI is built on Python (version 2.7.5 or 3.5 or later), running on Mac, Windows, or Linux. The Python code makes calls to Oracle Cloud Infrastructure APIs to provide the functionality implemented for the various services. These are REST APIs that use HTTPS requests and responses. For more information, see [About the API](#).

**Installation:** See [Quickstart](#).

**Reference:** For help with a specific command, you can enter `help <command>` on the command line or view the [Command Line Reference](#). This reference is derived from the APIs and help text in the Python source code.

### Requirements

To install and use the CLI, you must have:

- An Oracle Cloud Infrastructure account
- A user created in that account, in a group with a policy that grants the desired permissions. This account user can be you, another person, or a system that calls the

API. For an example of how to set up a new user, group, compartment, and policy, see [Adding Users](#). For a list of other typical Oracle Cloud Infrastructure policies, see [Common Policies](#).

- A keypair used for signing API requests, with the public key uploaded to Oracle. Only the user calling the API should possess the private key. See [Configuring the CLI](#).



### Note

To use the CLI without a keypair, you can use token-based authentication. For more information, see [Token-based Authentication for the CLI](#).

- Python version 2.7.5 or 3.5 or later, running on Mac, Windows, or Linux. Note that if you use the CLI Installer and do not have Python on your machine, the Installer offers to automatically install Python for you. If you already have Python installed on your machine, you can use the `python --version` command to find out which version is installed.
- If you require FIPS-compliance, see [Using FIPS-validated Libraries](#).

### Services Supported

- Analytics Cloud
- Announcements
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Compute Work Requests
- Content and Experience
- Core Services (Networking, Compute, Block Volume)

## CHAPTER 34 Developer Tools

---

- Data Transfer
- Database
- Digital Assistant
- DNS
- Email Delivery
- Events
- Functions
- File Storage
- Health Checks
- IAM
- Integration
- Key Management
- Limits
- Load Balancing
- Monitoring
- Notifications
- Object Storage
- Quotas
- Resource Manager
- Search
- Streaming
- Web Application Acceleration and Security

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the CLI is released, subscribe to the [Atom feed](#).

#### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Quickstart

Using the installer script and the setup command is the fastest way to get up and running with the CLI.



#### Warning

Oracle recommends that you avoid using string values that include confidential information.

#### INSTALLING THE CLI

The installer script automatically installs the CLI and its dependencies, Python and virtualenv. Before running the installer, be sure you meet the [Requirements](#).

*MACOS, LINUX, AND UNIX*

1. Open a terminal.
2. To run the installer script, run the following command.

## CHAPTER 34 Developer Tools

---

```
bash -c "$(curl -L https://raw.githubusercontent.com/oracle/oci-cli/master/scripts/install/install.sh)"
```

3. Respond to the [Installation Script Prompts](#).

### WINDOWS

1. Open the PowerShell console using the **Run as Administrator** option.
2. The installer enables auto-complete by installing and running a script. To allow this script to run, you must enable the RemoteSigned execution policy. To configure the remote execution policy for PowerShell, run the following command.

```
Set-ExecutionPolicy RemoteSigned
```

3. To run the installer script, run the following command.

```
powershell -NoProfile -ExecutionPolicy Bypass -Command "iex ((New-Object System.Net.WebClient).DownloadString('https://raw.githubusercontent.com/oracle/oci-cli/master/scripts/install/install.ps1'))"
```

4. Respond to the [Installation Script Prompts](#).

### INSTALLATION SCRIPT PROMPTS

The installation script prompts you for the following information.

- If you do not have a compatible version of Python installed:
  - Windows and Linux: You are prompted to provide a location for installing the binaries and executables. The script will install Python for you.
  - MacOS: You are notified that your version of Python is incompatible. You must upgrade before you can proceed with the installation. The script will not install Python for you.
- When prompted to upgrade the CLI to the newest version, respond with **Y** to overwrite an existing installation.
- When prompted to update your PATH, respond with **Y** to be able to invoke the CLI without providing the full path to the executable. This will add oci.exe to your PATH.

### SETTING UP THE CONFIG FILE

Before using the CLI, you must create a config file that contains the required credentials for working with Oracle Cloud Infrastructure. You can create this file using a setup dialog or manually using a text editor.

#### *USE THE SETUP DIALOG*

To have the CLI walk you through the first-time setup process, use the `oci setup config` command. The command prompts you for the information required for the config file and the API public/private keys. The setup dialog generates an API key pair and creates the config file.

For more information about how to find the required information, see:

- [Where to Get the Tenancy's OCID and User's OCID](#)
- [Regions and Availability Domains](#)

#### *MANUAL SETUP*

If you want to set up the API public/private keys yourself and write your own config file, see [SDK and Tool Configuration](#).



#### **Tip**

Use the `oci setup keys` command to generate a key pair to include in the config file.

### NEXT STEPS

- For details on starting a session, see [Starting a CLI Session](#)
- [Getting Started with the Command Line Interface](#) provides an end-to-end walk-through of using the CLI to launch an instance.

### Manual Installation

Instead of using the installer script as described in the [Quickstart](#), you can manually install the CLI and its dependencies. Before proceeding, be sure you meet the [Requirements](#).



#### Warning

Oracle recommends that you avoid using string values that include confidential information.

### STEP 1: INSTALLING PYTHON

Python installation instructions vary for each operating system.

## Windows and Windows Server 2008 R2

### WINDOWS

Install Python from the [Python Windows downloads](#) page. During installation, choose to add Python to the PATH and/or environment variables (depending on the prompt).

#### *WINDOWS SERVER 2008 R2*

To install Python on this version of Windows Server you must install Python 2.7, or install Windows 2008 R2 Service Pack 1 (SP1) to the instance before installing later versions of Python.

## Oracle Linux 7.3 and Oracle Linux 6

### ORACLE LINUX

Some versions of Oracle Linux come with incompatible versions of Python, and may require additional components to install the CLI.

## CHAPTER 34 Developer Tools

---

### *ORACLE LINUX 7.3*

Before you install the CLI, run the following command on a new Oracle Linux 7.3 image.

```
sudo yum install gcc libffi-devel python-devel openssl-devel
sudo easy_install pip
```

### *ORACLE LINUX 6*

On Oracle Linux 6, a newer version of Python is usually required. You can install a newer version alongside the existing version by downloading from [Python](#), and then install the CLI in a virtual environment that uses the new version. To install the new version of Python, run the following commands.

```
sudo yum install gcc libffi-devel python-devel openssl-devel
sudo easy_install pip
curl -O https://www.python.org/ftp/python/3.6.0/Python-3.6.0.tgz
tar -xvzf Python-3.6.0.tgz
cd Python-3.6.0
./configure
make
sudo make install
```

## CentOS 6 and CentOS 7

### **CENTOS 6 AND CENTOS 7**

Before you install the CLI, run the following commands on a new CentOS image.

```
sudo yum install gcc libffi-devel python-devel openssl-devel
sudo easy_install pip
```

## Ubuntu 16.04

### **UBUNTU 16.04**

Python should come installed. However, if you need to install Python, use the OS's normal package manager.

## CHAPTER 34 Developer Tools

---

To install the additional required components, run the following commands.

```
sudo apt-get update
sudo apt-get install build-essential libssl-dev libffi-dev python-dev
sudo apt-get install python3-pippip3 install --upgrade pip
```

### STEP 2: INSTALLING AND CONFIGURING VIRTUALENV

`virtualenv` is a virtual environment builder that lets you create isolated Python environments. For Linux users, `virtualenv` is usually in a separate package from the main Python package.

Download the software and documentation from:

- Software: [GitHub](#) or [PyPI](#).
- Documentation: [Docs Virtualenv](#)

#### *INSTALLING AND CONFIGURING VIRTUALENV*

After Python is installed, install and configure `virtualenv`.

1. To set up the environment for Python 2, use the following command.

```
pip install virtualenv
```

2. To set up the environment for Python 3, use the following command.

```
pip3 install virtualenv
```

`virtualenv` typically installs to your Python directory. For example:

```
/usr/local/share/python3/virtualenv
```

3. (Optional) To create a directory for storing your virtual environments, run the following command.

```
mkdir -p myvirtualspaces/virtualenvs
```

4. To create a new virtual environment without any packages, run the following command.

```
virtualenv myvirtualspaces/virtualenvs/cli-testing --no-site-packages
```

## CHAPTER 34 Developer Tools

---

If you're installing a newer version of Python to run alongside an existing version, you can create a virtual environment that uses the new version.

To reference the new version of Python, run the following command with the `-p` parameter.

```
virtualenv -p /usr/local/bin/python3.6 cli-testing
```

### STEP 3: INSTALLING THE COMMAND LINE INTERFACE

You can download the CLI from [GitHub](#) or install the package from [Python Package Index \(PyPI\)](#).

To install using the GitHub download:

- Download and unzip **oci-cli.zip**.
- Run the following command.

```
pip install oci_cli-*-py2.py3-none-any.whl
```

To install using PyPI, run the following command.

```
pip install oci-cli
```

For information on how to start a CLI session, see [Starting a CLI Session](#).

### INSTALLING WITHOUT A VIRTUAL ENVIRONMENT

We do not recommend installing the CLI in your system-wide Python and suggest that instead you install the CLI using the installer or virtual environment.

In cases where you are trying to install the CLI in your system-wide Python using the latest pip version, you might encounter conflicts with some `distutils` installed packages. Following is an example error message when this occurs:

```
sudo pip install oci-cli
...
...
Cannot uninstall 'requests'. It is a distutils installed project and thus we cannot accurately determine which files belong to it which would lead to only a partial uninstall.
```

Another option is to install the CLI for the user using the following command, although this approach is not supported:

```
pip install --user oci-cli
```

### Token-based Authentication for the CLI

Token-based authentication for the CLI allows customers to authenticate their session interactively, then use the CLI for a single session without an API signing key. This enables customers using an identity provider that is not SCIM-supported to use a federated user account with the CLI and SDKs.

#### REQUIREMENTS

The requirements are the same as those listed for the CLI in [Requirements](#), except that instead of a SSH keypair, you need a web browser for the authentication process.

#### STARTING A TOKEN-BASED CLI SESSION

To use token-based authentication for the CLI on a computer with a web browser:

1. In the CLI, run the following command. This will launch a web browser.

```
oci session authenticate
```

2. In the browser, enter your user credentials. This authentication information is saved to the `.config` file.

#### VALIDATING A TOKEN

To verify that a token is valid, run the following command:

```
oci session validate --config-file <path_to_config_file> --profile <profile_name> --auth security_token
```

You should receive a message showing the expiration date for the session. If you receive an error, check your profile settings.

#### REFRESHING A TOKEN

The default token TTL is set to 1 hour before it expires and can be refreshed within the validity period up to 24 hours.

To refresh the token, run the following command:

```
oci session refresh --profile <profile_name>
```

## CHAPTER 34 Developer Tools

---

### STARTING A TOKEN-BASED CLI SESSION WITHOUT A BROWSER

To use token-based authentication for the CLI on a computer without a web browser, you must export a session from a web-enabled computer, then import it to the computer without a web browser.

#### EXPORTING FROM SOURCE COMPUTER

On the source computer with the browser:

1. In the CLI, run the following command:

```
oci session authenticate
```

2. Enter the user credentials you wish to use on the target computer.
3. To export a zip file, run the following command:

```
oci session export --profile <profile_name> --output-file <output_filename>
```

To verify the export, see [Validating a Token](#).

#### IMPORTING TO TARGET COMPUTER

On the target computer without the browser, run the following command in the CLI,:

```
oci session import --session-archive <path_to_exported_zip>
```

You can test the import by running the following:

```
oci iam region list --config-file <path_to_config_file> --profile <profile_name> --auth security_token
```

It should return a list of regions. Successful execution of this command verifies that the token authentication is working as expected.

### RUNNING SCRIPTS ON A COMPUTER WITHOUT A BROWSER

After importing the authentication to the target computer, you can run the CLI and SDKs by using the following settings.

#### FOR CLI

To run scripts on the CLI, append the following suffix:

```
--config-file <path_to_config_file> --profile <profile_name> --auth security_token
```

### FOR SDKS

To run SDKs on the target computer, you must read in the token file, then use it to initialize the `SecurityTokenSigner`.

After creating a token file as shown in [Starting a Token-based CLI Session](#), use the following process.



#### Note

These code samples demonstrate how to accomplish this using the Python SDK. For other SDKs, follow the same process, but adjust the syntax accordingly.

1. Read the token file from the `security_token_file` parameter of the `.config` file.

```
config = oci.config.from_file(profile_name='TokenDemo')
token_file = config['security_token_file']
token = None
with open(token_file, 'r') as f:
 token = f.read()
```

2. Read the private key specified by the `.config` file.

```
private_key = oci.signer.load_private_key_from_file(config['key_file'])
```

3. Create the initial SDK client which targets the user-specified region.

```
signer = oci.auth.signers.SecurityTokenSigner(token, private_key)
client = oci.identity.IdentityClient({'region': region}, signer=signer)
```

4. Make the identity request.

```
result = client.list_region_subscriptions(config['tenancy'])
```

### Configuration

You can use these optional configurations to extend CLI functionality. The CLI supports using a file for CLI-specific configurations. You can:

## CHAPTER 34 Developer Tools

---

- Specify a default profile.
- Set default values for command options so you don't have to type them into the command line.
- Define aliases for commands. For example, using "ls" as an alias for `list`.
- Define aliases for options. For example, using "--ad" as an alias for `--availability-domain`.
- Define named queries that are passed to the `--query` option instead of typing a JMESPath expression on the command line.

The default location and file name is `~/.oci/oci_cli_rc`. You can also explicitly specify this file with the `--cli-rc-file` option or by with the legacy `--defaults-file` option. For example:

```
Uses the file from ~/.oci/oci_cli_rc
oci os bucket list

Uses a custom file
oci os bucket list --cli-rc-file path/to/my/cli/rc/file
```

To set up an `oci-cli-rc` file, run the following command.

```
oci setup oci-cli-rc --file path/to/target/file
```

This command creates the file you specify that includes examples of default command aliases, parameter aliases, and named queries.



### Note

If you are using Windows, you should use backslash as the directory separator in pathnames, instead of the forward slash.

### SPECIFYING A DEFAULT PROFILE

Specify a default profile in the `OCI_CLI_SETTINGS` section of the CLI configuration file. The next example shows how to specify a default profile named IAD. The CLI looks for a profile

## CHAPTER 34 Developer Tools

---

named IAD in your `~/.oci/config` file, or any other file that you specify using the `--config-file` option.

```
[OCI_CLI_SETTINGS]
default_profile=IAD
```

You can also specify a default value for the `--profile` option using the `OCI_CLI_PROFILE` environment variable.

If a default profile value has been specified in multiple locations, the order of precedence is:

1. The value specified in the `--profile` option.
2. The value specified in the `OCI_CLI_PROFILE` environment variable.
3. The value specified in the `default_profile` field in the `OCI_CLI_SETTINGS` section of the CLI configuration file.

### SPECIFYING DEFAULT VALUES

The CLI supports using a default values file so that you don't have to keep typing them into the command line. For example, instead of typing in a `--compartment-id` on each launch instance command or having to keep specifying the `--namespace` when using Object Storage commands. You can put this information in a default values file.

Default values can be applied at different levels, from general to specific:

- Globally, across all the CLI commands.
- To a particular service, such as Compute or Object Storage.
- To a specific group, such as commands related to exporting images.
- To a specific command.

Default values are treated hierarchically, with specific values having a higher order of precedence than general values. For example, if there is a globally defined value for `compartment-id` and a specific `compartment-id` defined for the `compute instance launch` command, the CLI uses the value for the `compute instance launch` instead of the global default.

## CHAPTER 34 Developer Tools

---

### DEFAULT VALUES FILE NAME AND LOCATION

When you start the CLI, the program looks for the default values file in `~/.oci/oci_cli_rc`. You can also specify a different file and location by using the `--cli-rc-file` option, as illustrated by the following:

```
Uses the default values file from ~/.oci/oci_cli_rc
oci os bucket list

Uses a custom default values file
oci os bucket list --cli-rc-file path/to/my/cli/custom-oci-cli-rc-file
```

### COMMAND VALUE PRIORITY

If a value is provided on the command line also exists in `--cli-rc-file`, the value from the command line has priority. For a command with options that take multiple values, the values are taken entirely from the command line or from `--cli-rc-file`. The 2 sources aren't merged.

### DEFAULTS VALUE FILE SYNTAX

The `--cli-rc-file` file can be divided into different sections with one or more keys per section.

## Sections

In the next example, the file has two sections, with a key in each section. To specify which section to use, you use the `--profile` option in the CLI.

```
[DEFAULT]
compartment-id = ocidl.compartment.oc1..aaaaaaaa15zx25nzpgeyqd3gzijdlg3ieqeyrggnx7il26astxxhqoljnhwa
[ANOTHER_SECTION]
compartment-id = ocidl.compartment.oc1..aaaaaaaa13gzieyqdrngnx7xil26astxxhqol2pgjjdlieqeyg35nz5znhwa
```

## Keys

Keys are named after command line options, but do not use a leading double hyphen (`--`). For example, the key for `--image-id` is `image-id`. You can specify keys for single values, multiple values, and flags.

- **Keys for Single Values.** The next example shows how to specify key values at different levels, and with different scope.

```
[DEFAULT]
Defines a global default for bucket-name
bucket-name = my-global-default-bucket-name

Defines a default for bucket-name, which applies to all 'compute' commands
compute.bucket-name = bucket-name-for-image-import-export

Defines a default for bucket-name, which applies to all 'os object' commands (e.g., os object
get)
os.object.bucket-name = bucket-name-for-object-commands

Defines a default for bucket-name, for the 'os object multipart list' command
os.object.multipart.list.bucket-name = bucket-name-for-multipart-list
```

- **Keys for Multiple Values.** Some options, such as `--include` and `--exclude` on the `oci os object bulk-upload` command can be specified more than once. For example:

```
oci os object bulk-upload -ns my-namespace -bn my-bucket --src-dir my-directory --include *.txt -
-include *.png
```

The next example shows how you would enter the `--include` values in the `--cli-rc-file` file

```
[DEFAULT]
os.object.bulk-upload.include =
 *.txt
 *.png
```

In the previous example, one value is given for each line and each line must be indented underneath its key. You can use tabs or spaces and the amount of indentation doesn't matter. You can also put a value on the same line as the key, add more values on the following lines, and use a path statement for a value. For example:

```
[DEFAULT]
os.object.bulk-upload.include = *.pdf
 *.txt
 *.png
 my-subfolder/*.tiff
```

## CHAPTER 34 Developer Tools

---

- **Keys for Flags.** Some command options are flags, like `--force`, which uses a Boolean value. To set a flag for the `--force` option, use the following command.

```
os.object.delete.force=true
```

### SPECIFYING COMMAND ALIASES

Specify named queries in the `OCI_CLI_COMMAND_ALIASES` section of the CLI configuration file. There are two types of aliases, global aliases and command sequence aliases. The following example shows each type of alias.

```
[OCI_CLI_COMMAND_ALIASES]
This is a global alias that lets you use "ls" instead of "list" for any list command in the CLI.
#
ls = list

Command examples:
oci os object ls or oci os compute ls

This is a command sequence alias that lets you use "oci os object rm" instead of "oci os
object delete".
<alias> = <dot-separated sequence of groups and sub-groups>.<command or group to alias>
#
rm = os.object.delete

Command example:
<alias> = rm, <sequence of groups and sub-groups> = os object, <command or group to alias> = delete
```

If you want to define default values for options in your CLI configuration file, you can use the alias names you have defined. For example, if you have `-ls` as an alias for `--list`, you can define a default for an availability domain when listing instances by using the following command.

```
[DEFAULT]
compute.instance.ls.compartment-
id=ocid1.compartment.oc1..aaaaaaaa15zx25nznpgyqd3gzijd1g3ieqeyrggnx7il26astxxhq1jnhwa
```

### SPECIFYING OPTION ALIASES

Specify option aliases in the `OCI_CLI_PARAM_ALIASES` section of the CLI configuration file. Option aliases are applied globally. The following example shows some aliases for command

options.

```
[OCI_CLI_PARAM_ALIASES]
Option aliases either start with a double hyphen (--) or are a single hyphen (-) followed by a
single letter. For example: --example-alias, -e
#
--ad = --availability-domain
--dn = --display-name
--egress-rules = --egress-security-rules
--ingress-rules = --ingress-security-rules
```

If you want to define default values for options in your CLI configuration file, you can use the alias names you have defined. For example, if you have `-ad` as an alias for `--availability-domain`, you can define a default for an availability domain when listing instances by using the following command.

```
[DEFAULT]
compute.instance.list.ad=xyx:PHX-AD-1
```

### SPECIFYING NAMED QUERIES

If you use the `--query` parameter to filter or manipulate output, you can define named queries instead of using a JMESPath expression on the command line.

Specify named queries in the `OCI_CLI_CANNED_QUERIES` section of the CLI configuration file.

### Examples of Named Queries

```
[OCI_CLI_CANNED_QUERIES]
For list results, this gets the ID and display-name of each item in the list.
Note that when the names of attributes have dashes in them they need to be surrounded
with double quotes. This query knows to look for a list because of the [*] syntax

get_id_and_display_name_from_list=data[*].{id: id, "display-name": "display-name"}

get_id_and_display_name_from_single_result=data.{id: id, "display-name": "display-name"}

Retrieves a comma separated string, for example:
ocid1.instance.oc1.phx.xyz....,cli_test_instance_675195,RUNNING
#
get_id_display_name_and_lifecycle_state_from_single_result_as_csv=data.[id, "display-name", "lifecycle-
```

## CHAPTER 34 Developer Tools

---

```
state"] | join(`,`, @)

Retrieves comma separated strings from a list of results
#
get_id_display_name_and_lifecycle_state_from_list_as_csv=data[*].[join(`,`, [id, "display-name",
"lifecycle-state"])][]

Filters where the display name contains some text
#
filter_by_display_name_contains_text=data[?contains("display-name", `your_text_here`)]

Filters where the display name contains some text and pull out certain attributes(id and time-
created)
#
filter_by_display_name_contains_text_and_get_attributes=data[?contains("display-name", `your_text_
here`)].{id: id, timeCreated: "time-created"}

Get the top 5 results from a list operation
#
get_top_5_results=data[:5]

Get the last 2 results from a list operation
#
get_last_2_results=data[-2:]
```

You can reference any of these queries using this syntax: `query://<query name>`.

For example, to get id and display name from a list, run the following command.

```
oci compute instance list -c $C --query query://get_id_and_display_name_from_list
```

### ENABLING AUTO-COMPLETE

If you used the CLI installer, you don't have to configure auto-complete because it's enabled automatically.

To enable auto-complete (tab completion) for a manual CLI installation, run the following command.

```
oci setup autocomplete
```

To enable auto-complete on a session by session basis, run the following command.

```
eval "$(_OCI_COMPLETE=source oci)"
```



### Note

#### *Support for Auto-complete on Windows*

Auto-complete on Windows is only supported if you're using PowerShell. A script runs to enable this feature. However, you must change the PowerShell execution policy to RemoteSigned. To configure this policy, run the following command at the PowerShell command line.

```
Set-ExecutionPolicy RemoteSigned
```

### USING FIPS-VALIDATED LIBRARIES

The CLI can be configured to use FIPS-validated libraries on Linux. The CLI is built on the Python SDK and leverages operating system level cryptographic libraries.

#### CONFIGURING THE ENVIRONMENT

1. Verify the installed version of OpenSSL is FIPS-compliant. Run the following command:

```
openssl version
```

If "fips" is not part of the version name, you should upgrade OpenSSL to a FIPS-compliant version. You can download the latest versions of OpenSSL at:

<https://www.openssl.org/source/>

2. Determine the location of the FIPS-compliant version of libcrypto:

```
ls -l /usr/lib64/libcrypto*
```

3. Set the environment variable OCI\_CLI\_FIPS\_LIBCRYPTO\_FILE to the location of libcrypto:

```
export OCI_CLI_FIPS_LIBCRYPTO_FILE=</path/to/libcrypto.x.x.x>
```

## CHAPTER 34 Developer Tools

---

If you do not want to run this command at the start of every session, you can add it to your `.bashrc` or `.bash_profile` file.

You can confirm that the environment variable is set properly with this command:

```
set | grep OCI_CLI_FIPS_LIBCRYPTO_FILE
```

You can now proceed to the standard installation process outlined in [Quickstart](#)

### VERIFYING THE CONFIGURATION

To verify that the CLI is using the library that you specified during [Configuring the Environment](#), execute the following commands in Python. Be sure to do so in the same environment that the CLI uses.

```
import ssl
ssl.FIPS_mode()
```

This should return 1, indicating that SSL is using the library specified by the `OCI_CLI_FIPS_LIBCRYPTO_FILE` environment variable.

### Using the CLI

This topic describes how to use the CLI to access Oracle Cloud Infrastructure and carry out service-related tasks. This topic assumes that you have [configured the CLI](#) and are ready to start using it.



#### Tip

[Getting Started with the Command Line Interface](#) provides an end-to-end walk-through of using the CLI to launch an instance.

### STARTING A CLI SESSION

#### MACOS, LINUX, AND UNIX

To start a CLI session, run the following commands.

## CHAPTER 34 Developer Tools

---

1. Open a terminal.
2. Change the working directory.

```
cd myvirtualspaces/virtualenvs/cli-testing/bin
```

3. Run the activate batch file.

```
source activate
```

To stop using the CLI, run the following command in a terminal.

```
deactivate
```

### WINDOWS

To start a CLI session, run the following commands.

1. Open the Command Prompt using the **Run as administrator** option.
2. Change the working directory.

```
cd myvirtualspaces/virtualenvs/cli-testing/Scripts
```

3. Run the activate batch file.

```
activate
```

To stop using the CLI, run the following command from the command line.

```
deactivate
```



### Warning

Avoid entering confidential information when providing resource names, descriptions, or other values that may expose sensitive information.

### COMMAND LINE SYNTAX

Most commands must specify a service, followed by a resource type and then an action. The basic command line syntax is:

```
oci <service> <type> <action> <options>
```

## CHAPTER 34 Developer Tools

---

For example, this syntax is applied as follows:

- `compute` is the `<service>`
- `instance` is the resource `<type>`
- `launch` is the `<action>`, and
- the rest of the command string consists of `<options>`.

The following command to launch an instance shows a typical command line construct.

```
oci compute instance launch --availability-domain "EMIr:PHX-AD-1" -c
ocid1.compartment.oc1..aaaaaaaa13gzijdlieqeyg35nz5zxil26astxxhqol2pgeyqdrqgnx7jnhwa --shape
"VM.Standard1.1" --display-name "Instance 1 for sandbox" --image-id
ocid1.image.oc1.phx.aaaaaaaqutj4qjxihpl4mboabsa27mrpusygv6gurp47kat5z7vljmq3puq --subnet-id
ocid1.subnet.oc1.phx.aaaaaaaaypsr25bzjmjyn6xwgkcrqxd3dbhiha6lodzus3gafscirbhj5bpa
```



### Warning

In the previous example, you can provide a friendly name for the instance using the `--display-name` option. Avoid entering confidential information when providing resource names or descriptions.

### BASIC EXAMPLES

This section provides examples of basic operations using the CLI.



### Note

#### *Using Environment Variables for OCIDs*

Several of the CLI examples use environment variables for OCIDs, such as:

- \$T for a tenancy OCID
- \$C for a compartment OCID

For example:

```
T=ocid1.tenancy.oc1..aaaaaaaaba3pv6wm2ytdrwx32uzr4h25vkr4jqae5f15p2b2qstifsfdsq

C=ocid1.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5fldrwpp2jojdcypxfga
```

To get a namespace, run the following command.

```
oci os ns get
```

To list compartments, run the following command.

```
oci iam compartment list -c $T
```

To get a list of buckets, run the following command.

```
oci os bucket list -ns mynamespace --compartment-id $C
```

To list users and limit the output, run the following command.

```
oci iam user list --compartment-id $T --limit 5
```

To add a user to a group, run the following command.

```
oci iam group add-user --user-id
ocid1.user.oc1..aaabcaaaakkhhtmgvhvqqq7rgvzwuj3drwmtlsgz6sbfo7y4uc5sprzli377q --group-id
ocid1.group.oc1..aaabcaaa66plootq6uwwxhfdw21sdqtgeb6l4pjsv5eeuexrauuj35b7b
```

### GETTING HELP WITH COMMANDS

You can get help for any command using `--help`, `-h`, or `-?`. For example:

## CHAPTER 34 Developer Tools

---

```
oci --help
```

```
oci os bucket -h
```

```
oci os bucket create -?
```

### *VIEWING ALL THE CLI HELP*

You can view the [command line help](#).

### **DETERMINING THE INSTALLED VERSION OF THE CLI**

To get the installed version of the CLI, run the following command.

```
oci --version
```

### **USING DATES AND TIMES IN CLI COMMANDS**

The CLI supports the following accepted date formats.

- UTC with milliseconds

```
Format: YYYY-MM-DDTHH:mm:ss.sssTZD, Example: 2017-09-15T20:30:00.123Z
```

- UTC without milliseconds

```
Format: YYYY-MM-DDTHH:mm:ssTZD, Example: 2017-09-15T20:30:00Z
```

- UTC with minute precision

```
Format: YYYY-MM-DDTHH:mmTZD, Example: 2017-09-15T20:30Z
```

- Timezone with milliseconds

```
Format: YYYY-MM-DDTHH:mm:ss.sssTZD, Example: 2017-09-15T12:30:00.456-08:00
```

- Timezone without milliseconds

```
Format: YYYY-MM-DDTHH:mm:ssTZD, Example: 2017-09-15T12:30:00-08:00
```

- Timezone with offset with minute precision

```
Format: YYYY-MM-DDTHH:mmTZD, Example: 2017-09-15T12:35-08:00
```

- Date Only (This date will be taken as midnight UTC of that day)

```
Format: YYYY-MM-DD, Example: 2017-09-15
```

- Epoch seconds

Example: 1412195400



### Note

In our datetime formats, the `T` can be replaced with a space. For example, both `"2017-09-15 20:30:00.123Z"` and `2017-09-15T20:30:00.123Z` are acceptable. (Note that if you do not include the `T`, you must wrap the value in quotes.) We also support time zones with and without the colon. Both `+10:00` and `+1000` are acceptable.

## MANAGING CLI INPUT AND OUTPUT

The CLI provides several options for managing command input and output.

### PASSING COMPLEX INPUT

Complex input, such as arrays and objects with more than one value, are passed in JSON format and can be provided as a string at the command line, as a file, or as a command line string and as a file.

### MacOS, Linux, or Unix

The following command shows how to pass two values for the `--metadata` object.

```
oci os bucket create -ns mynamespace --name mybucket --metadata '{"key1":"value1","key2":"value2"}' --compartment-id ocid1.compartment.oc1..aaaaaaaarhifmvrvuqtys5q66rck6copzqck3ukc5fldrwpp2jojdcypxfga
```

### Windows

On Windows, to pass complex input to the CLI as a JSON string, you must enclose the entire block in double quotes. Inside the block, each double quote for the key and value strings must be escaped with a backslash (`\`) character.

The following command shows how to pass two values for the `--metadata` object on Windows.

## CHAPTER 34 Developer Tools

```
oci os bucket create -ns mynamespace --name mybucket --metadata "{\
\"key1\": \"value1\", \"key2\": \"value2\"}" --compartment-id
ocidl.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5f1drwpp2jobjdcypxfga
```



### Note

#### *JSON Errors*

The error message "Parameter '<PARAMETER NAME>' must be in JSON format." indicates that the value you passed for the parameter with name "PARAMETER NAME" was not valid JSON. This error is typically a result of the JSON string not being escaped correctly.

For more information about using JSON strings, see [Advanced JSON Options](#)

#### *FORMAT OUTPUT AS A TABLE*

By default, all responses to a command are returned in JSON format. For example, a response like the following is returned when you issue the command to get a list of regions.

```
{
 "data": [
 {
 "key": "FRA",
 "name": "eu-frankfurt-1"
 },
 {
 "key": "IAD",
 "name": "us-ashburn-1"
 },
 {
 "key": "ICN",
 "name": "ap-seoul-1"
 },
 {
 "key": "PHX",
 "name": "us-phoenix-1"
 },
 {
```

## CHAPTER 34 Developer Tools

```
 "key": "LHR",
 "name": "uk-london-1"
 },
 {
 "key": "NRT",
 "name": "ap-tokyo-1"
 },
 {
 "key": "YYZ",
 "name": "ca-toronto-1"
 }
]
```

In some cases, readability can become an issue, which is easily resolved by formatting a response as a table. To get a response to a command formatted as a table, run the following command.

```
oci iam region list --output table
```

The following sample list of regions is returned as a two column table.

```
+-----+-----+
| key | name |
+-----+-----+
| FRA | eu-frankfurt-1 |
| IAD | us-ashburn-1 |
| ICN | ap-seoul-1 |
| PHX | us-phoenix-1 |
| NRT | ap-tokyo-1 |
| LHR | uk-london-1 |
| YYZ | ca-toronto-1 |
+-----+-----+
```

### *FILTER OUTPUT*

You can filter output using the JMESPath query option for JSON. Filtering is very useful when dealing with large amounts of output. For example, run the following command with the output table option to get a list of images.

```
oci compute image list -c
ocid1.compartment.oc1..aaaaaaaapxgklgmujxjzx2ypptfjrcieq7rrob2u2zbesh3w1afsgthhgtea --output table
```

## CHAPTER 34 Developer Tools

The image information is returned in table format, but too much data is returned, which overflows the width of the terminal. In addition, you might not need all the information that's returned.

```
| base-image-id | compartment-id | create-image-allowed | display-name
 | id | lifecycle-state | operating-system | operating-system-version | time-created
 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| None | | True | | Windows-Server-2012-R2-Standard-Edition-VM-2017.07.25-0 | ocid
1.image.oc1.phx.aaaaaaab2xgy6bijtudhsgsbgn6zwfqnkdb2bp4l4qap7e4mehv6bv3qca | AVAILABLE | Windows
 | Serv
er 2012 R2 Standard | 2017-07-25T23:59:59.311000+00:00 |
| None | | None | | True | | Windows-Server-2012-R2-Standard-Edition-VM-
2017.04.03-0 | ocid
1.image.oc1.phx.aaaaaaa53cliasgvqmutflwqkafbro2y4ywjebci5szc4eus5byy2e2b7ua | AVAILABLE | Windows
 | Serv
er 2012 R2 Standard | 2017-04-03T19:42:22.938000+00:00 |
| None | | None | | True | | Windows-Server-2012-R2-Standard-Edition-BM-
2017.07.25-0 | ocid
1.image.oc1.phx.aaaaaaaadcegaay43eux6uap55fhp6lqagh37xgocscktwm2yr7q14pcykxq | AVAILABLE | Windows
 | Serv
er 2012 R2 Standard | 2017-07-25T20:55:37.937000+00:00 |
| None | | None | | True | | Windows-Server-2012-R2-Standard-Edition-BM-
2017.04.13-0 | ocid1.image.oc1.phx.aaaaaaa7xgecq2kt7tikqfmrshu6gwukoc3lcnf2iqtwmjyarlprp6j6lna |
AVAILABLE | Windows | Serv
er 2012 R2 Standard | 2017-04-13T17:36:50.840000+00:00 |
| None | | None | | True | | Windows-Server-2008-R2-Standard-Edition-VM-
2017.08.03-0 | ocid
1.image.oc1.phx.aaaaaaaajmyrf52wf2blf7jd7y2dcrjvg6dyulfyp7d3r3oarc5ayka5liq | AVAILABLE | Windows
 | Serv
er 2008 R2 Standard | 2017-07-27T18:19:06.976000+00:00 |
| None | | None | | True | | Oracle-Linux-7.4-2017.09.29-0
 | ocid
1.image.oc1.phx.aaaaaaa3g2xpzlbrrdknqcjtzv2tvxcofj55vdcmpxdlbohmtt7encpana | AVAILABLE | Oracle
Linux | 7.4
| 2017-10-05T22:36:17.246000+00:00 |
| None | | None | | True | | Oracle-Linux-7.4-2017.08.25-1
 | ocid
1.image.oc1.phx.aaaaaaaajan2cd2g65tphaiegiz4lbs422rdc73okcu7dt2uya6p5szywsa | AVAILABLE | Oracle
Linux | 7.4
```

## CHAPTER 34 Developer Tools

```
| 2017-09-11T23:12:18.644000+00:00 |
| None | None | True | Oracle-Linux-7.4-2017.08.25-0
| ocid
1.image.oc1.phx.aaaaaaaabifl2bmaygtu4riw3vcuow15cqwdzdqzwndqneoybcfcn2pgyc6a | AVAILABLE | Oracle
Linux | 7.4 | 2017-08-25T01:21:37.176000+00:00 |
```

You can limit the amount of data returned by combining the `--query` option with `--output table` to get the information you want from a command.

To get filtered image information returned in a table format, run the following command.

```
oci compute image list -c
ocid1.compartment.oc1..aaaaaaaapxgklgmujxjzx2ypptfjrcieq7rrob2u2zbesh3wlafsgthhgtea --output table --
query "data [*].{ImageName:\\"display-name\\", OCID:id}"
```

The previous command returns the following image information, formatted as a two column table.

```
+-----+-----+
| ImageName | OCID |
+-----+-----+
| Windows-Server-2012-R2-Standard-Edition-VM-2017.07.25-0 | ocid1.image.oc1.phx.aaaaaaaab2xgy6bijtudhsbsgns6zwfqnkdb2bp414qap7e4mehv6bv3qca |
| Windows-Server-2012-R2-Standard-Edition-VM-2017.04.03-0 | ocid1.image.oc1.phx.aaaaaaa53cliasgvqmutflwqkafbro2y4ywjebci5szc4eus5byy2e2b7ua |
| Windows-Server-2012-R2-Standard-Edition-BM-2017.07.25-0 | ocid1.image.oc1.phx.aaaaaaaadcegaay43eux6uap55fhp6lqagh37xgocscctwm2yr7q14pcykxq |
| Windows-Server-2012-R2-Standard-Edition-BM-2017.04.13-0 | ocid1.image.oc1.phx.aaaaaaa7xgecq2kt7tikqfrmshu6gwukoc3lcnf2iqtwmjarlprp6j6lna |
| Windows-Server-2008-R2-Standard-Edition-VM-2017.08.03-0 | ocid1.image.oc1.phx.aaaaaaaajmyrf52wf2blf7jd7y2dcrjvg6dyulfyp7d3r3oarc5ayka5liq |
| Oracle-Linux-7.4-2017.09.29-0 | ocid1.image.oc1.phx.aaaaaaa3g2xpz1brrdknqjtzv2tvxcofjc55vdcmpxdlbohmtt7encpana |
| Oracle-Linux-7.4-2017.08.25-1 | ocid1.image.oc1.phx.aaaaaaaajan2cd2g65tphaiegiz4lbs422rdc73okcu7dt2uya6p5szywsa |
| Oracle-Linux-7.4-2017.08.25-0 | ocid1.image.oc1.phx.aaaaaaaabifl2bmaygtu4riw3vcuow15cqwdzdqzwndqneoybcfcn2pgyc6a |
| Oracle-Linux-7.3-2017.07.17-1 | ocid1.image.oc1.phx.aaaaaaa7jvfm572d4ehcgh3ijapvhr52voel33ispumnygi3k17mph55ha |
| Oracle-Linux-7.3-2017.07.17-0 |
```

## CHAPTER 34 Developer Tools

```
ocidl.image.oc1.phx.aaaaaaa5yu6pw3riqtuhxzov7fdngi4tsteganmao54nq3pyxu3hxcuzmoa |
| Oracle-Linux-6.9-2017.09.29-0 |
ocidl.image.oc1.phx.aaaaaaa2d243dmn6mj53zieyap5bdvtq7xfmr5kg5xulrldbzdavaaoj6a |
| Oracle-Linux-6.9-2017.08.25-0 |
ocidl.image.oc1.phx.aaaaaaaavlwrctgz2mx6c4q4gg4gwwvibx6g7xqkoe3tbbwjnifybwmexpnq |
| Oracle-Linux-6.9-2017.07.17-0 |
ocidl.image.oc1.phx.aaaaaaa3s4v5eamndtyghbo4bj2mhobkwjwbz3eowyy5cebmrsoxvoopixa |
| CentOS-7-2017.09.14-0 |
ocidl.image.oc1.phx.aaaaaaaauqtvzqhplzuyesb5tctig6qrwoavpnfiwdkvuyun7z646z72ahcq |
| CentOS-7-2017.07.17-0 |
ocidl.image.oc1.phx.aaaaaaaahmts5c5nktcnqsu6ppom72d7dnvkmqsoavpsiklamn7qd3a7szq |
| CentOS-7-2017.04.18-0 |
ocidl.image.oc1.phx.aaaaaaaamx6ta37ux1tor6n5lxfgd5lkb3lwmqurlpn2x4dz5ockekiuea |
| CentOS-6.9-2017.09.14-0 |
ocidl.image.oc1.phx.aaaaaaaagedr7qsbpxjylietj7dy2r4xoq6p65v3i6y4simkhgyww2ibzxq |
| CentOS-6.9-2017.07.17-0 |
ocidl.image.oc1.phx.aaaaaaaalm3mr4lpsnzjev2nzmkmhpiy7yxu3456qyg7r4nvjieslp4yngtq |
| CentOS-6.8-2017.06.13-0 |
ocidl.image.oc1.phx.aaaaaaaauk5k4km4epm7fxj5ifuylvnyjfkmlukqcg25clayx3ucuzqizjbia |
| Canonical-Ubuntu-16.04-2017.08.22-0 |
ocidl.image.oc1.phx.aaaaaaaalzhdvphf77ggvqo2apmve7o4s4jo77rluaf456qdzrtwmkq2xhra |
| Canonical-Ubuntu-16.04-2017.06.28-0 |
ocidl.image.oc1.phx.aaaaaaaak2idogwetkehtdvo7m673ojuucpfxybd3ehun7izzgjq4c4gga |
| Canonical-Ubuntu-16.04-2017.05.16-0 |
ocidl.image.oc1.phx.aaaaaaaee3a3oedsmmwsqu4dsrzntekefgq7vosngn4r6u6n5mis7dwpxxpa |
+-----+
-----+
```

For more information about the JMESPath query language for JSON, see [JMESPath](#).

### ADVANCED JSON OPTIONS

You can get the correct JSON format for command options and commands.

- For a command option, use `--generate-param-json-input` and specify the command option that you want to get the JSON for. To generate the JSON for creating or updating a security rule, run the following command.

```
oci network security-list create --generate-param-json-input ingress-security-rules
```

### Response from the Command

```
[
 {
 "icmpOptions": {
 "code": 0,
 "type": 0
 },
 "isStateless": true,
 "protocol": "string",
 "source": "string",
 "tcpOptions": {
 "destinationPortRange": {
 "max": 0,
 "min": 0
 },
 "sourcePortRange": {
 "max": 0,
 "min": 0
 }
 },
 "udpOptions": {
 "destinationPortRange": {
 "max": 0,
 "min": 0
 },
 "sourcePortRange": {
 "max": 0,
 "min": 0
 }
 }
 },
 {
 "icmpOptions": {
 "code": 0,
 "type": 0
 },
 "isStateless": true,
 "protocol": "string",
```

```
"source": "string",
"tcpOptions": {
 "destinationPortRange": {
 "max": 0,
 "min": 0
 },
 "sourcePortRange": {
 "max": 0,
 "min": 0
 }
},
"udpOptions": {
 "destinationPortRange": {
 "max": 0,
 "min": 0
 },
 "sourcePortRange": {
 "max": 0,
 "min": 0
 }
}
}
]
```

- For an entire command, use `--generate-full-command-json-input`. To generate the JSON for launching an instance, run the following command.

```
oci compute instance launch --generate-full-command-json-input
```

### Response from the Command

```
{
 "assignPublicIp": true,
 "availabilityDomain": "string",
 "compartmentId": "string",
 "displayName": "string",
 "extendedMetadata": {
 "string1": {
 "string1": "string",
 "string2": "string"
 }
 }
}
```

```
 },
 "string2": {
 "string1": "string",
 "string2": "string"
 }
 },
 "hostnameLabel": "string",
 "imageId": "string",
 "metadata": {
 "string1": "string",
 "string2": "string"
 },
 "privateIp": "string",
 "shape": "string",
 "skipSourceDestCheck": true,
 "subnetId": "string",
 "vnicDisplayName": "string"
}
```

### *ORDER OF PRECEDENCE FOR JSON INPUT*

The CLI supports combining arguments on the command line with file input. However, if the same values are provided in a file and on the command line, the command line takes precedence.

### *USING A JSON FILE FOR COMPLEX INPUT*

You can pass complex input from a file by referencing it from the command line. For Windows users, this removes the requirement of having to escape JSON text. You provide a path to the file using the `file://` prefix.

## **Path Types**

Using `testfile.json` as an example, the following types of paths are supported.

- Relative paths from the same directory, for example: `file://testfile.json` and `file://relative/path/to/testfile.json`
- Absolute paths on Linux, MacOS or Unix, for example:

file:///absolute/path/to/testfile.json

- Full file paths on Windows, for example: file://C:\path\to\testfile.json



### Note

#### *File Path Expansions*

File path expansions, such as "~/", "./", and "../", are supported. On Windows, the "~/" expression expands to your user directory, which is stored in the %USERPROFILE% environment variable. Using environment variables in paths is also supported.

## File Locations

The following file locations are supported.

- Your home directory.

```
oci os bucket create -ns mynamespace --name mybucket --compartment-id
ocidl.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5fldrwpp2jojdscyxfga --metadata
file://~/testfile.json
```

- The current directory.

```
oci os bucket create -ns mynamespace --name mybucket --compartment-id
ocidl.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5fldrwpp2jojdscyxfga --metadata
file://testfile.json
```

- The /tmp directory (Linux, Unix, or MacOS).

```
oci os bucket create -ns mynamespace --name mybucket --compartment-id
ocidl.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5fldrwpp2jojdscyxfga --metadata
file:///tmp/testfile.json
```

- The C:\temp directory (Windows).

```
oci os bucket create -ns mynamespace --name mybucket --compartment-id
ocidl.compartment.oc1..aaaaaaaarhifmvrvuqtye5q66rck6copzqck3ukc5fldrwpp2jojdscyxfga --metadata
file://C:\temp\testfile.json
```

### EXAMPLES OF USING A JSON FILE AS INPUT

The examples in this section use JSON that's generated for a command option and an entire command. The JSON is saved in a file, edited, and then used as command line input.

#### Use File Input for a Command Option

This end-to-end example shows how to generate the JSON for a security list id option used to create a subnet. The JSON is saved in a file, edited, and then used as command line input.

#### Use a JSON File as Input for a Security List Option

1. To generate the JSON for the `security-list-ids` option, run the following command.

```
oci network subnet create --generate-param-json-input security-list-ids
```

2. Create a file and add the following content, which was returned in step 1. This content doesn't have to be escaped or on a single line, it just has to contain valid JSON.

```
[
 "string",
 "string"
]
```

3. Edit the file and replace the "string" values with values, as shown in the following example.

```
[
 "ocid1.securitylist.oc1.phx.aaaaaaaaw7c62ybv4676muq5tdrwup3v2maiquhbk4sf75tjcf5dm6kvlq",
 "ocid1.securitylist.oc1.phx.aaaaaaa7snx4jh5drwo2h33rwcqev6elir55hnrhi2yfdjfon5rcqk4g"
]
```

4. Save the file as "security-list.json".
5. To create the subnet using "security-list.json" as input, run the following command.

```
oci network subnet create --vcn-id
ocid1.vcn.oc1.phx.aaaaaaa6wmuahgxejkv7ukyruqdrwlmrumt16vyisxxxavagiqw2eeet2sa -c
ocid1.compartment.oc1..aaaaaaa13gzijdliehxhql2rggndrwyg35nz5zxl126astpgeyq7jnhwa --
availability-domain "EMIr:PHX-AD-1" --display-name TESTSUB --dns-label "testinstances" --cidr-
block "10.0.0.0/16" --security-list-ids file://security-list.json
```

### Use File Input for an Entire Command

This end-to-end example shows how to generate the JSON to create a virtual cloud network (VCN). The JSON is saved in a file, edited, and then used as command line input.

### Use a JSON File as Input to Create a VCN

1. To generate the JSON needed to create a VCN, run the following command.

```
oci network vcn create --generate-full-command-json-input
```

2. Create a file and add the following content, which was returned in step 1. This content doesn't have to be escaped or on a single line, it just has to contain valid JSON.

```
{
 "cidrBlock": "string",
 "compartmentId": "string",
 "displayName": "string",
 "dnsLabel": "string"
}
```

3. Edit the file and replace the "string" values with values, as shown in the following example.

```
{
 "cidrBlock": "10.0.0.0/16",
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaa13gzijdljedxxhqol2rggndrwyg35nz5zxi126astpgeyq7jnhwa",
 "displayName": "TestVCN",
 "dnsLabel": "testdns"
}
```

4. Save the file and name it "create-vcn.json"
5. To create the VCN using "create-vcn.json" as input, run the following command.

```
oci network vcn create --from-json file://create-vcn.json
```

### ADVANCED EXAMPLES

The following examples show how you can use the CLI to complete complex tasks in Oracle Cloud Infrastructure.

## CHAPTER 34 Developer Tools

---

### WORKING WITH OBJECT STORAGE

You can use the CLI for several object operations with the Object Storage service.

#### Uploading and Downloading Files

Objects can be uploaded from a file or from the command line (STDIN), and can be downloaded to a file or to the command line (STDOUT).

Upload an object:

```
oci os object put -ns mynamespace -bn mybucket --name myfile.txt --file /Users/me/myfile.txt --metadata '{"key1":"value1","key2":"value2"}
```

Upload object contents from the command line (STDIN):

```
oci os object put -ns mynamespace -bn mybucket --name myfile.txt --file <--'object content'
```

Download an object:

```
oci os object get -ns mynamespace -bn mybucket --name myfile.txt --file /Users/me/myfile.txt
```

Print object contents to the command line (STDOUT):

```
oci os object get -ns mynamespace -bn mybucket --name myfile.txt --file -
```

#### Bulk Operations in Object Storage

The CLI supports the following bulk operations in Object Storage:

- Uploading files in a directory and all its subdirectories to a bucket

```
Upload all the files in a directory.
oci os object bulk-upload -ns mynamespace -bn mybucket --src-dir path/to/upload/directory
```

- Downloading all objects, or all the objects that match a specified prefix, in a bucket

```
Download all the objects.
oci os object bulk-download -ns mynamespace -bn mybucket --download-dir
path/to/download/directory

Download all the objects that match the specified prefix.
```

## CHAPTER 34 Developer Tools

---

```
oci os object bulk-download -ns mynamespace -bn mybucket --download-dir
path/to/download/directory --prefix myprefix
```

- Deleting all objects, or all the objects that match a specified prefix, in a bucket

```
Delete all the objects.
oci os object bulk-delete -ns mynamespace -bn mybucket

Delete objects that match the specified prefix.
oci os object bulk-delete -ns mynamespace -bn mybucket --prefix myprefix
```

Bulk operations support several options that let you:

- Overwrite or skip files/objects using `--overwrite` or `--no-overwrite`. (**Note:** If you pass neither of these options you are prompted for confirmation every time there is something to overwrite.)
- Limit delete, upload, or download operations using `--prefix` and/or `--delimiter`
- Preview a bulk deletion with `--dry-run`

To get more information about the commands for bulk operations, run the following help commands:

```
bulk-upload
oci os object bulk-upload -h

bulk-download
oci os object bulk-download -h

bulk-delete
oci os object bulk-delete -h
```

### Multipart Operations in Object Storage

Multipart operations for Object Storage include object uploads and downloads.

#### MULTIPART UPLOADS

Large files can be uploaded to Object Storage in multiple parts to speed up the upload. By default, files larger than 128 MiB are uploaded using multipart operations. You can override this default by using the `--no-multipart` option.

You can configure the following options for the `oci os object put` command:

- `--no-multipart` overrides an automatic multipart upload if the object is larger than 128 MiB. The object is uploaded as a single part, regardless of size.
- `--part-size` in MiB, to use in a multipart operation. The default part size is 128 MiB and a part size that you specify must be greater than 10 MiB. If the object is larger than the `--part-size`, it is uploaded in multiple parts.
- `--parallel-upload-count`, to specify the number of parallel operations to perform. You can use this value to balance resources and upload times. A higher value may improve times but consume more system resources and network bandwidth. The default value is 10.

The `--resume-put` command allows you to resume a large file upload in cases where the upload was interrupted.



### Note

#### *Multipart Uploads from STDIN*

Objects uploaded from STDIN are uploaded in multiple parts. If the object content is smaller than 10 MiB, the upload is only 1 part, and the MultipartUpload API is used for the upload. Specifying `--no-multipart` when uploading from STDIN will result in an error.

The following example shows the command for a multipart upload if the object is larger than 200 MiB.

```
oci os object put -ns my-namespace -bn my-bucket --file path/to/large/file --part-size 200
```

For more information about multipart uploads, see [Using Multipart Uploads](#).

### MULTIPART DOWNLOADS

Large files can be downloaded from Object Storage in multiple parts to speed up the download.

You can configure the following options for the `oci os object get` command:

- `--multipart-download-threshold` lets you specify the size, in MiB at which an object should be downloaded in multiple parts. This size must be at least 128 MiB.
- `--part-size`, in MiB, to use for a download part. This gives you the flexibility to use more (smaller size) or fewer (larger size) parts as appropriate for your requirements. For example, compute power and network bandwidth. The default minimum part size is 120 MiB.
- `--parallel-download-count` lets you specify how many parts are downloaded at the same time. A higher value may improve times but consume more system resources and network bandwidth. The default value is 10.

The following example shows the command to download any object with a size greater than 500 MiB. The object is downloaded in 128 MiB parts.

```
oci os object get -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 500 --part-size 128
```

### Upgrading the CLI

If you installed the CLI manually, use one of the following commands to upgrade the CLI.

- To upgrade a standard installation, run the following command.

```
pip install oci-cli --upgrade
```

- To upgrade a standard virtualenv installation, run the following command.

```
cli-testing/bin/pip install oci-cli --upgrade
```

If you installed the CLI using the install script, use the following process to upgrade the CLI:

- Run the install script and specify the same install directory.
- When prompted, reply **Y** to overwrite the existing installation.

### Uninstalling the CLI

*FOR MANUAL INSTALLATIONS*

If you manually installed the CLI using pip, run the following command:

## CHAPTER 34 Developer Tools

---

```
pip uninstall oci-cli
```

If you manually installed the CLI in a virtualenv, run the following command:

```
<path/to/virtualenv>/bin/pip uninstall oci-cli
```

### *FOR SCRIPT INSTALLATIONS*

If you used the install script and the default installation location, you should delete the following directories.

On Windows:

- %USERPROFILE%/lib
- %USERPROFILE%/bin

On Mac:

- \$HOME/lib
- \$HOME/bin

If you used the install script, but installed to a custom location, you should delete the directories at that location.

### **UNINSTALLING PYTHON**

The script also installs Python as a dependency if it was not already installed. In Windows 10, you can uninstall Python in Control Panel or at the command line.

#### *USING CONTROL PANEL*

To uninstall Python in Control Panel, select **Programs and Features**. Right-click Python and select **Uninstall**.

For more information, see [Repair or remove programs in Windows 10](#).

#### *USING THE COMMAND LINE*

To uninstall Python at the command line, run the following command:

```
msiexec /x python<version>.msi
```

If you do not have the MSI file, you can also use the package or product code. For more information, see [Using the Windows Installer](#).

### Troubleshooting the CLI

This topic describes how to resolve issues that you might encounter when installing Python or the CLI, or when using the CLI.

#### SERVICE ERRORS

Any operation resulting in a service error causes an error of type "ServiceError" to be returned by the CLI. For information about common service errors that Oracle Cloud Infrastructure returns, see [API Errors](#).

#### ORACLE LINUX PERMISSIONS ISSUES

On Oracle Linux 7.3, if you encounter permission issues when running `pip install`, you might need to use `sudo`.

#### OCI COMMAND NOT FOUND

If the `oci` command isn't found, this can be caused by one of the following reasons:

- `pip` installed the package to a different virtual environment than your active one.
- You switched to a different active virtual environment after you installed the CLI.

To determine where the CLI is installed, run the `which pip` and `which oci` commands.

#### WHEEL FILE WON'T INSTALL

If the wheel file won't install, verify that `pip` is up to date. To update `pip`, run the `pip install -U pip` command. Try to install the wheel again.

#### WINDOWS ISSUES

If the `oci` command isn't found, make sure that the `oci.exe` location is in your path (for example, the `Scripts` directory in your Python installation).

## CHAPTER 34 Developer Tools

---

### CONTACT INFORMATION

If you want to contribute ideas, report a bug, get notified about updates, have questions, or want to give feedback, use one of the following links.

#### CONTRIBUTIONS

Got a fix for a bug, or a new feature you'd like to contribute? The CLI is open source and accepting pull requests on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the CLI is released, subscribe to the [Atom](#) feed.

#### QUESTIONS OR FEEDBACK

Ways to get in touch:

- [GitHub](#): To file bugs and feature requests only.
- [Stack Overflow](#): Use the [oracle-cloud-infrastructure](#) and [oci-cli](#) tags in your post.
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

## Other Tools and Plug-ins

### Tools and Plug-ins for SDKs and CLI

Oracle Cloud Infrastructure provides additional developer tools for automating processes and facilitating development.

**Toolkit for Eclipse** - The toolkit is an open source plug-in for the Eclipse Integrated Development Environment (IDE) that enables Java developers to code and deploy applications more quickly and efficiently.

- **Documentation:** [Toolkit for Eclipse](#)
- **Download:** To build and install the Toolkit, clone the [GitHub repository](#) then follow instructions in Getting Started with Toolkit for Eclipse.

**HDFS Connector for Object Storage** - Read and write data with your Apache Hadoop application to and from the Oracle Cloud Infrastructure Object Storage service. Building the connector relies on Maven artifacts that are provided by the SDK for Java.

- **Documentation:** [Properties](#)
- **Download:** [GitHub](#)

### DevOps Tools and Plug-ins

Oracle Cloud Infrastructure provides a number of DevOps tools and plug-ins for working with Oracle Cloud Infrastructure services. These can simplify provisioning and managing infrastructure or enable automated testing and continuous delivery.

**Terraform Provider** - Manage "infrastructure as code" with this component that connects Terraform to a given Oracle Cloud Infrastructure service.

- **Documentation:** [Terraform Provider](#)
- **Download:** [GitHub](#)

**Ansible Modules** - Automate provisioning and configuring of Oracle Cloud Infrastructure resources, such as Compute, Load Balancing, and Database services.

- **Documentation:** [Ansible Modules](#)
- **Download:** [GitHub](#)

**Chef Knife Plug-in** - Manage Oracle Cloud Infrastructure resources with Chef Knife, a command line tool that provides an interface between a local chef-repo and the Chef server.

- **Documentation:** [Chef Knife Plug-in](#)
- **Download:** [GitHub](#)

**Compute Jenkins Plug-in** - Bring up and down services or nodes as required to serve Jenkins Build Jobs and dynamically allocate Oracle Cloud Infrastructure resources for continuous integration tasks.

- **Documentation:** [Compute Jenkins Plug-in](#)
- **Download:** [GitHub](#)

**Grafana Plug-in** - Visualize metrics from the Monitoring service in your Grafana instance.

- **Documentation:** [Grafana Plug-in](#)
- **Download:** [GitHub](#)

**Terraform Kubernetes Installer** - Provision and configure the resources needed to run a highly available and configurable Kubernetes cluster.

- **Download:** [GitHub](#)

**Kubernetes Volume Provisioner** - Enable dynamic provisioning of storage resources when running Kubernetes on Oracle Cloud Infrastructure.

- **Download:** [GitHub](#)

### DevOps Integrations

- [Jenkins X Integration](#): Create a new Kubernetes cluster on Oracle Cloud Infrastructure Container Engine for Kubernetes.
- [Packer Integration](#): Create reusable custom images.

### Other Services and Features for DevOps

Oracle Cloud Infrastructure provides other services and features relevant to DevOps professionals.

- [Container Engine for Kubernetes \(OKE\)](#)  
Reliably build, deploy, and manage cloud-native containerized applications. You specify

the compute resources that your applications require, and Container Engine for Kubernetes provisions them on Oracle Cloud Infrastructure in an existing tenancy.

- [Oracle Cloud Infrastructure Registry](#)  
Store, share, and manage development artifacts like Docker images. As Oracle Cloud Infrastructure Registry is managed by Oracle, your applications are deployed reliably and you don't have to deal with operational issues.

### Toolkit for Eclipse

The Oracle Cloud Infrastructure Toolkit for Eclipse is an open source plug-in for the Eclipse Integrated Development Environment (IDE). The toolkit provides a set of features that help developers connect to Oracle Cloud Infrastructure from within Eclipse. For example, by using the toolkit, you can deploy a project to the cloud by using the Object Storage feature to upload multiple files in one click. The Compute feature enables you to start a compute instance or restart it if needed. You can also switch between multiple accounts and regions from the Eclipse IDE.

**Download:** To install the Toolkit, download the `com.oracle.oci.eclipse.zip` toolkit from [the releases section on GitHub](#), then follow the instructions in [Getting Started with Toolkit for Eclipse](#).

### Requirements

To use the Oracle Cloud Infrastructure Toolkit for Eclipse, you must have the following:

- An Oracle Cloud Infrastructure account
- A user created in that account, in a group with a policy that grants the desired permissions. This can be a user for yourself, or another person/system that needs to call the API. For an example of how to set up a new user, group, compartment, and policy, see [Adding Users](#). For a list of typical policies you may want to use, see [Common Policies](#).
- A key pair used for signing API requests, with the public key uploaded to Oracle. For more information on generating and uploading keys, see [Required Keys and OCIDs](#).
- Java 8

## CHAPTER 34 Developer Tools

---

- Maven
- [SDK for Java](#)
- Eclipse IDE for Java Developers 4.3 or later

### Services Supported

- Compute
- Object Storage

### Contact Us

#### CONTRIBUTIONS

Got a fix for a bug or a new feature you'd like to contribute? The plug-in is open source and accepting pull requests on [GitHub](#).

#### NOTIFICATIONS

To be notified when a new version of the toolkit is released, subscribe to the [Atom feed](#).

#### QUESTIONS OR FEEDBACK

- [GitHub Issues](#): To file bugs and feature requests only
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Getting Started with Toolkit for Eclipse

#### DOWNLOADING THE TOOLKIT

You can download the `com.oracle.oci.eclipse.zip` toolkit from [the releases section on GitHub](#).

#### INSTALLING THE TOOLKIT

After building the toolkit, launch the Eclipse IDE.

1. From the top navigation bar, select **Help > Install New Software...**
2. In Install dialog, click **Add...**
3. In the Add Repository dialog, click **Archive...**
4. In the right pane of the Repository Archive window, select the zip file containing the toolkit. Click **Open**.
5. In the Add Repository dialog click **Add**.
6. In the Available Software dialog, select **Oracle Cloud Infrastructure Toolkit for Eclipse**, then click **Next**.
7. In the Install Details dialog, click **Finish**.

### CONFIGURING THE TOOLKIT

#### *ORACLE CLOUD INFRASTRUCTURE PREFERENCES*

Before you can use the toolkit, you must configure the Oracle Cloud Infrastructure Preferences in the Eclipse IDE. This process will provide the necessary identifiers and credentials so the toolkit can connect to your Oracle Cloud Infrastructure account. For more information, see [Required Keys and OCIDs](#).

1. From the top navigation bar, select **Preferences > Oracle Cloud Infrastructure Preferences** .
2. For **Profile Name**, provide a short descriptive name.
3. From the **Region** dropdown, select your region.
4. Enter your **User OCID** and **Tenancy OCID**. For information on how to locate this information, see [Where to Get the Tenancy's OCID and User's OCID](#).
5. For **Key File**, click **Browse** and select the appropriate file. For more information, see [How to Generate an API Signing Key](#).
6. Enter the **Fingerprint** for the Key File. For more information, see [How to Get the Key's Fingerprint](#).
7. Enter the **Passphrase**, if you created one for the key pair. If not, leave this field blank.

8. Click **Save Profile**.
9. Click **Apply and Close**.

### *PROXY SETTINGS*

If you are on a network that uses a proxy to connect to the internet, you must configure Eclipse proxy settings. For more information, see Network Connections in the [Eclipse IDE Documentation](#).

### **UNINSTALLING THE TOOLKIT**

Launch the Eclipse IDE.

1. From the top navigation bar, select **Help > About Eclipse IDE**
2. Click **Installation Details**.
3. In the Installation Details window, select the Installed Software tab.
4. Select Oracle Cloud Infrastructure Toolkit for Eclipse and click **Uninstall...**
5. In the Uninstall dialog, confirm the items to be uninstalled then click **Finish**.
6. Click **Restart**.

### **Using Toolkit for Eclipse**

After configuring the [Oracle Cloud Infrastructure Preferences](#), you can connect to your tenancy via Eclipse and use the Oracle Cloud Infrastructure Explorer to view and update your resources. You can also switch between different accounts saved in the profile.

To change the region, click the region icon in the Explorer navigation bar and select from the dropdown.

To change the compartment, click the compartment icon and select from the dropdown.

### **USING THE TOOLKIT FOR ECLIPSE WITH COMPUTE INSTANCES**

In the Oracle Cloud Infrastructure Explorer, double-click **Compute** to view available resources.

Double-click **Instances** to display a list of instances. You can right-click each one in the list to start, stop, or reboot it.

You can also double-click **Block Volumes** to view a list of block volumes.

### USING THE TOOLKIT FOR ECLIPSE WITH OBJECT STORAGE

In the Oracle Cloud Infrastructure Explorer, double-click Object Storage to view available resources.

#### *WORKING WITH BUCKETS*

To create a bucket, right-click **Object Storage** and select **Create New Bucket**.

To view content and details, double-click or right-click the bucket and select **Open Bucket**.

To delete a bucket, view the bucket's details and click **Delete Bucket**.

#### *WORKING WITH OBJECTS*

To upload an object, select the destination bucket and view its details. Right-click the Object list and select **Upload Object**. You can also drag one or more files to the Object list to upload.

To download one or more objects, right-click and select **Download Object**.

To delete one or more objects, right-click and select **Delete Object**.

## HDFS Connector for Object Storage

The Hadoop Distributed File System (HDFS) Connector lets your Apache Hadoop application read and write data to and from the Oracle Cloud Infrastructure Object Storage service.

This SDK and sample is dual-licensed under the Universal Permissive License 1.0 and the Apache License 2.0; third-party content is separately licensed as described in the code.

- **Services supported:** Object Storage
- **Download:** [GitHub](#) or [Maven](#)
- **API Documentation:** [HDFS Connector API Reference](#)

### Requirements

To use the HDFS connector, you must have:

- An Oracle Cloud Infrastructure account.
- A user created in that account, in a group with a policy that grants the desired permissions for any bucket you want to use. This can be a user for yourself, or another person/system that needs to call the API. For an example of how to set up a new user, group, compartment, and policy, see [Adding Users](#). For a basic Object Storage policy, see [Let Object Storage admins manage buckets and objects](#).
- Java 8
- A TTL value of 60. For more information, see [Configuring JVM TTL for DNS Name Lookups](#).

### CREDENTIALS AND PASSWORDS

If you use an encrypted PEM file for credentials, the passphrase will be read from configuration using the `getPassword` Hadoop Configuration method. The `getPassword` option checks for a password in a registered security provider. If the security provider doesn't contain the requested key, it will fallback to reading the plaintext passphrase directly from the configuration file.

### CONFIGURING JVM TTL FOR DNS NAME LOOKUPS

The Java Virtual Machine (JVM) caches DNS responses from lookups for a set amount of time, called *time-to-live* (TTL). This ensures faster response time in code that requires frequent name resolution.

The JVM uses the [networkaddress.cache.ttl](#) property to specify the caching policy for DNS name lookups. The value is an integer that represents the number of seconds to cache the successful lookup. The default value for many JVMs, `-1`, indicates that the lookup should be cached forever.

Because resources in Oracle Cloud Infrastructure use DNS names that can change, we recommend that you change the the TTL value to 60 seconds. This ensures that the new IP address for the resource is returned on next DNS query. You can change this value globally or specifically for your application:

## CHAPTER 34 Developer Tools

---

- To set TTL globally for all applications using the JVM, add the following in the `$JAVA_HOME/jre/lib/security/java.security` file:

```
networkaddress.cache.ttl=60
```

- To set TTL only for your application, set the following in your application's initialization code:

```
java.security.Security.setProperty("networkaddress.cache.ttl" , "60");
```

### Installation

Copy the bundled jars from `lib` and `third-party/lib` to each node of the Hadoop cluster so that they are included in Hadoop's CLASSPATH.

### SDK FOR JAVA AND MAVEN ARTIFACTS

Building an HDFS connector relies on Maven artifacts that are provided by the SDK for Java. To obtain the artifacts, you must [download the SDK for Java](#) and build it locally. You can then build the HDFS connector.



#### Important

The SDK for Java version that you download from the [Oracle Releases page](#) must match the HDFS connector version, which you can find in the `hdfs-connector/pom.xml` file in the dependency tag block that has the `groupId` attribute.

### HDFS CONNECTOR AND MAVEN ARTIFACTS

The HDFS Connector is available on [Maven Central](#) and [JCenter](#).

To use the HDFS Connector in your project, import the following project dependency. For example:

```
<dependency>
 <groupId>com.oracle.oci.sdk</groupId>
 <artifactId>oci-hdfs-connector</artifactId>
```

## CHAPTER 34 Developer Tools

---

```
<!-- Replace the version below with your required version -->
<version>2.9.2.0</version>
</dependency>
```

### PROPERTIES

You can set the following HDFS Connector properties in the `core-site.xml` file. The [BmcProperties](#) page lists additional properties that you can configure for a connection to Object Storage.

Property	Description	Type	Required
<code>fs.oci.client.hostname</code>	URL of the host endpoint.  For example, <code>https://www.foo.com.</code>	String	Yes
<code>fs.oci.client.auth.tenantId</code>	OCID of your tenancy.  To get the value, see <a href="#">Required Keys and OCIDs</a> .	String	Yes
<code>fs.oci.client.auth.userId</code>	OCID of the user calling the API.  To get the value, see <a href="#">Required Keys and OCIDs</a> .	String	Yes

## CHAPTER 34 Developer Tools

Property	Description	Type	Required
<code>fs.oci.client.auth.fingerprint</code>	Fingerprint for the key pair being used.  To get the value, see <a href="#">Required Keys and OCIDs</a> .	String	Yes, unless you provide a custom authenticator.
<code>fs.oci.client.auth.pemfilepath</code>	Full path and file name of the private key used for authentication. The file should be on the local file system.	String	Yes, unless you provide a custom authenticator
<code>fs.oci.client.auth.passphrase</code>	Passphrase used for the key, if it is encrypted.	String	Only if your key is encrypted

You can specify that a property value applies to a specific bucket by appending `.<bucket_name>.<namespace_name>` to the property name.

This example shows how properties can be configured in a `core-site.xml` file (the OCIDs are shortened for brevity):

```
<configuration>
...
 <property>
 <name>fs.oci.client.hostname</name>
 <value>https://objectstorage.us-ashburn-1.oraclecloud.com</value>
 </property>
 <property>
 <name>fs.oci.client.hostname.myBucket.myNamespace</name>
 <value>https://objectstorage.phoenix-1.oraclecloud.com</value><!-- Use Phoenix for
myBucket@myNamespace -->
 </property>
 <property>
 <name>fs.oci.client.auth.tenantId</name>
 <value>ocid1.tenancy.oc1..aaaaaaaaba3pv6wkr4j...stifsfdsq</value>
```

```
</property>
<property>
 <name>fs.oci.client.auth.userId</name>
 <value>ocidl.user.oc1..aaaaaaaat5nvwcnazjc...aqw3rynjq</value>
</property>
<property>
 <name>fs.oci.client.auth.fingerprint</name>
 <value>20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34</value>
</property>
<property>
 <name>fs.oci.client.auth.pemfilepath</name>
 <value>~/oci/oci_api_key.pem</value>
</property>
...
</configuration>
```

### Using Instance Principals for Authentication

Oracle provides *instance principals* so that you no longer need to configure user credentials or provide PEM files on services running on instances. Each of these instances has its own identity and authenticates by using certificates added to the instance by instance principals.

To use instance principals authentication with the HDFS connector, simply provide the property `fs.oci.client.custom.authenticator` and set the value to `com.oracle.bmc.hdfs.auth.InstancePrincipalsCustomAuthenticator`.

Because using instance principals provides your instance with a custom authenticator, it is no longer necessary to configure the following properties:

- `fs.oci.client.auth.tenantId`
- `fs.oci.client.auth.userId`
- `fs.oci.client.auth.fingerprint`
- `fs.oci.client.auth.pemfilepath`
- `fs.oci.client.auth.passphrase`

The following example code illustrates using instance principals for authentication with the HDFS connector:

## CHAPTER 34 Developer Tools

```
<?xml version="1.0"?>
<configuration>
 <property>
 <name>fs.oci.client.hostname</name>
 <value>https://objectstorage.us-phoenix-1.oraclecloud.com</value>
 </property>
 <property>
 <name>fs.oci.client.custom.authenticator</name>
 <value>com.oracle.bmc.hdfs.auth.InstancePrincipalsCustomAuthenticator</value>
 </property>
</configuration>
```

For more information about instance principals, see [Announcing Instance Principals for Identity and Access Management](#).

### Configuring a HTTP Proxy

You can set the following optional properties in the `core-site.xml` file to configure a HTTP proxy:

Property	Description	Type	Required
<code>fs.oci.client.proxy.uri</code>	The URI of the proxy endpoint.  For example, <code>http://proxy.mydomain.com:80</code> .	String	No
<code>fs.oci.client.proxy.username</code>	The username to authenticate with the proxy.	String	No
<code>fs.oci.client.proxy.password</code>	The password to authenticate with the proxy.	String	No
<code>fs.oci.client.multipart.allowed</code>	Enables the upload manager to support multipart uploads	Boolean	No

Property	Description	Type	Required
<code>fs.oci.client.multipart.minobjects ize.mb</code>	Specifies the minimum object size in mebibytes in order to use the upload manager.	Integer	No
<code>fs.oci.client.multipart.partsize.m b</code>	Specifies the part size in mebibytes for the upload manager.	Integer	No

**Note**

Configuring a proxy enables use of the `ApacheConnectorProvider` when making connections to Object Storage. It buffers requests into memory and can impact memory utilization when uploading large objects. It is recommended to enable multipart uploads and adjust the multipart properties to manage memory consumption.

**Large Object Uploads**

Large objects are uploaded to Object Storage using multipart uploads. The file is split into smaller parts that are uploaded in parallel, which reduces upload times. This also enables the HDFS connector to retry uploads of failed parts instead of failing the entire upload. However, uploads may transiently fail, and the connector will attempt to abort partially uploaded files. Because these files accumulate (and you will be charged for storage), list the uploads periodically and then after a certain number of days abort them manually using the SDK for Java.

Information about using the Object Storage API for managing multipart uploads can be found in [API Documentation](#).



### Note

If you prefer not to use multipart uploads, you can disable them by setting the `fs.oci.client.multipart.allowed` property to `false`.

## Best Practices

The following sections contain best practices to optimize usage and performance.

### DIRECTORY NAMES

There are no actual directories in Object Storage. Directory grouping is a function of naming convention, where objects use / delimiters in their names. For example, an object named `a/example.json` implies there is a directory named `a`. However, if that object is deleted, the `a` directory is also deleted implicitly. To preserve filesystem semantics where the directory can exist without the presence of any files, the HDFS connector creates an actual object whose name ends in / with a path that represents the directory, (e.g., create an object named `a/`). Now, deleting `a/example.json` doesn't affect the existence of the `a` directory, because the `a/` object maintains its presence. However, it's entirely possible that somebody could delete that `a/` object without deleting the files/directories beneath it. The HDFS connector will only delete the folder object if there are no objects beneath that path. The folder object itself is zero bytes.

### INCONSISTENT FILESYSTEM

Deleting a directory means deleting all objects that start with the prefix representing that directory. HDFS allows you to query for the file status of a file or a directory. The file status of a directory is implemented by verifying that the folder object for that directory exists. However, it's possible that the folder object has been deleted, but some of the objects with that prefix still exist. For example, in a situation with these objects:

## CHAPTER 34 Developer Tools

---

- `a/b/example.json`
- `a/b/file.json`
- `a/b/`

HDFS would know that directory `/a/b/` exists and is a directory, and scanning it would result in `example.json` and `file.json`. However, if object `a/b/` was deleted, the filesystem would appear to be in an inconsistent state. You could query it for all files in directory `/a/b/` and find the two entries, but querying for the status of the actual `/a/b/` directory would result in an exception because the directory doesn't exist. The HDFS connector does not attempt to fix up the state of the filesystem.

### FILE CREATION

Object Storage supports objects that can be many gigabytes in size. Creating files will normally be done by writing to a temp file and then uploading the contents of the file when the stream is closed. The temp space must be large enough to handle multiple uploads. The temp directory used is controlled by the `hadoop.tmp.dir` configuration property.

### READ/SEEK SUPPORT

When in-memory buffers are enabled (`fs.oci.io.read.inmemory`), seek is fully supported because the entire file is buffered into a byte array. When in-memory buffer is not enabled (likely because object sizes are large), seek is implemented by closing the stream and making a new range request starting at the specified offset.

### DIRECTORY LISTING

Listing a directory is essentially a List bucket operation with a prefix and delimiter specified. To create an HDFS FileStatus instance for each key, the connector performs an additional HEAD request to get ObjectMetadata for each individual key. This will be required until Object Storage supports richer list operation data.

### URI Format for Filesystems and Files

HDFS filesystems and files are referenced through URIs. The scheme specifies the type of filesystem, and the remaining part of the URI is largely free for the filesystem

implementation to interpret as it wants.

Because Object Storage is an object store, its ability to name objects as if they were files in a filesystem is used to mimic an actual filesystem.

### Root

The root of Object Storage filesystem is denoted by a path where the authority component includes the bucket name and the namespace name, as shown:



#### Note

In the examples, "MyBucket" and "MyNamespace" are placeholders and should be replaced with appropriate values.

```
oci://MyBucket@MyNamespace/
```

This is always the root of the filesystem. The reason for using authority for both bucket and namespace is that HDFS only allows the authority portion to determine where the filesystem is; the path portion denotes just the path to the resource (so "oci//MyNamespace/MyBucket" won't work, for example). Note that the @ character is not a valid character for buckets or namespaces, and should allow the authority to be parsed correctly.

### Sub-directories

Sub-directories do not actually exist, but can be mimicked by creating objects with / characters. For example, two files named `a/b/c/example.json` and `a/b/d/path.json` would appear as if they were in a common directory `a/b`. This would be achieved by using the Object Storage prefix- and delimiter-based querying. In the given example, referencing a sub-directory as a URI would be:

```
oci://MyBucket@MyNamespace/a/b/
```

## CHAPTER 34 Developer Tools

---

### OBJECTS/FILES

An object named `a/b/c/example.json` is referenced as:

```
oci://MyBucket@MyNamespace/a/b/c/example.json
```

### Logging

Logging in the connector is done through [SLF4J](#). SLF4J is a logging abstraction that allows the use of a user-supplied logging library (e.g., log4j). For more information, see the [SLF4J manual](#).

The following example shows how to enable basic logging to standard output.

1. Download the SLF4J Simple binding jar: [SLF4J Simple Binding](#)
2. Add the jar to your classpath
3. Add the following VM arg to enable debug level logging (by default, info level is used): -  
`Dorg.slf4j.simpleLogger.defaultLogLevel=debug`

You can configure more advanced logging options by using the log4j binding.

### Sample Hadoop Job

[hadoop\\_sample\\_hdfs](#):

```
package com.oracle.oci.hadoop.example;

import java.io.ByteArrayOutputStream;
import java.io.IOException;
import java.net.URI;
import java.net.URISyntaxException;

import org.apache.commons.io.IOUtils;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.fs.FSDataInputStream;
import org.apache.hadoop.fs.FSDataOutputStream;
import org.apache.hadoop.fs.Path;
import org.apache.hadoop.io.IntWritable;
```

## CHAPTER 34 Developer Tools

---

```
import org.apache.hadoop.io.Text;
import org.apache.hadoop.mapreduce.Job;
import org.apache.hadoop.mapreduce.lib.input.FileInputFormat;
import org.apache.hadoop.mapreduce.lib.output.FileOutputFormat;
import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

import com.oracle.oci.hdfs.BmcFilesystem;

import lombok.RequiredArgsConstructor;

@RequiredArgsConstructor
public class SampleOracleBmcHadoopJob
{
 private static final String SAMPLE_JOB_PATH = "/samplehadoopjob";
 private static final String INPUT_FILE = SAMPLE_JOB_PATH + "/input.dat";
 private static final String OUTPUT_DIR = SAMPLE_JOB_PATH + "/output";

 // non-static since this is the runner class it needs to initialize after we set the properties
 private final Logger log = LoggerFactory.getLogger(SampleOracleBmcHadoopJob.class);

 /**
 * Runner for sample hadoop job. This expects 3 args: path to configuration file, Object Store
 namespace, Object
 * Store bucket. To run this, you must:
 *{@code
 *
 * *
 * Create a standard hadoop configuration file
 *
 * *
 * Create the bucket ahead of time.
 *
 * }
 *
 * This runner will create a test input file in a file '/samplehadoopjob/input.dat', and job results
 will be written
 * to '/samplehadoopjob/output'.
 *
 * @param args

```

## CHAPTER 34 Developer Tools

---

```
* 1) path to configuration file, 2) namespace, 3) bucket
* @throws Exception
*/
public static void main(final String[] args) throws Exception
{
 if (args.length != 3)
 {
 throw new IllegalArgumentException(
 "Must have 3 args: 1) path to config file, 2) object storage namespace, 3) object
storage bucket");
 }

 // redirect all logs to sysout
 System.setProperty("org.slf4j.simpleLogger.logFile", "System.out");
 System.setProperty("org.slf4j.simpleLogger.defaultLogLevel", "debug");

 final SampleOracleBmcHadoopJob job = new SampleOracleBmcHadoopJob(args[0], args[1], args[2]);
 System.exit(job.execute());
}

private final String configurationFilePath;
private final String namespace;
private final String bucket;

public int execute() throws IOException, ClassNotFoundException, InterruptedException,
URISyntaxException
{
 log.info("Creating hadoop configuration");
 final Configuration configuration = this.createConfiguration(this.configurationFilePath);

 final String authority = this.bucket + "@" + this.namespace;
 final String uri = "oci://" + authority;
 log.info("Using uri: {}", uri);

 log.info("Creating job inputs");
 this.setup(uri, configuration);

 log.info("Creating job");
 final Job job = this.createJob(configuration);

 final String in = uri + INPUT_FILE;
 final String out = uri + OUTPUT_DIR;
```

## CHAPTER 34 Developer Tools

---

```
log.info("Using input: {}", in);
log.info("Using output: {}", out);

FileInputFormat.addInputPath(job, new Path(in));
FileOutputFormat.setOutputPath(job, new Path(out));

log.info("Executing job...");
final int response = job.waitForCompletion(true) ? 0 : 1;

log.info("Attempting to read job results");
this.tryReadResult(uri, configuration);
return response;
}

private Configuration createConfiguration(final String configFile_path)
{
 final Configuration configuration = new Configuration();
 configuration.addResource(new Path(configFile_path));
 return configuration;
}

private void setup(final String uri, final Configuration configuration) throws IOException,
URISyntaxException
{
 try (final BmcFilesystem fs = new BmcFilesystem())
 {
 fs.initialize(new URI(uri), configuration);
 fs.delete(new Path(SAMPLE_JOB_PATH), true);
 final FSDataOutputStream output = fs.create(new Path(INPUT_FILE));
 output.writeChars("example\ncpath\ngak\ntest\nexample\ngak\n\ngak");
 output.close();
 }
}

private Job createJob(final Configuration configuration) throws IOException
{
 final Job job = Job.getInstance(configuration, "word count");
 job.setJarByClass(SampleOracleBmcHadoopJob.class);
 job.setMapperClass(SimpleMapper.class);
 job.setCombinerClass(SimpleReducer.class);
 job.setReducerClass(SimpleReducer.class);
 job.setOutputKeyClass(Text.class);
}
```

## CHAPTER 34 Developer Tools

---

```
 job.setOutputValueClass(IntWritable.class);
 return job;
 }

 private void tryReadResult(final String uri, final Configuration configuration)
 throws IOException, URISyntaxException
 {
 try (final BmcFilesystem fs = new BmcFilesystem())
 {
 fs.initialize(new URI(uri), configuration);
 // this should be the output file name, but that could change
 final FSDataInputStream input = fs.open(new Path(OUTPUT_DIR + "/part-r-00000"));

 final ByteArrayOutputStream baos = new ByteArrayOutputStream();
 IOUtils.copy(input, baos);
 log.info("\n====\n" + baos.toString() + "====");
 input.close();
 }
 }
}

package com.oracle.oci.hadoop.example;

import java.io.IOException;
import java.util.StringTokenizer;

import org.apache.hadoop.io.IntWritable;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.mapreduce.Mapper;

public class SimpleMapper extends Mapper
{
 private final static IntWritable one = new IntWritable(1);
 private final Text word = new Text();

 @Override
 public void map(final Object key, final Text value, final Context context) throws IOException,
 InterruptedException
 {
 final StringTokenizer itr = new StringTokenizer(value.toString());
 while (itr.hasMoreTokens())
```

## CHAPTER 34 Developer Tools

---

```
 {
 this.word.set(itr.nextToken());
 context.write(this.word, one);
 }
 }
}

package com.oracle.oci.hadoop.example;

import java.io.IOException;

import org.apache.hadoop.io.IntWritable;
import org.apache.hadoop.io.Text;
import org.apache.hadoop.mapreduce.Reducer;

public class SimpleReducer extends Reducer
{
 private final IntWritable result = new IntWritable();

 @Override
 public void reduce(final Text key, final Iterable values, final Context context)
 throws IOException, InterruptedException
 {
 int sum = 0;
 for (final IntWritable val : values)
 {
 sum += val.get();
 }
 this.result.set(sum);
 context.write(key, this.result);
 }
}
```

### Troubleshooting

This section contains troubleshooting information for the HDFS Connector.

### TROUBLESHOOTING SERVICE ERRORS

Any operation resulting in a service error will cause an exception of type `com.oracle.bmc.model.BmcException` to be thrown by the HDFS Connector. For information about common service errors returned by OCI, see [API Errors](#).

### JAVA ENCRYPTION KEY SIZE ERRORS

The HDFS Connector can only handle keys of 128 bit or lower key length. Users get "Invalid Key Exception" and "Illegal key size" errors when they use longer keys, such as AES256. Use one of the following workarounds to fix this issue:

- Use a 128 bit key, such as AES128.
- Install the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction from the following location:  
<http://www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html>

### Contributions

Got a fix for a bug, or a new feature you'd like to contribute? The SDK is open source and [accepting pull requests](#) on [GitHub](#).

### Notifications

If you wish to be notified when a new version of the HDFS connector is released, subscribe to the [Atom feed](#).

### Questions or Feedback

Ways to get in touch:

- [GitHub Issues](#): To file bugs and make feature requests
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) and [oci-hdfs-connector](#) tags in your post

- [Oracle Cloud Infrastructure Forum](#): Threads tagged with [Developer Tools](#)
- [My Oracle Support](#)

### Using the HDFS Connector with Spark

#### INTRODUCTION

This article provides a walk through that illustrates using the HDFS connector with the Spark application framework. For the walkthrough, we use the Oracle Linux 7.4 operating system, and we run Spark as a standalone on a single computer.

#### PREREQUISITES

Following are prerequisites for completing the walkthrough:

- You must have permission to launch a Compute instance. For guidance, see [Launching an Instance](#).
- You must be able to connect to the service instance that you've launched. For guidance, see [Connecting to an Instance](#).
- You must have the appropriate OCID, fingerprint, and private key for the Identity and Access Management (IAM) user that you will use to interact with an Object Storage. For guidance, see [SDK and Tool Configuration](#); see also [Resource Identifiers](#).
- You must have an Object Storage bucket that you can connect to.
- The IAM user must be able to read and write to that bucket using the Console.

#### USING SPARK

##### INSTALL SPARK AND DEPENDENCIES



#### Note

For the purpose of this example, install Spark into the current user's home directory. Note that for production scenarios, you would not do this.



### Note

Versions 2.7.7.0 and later no longer install all of the required third party dependencies. Required third party dependencies are bundled under the `third-party/lib` folder in the zip archive and should be installed manually.

1. Launch an instance of your Compute service. For guidance, see [Launching an Instance](#).
2. Ensure that your service instance has a public IP address so that you can connect using a Secure Shell (SSH) connection. For guidance, see [Instance Console Connections](#).
3. Connect to your service instance using an SSH connection.
4. Install Spark and its dependencies, Java and Scala, by using the code examples that follow.

```
We'll use wget to download some of the artifacts that need to be installed
sudo yum install wget

First install Java
sudo yum install java-1.8.0-openjdk.x86_64
export JAVA_HOME=/usr/lib/jvm/jre-1.8.0-openjdk
Should be something like: OpenJDK Runtime Environment (build 1.8.0_161-b14)
java -version

Then install Scala
wget https://downloads.lightbend.com/scala/2.12.4/scala-2.12.4.rpm
sudo yum install scala-2.12.4.rpm
Should be something like: Scala code runner version 2.12.4 -- Copyright 2002-2017, LAMP/EPFL and
Lightbend, Inc.
scala -version

Then download Spark
wget https://archive.apache.org/dist/spark/spark-2.2.1/spark-2.2.1-bin-hadoop2.7.tgz
tar xvf spark-2.2.1-bin-hadoop2.7.tgz
export SPARK_HOME=$HOME/spark-2.2.1-bin-hadoop2.7
export PATH=$PATH:$SPARK_HOME/bin
```

## CHAPTER 34 Developer Tools

---

```
Start a Spark master
cd $SPARK_HOME
./sbin/start-master.sh
```

### DOWNLOAD THE HDFS CONNECTOR AND CREATE CONFIGURATION FILES



#### Note

For the purposes of this example, place the JAR and key files in the current user's home directory. For production scenarios you would instead put these files in a common place that enforces the appropriate permissions (that is, readable by the user under which Spark and Hive are running).

Download the HDFS connector to the service instance and add the relevant configuration files by using the following code example. For additional information, see [HDFS Connector for Object Storage](#).

```
wget https://github.com/oracle/oci-hdfs-connector/releases/download/v2.9.2.1/oci-hdfs.zip
unzip oci-hdfs.zip -d oci-hdfs

cd $HOME
mkdir .oci
Create or copy your API key into the $HOME/.oci directory

cd $SPARK_HOME/conf
Create a core-site.xml (e.g. by transferring one you have, using vi etc.). Consult
https://docs.cloud.oracle.com/Content/API/SDKDocs/hdfsconnector.htm#two
for what this should look like

Create a spark-defaults.conf file from the template
cp spark-defaults.conf.template spark-defaults.conf
```

In the `spark-defaults.conf` file, add the following at the bottom:

```
spark.sql.hive.metastore.sharedPrefixes= shaded.oracle,com.oracle.bmc
```

## CHAPTER 34 Developer Tools

---

### PREPARE DATA

For testing data, we will use the MovieLens data set.

1. Download the latest data set at <https://grouplens.org/datasets/movielens/latest/>. Be sure to download the "Small" data set.
2. Unzip the download file.
3. Upload the `movies.csv` file to your Object Storage bucket.

### TEST USING THE SPARK SHELL

With the data ready, we can now launch the Spark shell and test it using a sample command:

```
cd $SPARK_HOME
./bin/spark-shell

scala> sc.wholeTextFiles("oci://PipedUploadTest@sampletenancy/")
java.io.IOException: No FileSystem for scheme: oci
```

You receive an error at this point because the `oci://` file system schema is not available. We need to reference the JAR file before starting the Spark shell. Here's an example for doing so:

```
./bin/spark-shell --jars $HOME/oci-hdfs/lib/oci-hdfs-full-1.2.7.jar

scala> sc.wholeTextFiles("oci://PipedUploadTest@sampletenancy/")
res0: org.apache.spark.rdd.RDD[(String, String)] = oci://PipedUploadTest@sampletenancy/
MapPartitionsRDD[1] at wholeTextFiles at <console>:25

scala> sc.textFile("oci://PipedUploadTest@sampletenancy/movies.csv").take(20).foreach(println)
movieId,title,genres
1,Toy Story (1995),Adventure|Animation|Children|Comedy|Fantasy
2,Jumanji (1995),Adventure|Children|Fantasy
3,Grumpier Old Men (1995),Comedy|Romance
4,Waiting to Exhale (1995),Comedy|Drama|Romance
5,Father of the Bride Part II (1995),Comedy
6,Heat (1995),Action|Crime|Thriller
7,Sabrina (1995),Comedy|Romance
8,Tom and Huck (1995),Adventure|Children
9,Sudden Death (1995),Action
10,GoldenEye (1995),Action|Adventure|Thriller
11,"American President, The (1995)",Comedy|Drama|Romance
12,Dracula: Dead and Loving It (1995),Comedy|Horror
```

## CHAPTER 34 Developer Tools

---

```
13,Balto (1995),Adventure|Animation|Children
14,Nixon (1995),Drama
15,Cutthroat Island (1995),Action|Adventure|Romance
16,Casino (1995),Crime|Drama
17,Sense and Sensibility (1995),Drama|Romance
18,Four Rooms (1995),Comedy
19,Ace Ventura: When Nature Calls (1995),Comedy
```

The command is successful so we are able to connect to Object Storage. Note that if you do not wish to pass the `--jars` argument each time the command executes, you can instead copy the `oci-hdfs-full` JAR file into the `$$SPARK_HOME/jars` directory.

### *START THE SPARK THRIFT SERVER*

Start the Spark Thrift Server on port 10015 and use the Beeline command line tool to establish a JDBC connection and then run a basic query, as shown here:

```
cd $$SPARK_HOME
./sbin/start-thriftserver.sh --hiveconf hive.server2.thrift.port=10015
```

Once the Spark server is running, we can launch Beeline, as shown here:

```
cd $$SPARK_HOME
./bin/beeline
Beeline version 1.2.1.spark2 by Apache Hive
beeline>
```

Next, connect to the server, as shown here:



### **Note**

For the purposes of this example, we have not configured any security, so any user name and password will be accepted. For production scenarios you would not do this.

```
beeline> !connect jdbc:hive2://localhost:10015 testuser testpass
Connecting to jdbc:hive2://localhost:10015
log4j:WARN No appenders could be found for logger (org.apache.hive.jdbc.Utils).
log4j:WARN Please initialize the log4j system properly.
```

## CHAPTER 34 Developer Tools

```
log4j:WARN See http://logging.apache.org/log4j/1.2/faq.html#noconfig for more info.
Connected to: Spark SQL (version 2.2.1)
Driver: Hive JDBC (version 1.2.1.spark2)
Transaction isolation: TRANSACTION_REPEATABLE_READ
0: jdbc:hive2://localhost:10015>
```

If we now check to see what tables exist, we see the following:

```
0: jdbc:hive2://localhost:10015> show tables;
+-----+-----+-----+---+
| database | tableName | isTemporary |
+-----+-----+-----+---+
+-----+-----+-----+---+
No rows selected (0.724 seconds)
```

None exist presently; however, we can create a table and link it to the `movies.csv` file that we downloaded and placed in the Object Storage bucket, as shown here:

```
0: jdbc:hive2://localhost:10015> create table test_table (movieId integer, title string, genres
string) using csv options (path "oci://myBucket@myTenant/movies.csv", header "true", delimiter ",");

0: jdbc:hive2://localhost:10015> describe formatted test_table;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
----+---+
| col_name | data_type | |
comment |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
----+---+
| movieId | int | | NULL
| | | |
| title | string | | NULL
| | | |
| genres | string | | NULL
| | | |
| | | |
| # Detailed Table Information | | |
| | | |
| Database | default | |
| | | |
| Table | test_table | |
| | | |
| Owner | opc | |
```

## CHAPTER 34 Developer Tools

```
|
| Created | Thu Mar 01 20:45:18 GMT 2018 |
|
| Last Access | Thu Jan 01 00:00:00 GMT 1970 |
|
| Type | EXTERNAL |
|
| Provider | csv |
|
| Table Properties | [transient_lastDdlTime=1519937118] |
|
| Location | oci://PipedUploadTest@sampletenancy/movies.csv |
| Serde Library | org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe |
|
| InputFormat | org.apache.hadoop.mapred.SequenceFileInputFormat |
|
| OutputFormat | org.apache.hadoop.hive ql.io.HiveSequenceFileOutputFormat |
|
| Storage Properties | [delimiter=,, header=true, serialization.format=1] |
|
+-----+-----+-----+-----+
----+---+
```

Note that the table stores its data externally in Object Storage and the data can be accessed using the HDFS connector (the `oci://` file system scheme). Now that we have a table, we can query it:

```
0: jdbc:hive2://localhost:10015> select * from test_table limit 10;
+-----+-----+-----+-----+
| movieId | title | genres |
+-----+-----+-----+-----+
| 1 | Toy Story (1995) | Adventure|Animation|Children|Comedy|Fantasy |
| 2 | Jumanji (1995) | Adventure|Children|Fantasy |
| 3 | Grumpier Old Men (1995) | Comedy|Romance |
| 4 | Waiting to Exhale (1995) | Comedy|Drama|Romance |
| 5 | Father of the Bride Part II (1995) | Comedy |
| 6 | Heat (1995) | Action|Crime|Thriller |
| 7 | Sabrina (1995) | Comedy|Romance |
| 8 | Tom and Huck (1995) | Adventure|Children |
| 9 | Sudden Death (1995) | Action |
| 10 | GoldenEye (1995) | Action|Adventure|Thriller |
+-----+-----+-----+-----+
```

### FOR MORE INFORMATION

- [HDFS Connector for Object Storage](#)
- [Overview of Object Storage](#)
- [Apache Spark](#)

## Terraform Provider

This topic provides information about installing, configuring, and using the Terraform provider with Oracle Cloud Infrastructure.

Terraform is a tool that allows you to programmatically manage, version, and persist your IT infrastructure as "infrastructure as code." Terraform uses declarative syntax to describe your infrastructure and then persist it in configuration files that can be shared, reviewed, edited, versioned, preserved, and reused.

The Oracle Cloud Infrastructure Terraform provider is a component that connects Terraform to the service infrastructure that you wish to manage.

- **Services supported:**

- Analytics Cloud
- Audit
- Budgets
- Container Engine for Kubernetes
- Compute Autoscaling
- Content and Experience
- Core Services (Networking, Compute, Block Volume)
- Database
- Digital Assistant
- DNS Service
- Email Delivery

- Events
  - FastConnect
  - File Storage
  - Functions
  - Health Checks
  - IAM
  - Integration Cloud
  - Key Management
  - Limits
  - Load Balancing
  - Monitoring
  - Notifications
  - Object Storage
  - Streaming
  - Web Application Firewall (WAF)
- **Licensing:** This provider and sample is licensed under the Mozilla Public License 2.0; third-party content is separately licensed as described in the code.
  - **Documentation:** [Oracle Cloud Infrastructure Provider](#)



### Note

For troubleshooting, see [Terraform state drift with tag defaults and tags for secondary resources](#) for a known issue with tags related to Terraform.

### Contributions

Got a fix for a bug, or a new feature you'd like to contribute? The Terraform provider is open source and accepting pull requests on [GitHub](#).

### Notifications

To be notified when a new version of the Terraform provider is released, subscribe to the [Atom feed](#).

### Questions or Feedback

Ways you can get in touch:

- [GitHub](#): To file bugs and feature requests only.
- [Stack Overflow](#): Please use the [oci-terraform](#) and [oracle-cloud-infrastructure](#) tags in your post.
- [Developer Tools section](#) of the Oracle Cloud forums.

### Getting Started with the Terraform Provider

This topic provides instructions for downloading and installing both Terraform and the Oracle Cloud Infrastructure Terraform provider, and provides a brief introduction to the key concepts for understanding and using the Oracle Cloud Infrastructure Terraform provider.

#### TERRAFORM OVERVIEW

Terraform is "infrastructure-as-code" software that allows you to define your infrastructure resources in files that you can persist, version, and share. These files describe the steps required to provision your infrastructure and maintain its desired state; it then executes these steps and builds out the described infrastructure.

Terraform's configuration and execution building blocks are modules, which are self-contained configuration packages. You can use these modules to organize your code and to create reusable components. HashiCorp, the developer of Terraform, provides a library of open-source Terraform modules "out of the box" to support many common tasks.

## CHAPTER 34 Developer Tools

---

To use Terraform for Oracle Cloud Infrastructure, you must download two components - Terraform from HashiCorp, and then the Oracle Cloud Infrastructure Terraform provider.

### DOWNLOAD AND INSTALL TERRAFORM

Download Terraform from the [HashiCorp download page](#). Ensure that you download the correct binary file for your system.



#### Important

The Oracle Cloud Infrastructure Terraform provider version 2.2.0 and greater requires Terraform version 0.10.1 or greater.

### DOWNLOAD THE ORACLE CLOUD INFRASTRUCTURE TERRAFORM PROVIDER

#### *PREREQUISITES FOR INSTALLING AND USING THE TERRAFORM PROVIDER*

- An Oracle Cloud Infrastructure account that has user credentials sufficient to execute a Terraform plan.
- A user in that account.
- Required keys and Oracle Cloud Infrastructure IDs (OCIDs). For guidance, see "Required Keys and OCIDs" in the *Oracle Cloud Infrastructure User Guide*.
- The correct Terraform binary file for your operating system (version 0.10.1 or greater).

#### *INSTALLING AND CONFIGURING THE TERRAFORM PROVIDER*

- For guidance on installing or on upgrading a previous version of the Oracle Cloud Infrastructure Terraform provider, see [Terraform Provider Version 3](#).
- For guidance on setting up the Terraform provider, see [Oracle Cloud Infrastructure Provider](#).

### FOR MORE INFORMATION

- [GitHub](#)
- [Hashicorp Terraform Documentation](#)
- [Creating Terraform Modules](#)
- [Terraform Configurations](#)
- [Terraform Configuration Syntax](#)

### Writing Terraform Configurations

#### OVERVIEW

Using Terraform, you can describe your Oracle Cloud Infrastructure using the HashiCorp Configuration Language format (HCL) in Terraform configuration files (see [Configuration Syntax](#)). Terraform configuration files can use either of two formats: Terraform domain-specific language (HashiCorp Configuration Language format [HCL]), which is the recommended approach, or JSON format if the files need to be machine-readable. Configuration files that use the HCL format end with the `.tf` file extension; those using JSON format end with the `.tf.json` file extension. The Terraform format is human-readable, while the JSON format is machine readable.

Use Terraform configurations to define your Oracle Cloud Infrastructure resources, variable definitions, data sources, and a great deal more. Terraform, then, converts your Oracle Cloud Infrastructure configurations into a set of API calls against Oracle Cloud Infrastructure API endpoints. The key to writing Terraform configuration is understanding how to abstract the desired infrastructure conceptually into [Terraform configuration syntax](#).



### Important

While the Oracle Cloud Infrastructure API uses camelCase extensively, Terraform does not support camelCase in configuration files. For this reason, in the configurations you see underscores rather than capitalization as separators. For example, where the API uses `availabilityDomain`, the Terraform configuration uses `availability_domain`.

### CONFIGURATION FILE REQUIREMENTS

Terraform configuration (`.tf`) files have specific requirements, depending on the components that are defined in the file. For example, you might have your Terraform provider defined in one file (`provider.tf`), your variables defined in another (`variables.tf`), your data sources defined in yet another.



### Note

We provide a great many example configuration files in the [Terraform Provider Examples](#) on our Oracle Cloud Infrastructure GitHub.

### PROVIDER DEFINITIONS

The following example using Terraform syntax illustrates the requirements for an Oracle Cloud Infrastructure Terraform provider definition, and also shows associated variable definitions. The provider definition relies on variables so that the configuration file itself does not contain sensitive data. Including sensitive data creates a security risk when exchanging or sharing configuration files.

```
variable "tenancy_ocid" {}
variable "user_ocid" {}
variable "fingerprint" {}
```

## CHAPTER 34 Developer Tools

---

```
variable "private_key_path" {}
variable "region" {}

provider "oci" {
 tenancy_ocid = "${var.tenancy_ocid}"
 user_ocid = "${var.user_ocid}"
 fingerprint = "${var.fingerprint}"
 private_key_path = "${var.private_key_path}"
 region = "${var.region}"
}
```

The `region` attribute specifies the geographical region in which your provider resources are created. To target multiple regions in a single configuration, you simply create a provider definition for each region and then differentiate by using a provider alias, as shown in the following example. Notice that only one provider, named "oci" is defined, and yet the `oci` provider definition is entered twice, once for the `us-phoenix-1` region (with the alias "phx"), and once for the region `us-ashburn-1` (with the alias "iad").

```
variable "tenancy_ocid" {}
variable "user_ocid" {}
variable "fingerprint" {}
variable "private_key_path" {}
variable "compartment_ocid" {}

provider "oci" {
 region = "us-phoenix-1"
 alias = "phx"
 tenancy_ocid = "${var.tenancy_ocid}"
 user_ocid = "${var.user_ocid}"
 fingerprint = "${var.fingerprint}"
 private_key_path = "${var.private_key_path}"
}

provider "oci" {
 region = "us-ashburn-1"
 alias = "iad"
 tenancy_ocid = "${var.tenancy_ocid}"
 user_ocid = "${var.user_ocid}"
 fingerprint = "${var.fingerprint}"
 private_key_path = "${var.private_key_path}"
}
```

For more information, see [Provider Configuration](#).

### VARIABLE DEFINITIONS

Variables in Terraform represent parameters for Terraform modules. In variable definitions, each block configures a single input variable, and each definition can take any or all of three optional arguments:

- `type` (optional): Defines the variable type as one of three allowed values: `string`, `list`, and `map`. If this argument is not used, the variable type is inferred based on `default`. If no `default` is provided, the type is assumed to be `string`.
- `default` (optional): Sets the default value for the variable. If no default value is provided, the caller must provide a value or Terraform throws an error.
- `description` (optional): A human-readable description of the variable.

Following are examples of several variable definitions. Some definitions include optional parameters.

```
variable "tenancy_ocid" {}
variable "user_ocid" {}
variable "fingerprint" {}
variable "private_key_path" {}
variable "region" {}

variable "AD" {
 default = "1"
 description = "Availability Domain"
}

variable "CPUCoreCount" {
 default = "2"
 type = "string"
}
```

For more information, see [Input Variable Configuration](#). See also [Input Variables](#).

### OUTPUT CONFIGURATION

Output variables provide a means to support Terraform end-user queries. This allows users to extract meaningful data from among the potentially massive amount of data associated with a complex infrastructure. For example, you might be interested only in a handful of key values

## CHAPTER 34 Developer Tools

---

at any given time and defining output variables allows you to extract exactly the information that you need.

Following is a simple example in which only a few output variables (instance IP addresses and boot volume IDs) are defined:

```
Output the private and public IPs of the instance
output "InstancePrivateIPs" {
 value = ["${oci_core_instance.TFInstance.*.private_ip}"]
}

output "InstancePublicIPs" {
 value = ["${oci_core_instance.TFInstance.*.public_ip}"]
}

Output the boot volume IDs of the instance
output "BootVolumeIDs" {
 value = ["${oci_core_instance.TFInstance.*.boot_volume_id}"]
}
```

For more information, see [Output Variables](#). See also [Output Configuration](#).

### RESOURCE CONFIGURATION

Resources are components of your Oracle Cloud Infrastructure. These resources include everything from low-level components such as physical and virtual servers, to higher-level components such as email and database providers, your DNS record.

### Declaring Resources

Following is a simple example of a resource definition that illustrates their basic structure.

```
resource "oci_core_virtual_network" "vcn1" {
 cidr_block = "10.0.0.0/16"
 dns_label = "vcn1"
 compartment_id = "${var.compartment_ocid}"
 display_name = "vcn1"
}
```

The resource declaration on the first line of the example uses the keyword "resource" and takes two parameters, resource `type` and resource `name` ("oci\_core\_virtual\_network" and "vcn1" in the example). Inside the code block, then, is the resource configuration.

For more information, see [Resource Configuration](#).

### Referencing Resources in Another Stack

You can reference resources that exist in other stacks. The Terraform [remote\\_state](#) data source allows you to read output variables from state files.

For example, when writing a Terraform configuration for a new web application, you can make the web application use the subnet previously created from your network stack, as long as the required subnet values were [output](#) in the network stack state file. In the Terraform configuration for your new web application, do the following:

- Pull the state file of the existing network stack into the context of your current Terraform configuration.
- Load the pulled state file to a data source for remote state files.
- Populate the subnet data source in your current configuration with values from the relevant output variables of the referenced state file.
- Optionally print the identifying information for the populated data source to confirm expected values.



#### Note

In addition to permissions required for Resource Manager operations, you'll need appropriate permissions for resource types you're referencing, in the compartment that you're referencing them. In this example, you need read permissions for network resources in the compartment where they're located.

The following Terraform configuration excerpt references a subnet in another stack.

```
The following example assumes that the source stack (defined by `stack_id`) has output a value named `subnet_id`
Terraform v0.12 is assumed
```

## CHAPTER 34 Developer Tools

---

```
variable "stack_id" {}

Pull the state file of the existing Resource Manager stack (the network stack) into this context
data "oci_resourcemanager_stack_tf_state" "stack1_tf_state" {
 stack_id = "${var.stack_id}"
 local_path = "stack1.tfstate"
}

Load the pulled state file into a remote state data source
data "terraform_remote_state" "external_stack_remote_state" {
 backend = "local"
 config = {
 path = "${data.oci_resourcemanager_stack_tf_state.stack1_tf_state.local_path}"
 }
}

Populate a data source in this configuration using a value from the remote state data source
data "oci_core_subnet" "subnet1" {
 subnet_id = "${data.terraform_remote_state.external_stack_remote_state.outputs.subnet_id}"
}

Print the values of the populated data source
output "print-subnet1" {
 value = "${data.oci_core_subnet.subnet1}"
}
```

### *DATA SOURCE CONFIGURATION*

Data sources represent read-only views of existing infrastructure intended for semantic use in Terraform configurations. Following is a simple example of a data source configuration to illustrate its basic structure:

```
Gets a list of Availability Domains
data "oci_identity_availability_domains" "ADs" {
 compartment_id = "${var.tenancy_ocid}"
}

Get DB node list
data "oci_database_db_nodes" "DBNodeList" {
 compartment_id = "${var.compartment_ocid}"
 db_system_id = "${oci_database_db_system.TFDBNode.id}"
}
```

## CHAPTER 34 Developer Tools

---

```
}

Get DB node details
data "oci_database_db_node" "DBNodeDetails" {
 db_node_id = "${lookup(data.oci_database_db_nodes.DBNodeList.db_nodes[0], "id")}"
}

Gets the OCID of the first (default) vNIC
data "oci_core_vnic" "DBNodeVnic" {
 vnic_id = "${data.oci_database_db_node.DBNodeDetails.vnic_id}"
}
```

For more information, see [Data Source Configuration](#).

### ENABLING INSTANCE PRINCIPAL AUTHORIZATION

Instance principal authorization allows your provider to make API calls from an Oracle Cloud Infrastructure compute instance without needing the `tenancy_ocid`, `user_ocid`, `fingerprint`, and `private_key_path` attributes in your provider definition.



#### Note

Instance principle authorization applies only to instances that are running in the Oracle Cloud Infrastructure.

To enable instance principal authorization for Oracle Cloud Infrastructure Terraform providers, set the `auth` attribute to "InstancePrincipal" in your provider definition, as shown in the following example:

```
variable "region" {}

provider "oci" {
 auth = "InstancePrincipal"
 region = "${var.region}"
}
```

### EXAMPLE TERRAFORM PROVIDERS

To see examples of the Oracle Cloud Infrastructure Terraform provider, see [Terraform Provider Examples](#). Several examples are provided and are grouped by service, including Compute, Database, Networking, Load Balancing, and several others.

### FOR MORE INFORMATION

- [Configuration](#)
- [Configuration Syntax](#)
- [Terraform Provider Examples](#)

### Using the Object Store for Terraform State Files

You can store [Terraform state files](#) in the Oracle Cloud Infrastructure Object Storage. Doing so requires that you configure a backend using one of the Terraform backend types.

Terraform supports various backend types to allow flexibility in how state files are loaded into Terraform. (For more information, see [Terraform Backend Types](#).) For our purposes, we address two of these approaches:

- Using an HTTP remote state backend
- Using an S3-compatible remote state backend

### USING AN HTTP BACKEND

Using the [HTTP backend type](#) allows you to store state using a simple REST client. With the HTTP backend type, you can easily fetch, update, and purge state using the `HTTP GET`, `POST`, and `DELETE` methods.

To configure the HTTP backend to store your Oracle Cloud Infrastructure Terraform state files, do the following:

#### *CREATE A PRE-AUTHENTICATED REQUEST*

Creating a pre-authenticated request in Oracle Object Storage enables accessing a bucket or object in the Oracle Cloud Infrastructure without needing to provide credentials. To do so, you must create a pre-authenticated request that has read/write permissions to the object store

where you intend to save the Terraform state file. You can do so in any of three ways: by using the Console UI, by using the command line interface (CLI), or by using the REST APIs.



### Note

A state file must exist in the bucket before you create the pre-authenticated request. This file can be an existing state file, or an empty file for the initial state.

For guidance, see [Using Pre-Authenticated Requests](#).

#### UPLOAD EXISTING STATE

If you have an existing state file, you can upload it using Curl to make an `HTTP PUT` request to the object store URL, as shown here:

```
curl -X PUT -H "Content-Type: text/plain" --data-binary "@path/to/local/tfstate" http://<prefix>/<my-access-uri>
```

#### CONFIGURE HTTP AS A TERRAFORM BACKEND

The [HTTP backend type](#) stores state using a simple REST client and allows you to easily fetch, update, and purge state using the `HTTP GET`, `POST`, and `DELETE` methods.

The access URI for addressing Oracle Cloud Infrastructure Terraform configurations must be of the form : `https://objectstorage.us-phoenix-1.oraclecloud.com/my-access-uri` (where region and access URI are specific to you).

For more example configuration and state files that reference code, and a summary of configuration variables, see [Standard Backends: HTTP](#).

Following is an example Terraform configuration. The region in the URL can be something other than the Phoenix region.

```
terraform {
 backend "http" {
 address = "https://objectstorage.us-phoenix-1.oraclecloud.com/<my-access-uri>" update_method =
"PUT" }
}
```

## CHAPTER 34 Developer Tools

---

### REINITIALIZE TERRAFORM

Finally, you must reinitialize Terraform and then run the `apply` command, as shown following.

```
terraform init
terraform apply
```

After completing these steps,, you are able to use Oracle Cloud Infrastructure as the backend for storing Terraform state files.

### USING AN S3-COMPATIBLE BACKEND

Configuring the S3-compatible backend requires that the account be enabled with S3 authentication keys, which are set on a per-user basis.

1. In the Console, open the navigation menu, then, under **Governance and Administration**, navigate to **Identity**, then **Users**. Under **User Details**, click **Customer Secret Key**. For guidance, see [Working with Amazon S3 Compatibility API Keys](#).
2. Set the location for the credentials file. The default location is `~/.aws/credentials`. You can set an alternate location by using the S3 backend `shared_credentials_file` option.



#### Warning

Never set the `access_key` and the `secret_key` attributes in the same Terraform backend configuration, since this creates a security risk.

3. Configure the `[default]` entry in the credentials file with the appropriate object storage credentials. The file can contain any number of credential profiles. If you provide a different profile name, you must also update the backend `profile` option in your Terraform configuration file.

Following is an example of Object Storage credentials:

```
[default]
aws_access_key_
id=ocid1.credential.oc1..aaaaaaaasbmhehdmevolvqwtbdjgwfsvxjsgxgipdbph7odn2brgurgsytca
aws_secret_access_key=mSTdaWhlbWj3ty4JZX1m0NUZV52x1ImWjayJLJ6OH9A=
```

Where `aws_access_key_id` and `aws_secret_access_key` are user-specific values provided from the Console. The key values provided in the example are not valid and provided as examples only.

4. Set the object storage `endpoint` value in the following format:

```
https://{namespace}.compat.objectstorage.{region}.oraclecloud.com
```

Following is a full example of an Object Storage backend configuration:

```
terraform {
 backend "s3" {
 bucket = "terraform-state"
 key = "terraform.tfstate"
 region = "us-phoenix-1"
 endpoint = "https://acme.compat.objectstorage.us-phoenix-1.oraclecloud.com"

 skip_region_validation = true
 skip_credentials_validation = true
 skip_requesting_account_id = true
 skip_get_ec2_platforms = true
 skip_metadata_api_check = true
 force_path_style = true
 }
}
```



### Note

The S3 backend configuration can also be used for the `terraform_remote_state` data source to enable sharing state across Terraform projects.

Once you have configured the backend, you must run `terraform init` to finish the setup. If you already have an existing `terraform.tfstate` file, then Terraform prompts you to confirm that the current state file is the one to upload to the remote state.

### FOR MORE INFORMATION

- [Using Pre-Authenticated Requests](#)
- [State Files](#)
- [Terraform Backend Types](#)

### Terraform Provider Best Practices

Following are recommended best practices for writing configurations for the Oracle Cloud Infrastructure Terraform provider.

- [Referencing Images](#)
- [Availability Domains](#)

### REFERENCING IMAGES

When launching Compute instances, your Terraform configuration should use the same image every time you run a Terraform Apply job.

To ensure this, specify the image OCID directly, rather than locating it using the `oci_core_image` data source. This is because the `oci_core_image` data source calls into the `ListImages` API, whose return values can change over time because over time images are added and older ones deleted. For a list of Oracle-Provided images and their OCIDs, see [Oracle-Provided Images](#). For more information, see [Results of `oci\_core\_images` will change over time for Oracle-provided images](#).

We recommend the following pattern for specifying an image for a given region:

```
variable "image_id" {
 type = "map"
 default = {
 // See https://docs.cloud.oracle.com/iaas/images/
 // Oracle-provided image "Oracle-Linux-7.4-2018.02.21-1"
 us-phoenix-1 = "ocid1.image.oc1.phx.aaaaaaaaaupbfz5f5hdvejulmalhyb6goieolullgkpumorbvxlwkaowg1slq"
 us-ashburn-1 = "ocid1.image.oc1.iad.aaaaaaaaj1w3xfie2t5t52uegyhiq2npx7bqyu4uvi2zyu3w3mqayc2bxmaa"
 eu-frankfurt-1 = "ocid1.image.oc1.eu-frankfurt-
1.aaaaaaaa7d3fsb6272srnftyi4dphdghfj6gurbqhm6ileds7ba3m2gltxq"
 uk-london-1 = "ocid1.image.oc1.uk-london-
1.aaaaaaaa6h6gj6v4n56mqrbrgnoskq63blyv2752g36zerymy63cfkojiiq"
```

## CHAPTER 34 Developer Tools

---

```
}
}
```

A Compute instance can use this in the following way:

```
resource "oci_core_instance" "TFInstance" {
 image = "${var.image_id[var.region]}"
 ...
}
```

### AVAILABILITY DOMAINS

With respect to Availability Domains, we caution against a common pattern, as shown here:

```
// Get all availability domains for the region
data "oci_identity_availability_domains" "ads" {
 compartment_id = "${var.tenancy_ocid}"
}

// Then either use it to get a single AD name based on the index:
resource "oci_core_instance" "nat" {
 availability_domain = "${lookup(data.oci_identity_availability_domains.ads.availability_domains
[var.nat_instance_ad], "name")}"
 ...
}

// Or iterate through all the ADs:
resource "oci_core_subnet" "nat" {
 count = "${length(data.oci_identity_availability_domains.ads.availability_domains)}"
 availability_domain = "${lookup(data.oci_identity_availability_domains.ad.availability_domains
[count.index], "name")}"
 ...
}
```

The recommendation, then, is to explicitly list the Availability Domain names for the regions in your configuration. To do so, use a variable that you have defined as follows:

```
variable "ad_list" {
 type = "list"
}
```

You can then use the variable as shown here:

```
// Index:
resource "oci_core_instance" "nat" {
```

## CHAPTER 34 Developer Tools

---

```
availability_domain = "${var.ad_list[var.nat_instance_ad_index]}"
...
}

// Or iterate through all the ADs:
resource "oci_core_subnet" "nat" {
 count = "${length(var.ad_list)}"
 availability_domain = "${var.ad_list[count.index]}"
 ...
}
```

You can then set the `ad_list` variable directly by using the availability domain names for your tenant and region, as shown here:

```
variable "ad_list" {
 type = "list"
 default = ["kIdk:PHX-AD-1", "kIdk:PHX-AD-2", "kIdk:PHX-AD-3"]
}
```

The advantage of using this method is that it gives you control over your availability domain usage and prevents unexpected changes over time. However, this approach is problematic when configurations are shared between tenancies and regions, since availability domain names are tenancy- and region-specific.

A convenient alternative is to instead set the `ad_list` value by using the `oci_identity_availability_domains` data source. You should do this in the configuration, then pass them into the modules. This effectively centralizes the list of ADs, making it is easy to switch to an explicit list later, should that become necessary: Note that the modules themselves should not use the `oci_identity_availability_domains` data source.

```
data "oci_identity_availability_domains" "ad" {
 compartment_id = "${var.tenancy_ocid}"
}

data "template_file" "ad_names" {
 count = "${length(data.oci_identity_availability_domains.ad.availability_domains)}"
 template = "${lookup(data.oci_identity_availability_domains.ad.availability_domains[count.index],
"name")}"
}

module "ssm_network" {
 ad_list = "${data.template_file.ad_names.*.rendered}"
}
```

```
...
}
```

### Ansible Modules

This topic provides information about installing, configuring, and using Ansible and the Oracle Cloud Infrastructure Ansible modules.

Oracle supports the use of Ansible for cloud infrastructure provisioning, orchestration, and configuration management. Ansible allows you to automate configuring and provisioning your cloud infrastructure, deploying and updating software assets, and orchestrating your complex operational processes.

What enables orchestrating, provisioning, and configuration management tasks are the Ansible modules for Oracle Cloud Infrastructure. Ansible provides a library of these Ansible modules "out of the box" for managing common tasks, and libraries of custom modules from cloud providers like AWS and Azure (see the [Module Index](#)). Oracle also provides a library of Ansible cloud modules that support provisioning and managing Oracle Cloud Infrastructure services.

Ansible [playbooks](#) automate configuration, deployment, and orchestration tasks using a declarative language ([YAML](#)), which allows you to describe infrastructure configuration, deployment policy, and orchestrating complex process steps, either synchronously or asynchronously. Ansible playbooks can be thought of as automation instruction manuals; Ansible modules, then, are your task execution tools.

Ansible modules allow you to author Ansible playbooks that enable automating the provisioning and configuring of Oracle Cloud Infrastructure services and resources, such as Compute, Load Balancing, Database, and other Oracle Cloud Infrastructure services.

- **Services supported:**

- Block Volume
  - Compute
  - Container Engine for Kubernetes
  - Database
  - DNS
  - Email Delivery
  - File Storage
  - IAM
  - Load Balancing
  - Networking
  - Object Storage
  - Search
  - WAF
- **Licensing:** Copyright © 2018, Oracle and/or its affiliates. This software is made available to you under the terms of the GPL 3.0 license or the Apache 2.0 license. See LICENSE.TXT for details.
  - **Download:** You can download the Oracle Cloud Infrastructure Ansible Module from the Oracle GitHub repository. Please follow guidance in [Getting Started with Ansible for Oracle Cloud Infrastructure](#), in the section "Installing the Oracle Cloud Infrastructure Ansible Modules."
  - **Documentation:** [Getting Started with Ansible for Oracle Cloud Infrastructure](#)

### Requirements

To use Ansible, you must have the following prerequisites on your controller computer, that is, the computer from which Ansible playbooks are executed. For more information, see the [Ansible Installation Guide](#).

- An Oracle Cloud Infrastructure account.
- A user created in that account, in a security group with a policy that grants the necessary permissions for working with resources in those compartments. For guidance, see [How Policies Work](#).
- The necessary credentials and OCID information.
- The Oracle Cloud Infrastructure Python SDK. To download and install the SDK, see the topic [Python SDK](#).
- Install the Ansible modules by following guidance in [Getting Started with Ansible for Oracle Cloud Infrastructure](#), in the section "Installing the Oracle Cloud Infrastructure Ansible Modules."

### Ansible Key Concepts

#### MODULES

Modules represent discrete provisioning tasks or operations that you can invoke individually from the command line, or else run individually or in sequence from a playbook. Ansible provides a large library of out-of-box modules that are listed in the [Module Index](#). Other Ansible providers like Microsoft Azure and Amazon Web Services (AWS) also provide libraries of Ansible modules. Oracle also provides a library of Ansible modules that interact with services. For more information, see [Working with Modules](#).

#### PLAYBOOKS

Playbooks provide a declarative language that allows you to create and persist your Ansible cloud infrastructure provisioning tasks. Playbooks are coded sets of instructions that you create to manage cloud infrastructure provisioning, and more advanced processes. For more information, see [Working with Playbooks](#). See also a set of available [Example Playbooks](#).

#### ROLES

Ansible roles are units of organization that allows you to abstract configuration, orchestration, and provisioning tasks into roles that you can save and share among playbooks and other users, and that are useful for organizing functionality in playbooks. In a sense, playbooks are

minimalist playbooks that encapsulate common configuration steps so you can share them across users or playbooks. For more information, see [Ansible Roles](#).

### INVENTORY

Ansible inventory is a means for describing hosts and groups. The inventory can be static (as a simple .ini file), or dynamic, where a look-up script assembles an up-to-date infrastructure inventory. For more information about Ansible inventory, see [Working with Inventory](#).



#### Important

When using Ansible to work with Oracle Cloud Infrastructure hosts, we recommend using the Oracle dynamic inventory script to obtain a dynamic inventory list. For more information, see [Using the Dynamic Inventory Script](#).

### Notifications

To be notified when new Ansible modules are released, subscribe to the [Ansible Atom Feed](#).

### Questions or Feedback

Ways to get in touch:

- [GitHub](#): To file bugs and feature requests only.
- [Stack Overflow](#): Use the [oci-ansible](#) and [oracle-cloud-infrastructure](#) tags in your post.
- [Developer Tools section](#) of the Oracle Cloud forums.

### Getting Started with Ansible for Oracle Cloud Infrastructure

This topic discusses how to get started with downloading and using Ansible with the Oracle Cloud Infrastructure. There are four initial steps for getting started with Ansible:

## CHAPTER 34 Developer Tools

---

- Ensure that you have all of the prerequisites
- Download and install the Oracle Cloud Infrastructure Python SDK
- Download and install Ansible
- Download and install the Ansible modules for Oracle Cloud Infrastructure

### PREREQUISITES FOR USING ANSIBLE FOR ORACLE CLOUD INFRASTRUCTURE

- You must have an Oracle Cloud Infrastructure account.
- Create a user in that account, in a security group with a policy that grants necessary permissions for working with resources in the account compartments.
- You must have the necessary credentials and OCID information.

### INSTALLING THE ORACLE CLOUD INFRASTRUCTURE PYTHON SDK

1. Download and install the Python SDK by following instructions in the topic, [Python SDK](#). For additional guidance, see [Downloading and Installing the SDK](#).
2. After installing the Python SDK, you must configure it using instructions in the topic [Configuring the SDK](#).

### INSTALLING AND CONFIGURING ANSIBLE

- To install Ansible, follow the instructions provided in the [Ansible Installation Guide](#).
- For guidance configuring Ansible, see [Configuring Ansible](#).

### INSTALLING ORACLE CLOUD INFRASTRUCTURE ANSIBLE MODULES ON AN ORACLE LINUX IMAGE

Oracle Cloud Infrastructure Ansible Modules come pre-installed on the [Oracle Cloud Developer image](#).



#### Note

For more information on Oracle Cloud Infrastructure-provided images, see [Oracle-Provided Images](#).

## CHAPTER 34 Developer Tools

---

To install Oracle Cloud Infrastructure Ansible Modules on an Oracle Linux image:

1. `$ yum install oci-ansible-modules`

### INSTALLING ORACLE CLOUD INFRASTRUCTURE ANSIBLE MODULES ON A NON-ORACLE LINUX IMAGE

1. `$ git clone https://github.com/oracle/oci-ansible-modules.git`
2. `$ cd oci-ansible-modules`
3. Run one of the following commands:
  - a. If Ansible is installed as a user:  
`$ ./install.py`
  - b. If Ansible is installed as root:  
`$ sudo ./install.py`

### SAMPLE ANSIBLE MODULES

Sample modules are available in the Oracle Cloud Infrastructure Ansible Module GitHub project. The samples library is updated regularly with the addition of new samples. You can access the samples at <https://github.com/oracle/oci-ansible-modules>.

### WRITING A SAMPLE PLAYBOOK

You can now write a sample playbook that uses Ansible modules. Following is an example playbook (named `list_buckets.yml`) that uses the `oci_bucket_facts` module to fetch all of the facts pertaining to all of the buckets in your compartment.

```

- name : List summary of existing buckets in OCI object storage
 connection: local
 hosts: localhost
 tasks:
 - name: List bucket facts
 oci_bucket_facts:
 namespace_name: '<yournamespace>'
```

## CHAPTER 34 Developer Tools

---

```
 compartment_id: '<yourcompartmentocid>'
 register: result
- name: Dump result
 debug:
 msg: '{{result}}'
```

### EXECUTING THE PLAYBOOK

Execute the Ansible playbook using Python by invoking this command:

```
$ ansible-playbook list_buckets.yml
```

### HOW TO OBTAIN MODULE DOCUMENTATION

To obtain access to detailed information about using Ansible modules in the CLI, including documentation of a module's configurable options, samples, return values, and so forth, use the `ansible-doc` command on the module's name. For example, to get the documentation for the `oci_bucket_facts` module, execute the following command:

```
$ ansible-doc oci_bucket_facts
```

Documentation of the Oracle Cloud Infrastructure Ansible modules is also available in the [Cloud Modules](#) page of the [Oracle Cloud Infrastructure Ansible Modules](#) site.

### CONFIGURING AUTHENTICATION

When creating and configuring Oracle Cloud Infrastructure resources, Ansible modules use authentication information that is outlined in [SDK and CLI Configuration File](#).



### Warning

IAM credentials that are referenced in Oracle Cloud Infrastructure SDK configuration files grant access to Oracle Cloud Infrastructure resources. Because of this, it is important to secure the credentials to prevent unauthorized access to these resources. To secure the credentials on the controller node where your Ansible playbooks run, follow guidelines outlined in the document [Securing IAM](#) (see section entitled "IAM Credentials").

Ansible modules permit you to override authentication information specified in the SDK configuration file by using module options and environment variables. Documentation for this is provided internally, as described in the preceding section, **How to Obtain Module Documentation**. However, using environment variables and Ansible module options to override authentication information must be avoided in production scenarios.

We recommend using Oracle Cloud Infrastructure SDK configuration files to specify authentication information. Use the "profiles" feature in the SDK configuration file to support different users. When distributing roles that use Ansible modules, ensure that no IAM credentials are included with the roles.

#### FOR MORE INFORMATION

- [Sample Ansible Playbooks](#)
- [Get Started with Ansible](#)
- [How Ansible Works](#)
- [Ansible for DevOps](#)

#### Using the Dynamic Inventory Script

Ansible tracks configuration resources by preserving lists, called inventory lists, as simple list files (also sometimes called a *hostfile*). These inventory files can be either simple static lists,

or they can be dynamic lists that automatically update when inventory resources are added, deleted, or moved. For more information about Ansible inventory files, see [Working with Inventory](#). See also, [Working with Dynamic Inventory](#).

When using Ansible to work with hosts that you have provisioned in Oracle Cloud Infrastructure, static inventory lists can cause problems because Compute instances are added and deleted over time. They can also be affected by external tools such as Terraform, or by the Oracle Cloud Infrastructure SDKs.

Having up-to-date and accurate inventory lists is essential for running Ansible playbooks. Oracle Cloud Infrastructure provides you with a script that you can download and run to ensure that your instance inventory list is always up-to-date. The script ensures that you always have the current set of Oracle Cloud Infrastructure compute instances available to your playbooks.

### DOWNLOAD AND CONFIGURE THE DYNAMIC INVENTORY SCRIPT

Download the dynamic inventory script (`oci_inventory.py`) and the default configuration file (`oci_inventory.ini`) from the following:

- [https://github.com/oracle/oci-ansible-modules/blob/master/inventory-script/oci\\_inventory.py](https://github.com/oracle/oci-ansible-modules/blob/master/inventory-script/oci_inventory.py)
- [https://github.com/oracle/oci-ansible-modules/blob/master/inventory-script/oci\\_inventory.ini](https://github.com/oracle/oci-ansible-modules/blob/master/inventory-script/oci_inventory.ini)



#### Important

Before you can use the script, ensure that you have a valid Oracle Cloud Infrastructure configuration. For guidance configuring `~/.oci/config`, see [SDK and CLI Configuration File](#).

#### SCRIPT AND CONFIGURATION DETAILS

The Python script, `oci_inventory.py`, uses the Oracle Cloud Infrastructure [Python SDK](#) to query compute instances in your Oracle Cloud Infrastructure tenancy and then uses this

## CHAPTER 34 Developer Tools

---

information to create a dynamic inventory that you can use with your Ansible playbooks. Use arguments in the Python script to control the configuration profile and the tenancy compartment that you query against.

You can use the configuration file, `oci_inventory.ini`, to control how inventory details are cached, and to control which Oracle Cloud Infrastructure profile to use. This file allows you to modify host names, and to control how hosts are named in the inventory list that is generated.

### *ARGUMENTS AND ENVIRONMENT VARIABLES*

For a complete list of command-line arguments and environment variables that the script accepts, see [Dynamic Inventory Script](#).

### *ORDER OF PRECEDENCE*

Following is the order of precedence for configurations that are used by the dynamic inventory script:

1. Command-line arguments.
2. Environment variables.
3. Configuration settings in the selected `profile` in your Oracle Cloud Infrastructure configuration file.



#### **Note**

The default configuration file used by the inventory script is `./oci_inventory.ini`. The Oracle Cloud Infrastructure SDK configuration file, however, defaults to `~/.oci/config`. The script reads the `DEFAULT` profile from the `config` file if no profile name is specified.



### Important

By default, the inventory is generated for all the compartments in the tenancy. You must have `COMPARTMENT_INSPECT` permission on the root compartment for this script to be able to access all compartments. However, when `compartment_ocid` is specified, the inventory is generated for only the specific compartment, so you only need `COMPARTMENT_INSPECT` permission on the specified compartment.

The inventory list that is generated by the dynamic inventory script is grouped using the following attributes:

- The region in which the compute instance resides.
- The name of the compartment the compute instance belongs to.
- The Availability Domain the compute instance is in.
- The `vcn_id` of the vcn the compute instance is in.
- The `subnet_id` of the subnet the compute instance is in.
- The `security_list_ids` of the subnet the compute instance is in.
- The `image_id` of the image used to launch the compute instance.
- Shape of the compute instance.
- The instance's free-form tags, with the group name set to `tag_<tag_name>=<tag_value>`.
- The instance's defined tags, with the group name set to `<tag_namespace>#<tag_name>=<tag_value>`.
- Oracle Cloud Infrastructure compute instance metadata (key-value pairs), with the group name set to `<metadata-key>=<metadata-value>`.
- Oracle Cloud Infrastructure compute instance extended metadata (key-value pairs), with the group name set to `<metadata-key>=<metadata-value>`.



### Important

By default, all non-alphanumeric characters in group names and host names are replaced with an underscore (`_`) when the inventory is generated *except* hash (`#`), equals (`=`), period (`.`) and dash (`-`). This allows you to use these names as Ansible group names. To disable this default substitution, set `sanitize_names` to `False` in the dynamic inventory settings file, whose default location is `./oci_inventory.ini`). To also replace the dash (`-`) when `sanitize_names` is `True`, set `replace_dash_in_names` to `True` in the settings file.

## HOW TO USE DYNAMIC INVENTORY

### USING DYNAMIC INVENTORY DURING PLAYBOOK EXECUTION

To use your dynamic inventory, first ensure that you have a correct Oracle Cloud Infrastructure SDK configuration file. Optionally, you can also have an `oci_inventory.ini` file.

Invoke `ansible-playbook` using the following command:

```
$ ansible-playbook -i <path-to-inventory-file>/oci_inventory.py <your-playbook-using-the-generated-inventory>
```

Alternatively, use the `ANSIBLE_HOSTS` environment variable by using the following command:

```
$ ANSIBLE_HOSTS=<path-to-inventory-file>/oci_inventory.py ansible-playbook <your-playbook-using-the-generated-inventory>
```

### DISABLING THE CACHE AND FETCHING THE LATEST INVENTORY LIST

If you are running the dynamic inventory script in a stand-alone manner, you can ignore the cached inventory list and fetch the most current inventory list by using the `--refresh` (`-r`) argument, as shown in the following example:

```
$ \
```

## CHAPTER 34 Developer Tools

---

If instead you use the inventory script during an Ansible playbook invocation, set the `OCI_CACHE_MAX_AGE` environment variable to 0 (zero) to ignore the cached list, and fetch the latest compute instance, as shown in the following example:

```
$ OCI_CACHE_MAX_AGE=0 ansible-playbook -i <path-to-inventory-file>/oci_inventory.py <your-playbook-using-the-generated-inventory>
```

### *DEBUGGING THE INVENTORY LIST*

To examine the inventory list, run the dynamic inventory script using the `--list` argument, as shown here:

```
$ \
```

If you wish to print additional debug information to `STDERR`, use the `--debug` argument, as shown:

```
$ \
```

### *RETRIEVE INFORMATION ABOUT A HOST*

You can configure the inventory script to provide information about a specified host by using the `--host` argument and providing the host's IP address, as shown:

```
$ \
```

The command returns the following set of variables and values:

```
{
 "availability_domain": "IwGV:US-ASHBURN-AD-1",
 "compartment_id": "ocidl.compartment.oc1..<xxxxxEXAMPLExxxxx>",
 "defined_tags": {},
 "display_name": "ansible-test-instance",
 "extended_metadata": {},
 "freeform_tags": {},
 "id": "ocidl.instance.oc1.iad.<xxxxxEXAMPLExxxxx>",
 "image_id": "ocidl.image.oc1.iad.<xxxxxEXAMPLExxxxx>",
 "ipxe_script": null,
 "launch_mode": "CUSTOM",
 "launch_options": {
 "boot_volume_type": "ISCSI",
 "firmware": "UEFI_64",
 "network_type": "VFIO",
 "remote_data_volume_type": "ISCSI"
 }
}
```

```
},
"lifecycle_state": "AVAILABLE",
"metadata": {
 "baz": "quux",
 "foo": "bar"
},
"region": "iad",
"shape": "VM.Standard1.1",
"source_details": {
 "image_id": "ocidl.image.oc1.iad.<xxxxxEXAMPLExxxxx>",
 "source_type": "image"
},
"time_created": "2018-01-16T12:13:35.336000+00:00"
}
```

### TROUBLESHOOTING THE DYNAMIC INVENTORY SCRIPT

If the inventory list generated by the inventory script does not include all of the compute instances in your tenancy, review the following information.

#### *USER PERMISSIONS*

Ensure that the user OCID (specified using either the `OCI_USER` environment variable, or the `profile` section in your SDK configuration file) has the policy permissions to list the compute instances. To see a list of permissions for API operations, see [Details for the Core Services](#).

The inventory script makes API calls for the following operations:

- ListCompartments
- ListVNICAttachments
- GetSubnet
- GetVCN
- GetVNIC
- GetInstance

#### *HOSTNAME FORMAT*

The default for `OCI_HOSTNAME_FORMAT` is `public_ip`. The dynamic inventory generated using `OCI_HOSTNAME_FORMAT` set to `public_ip` contains only compute instances that have a public

IP address. This can be useful in cases where the Ansible controller node is outside the VCN, since Ansible can only reach instances that have public IP addresses.

However, if running Ansible in a compute instance within your VCN, and it has access to all of the subnets within your VCN, including compute instances that have private IP addresses, you must set the `OCI_HOSTNAME_FORMAT` to `private_ip` to list compute instances using their private IP addresses.

### Sample Ansible Playbooks

Provided here is a catalog of sample Ansible playbooks for Oracle Cloud Infrastructure that illustrate how to carry out common infrastructure provisioning and configuration tasks. The samples are organized in groups associated with Oracle Cloud Infrastructure services:

- Block Volume
- Compute
- Container Engine for Kubernetes
- Database
- File Storage
- IAM
- Load Balancing
- Object Storage
- Deployment Solution (MongoDB)

You find a brief description of each playbook in the sections that follow, along with links to each sample on the Oracle GitHub repository. Begin by reviewing the `Readme.md` file that you find in each playbook's root directory.

### **BLOCK VOLUME**

#### *ATTACH A BLOCK VOLUME TO A COMPUTE INSTANCE*

This sample playbook shows how to attach a block volume to a compute instance using the iSCSI volume attachment type, and then connect it to the compute instance using `iscsiadm`. The sample shows how to do the following:

- Generate a temporary, host-specific SSH key pair.
- Specify the public key from the key pair for connecting to the instance, and then launch the instance.
- Create a new Block Volume for the instance, attach the volume to the instance, and specify `iSCSI` as the volume attachment type.
- Connect to and then mount the volume from the compute instance by executing `iscsiadm` commands over SSH using an Ansible module.

[Go to the sample on Oracle GitHub.](#)

### **COMPUTE**

#### *LAUNCH A COMPUTE INSTANCE*

This sample playbook shows how to launch a public Compute instance and then access the instance from an Ansible module over an SSH connection. The sample illustrates how to do the following:

- Generate a temporary, host-specific SSH key pair.
- Specify the public key from the key pair for connecting to the instance, and then launch the instance.
- Connect to the newly launched instance using SSH.

[Go to the sample on Oracle GitHub.](#)

#### *USE NAT TO ENABLE INTERNET ACCESS FROM A COMPUTE INSTANCE*

This sample playbook shows how to enable internet access from a Compute instance in a private subnet. The example uses a network address translation (NAT) instance in a public

subnet through an Ansible module. The sample illustrates how to do the following:



### Note

For guidance setting up the network topology to support a NAT instance, see the white paper [NAT Instance Configuration: Enabling Internet Access for Private Subnets](#). See also [Tutorial: Automatically Set Up a NAT Instance in Oracle Cloud Infrastructure with Terraform](#).

- Set up the applicable network topology, including creating the VCN, internet gateway, public and private subnets, and required security lists and route rules.
- Provision a NAT instance in the public subnet, and a private instance in the private subnet.
- Enable outbound internet access for the private instance through the NAT instance on the public subnet.

[Go to the sample on Oracle GitHub.](#)

*ENABLE INTERNET ACCESS FROM A COMPUTE INSTANCE USING THE ORACLE CLOUD INFRASTRUCTURE NAT GATEWAY*

This sample is similar to the previous sample except that while the previous sample configures a compute instance to operate as a NAT Gateway, the present sample employs the Oracle Cloud Infrastructure NAT Gateway service.



### Note

For more information about the Oracle Cloud Infrastructure NAT Gateway service, see [NAT Gateway](#). For a blog post discussing how to use the Oracle Cloud Infrastructure NAT Gateway, see [Access Resources on the Public Internet Through an Oracle Cloud Infrastructure NAT Gateway](#).

The sample show how to complete the following:

- Set up the VCN, the NAT Gateway, the Internet Gateway, the public and private subnets, and the necessary security lists and route rules.
- Provision a bastion instance in the public subnet and a private instance in the private subnet.

Once set up, the private instance will have outbound Internet access through the Oracle Cloud Infrastructure NAT Gateway, and will be accessible using SSH from the bastion instance.

[Go to the sample on Oracle GitHub](#)

### *CREATE AN INSTANCE POOL*

This sample shows how to manage your Compute instances using resources such as instance configurations and instance pools that are provided using Oracle Cloud Infrastructure Ansible modules. Instance pools help you create and provision multiple Compute instances within the same region based on a single instance configuration.

The sample illustrated completing the following tasks:

- Generate a temporary, host-specific SSH key pair.
- From the SSH key pair, specify the public key for connecting to the instance during launch.
- Create an instance configuration that defines settings for creating a Compute instance as part of the instance pool. The configuration provides details such as base image, shape, and metadata.
- Demonstrates how the Compute instances based on the instance configuration can be launched using instance pools.
- Connect to one of the Compute instances using SSH.

[Go to the sample on Oracle GitHub](#)

### *CREATE INSTANCE CONSOLE CONNECTIONS AND CAPTURE CONSOLE HISTORY*

This sample shows you can create VNC and serial console connections to a Compute instance, and how you can fetch and capture the serial Console data from the instance. For more

information about Console connections, see [Instance Console Connections](#).

This sample illustrates completing the following tasks:

- Generate a temporary SSH key pair for the serial Console connection.
- Create an instance Console connection for a Compute instance.
- Capture serial Console data for a Compute instance, and then save the data to a local machine so you can troubleshoot and debug issues.

[Go to the sample on Oracle GitHub](#)

### *ACCESS OBJECT STORAGE FROM A PRIVATE INSTANCE USING SERVICE GATEWAY*

This sample playbook shows how you can enable private access to an Object Storage from a Compute instance using a service gateway.



#### **Note**

For more information about service gateways, see [Access to Oracle Services: Service Gateway](#). To read a blog post discussing how to connect Compute instances using the service gateway, see [Connect Private Instances with Oracle Services Through an Oracle Cloud Infrastructure Service Gateway](#).

The sample shows how to complete the following tasks:

- Set up a user, group, and policies required for managing buckets.
- Create and upload the required API keys to the user.
- Set up the VCN, the NAT gateway, the Internet gateway, the public and private subnets, as well as the required security lists and route tables. Note that a bastion instance is provisioned in the public subnet, and a private instance is provisioned in the private subnet.
- Provision a Compute instance in the private subnet,

## CHAPTER 34 Developer Tools

---

- Install the Oracle Cloud Infrastructure command line interface (CLI) and configure the CLI using the cloud init script.
- Disable the NAT gateway to restrict public access to the private instance.
- Create a bucket from the private instance using the Oracle Cloud Infrastructure CLI, then verify that the bucket is created.

Following this setup, the private instance will have private access to Object Storage.

[Go to the sample on Oracle GitHub](#)

### CONTAINER ENGINE FOR KUBERNETES

#### *CREATE A CLUSTER USING CONTAINER ENGINE FOR KUBERNETES*

This sample playbook uses Container Engine for Kubernetes (OKE) to create a cluster and deploys a sample application on the cluster. This sample complements an existing example, [Creating a Cluster with Oracle Cloud Infrastructure Container Engine for Kubernetes](#).

This sample illustrates how to do the following:

- Creates and configures a VCN and related resources required for setting up an OKE cluster.
- Creates a cluster.
- Creates a node pool.
- Downloads the kubeconfig file for the cluster.
- Deploys a sample application on the cluster.
- Verifies a successful deployment.

[Go to the sample on Oracle GitHub](#).

### DATABASE

#### *BARE METAL/VM DATABASE PROVISIONING*

This sample playbook shows how to retrieve the public and private IP addresses of a database system node so that you can access it through an Ansible module. The sample illustrates how

to do the following:

- Collect database node VNIC information for a specified database.
- Extract public and private IP addresses of the database node from the VNIC.

[Go to the sample on Oracle GitHub.](#)

### *AUTONOMOUS DATA WAREHOUSE*

This sample playbook shows how to create an Autonomous Data Warehouse and manage its lifecycle. The sample shows how to do the following:

- Set up an Autonomous Data Warehouse.
- List all of the Autonomous Data Warehouse instances available in a compartment, filtered by display name.
- Get the "facts" for a specified Autonomous Data Warehouse.
- Stop and start an Autonomous Data Warehouse instance.
- Delete an Autonomous Data Warehouse instance.

[Go to the sample on Oracle GitHub.](#)

### *AUTONOMOUS TRANSACTION PROCESSING*

This sample playbook shows how to create an Autonomous Transaction Processing database and manage its lifecycle. The sample shows how to do the following:

- Set up an Autonomous Transaction Processing database instance.
- List all of the Autonomous Transaction Processing instances in a compartment, filtered by display name.
- Get the "facts" for a specified Autonomous Transaction Processing instance.
- Delete an Autonomous Transaction Processing database instance.

[Go to the sample on Oracle GitHub.](#)

### FILE STORAGE

#### *CREATE AND MOUNT A FILE SYSTEM*

The sample shows how to create a file system that you can access through an Oracle Cloud Infrastructure Compute instance using Ansible cloud modules. The sample illustrates completing the following tasks:

- Creates network dependencies like VCNs, subnets, and so forth, as well as a security list that is configured as required by the File Systems service.
- Generates the certificates required by the Compute instances.
- Demonstrates how to create File Storage service components, such as mount targets, file systems, exports, and snapshots.
- Demonstrates how to mount the file system using a Compute instance, and how to then access the file system content from a different Compute instance.

[Go to the sample on Oracle GitHub](#)

#### *MULTIPLE FILE SYSTEMS WITH MOUNT TARGETS*

This sample shows how one you can export one file system using two different export paths located on two different mount targets. It also shows how a single mount target can export paths from two different file systems. The sample illustrates completing the following tasks:

- Creates network dependencies like VCNs, subnets, and so forth, as well as a security list that is configured as required by the File Systems service.
- Generates the certificates required by the Compute instances.
- Demonstrates how to create File Storage service components, such as mount targets, file systems, exports, and snapshots.
- Demonstrates how one file system can be exported onto two different mount targets.
- Demonstrates how a single mount target can export paths from two different file systems.
- Demonstrates how to mount the file system using an Oracle Cloud Infrastructure Compute instance.

[Go to the sample on Oracle GitHub](#)

### IAM

*USE ANSIBLE MODULES TO PERFORM IAM TASKS*

This sample shows how to perform basic identity and access management (IAM) tasks using Ansible modules. The sample also shows how to execute an Ansible playbook or execute individual tasks as a different user. The sample illustrates how to do the following:

- Create groups (**ObjectReaders** and **ObjectWriters**).
- Create an IAM policy that enables the following:
  - **ObjectReaders** to list and read buckets and objects.
  - **ObjectWriters** to create, update, list, and read buckets and objects in a specified compartment.
  - Assign the policy to groups.
- Create users (**alice** and **bob**) and then do the following:
  - Add **alice** to the **ObjectWriters** group.
  - Add **bob** to the **ObjectReaders** group.
  - Run as **alice** to create a bucket, then upload objects to the bucket.
  - Run as **bob** to list all objects in a bucket.

[Go to the sample on Oracle GitHub.](#)

### LOAD BALANCING

This sample playbook shows how to create a public load balancer using an Ansible module. The sample illustrates the following:

- Generating network-related artifacts, such as subnets and VCNs, for example.
- Generating the required certificates for the load balancer.
- Using an Ansible playbook to create a public load balancer.

[Go to the sample on Oracle GitHub.](#)

## CHAPTER 34 Developer Tools

---

### OBJECT STORAGE

#### *LIST OBJECTS AND BUCKETS*

This sample playbook shows how to list all objects and buckets in a namespace.

[Go to the sample on Oracle GitHub.](#)

#### *DELETE OBJECTS*

This sample playbook shows how to delete objects created within a specified range of days from the specified buckets. You can also modify the sample so it deletes objects older than a specified number of days, which helps you prune old or unwanted objects that are stored in the service.

[Go to the sample on Oracle GitHub.](#)

### MONGODB DEPLOYMENT

#### *USE ANSIBLE MODULES TO DEPLOY A MONGODB DATABASE*

This sample playbook shows how to deploy a MongoDB database in the securely in the cloud using Ansible modules. The sample implements security measures using the Castle strategy ("defense in depth"), which is discussed in the article [Secure MongoDB on Oracle Bare Metal Cloud Services](#).

[Go to the sample on Oracle GitHub.](#)

### NETWORKING

#### *USE ANSIBLE MODULES TO PROVISION A VCN*

This sample playbook shows how to provision a virtual cloud network (VCN) with two private subnets in different availability domains, and an IPSec VPN. The sample provisions infrastructure resources that are illustrated in the knowledge base article [Scenario B: Private Subnet with a VPN](#). The sample provisions the following resources:

- A VCN.
- Two private subnets.

- A dynamic routing gateway (DRG).
- Customer premises equipment (CPE).
- An IPSec connection between the DRG and the CPE, and retrieves IPSec configuration information and status.

[Go to the sample on Oracle GitHub.](#)

### Chef Knife Plug-in

This topic provides information about installing, configuring, and using the Chef Knife Plug-in for Oracle Cloud Infrastructure.

- **Licensing:** This provider and sample is licensed under the Mozilla Public License 2.0; third-party content is separately licensed as described in the code.
- **Download:** [GitHub](#)
- **Documentation:** [README](#)

### Contributions

Got a fix for a bug, or a new feature you'd like to contribute? The Chef Knife Plug-in for Oracle Cloud Infrastructure is open source and accepting pull requests on [GitHub](#).

### Notifications

To be notified when a new version of the Chef Knife Plug-in for Oracle Cloud Infrastructure is released, [subscribe to the Atom feed](#).

### Questions or Feedback

- [GitHub](#): To file bugs and feature requests only.
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) tag in your post.

- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Compute Jenkins Plug-in

This topic provides information about installing, configuring, and using the Compute Jenkins plug-in for Oracle Cloud Infrastructure services.

- **Licensing:** The Oracle Cloud Infrastructure Compute Jenkins plug-in is dual-licensed under the Universal Permissive License (UPL) and the Apache License 2.0; third-party content is separately licensed as described in the code.
- **Download:** [GitHub](#)
- **Documentation:** [Oracle Cloud Infrastructure Compute Plug-in - Jenkins Wiki](#)

### Contributions

Got a fix for a bug, or a new feature you'd like to contribute? The Oracle Cloud Infrastructure Compute Jenkins plug-in is open source and accepting pull requests on [GitHub](#).

### Notifications

To be notified when a new version of the Oracle Cloud Infrastructure Compute Jenkins plug-in is released, [subscribe to the Atom feed](#).

### Questions or Feedback

- [GitHub](#): To file bugs and feature requests only.
- [Stack Overflow](#): Please use the [oracle-cloud-infrastructure](#) tag in your post.
- [Developer Tools section](#) of the Oracle Cloud forums
- [My Oracle Support](#)

### Grafana Plug-in

This topic provides instructions for installing, configuring, and using Oracle Cloud Infrastructure Data Source for Grafana, otherwise referenced as the Grafana Plug-in.

#### Grafana Plug-in Overview

Grafana is an open-source visualization and alerting tool that you can use for analytics and monitoring of time-series data (metrics). While metrics from [Oracle Cloud Infrastructure Monitoring](#) are visible in metrics charts through the Console, you can use Oracle Cloud Infrastructure Data Source for Grafana ("the Grafana Plug-in") to view metrics from resources across providers on a single Grafana dashboard.

#### Prerequisites for Using the Grafana Plug-in

- An Oracle Cloud Infrastructure account.
- A user in that account, in a security group with an IAM policy that grants necessary permissions for working with resources in the account compartments. The policy must give you access to the metric namespaces emitting metrics (such as Compute) as well as the related resources (such as a set of Compute instances). If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. Administrators: For common policies that give groups access to metrics, see [Let users view metric definitions in a compartment](#) and [Restrict user access to a specific metric namespace](#). To authorize resources, such as instances, to make API calls, add the resources to a [dynamic group](#) through its matching rules, and then create a policy that allows that dynamic group access to metrics. To allow access across compartments, create the policy in the tenancy. Because of the concept of policy inheritance, instances in the indicated dynamic group can then access metrics in any compartment. To reduce the scope of access to a particular compartment, specify that compartment instead of the tenancy. See [Let instances make API calls to access monitoring metrics in the tenancy](#).

- Compute instances: To emit metrics, Compute instances must be monitoring-enabled. OracleCloudAgent software installation may also be required. For more information, see [Enabling Monitoring for Compute Instances](#).
- Required keys and Oracle Cloud Infrastructure IDs (OCIDs). For guidance, see "Required Keys and OCIDs" in the *Oracle Cloud Infrastructure User Guide*.

### Download and Install the Grafana Plug-in

This section provides instructions for downloading and installing the Grafana Plug-in, either locally (Linux or Mac) or by Terraform script. Use Terraform script to provision a Compute instance that runs Grafana in Oracle Cloud Infrastructure.

Authentication for metric access depends on where Grafana is running. If you are running Grafana on a local machine outside Oracle Cloud Infrastructure, you must call the Monitoring API using the [Command Line Interface \(CLI\)](#). If you are running Grafana on the new Compute instance created by the Terraform script, use the script-generated dynamic group. If you are running Grafana on an Oracle Cloud Infrastructure Compute instance that you created, then you can add the instance to a [dynamic group](#) by configuring a matching rule. In all scenarios, you have the option of calling the API using the [CLI](#).

#### LOCAL: DOWNLOAD AND INSTALL THE PLUG-IN

### Local Installation Prerequisites

[Grafana installation](#). Version 3.0 or later required.

To install the Grafana Plug-in (Oracle Cloud Infrastructure Data Source for Grafana), see <https://grafana.com/plugins/oci-datasource/installation>. After installation, you're ready to configure the Grafana Plug-in.

#### TERRAFORM SCRIPT: CREATE A GRAFANA ENVIRONMENT

This section describes how to use Terraform scripts to create a Grafana environment on your virtual machine.

The Terraform scripts create a dynamic group using the name you specify, configure an IAM policy named "grafana\_policy," provision a Compute instance named "TFInstance0" (in

the compartment and VCN that you specify), and download and install both Grafana and the Grafana Plug-in onto the instance.

### Terraform Script Prerequisites

Before running the Terraform scripts, make sure you have the following:

- Terraform and Oracle Cloud Infrastructure Terraform Provider must both be installed on your development machine. See [Getting Started with the Terraform Provider](#).
- A Virtual Cloud Network (VCN) with [access to the Internet](#). For reference, see the [Linux Tutorial](#) for creating a VCN.
- Values for the following variables. You can pass the values to your `.bashrc` or `.bash_` profile, or note them somewhere for entering when needed. These variables are required for creating the Grafana environment and using the Grafana Plug-in.

#### Required variables

```
variable "tenancy_ocid" {}
variable "user_ocid" {}
variable "fingerprint" {}
variable "private_key_path" {}
variable "region" {}
variable "compartment_ocid" {}
variable "ssh_public_key" {}
variable "ssh_private_key" {}
variable "subnet_id" {}
variable "availability_domain" {}
variable "dynamic_group_name" {}
```



### Note

For help on getting the tenancy OCID and other values needed for these variables, see [Required Keys and OCIDs](#).

## To download and run the Terraform scripts

Make sure you satisfy the [prerequisites](#) before running the scripts.

1. Download the Terraform scripts using the following command.  

```
wget https://objectstorage.us-ashburn-1.oraclecloud.com/n/oracle-cloudnative/b/GrafanaTerraform/o/terraform_grafana.tar && tar -xvf terraform_grafana.tar
```
2. Change to the `terraform_grafana` directory using the following command.  

```
cd /terraform_grafana
```
3. Initialize a new Terraform configuration using the following command.  

```
terraform init
```

Initialization creates the files `terraform.tfstate` and `terraform.tfstate.backup`.
4. Generate an execution plan using the following command.  

```
terraform plan
```

Generating a plan can help you validate your Terraform script before you apply it to your environment.
5. Run the scripts using the following command, using the same availability domain and dynamic group variables passed in with `terraform plan`.  

```
terraform apply
```

On completion, the "Apply complete!" message appears. You're now ready to configure the Grafana Plug-in.

### Configure the Grafana Plug-in

This section describes how to add the Grafana Plug-in and set up a dashboard.

### To configure the Grafana Plug-in

#### Local installation



#### Note

When installed locally, the Grafana plug-in accesses metrics using calls to the Monitoring API.

1. Set up the [CLI](#) for accessing Oracle Cloud Infrastructure APIs. You'll need to access the Monitoring API for authentication.
2. Navigate to the Grafana homepage at the following URL.  
<http://localhost:3000>
3. Add the Grafana Plug-In (Oracle Cloud Infrastructure Data Source for Grafana). In addition to the steps below, see the Grafana instructions for adding data sources at [http://docs.grafana.org/guides/getting\\_started/](http://docs.grafana.org/guides/getting_started/).
  - a. In Grafana, on the Home Dashboard, click the gear icon on the left.
  - b. Click **Add data source**.
  - c. In the Filter text box, type: `oracle-oci-datasource`
  - d. In the filtered list, select **oracle-oci-datasource**.
  - e. In the Settings page, fill in your **Tenancy OCID**, **Default Region**, and **Environment**. For **Environment** choose **local**.

#### Environment options

**local** is for Grafana deployments outside Oracle Cloud Infrastructure.

**OCI Instance** is for Grafana deployments on Oracle Cloud Infrastructure resources.

The data source is now added, enabling you to set up a dashboard showing metrics from Oracle Cloud Infrastructure.

4. Set up a [dashboard](#) of the type "Graph."
  5. (Optional) To confirm access to metrics from Oracle Cloud Infrastructure, update your dashboard query ("Metrics") with a specific region, compartment, namespace, metric. You can also add one or more dimensions.
- Congratulations. You can now view your Oracle Cloud Infrastructure metrics in Grafana!

### Terraform environment



#### Note

When installed by Terraform script, the Grafana plug-in accesses metrics using the script-generated dynamic group.

1. Connect to the new instance "TFInstance0."  
For instructions, see [Connecting to an Instance](#).
2. Navigate to the Grafana homepage by running the following command, where *[Instance Public IP]* is the IP address of your new instance "TFInstance0."  

```
ssh opc@[Instance Public IP] -L 3000:localhost:3000
```



### Note

You can get the address from the list of instances in the Console. Click **Compute**, choose your **Compartment**, and then find your instance in the list. Alternatively, you can use the Core Services API [ListVnicAttachments](#) and [GetVnic](#) operations.

3. Add the Grafana Plug-In (Oracle Cloud Infrastructure Data Source for Grafana). In addition to the steps below, see the Grafana instructions for adding data sources at [http://docs.grafana.org/guides/getting\\_started/](http://docs.grafana.org/guides/getting_started/).
  - a. In Grafana, on the Home Dashboard, click the gear icon on the left.
  - b. Click **Add data source**.
  - c. In the Filter text box, type: `oracle-oci-datasource`
  - d. In the filtered list, select **oracle-oci-datasource**.
  - e. In the Settings page, fill in your **Tenancy OCID**, **Default Region**, and **Environment**. For **Environment** choose **OCI Instance**.

### Environment options

**local** is for Grafana deployments outside Oracle Cloud Infrastructure.

**OCI Instance** is for Grafana deployments on Oracle Cloud Infrastructure resources.

The data source is now added, enabling you to set up a dashboard showing metrics from Oracle Cloud Infrastructure.

4. Set up a [dashboard](#) of the type "Graph."
5. (Optional) To confirm access to metrics from Oracle Cloud Infrastructure, update your dashboard query ("Metrics") with a specific region, compartment, namespace, metric.

You can also add one or more dimensions.

Congratulations. You can now view your Oracle Cloud Infrastructure metrics in Grafana!

### Troubleshoot the Plug-In

If the dashboard query ("Metrics") fails to populate with options, or if you have other issues accessing metrics, then the IAM policy used to access metrics may be malformed, or it may not include all required matching rules for your dynamic group (Terraform script).

To resolve this issue, do the following.

- Review your IAM policy to ensure that it matches prerequisites.
- For Terraform, consider adding the following matching rule to your dynamic group:  
`matching_rule = "ANY {instance.compartment.id = '${var.compartment_ocid}'}"`
- If you updated your IAM policy, then restart the Grafana server and refresh the Grafana homepage.

### Terraform: Remove the environment

If you ran the Terraform script to create a Grafana environment on your virtual machine, you may want to remove the environment after your work is done. Follow the instructions in this section to do so.

#### To remove your Grafana environment

Run the following command, using the same availability domain and dynamic group variables passed in with `terraform apply`.

```
terraform destroy
```

## Tools Configuration

This general reference shows how to configure the SDKs and other developer tools to integrate with Oracle Cloud Infrastructure services.

- [Required Keys and OCIDs](#) - Details on identity and access management
- [SDK and CLI Configuration File](#) - Methods for providing configuration information when using the SDKs or CLI.

### Required Keys and OCIDs

Whether you're using an Oracle client (see [Software Development Kits and Command Line Interface](#)) or a client you built yourself, you need to do the following:

1. Create a user in IAM for the person or system who will be calling the API, and put that user in at least one IAM group with any desired permissions. See "Adding Users" in the *Oracle Cloud Infrastructure Getting Started Guide*. You can skip this if the user exists already.
2. Get these items:
  - RSA key pair **in PEM format** (minimum 2048 bits). See [How to Generate an API Signing Key](#).
  - Fingerprint of the public key. See [How to Get the Key's Fingerprint](#).
  - Tenancy's OCID and user's OCID. See [Where to Get the Tenancy's OCID and User's OCID](#).
3. Upload the public key from the key pair in the Console. See [How to Upload the Public Key](#).
4. If you're using one of the Oracle SDKs or tools, supply the required credentials listed above in either a configuration file or a config object in the code. See [SDK and CLI Configuration File](#). If you're instead building your own client, see [Request Signatures](#).



### Important

This key pair is **not** the SSH key that you use to access compute instances. See [Security Credentials](#).

Both the private key and public key must be in PEM format (not SSH-RSA format). The public key in PEM format looks something like this:

```
-----BEGIN PUBLIC KEY-----
MIIBIjANBgkqhkiG9w0BAQE...
...
-----END PUBLIC KEY-----
```

### How to Generate an API Signing Key

You can use the following [OpenSSL](#) commands to generate the key pair in the required PEM format. If you're using Windows, you'll need to install [Git Bash for Windows](#) and run the commands with that tool.

1. If you haven't already, create a `.oci` directory to store the credentials:

```
mkdir ~/.oci
```

2. Generate the private key with one of the following commands.

- Recommended: To generate the key, encrypted with a passphrase you provide when prompted:

```
openssl genrsa -out ~/.oci/oci_api_key.pem -aes128 2048
```

**Note:** For Windows, you may need to insert `-passout stdin` to be prompted for a passphrase. The prompt will just be the blinking cursor, with no text.

```
openssl genrsa -out ~/.oci/oci_api_key.pem -aes128 -passout stdin 2048
```

- To generate the key with no passphrase:

```
openssl genrsa -out ~/.oci/oci_api_key.pem 2048
```

## CHAPTER 34 Developer Tools

---

3. Ensure that only you can read the private key file:

```
chmod go-rwx ~/.oci/oci_api_key.pem
```

4. Generate the public key:

```
openssl rsa -pubout -in ~/.oci/oci_api_key.pem -out ~/.oci/oci_api_key_public.pem
```

**Note:** For Windows, if you generated the private key with a passphrase, you may need to insert `-passin stdin` to be prompted for the passphrase. The prompt will just be the blinking cursor, with no text.

```
openssl rsa -pubout -in ~/.oci/oci_api_key.pem -out ~/.oci/oci_api_key_public.pem -passin stdin
```

5. Copy the contents of the public key to the clipboard using `pbcopy`, `xclip` or a similar tool (you'll need to paste the value into the Console later). For example:

```
cat ~/.oci/oci_api_key_public.pem | pbcopy
```

Your API requests will be signed with your private key, and Oracle will use the public key to verify the authenticity of the request. You must upload the public key to IAM (instructions below).

### How to Get the Key's Fingerprint

You can get the key's fingerprint with the following OpenSSL command. If you're using Windows, you'll need to install [Git Bash for Windows](#) and run the command with that tool.

```
openssl rsa -pubout -outform DER -in ~/.oci/oci_api_key.pem | openssl md5 -c
```

When you upload the public key in the Console, the fingerprint is also automatically displayed there. It looks something like this: 12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:cd:ef

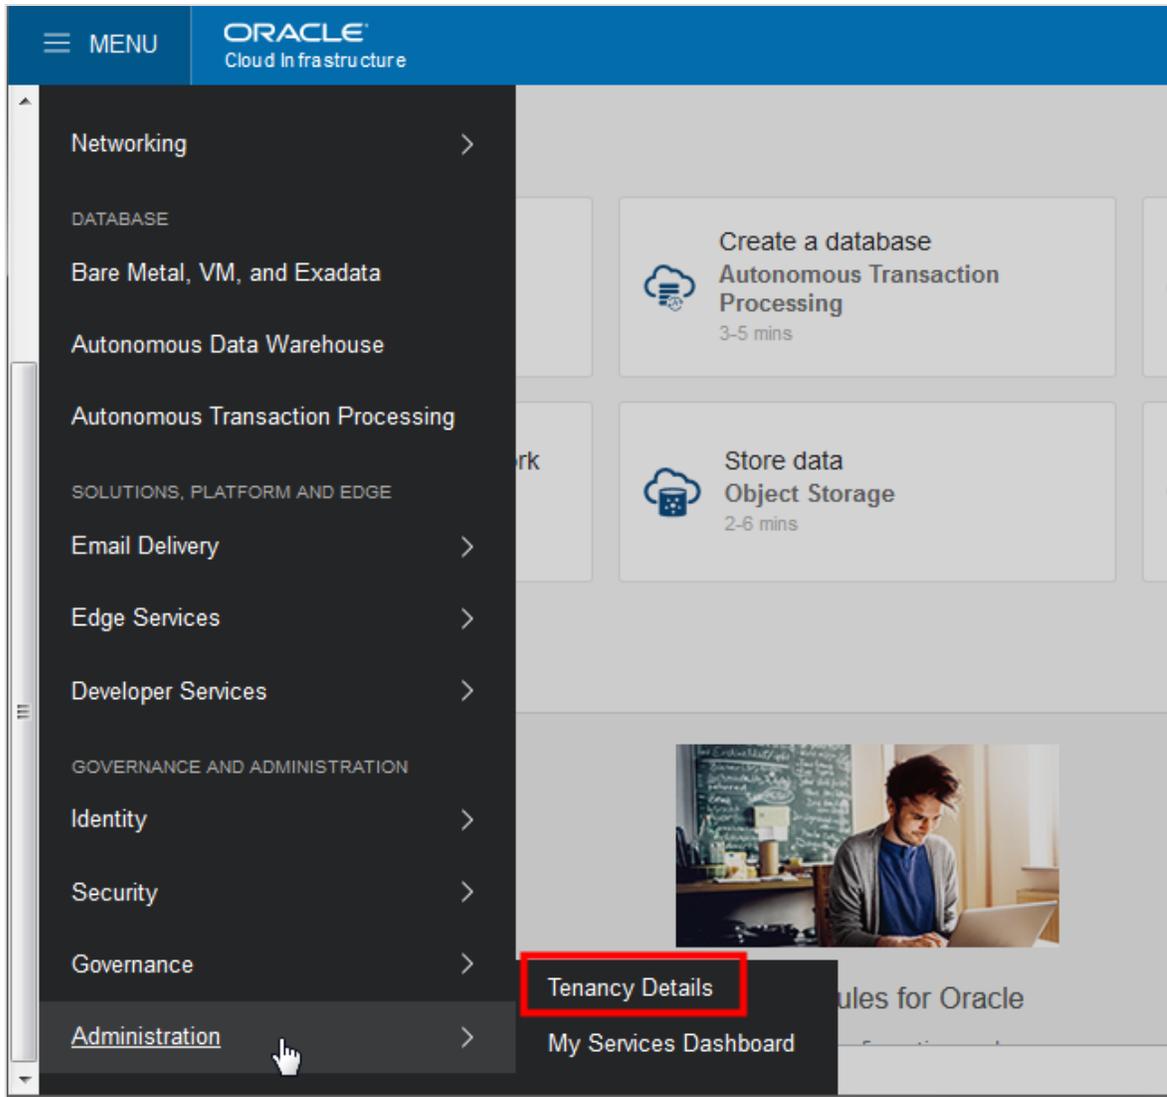
### Where to Get the Tenancy's OCID and User's OCID

Both OCIDs are in the Console, which is located at <https://console.us-ashburn-1.oraclecloud.com>. If you don't have a login and password for the Console, contact an administrator. If you're not familiar with OCIDs, see [Resource Identifiers](#).

### TENANCY'S OCID

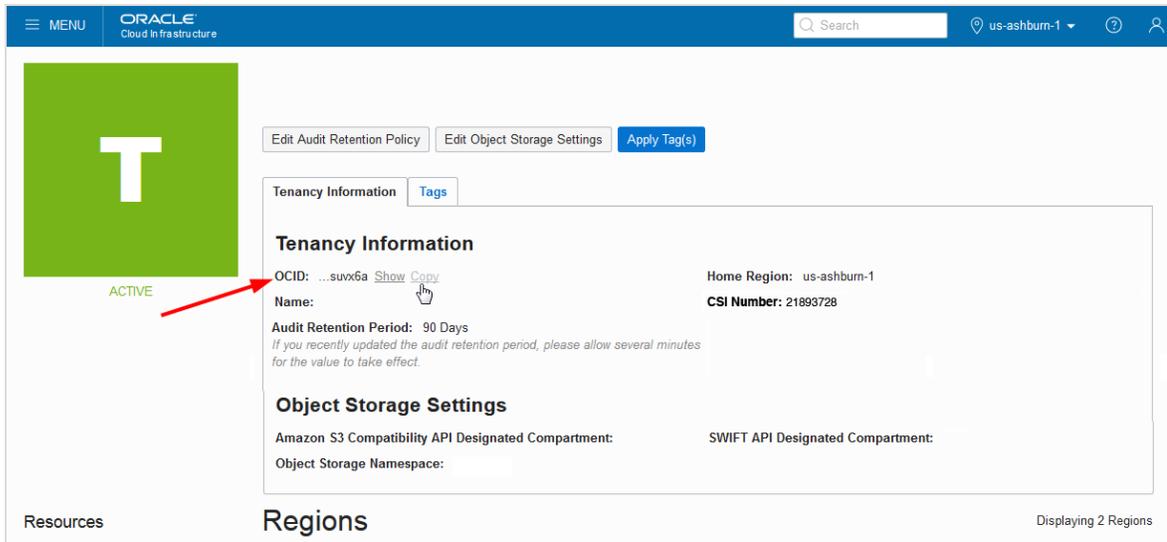
Get the tenancy OCID from the Oracle Cloud Infrastructure Console on the **Tenancy Details** page:

1. Open the navigation menu, under Governance and Administration, go to **Administration** and click **Tenancy Details**.



2. The tenancy OCID is shown under **Tenancy Information**. Click **Copy** to copy it to your clipboard.

## CHAPTER 34 Developer Tools



### USER'S OCID

Get the user's OCID in the Console on the page showing the user's details. To get to that page:

- If you're signed in as the user: Open the **Profile** menu (👤) and click **User Settings**.
- If you're an administrator doing this for another user: Open the navigation menu. Under **Governance and Administration**, go to **Identity** and click **Users**. Select the user from the list.

### How to Upload the Public Key

You can upload the PEM public key in the Console, located at <https://console.us-ashburn-1.oraclecloud.com>. If you don't have a login and password for the Console, contact an administrator.

1. Open the Console, and sign in.
2. View the details for the user who will be calling the API with the key pair:
  - If you're signed in as this user, click your username in the top-right corner of the Console, and then click **User Settings**.

- If you're an administrator doing this for another user, instead click **Identity**, click **Users**, and then select the user from the list.
3. Click **Add Public Key**.
  4. Paste the contents of the PEM public key in the dialog box and click **Add**.

The key's fingerprint is displayed (for example, 12:34:56:78:90:ab:cd:ef:12:34:56:78:90:ab:cd:ef).

Notice that after you've uploaded your first public key, you can also use the [UploadApiKey](#) API operation to upload additional keys. You can have up to three API key pairs per user. In an API request, you specify the key's fingerprint to indicate which key you're using to sign the request.

### SDK and CLI Configuration File

Oracle Cloud Infrastructure SDKs and CLI require basic configuration information, like user credentials and tenancy OCID. You can provide this information by:

- Using a configuration file
- Declaring a configuration at runtime

The SDKs fully support both options. Refer to the documentation for each SDK for information about the config object and any exceptions when using a configuration file:

- [SDK for Java Configuration](#)
- [Python SDK Configuration](#)
- [Ruby SDK Configuration](#)
- [Go SDK Configuration](#)

The CLI requires a configuration file. `--region` is the only configuration value that can be passed as a parameter for a CLI operation.

#### File Name and Location

The default configuration file name and location is `~/.oci/config`.

**Note**

On Windows, you can use PowerShell to create the folder with the following command: `mkdir ~/.oci`. File Explorer does not support creating folder names that start with a period.

**File Entries**

The following table lists the basic entries that are required for the configuration file, as well as where to get the required information.

Entry	Description and Where to Get the Value	Required?
<code>user</code>	OCID of the user calling the API. To get the value, see <a href="#">Required Keys and OCIDs</a> .  Example: <code>ocidl.user.oc1..aaaaaaaa65vwl75tewwm32rgqvm6i34unq</code> (shortened for brevity)	Yes
<code>fingerprint</code>	Fingerprint for the key pair being used. To get the value, see <a href="#">Required Keys and OCIDs</a> .  Example: <code>20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34</code>	Yes

## CHAPTER 34 Developer Tools

Entry	Description and Where to Get the Value	Required?
key_file	<p>Full path and filename of the private key.</p> <p><b>Important:</b> The key pair must be in PEM format. For instructions on generating a key pair in PEM format, see <a href="#">Required Keys and OCIDs</a>.</p> <p>If you encrypted the key with a passphrase, you must also include the <code>pass_phrase</code> entry in the config file.</p> <p>Example: <code>~/.oci/oci_api_key.pem</code></p>	Yes
pass_phrase	<p>Passphrase used for the key, if it is encrypted.</p> <p>Example: <code>examplephrase</code></p>	If key is encrypted
tenancy	<p>OCID of your tenancy. To get the value, see <a href="#">Required Keys and OCIDs</a>.</p> <p>Example: <code>ocid1.tenancy.oc1..aaaaaaaaba3pv6wuzr4h25vqstifsfdsq</code> (shortened for brevity)</p>	Yes
region	<p>An Oracle Cloud Infrastructure region. See <a href="#">Regions and Availability Domains</a>.</p> <p>Example: <code>us-ashburn-1</code></p>	Yes

### CUSTOM VALUES

Some Oracle Cloud Infrastructure SDKs support defining custom values in the configuration file. Refer to the documentation for each SDK for more information.



### Warning

Avoid saving confidential information in the configuration file.

### PROFILES AND INHERITANCE

You can create multiple profiles with different values for these entries, then you can specify which profile to load.

Some Oracle Cloud Infrastructure SDKs require a DEFAULT profile and support profile inheritance. This means that any value that isn't explicitly defined for a given profile is inherited from the DEFAULT profile. Refer to the documentation for each SDK for more information.

### Example Configuration

The following example shows key values in a configuration file and how to set profiles for a SDK that supports profile inheritance.

```
[DEFAULT]
user=ocidl.user.oc1..aaaaaaat5nvwcnA5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjg
fingerprint=20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34
key_file=~/.oci/oci_api_key.pem
tenancy=ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq
region=us-ashburn-1

[ADMIN_USER]
user=ocidl.user.oc1..aaaaaaa65vw17zut55hiavppn4nbfwyccuecuch5tewwm32rgqvm6i34unq
fingerprint=72:00:22:7f:d3:8b:47:a4:58:05:b8:95:84:31:dd:0e
key_file=keys/admin_key.pem
pass_phrase=mysecretphrase
```

## REST APIs

The Oracle Cloud Infrastructure APIs are typical REST APIs that use HTTPS requests and responses. This topic describes basic information about using the APIs.



### Warning

Oracle recommends that you avoid using string values that include confidential information in the Oracle Cloud Infrastructure API.

## API Reference and Endpoints

For links to the Oracle Cloud Infrastructure API reference and a list of the regional API endpoints, see [API Reference and Endpoints](#).

## API Version

The base path of the endpoint includes the desired API version (for example, 20160918). Here's an example for a POST request to create a new VCN in the Ashburn region:

```
POST https://iaas.us-ashburn-1.oraclecloud.com/20160918/vcns
```

## Request Signing Required

All Oracle Cloud Infrastructure API requests must be signed for authentication purposes. For information about the required credentials and how to sign the requests, see [Request Signatures](#).

## HTTPS and TLS 1.2 Required

All Oracle Cloud Infrastructure API requests must support HTTPS and SSL protocol TLS 1.2.

## Maximum Allowed Client Clock Skew

HTTP status code 401 (NotAuthenticated) is returned if the client's clock is skewed more than 5 minutes from the server's. To determine the server's clock time, use this curl command

## CHAPTER 34 Developer Tools

---

with the API endpoint:

```
curl -s --head <endpoint> | grep Date
```

For example:

```
curl -s --head https://iaas.us-phoenix-1.oraclecloud.com | grep Date
```

### Request and Response Format

The Oracle Cloud Infrastructure APIs use standard HTTP requests and responses. Each may contain Oracle-specific headers for pagination, entity tags (ETags), and so on as described elsewhere in this topic and in the API documentation.

Each response includes a unique Oracle-assigned request ID (for example, bb3f3275-f356-462a-93c4-bf40fb82bb02) in the `opc-request-id` response header. If you need to contact Oracle about a particular request, please provide this request ID.

Many of the API operations require JSON in the request body or return JSON in the response body. The specific contents of the JSON are described in the API documentation for the individual operation. Notice that the JSON is not wrapped or labeled according to the operation's name or the object's name or type.



#### Note

Make sure to set the `Content-Type` header to `application/json` in your POST and PUT requests that contain JSON in the body.

### Example CreateVcn Request

```
POST https://iaas.us-phoenix-1.oraclecloud.com/20160918/vcns
host: iaas.us-phoenix-1.oraclecloud.com
opc-retry-token: 239787fs987
Content-Type: application/json
HTTP headers required for authentication
Other HTTP request headers per the HTTP spec
```

## CHAPTER 34 Developer Tools

---

```
{
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaauwjnv47knr7uuuvqar5bshnspi6xoxsfebh3vy72fi4swgrkvuvq",
 "displayName": "Apex Virtual Cloud Network",
 "cidrBlock": "172.16.0.0/16"
}
```

### Example CreateVcn Response

200 OK

opc-request-id: 6c4d01a6-f764-4325-a3f8-720c8b5cae7b

```
{
 "id": "ocidl.vcn.oc1.phx.aaaaaaa4ex5pqjtkjhdb4h4gcnko7vx5uto5puj5noa5awznsqpwt3pqyq",
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaauwjnv47knr7uuuvqar5bshnspi6xoxsfebh3vy72fi4swgrkvuvq",
 "displayName": "Apex Virtual Cloud Network",
 "cidrBlock": "172.16.0.0/16"
 "defaultRouteTableId":
"ocidl.routetable.oc1.phx.aaaaaaaaba3pv6wkr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfsdq",
 "defaultSecurityListId":
"ocidl.securitylist.oc1.phx.aaaaaaaac6h4ckr3ncbxmvwinfvzxb7owu5hfzbvtu33kfe7hgscs5fjaq"
 "defaultDhcpOptionsId":
"ocidl.dhcpoptions.oc1.phx.aaaaaaaawglzn7s5sogyfznl25a4vxgu76c2hrgvzcd3psn6vcx331zmu2xa"
 "state": "PROVISIONING",
 "timeCreated": "2016-07-22T17:43:01.389+0000"
}
```

### Error Format

If a request results in an error, the response contains a standard HTTP response code with 4xx for client errors and 5xx for server errors. The body also includes JSON with an error code and a description of the error. For example:

```
{
 "code": "InvalidParameter",
 "message": "Description may not be empty; description size must be between 1 and 400"
}
```

For a list of common errors across all services, see [API Errors](#).

### Request Throttling

Oracle Cloud Infrastructure applies throttling to many API requests to prevent accidental or abusive use of resources. If you make too many requests too quickly, you might see some succeed and others fail. Oracle recommends that you implement an exponential back-off, starting from a few seconds to a maximum of 60 seconds. When a request fails due to throttling, the system returns response code 429 and the following error code and description:

```
{
 "code": "TooManyRequests",
 "message": "User-rate limit exceeded."
}
```

### Polling for Resource Status

Most Oracle Cloud Infrastructure resources, such as compute instances, have lifecycles. In many cases, you want your code to wait until a resource or work request reaches a specific state, or a timeout is exceeded, before taking further action.

You can poll a resource to determine its state. For example, when you call [GetInstance](#), the response body contains an [instance resource](#) that includes the `lifecycleState` attribute. You might want your code to wait until the instance's `lifecycleState` is `RUNNING` before proceeding.

Different resources take different amounts of time to transition between states. Therefore, the optimal frequency and duration parameters for a polling strategy can vary among resources. The Oracle Cloud Infrastructure SDK waiters use the following default strategy:

- Use an exponential back-off, starting from a few seconds to a maximum of 30 seconds between poll attempts.
- Poll up to 20 minutes, and then stop.

Or more information on waiters, see:

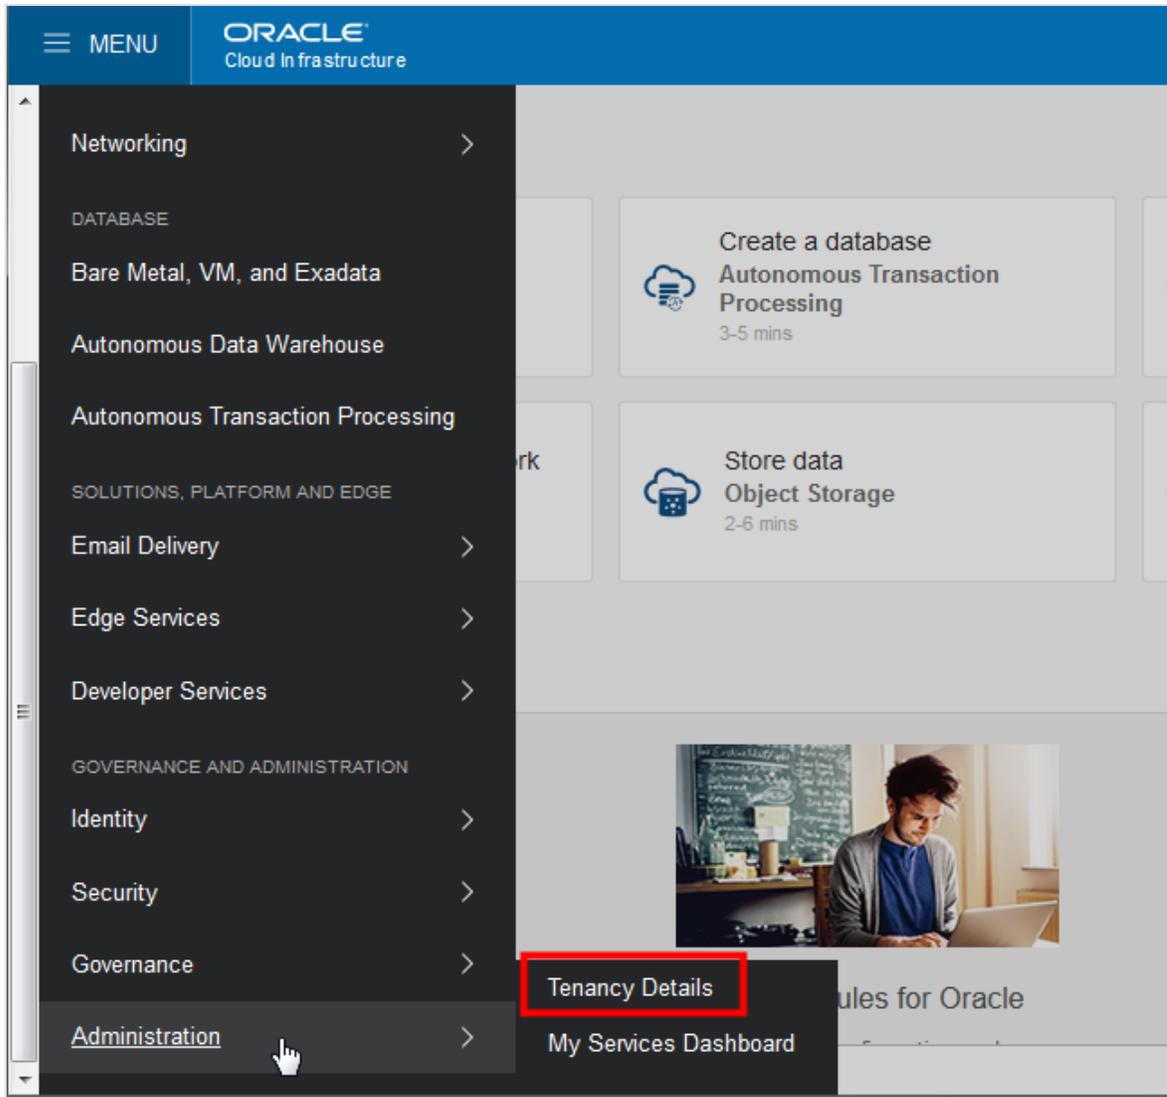
- [SDK for Java waiters documentation](#)
- [Ruby SDK waiters documentation](#)

### Where to Find Your Tenancy's OCID

If you use the API, you'll need your tenancy's OCID in order to sign the requests (see [Request Signatures](#)). You'll also need it for some of the IAM API operations. An OCID is an Oracle Cloud ID (see [Resource Identifiers](#)).

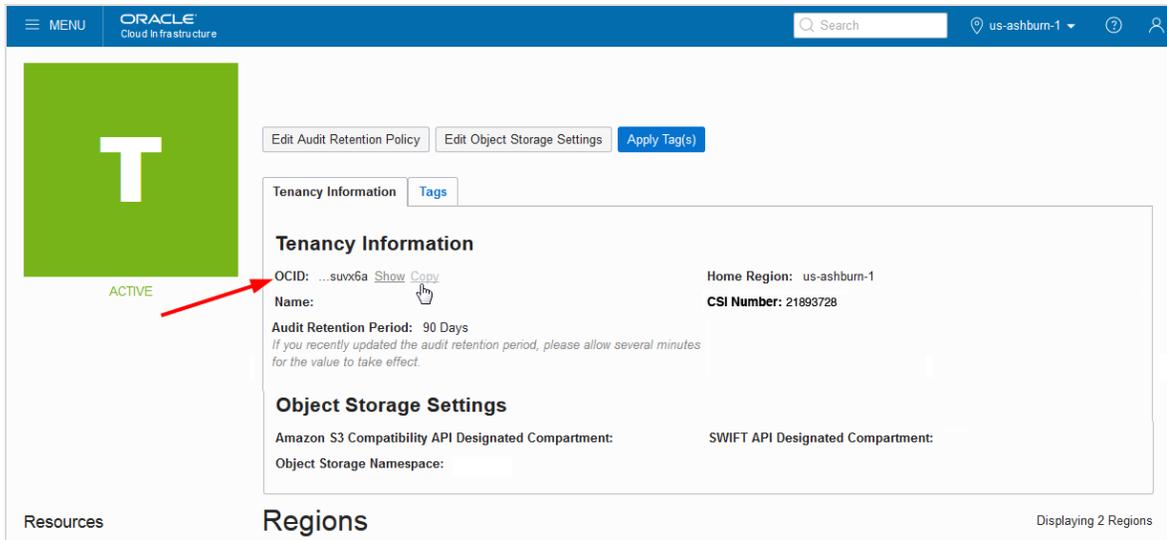
Get the tenancy OCID from the Oracle Cloud Infrastructure Console on the **Tenancy Details** page:

1. Open the navigation menu, under Governance and Administration, go to **Administration** and click **Tenancy Details**.



2. The tenancy OCID is shown under **Tenancy Information**. Click **Copy** to copy it to your clipboard.

## CHAPTER 34 Developer Tools



The tenancy OCID looks something like this (notice the word "tenancy" in it):  
`ocid1.tenancy.oc1..<unique_ID>`.

### List Pagination

Most List operations paginate results. For example, results are paginated for the [ListInstances](#) operation in the Core Services API. When you call a paginated List operation, the response indicates additional pages of results by including the `opc-next-page` header.



#### Note

A page can be empty even when more results remain. Any time the `opc-next-page` header appears, there are more list items to get. For more information about resource list control, see [Overview of Search](#).

### To get the next page of results

Make a new GET request against the same URL, modified by setting the page query parameter to the value from the `opc-next-page` header. Repeat this process until you get a response without an `opc-next-page` header. The absence of this header indicates that you have reached the last page of the list.



#### Note

For an alternative to writing pagination code, see the functions in the [pagination module](#) provided with the Python SDK.

### To get the previous page of results

(Available with some APIs.) Make a new GET request against the same URL, modified by setting the page query parameter to the value from the `opc-prev-page` header. Repeat this process until you get a response without an `opc-prev-page` header. The absence of this header indicates that you have reached the first page of the list.



#### Note

For an alternative to writing pagination code, see the functions in the [pagination module](#) provided with the Python SDK.

### To change the maximum number of results per page

In the GET request, set the `limit` to the number of items you want returned in the response.



### Note

The service will return no more than the number specified as `limit`, but might not return that exact number.

## Retry Token

For some operations you can provide a unique retry token (`opc-retry-token`) so the request can be retried in case of a timeout or server error without the risk of executing that same action again. The token expires after 24 hours, but can be invalidated before then due to conflicting operations (for example, if a resource has been deleted and purged from the system, then a retry of the original creation request may be rejected).

## ETags for Optimistic Concurrency Control

The API supports etags for the purposes of optimistic concurrency control. The GET and POST calls return an `etag` response header with a value you should store. When you later want to update or delete the resource, set the `if-match` header to the ETag you received for the resource. The resource will then be updated or deleted *only* if the ETag you provide matches the current value of that resource's ETag.

## Null vs. Empty Strings for Optional Parameters

If you send an empty string ("" ) as the value of an optional parameter, the API validates the value as normal (for example, checks against minimum and maximum allowed length, and so on). Often the minimum allowed length is 1, so an error would be returned. If you don't set the value (it's null), the API performs no validation, and some other action may occur. For example: if you don't set a value for the `displayName` when creating a new VCN object, the service will auto-generate a value.

### API Reference and Endpoints

Oracle Cloud Infrastructure has these APIs and corresponding regional endpoints:

#### Analytics Cloud API

##### [API reference](#)

- <https://analytics.ap-tokyo-1.ocp.oraclecloud.com>
- <https://analytics.us-ashburn-1.ocp.oraclecloud.com>

#### Announcements API

##### [API reference](#)

- <https://announcements.ap-seoul-1.oraclecloud.com>
- <https://announcements.ap-tokyo-1.oraclecloud.com>
- <https://announcements.ca-toronto-1.oraclecloud.com>
- <https://announcements.eu-frankfurt-1.oraclecloud.com>
- <https://announcements.uk-london-1.oraclecloud.com>
- <https://announcements.us-ashburn-1.oraclecloud.com>
- <https://announcements.us-phoenix-1.oraclecloud.com>

#### Audit API

##### [API reference](#)

- <https://audit.ap-mumbai-1.oraclecloud.com>
- <https://audit.ap-seoul-1.oraclecloud.com>
- <https://audit.ap-sydney-1.oraclecloud.com>
- <https://audit.ap-tokyo-1.oraclecloud.com>

- <https://audit.ca-toronto-1.oraclecloud.com>
- <https://audit.eu-frankfurt-1.oraclecloud.com>
- <https://audit.eu-zurich-1.oraclecloud.com>
- <https://audit.sa-saopaulo-1.oraclecloud.com>
- <https://audit.us-ashburn-1.oraclecloud.com>
- <https://audit.us-phoenix-1.oraclecloud.com>

### Budgets API

#### [API reference](#)

- <https://usage.ap-mumbai-1.oci.oraclecloud.com>
- <https://usage.ap-seoul-1.oci.oraclecloud.com>
- <https://usage.ap-sydney-1.oci.oraclecloud.com>
- <https://usage.ap-tokyo-1.oci.oraclecloud.com>
- <https://usage.ca-toronto-1.oci.oraclecloud.com>
- <https://usage.eu-frankfurt-1.oci.oraclecloud.com>
- <https://usage.eu-zurich-1.oci.oraclecloud.com>
- <https://usage.sa-saopaulo-1.oraclecloud.com>
- <https://usage.uk-london-1.oci.oraclecloud.com>
- <https://usage.us-ashburn-1.oci.oraclecloud.com>
- <https://usage.us-phoenix-1.oci.oraclecloud.com>

### Container Engine for Kubernetes API

#### [API reference](#)

- <https://containerengine.ap-mumbai-1.oraclecloud.com>
- <https://containerengine.ap-seoul-1.oraclecloud.com>

- <https://containerengine.ap-sydney-1.oraclecloud.com>
- <https://containerengine.ap-tokyo-1.oraclecloud.com>
- <https://containerengine.ca-toronto-1.oraclecloud.com>
- <https://containerengine.eu-frankfurt-1.oraclecloud.com>
- <https://containerengine.eu-zurich-1.oraclecloud.com>
- <https://containerengine.sa-saopaulo-1.oraclecloud.com>
- <https://containerengine.uk-london-1.oraclecloud.com>
- <https://containerengine.us-ashburn-1.oraclecloud.com>
- <https://containerengine.us-phoenix-1.oraclecloud.com>

### Core Services (covering Networking, Compute, and Block Volume)

The Networking, Compute, and Block Volume services are accessible with the following API:

#### Core Services API

##### [API reference](#)

- <https://iaas.ap-mumbai-1.oraclecloud.com>
- <https://iaas.ap-seoul-1.oraclecloud.com>
- <https://iaas.ap-sydney-1.oraclecloud.com>
- <https://iaas.ap-tokyo-1.oraclecloud.com>
- <https://iaas.ca-toronto-1.oraclecloud.com>
- <https://iaas.eu-frankfurt-1.oraclecloud.com>
- <https://iaas.eu-zurich-1.oraclecloud.com>
- <https://iaas.sa-saopaulo-1.oraclecloud.com>
- <https://iaas.uk-london-1.oraclecloud.com>

## CHAPTER 34 Developer Tools

---

- <https://iaas.us-ashburn-1.oraclecloud.com>
- <https://iaas.us-phoenix-1.oraclecloud.com>

You can manage Compute instances with the following API:

### Autoscaling API

#### [API reference](#)

- <https://autoscaling.ap-mumbai-1.oci.oraclecloud.com>
- <https://autoscaling.ap-seoul-1.oci.oraclecloud.com>
- <https://autoscaling.ap-sydney-1.oci.oraclecloud.com>
- <https://autoscaling.ap-tokyo-1.oci.oraclecloud.com>
- <https://autoscaling.ca-toronto-1.oci.oraclecloud.com>
- <https://autoscaling.eu-frankfurt-1.oci.oraclecloud.com>
- <https://autoscaling.eu-zurich-1.oci.oraclecloud.com>
- <https://autoscaling.sa-saopaulo-1.oci.oraclecloud.com>
- <https://autoscaling.uk-london-1.oci.oraclecloud.com>
- <https://autoscaling.us-ashburn-1.oci.oraclecloud.com>
- <https://autoscaling.us-phoenix-1.oci.oraclecloud.com>

You can track the progress of long-running Compute operations with the [Work Requests](#) API.

### Database API

#### [API reference](#)

- <https://database.ap-mumbai-1.oraclecloud.com>
- <https://database.ap-seoul-1.oraclecloud.com>

## CHAPTER 34 Developer Tools

---

- <https://database.ap-sydney-1.oraclecloud.com>
- <https://database.ap-tokyo-1.oraclecloud.com>
- <https://database.ca-toronto-1.oraclecloud.com>
- <https://database.eu-frankfurt-1.oraclecloud.com>
- <https://database.eu-zurich-1.oraclecloud.com>
- <https://database.sa-saopaulo-1.oraclecloud.com>
- <https://database.uk-london-1.oraclecloud.com>
- <https://database.us-ashburn-1.oraclecloud.com>
- <https://database.us-phoenix-1.oraclecloud.com>

You can track the progress of long-running Database operations with the [Work Requests](#) API.

### DNS API

#### [API reference](#)

- <https://dns.ap-mumbai-1.oraclecloud.com>
- <https://dns.ap-seoul-1.oraclecloud.com>
- <https://dns.ap-sydney-1.oraclecloud.com>
- <https://dns.ap-tokyo-1.oraclecloud.com>
- <https://dns.ca-toronto-1.oraclecloud.com>
- <https://dns.eu-frankfurt-1.oraclecloud.com>
- <https://dns.eu-zurich-1.oraclecloud.com>
- <https://dns.sa-saopaulo-1.oraclecloud.com>
- <https://dns.uk-london-1.oraclecloud.com>
- <https://dns.us-ashburn-1.oraclecloud.com>
- <https://dns.us-phoenix-1.oraclecloud.com>

### Email Delivery API

#### [API reference](#)

- <https://email.us-ashburn-1.oraclecloud.com>
- <https://email.us-phoenix-1.oraclecloud.com>
- <https://ctrl.email.uk-london-1.oci.oraclecloud.com>

### Events API

#### [API reference](#)

- <https://events.ap-mumbai-1.oraclecloud.com>
- <https://events.ap-seoul-1.oraclecloud.com>
- <https://events.ap-sydney-1.oraclecloud.com>
- <https://events.ap-tokyo-1.oraclecloud.com>
- <https://events.ap-toronto-1.oraclecloud.com>
- <https://events.eu-frankfurt-1.oraclecloud.com>
- <https://events.eu-zurich-1.oraclecloud.com>
- <https://events.sa-saopaulo-1.oraclecloud.com>
- <https://events.uk-london-1.oraclecloud.com>
- <https://events.us-ashburn-1.oraclecloud.com>
- <https://events.us-phoenix-1.oraclecloud.com>

### File Storage API

#### [API reference](#)

- <https://filestorage.ap-mumbai-1.oraclecloud.com>
- <https://filestorage.ap-seoul-1.oraclecloud.com>

- <https://filestorage.ap-sydney-1.oraclecloud.com>
- <https://filestorage.ap-tokyo-1.oraclecloud.com>
- <https://filestorage.ca-toronto-1.oraclecloud.com>
- <https://filestorage.eu-frankfurt-1.oraclecloud.com>
- <https://filestorage.eu-zurich-1.oraclecloud.com>
- <https://filestorage.sa-saopaulo-1.oraclecloud.com>
- <https://filestorage.uk-london-1.oraclecloud.com>
- <https://filestorage.us-ashburn-1.oraclecloud.com>
- <https://filestorage.us-phoenix-1.oraclecloud.com>

### Functions API

#### [API reference](#)

- <https://functions.ap-mumbai-1.oci.oraclecloud.com>
- <https://functions.ap-seoul-1.oci.oraclecloud.com>
- <https://functions.ap-sydney-1.oci.oraclecloud.com>
- <https://functions.ap-tokyo-1.oci.oraclecloud.com>
- <https://functions.ca-toronto-1.oci.oraclecloud.com>
- <https://functions.eu-frankfurt-1.oci.oraclecloud.com>
- <https://functions.eu-zurich-1.oci.oraclecloud.com>
- <https://functions.sa-saopaulo-1.oci.oraclecloud.com>
- <https://functions.uk-london-1.oci.oraclecloud.com>
- <https://functions.us-ashburn-1.oci.oraclecloud.com>
- <https://functions.us-phoenix-1.oci.oraclecloud.com>

These endpoints are shown in the preferred format. Note that an older format (`https://functions.<region-identifier>.oraclecloud.com`) is still supported, but is not preferred.

### Health Checks API

#### [API reference](#)

- `https://healthchecks.ap-mumbai-1.oraclecloud.com`
- `https://healthchecks.ap-seoul-1.oraclecloud.com`
- `https://healthchecks.ap-sydney-1.oraclecloud.com`
- `https://healthchecks.ap-tokyo-1.oraclecloud.com`
- `https://healthchecks.ca-toronto-1.oraclecloud.com`
- `https://healthchecks.eu-frankfurt-1.oraclecloud.com`
- `https://healthchecks.eu-zurich-1.oraclecloud.com`
- `https://healthchecks.sa-saopaulo-1.oraclecloud.com`
- `https://healthchecks.uk-london-1.oraclecloud.com`
- `https://healthchecks.us-ashburn-1.oraclecloud.com`
- `https://healthchecks.us-phoenix-1.oraclecloud.com`

### IAM API

#### [API reference](#)

- `https://identity.ap-mumbai-1.oraclecloud.com`
- `https://identity.ap-seoul-1.oraclecloud.com`
- `https://identity.ap-sydney-1.oraclecloud.com`
- `https://identity.ap-tokyo-1.oraclecloud.com`
- `https://identity.ca-toronto-1.oraclecloud.com`

- <https://identity.eu-frankfurt-1.oraclecloud.com>
- <https://identity.eu-zurich-1.oraclecloud.com>
- <https://identity.sa-saopaulo-1.oraclecloud.com>
- <https://identity.uk-london-1.oraclecloud.com>
- <https://identity.us-ashburn-1.oraclecloud.com>
- <https://identity.us-phoenix-1.oraclecloud.com>



### Note

*Use the Endpoint of Your Home Region for All IAM API Calls*

When you sign up for Oracle Cloud Infrastructure, Oracle creates a tenancy for you in one region. This is your home region. Your home region is where your IAM resources are defined. When you subscribe to a new region, your IAM resources are replicated in the new region, however, the master definitions reside in your home region and can only be changed there. Make all IAM API calls against your home region endpoint. The changes automatically replicate to all regions. If you try to make an IAM API call against a region that is not your home region, you will receive an error.

## Integration API

### Oracle Integration Service Instance REST API

Use this API to create, modify, delete, get information about an Oracle Integration instance, and list Oracle Integration instances in a compartment.

[Oracle Integration Service Instance REST API reference](#)

- <https://integration.ap-mumbai-1.ocp.oraclecloud.com>
- <https://integration.ap-seoul-1.ocp.oraclecloud.com>
- <https://integration.ap-sydney-1.ocp.oraclecloud.com>
- <https://integration.ap-tokyo-1.ocp.oraclecloud.com>
- <https://integration.ca-toronto-1.ocp.oraclecloud.com>
- <https://integration.eu-frankfurt-1.ocp.oraclecloud.com>
- <https://integration.uk-london-1.ocp.oraclecloud.com>
- <https://integration.us-ashburn-1.ocp.oraclecloud.com>
- <https://integration.us-phoenix-1.ocp.oraclecloud.com>

### Oracle Integration REST API

Use this API to integrate applications and automate business processes.

[Oracle Integration REST API Reference](#)

### Key Management API

[API reference](#)

- <https://kms.ap-mumbai-1.oraclecloud.com>
- <https://kms.ap-seoul-1.oraclecloud.com>
- <https://kms.ap-sydney-1.oraclecloud.com>
- <https://kms.ap-tokyo-1.oraclecloud.com>
- <https://kms.ca-toronto-1.oraclecloud.com>
- <https://kms.eu-frankfurt-1.oraclecloud.com>
- <https://kms.eu-zurich-1.oraclecloud.com>
- <https://kms.sa-saopaulo-1.oraclecloud.com>

- <https://kms.uk-london-1.oraclecloud.com>
- <https://kms.us-ashburn-1.oraclecloud.com>
- <https://kms.us-phoenix-1.oraclecloud.com>

In addition to these endpoints, each vault has a unique endpoint for create, update, and list operations for keys. This endpoint is referred to as the control plane URL or management endpoint. Each vault also has a unique endpoint for cryptographic operations. This endpoint is known as the data plane URL or the cryptographic endpoint.

### Load Balancing API

#### [API reference](#)

- <https://iaas.ap-mumbai-1.oraclecloud.com>
- <https://iaas.ap-seoul-1.oraclecloud.com>
- <https://iaas.ap-sydney-1.oraclecloud.com>
- <https://iaas.ap-tokyo-1.oraclecloud.com>
- <https://iaas.ca-toronto-1.oraclecloud.com>
- <https://iaas.eu-frankfurt-1.oraclecloud.com>
- <https://iaas.eu-zurich-1.oraclecloud.com>
- <https://iaas.sa-saopaulo-1.oraclecloud.com>
- <https://iaas.uk-london-1.oraclecloud.com>
- <https://iaas.us-ashburn-1.oraclecloud.com>
- <https://iaas.us-phoenix-1.oraclecloud.com>

### Marketplace API

#### [API reference](#)

- <https://marketplace.ap-mumbai-1.oci.oraclecloud.com>
- <https://marketplace.ap-seoul-1.oci.oraclecloud.com>
- <https://marketplace.ap-sydney-1.oci.oraclecloud.com>
- <https://marketplace.ap-tokyo-1.oci.oraclecloud.com>
- <https://marketplace.ca-toronto-1.oci.oraclecloud.com>
- <https://marketplace.eu-frankfurt-1.oci.oraclecloud.com>
- <https://marketplace.eu-zurich-1.oci.oraclecloud.com>
- <https://marketplace.sa-saopaulo-1.oci.oraclecloud.com>
- <https://marketplace.uk-london-1.oci.oraclecloud.com>
- <https://marketplace.us-ashburn-1.oci.oraclecloud.com>
- <https://marketplace.us-phoenix-1.oci.oraclecloud.com>

### Monitoring

#### [API reference](#)

#### [PostMetricData](#) operation:

- <https://telemetry-ingestion.ap-mumbai-1.oraclecloud.com>
- <https://telemetry-ingestion.ap-seoul-1.oraclecloud.com>
- <https://telemetry-ingestion.ap-sydney-1.oraclecloud.com>
- <https://telemetry-ingestion.ap-tokyo-1.oraclecloud.com>
- <https://telemetry-ingestion.ca-toronto-1.oraclecloud.com>
- <https://telemetry-ingestion.eu-frankfurt-1.oraclecloud.com>
- <https://telemetry-ingestion.eu-zurich-1.oraclecloud.com>
- <https://telemetry-ingestion.sa-saopaulo-1.oraclecloud.com>
- <https://telemetry-ingestion.uk-london-1.oraclecloud.com>

## CHAPTER 34 Developer Tools

---

- <https://telemetry-ingestion.us-ashburn-1.oraclecloud.com>
- <https://telemetry-ingestion.us-phoenix-1.oraclecloud.com>

All other operations:

- <https://telemetry.ap-mumbai-1.oraclecloud.com>
- <https://telemetry.ap-seoul-1.oraclecloud.com>
- <https://telemetry.ap-sydney-1.oraclecloud.com>
- <https://telemetry.ap-tokyo-1.oraclecloud.com>
- <https://telemetry.ca-toronto-1.oraclecloud.com>
- <https://telemetry.eu-frankfurt-1.oraclecloud.com>
- <https://telemetry.eu-zurich-1.oraclecloud.com>
- <https://telemetry.sa-saopaulo-1.oraclecloud.com>
- <https://telemetry.uk-london-1.oraclecloud.com>
- <https://telemetry.us-ashburn-1.oraclecloud.com>
- <https://telemetry.us-phoenix-1.oraclecloud.com>

## Notifications

### [API reference](#)

- <https://cp.notification.ap-mumbai-1.oraclecloud.com>
- <https://cp.notification.ap-seoul-1.oraclecloud.com>
- <https://cp.notification.ap-sydney-1.oraclecloud.com>
- <https://cp.notification.ap-tokyo-1.oraclecloud.com>
- <https://cp.notification.ca-toronto-1.oraclecloud.com>
- <https://cp.notification.eu-frankfurt-1.oraclecloud.com>
- <https://cp.notification.eu-zurich-1.oraclecloud.com>

- <https://cp.notification.sa-saopaulo-1.oraclecloud.com>
- <https://cp.notification.uk-london-1.oraclecloud.com>
- <https://cp.notification.us-ashburn-1.oraclecloud.com>
- [https://cp.notification.US West \(Phoenix\).oraclecloud.com](https://cp.notification.US West (Phoenix).oraclecloud.com)

### Object Storage and Archive Storage APIs

Both Object Storage and Archive Storage are accessible with the following APIs:

#### Object Storage API

##### [API reference](#)

- <https://objectstorage.ap-mumbai-1.oraclecloud.com>
- <https://objectstorage.ap-seoul-1.oraclecloud.com>
- <https://objectstorage.ap-sydney-1.oraclecloud.com>
- <https://objectstorage.ap-tokyo-1.oraclecloud.com>
- <https://objectstorage.ca-toronto-1.oraclecloud.com>
- <https://objectstorage.eu-frankfurt-1.oraclecloud.com>
- <https://objectstorage.eu-zurich-1.oraclecloud.com>
- <https://objectstorage.sa-saopaulo-1.oraclecloud.com>
- <https://objectstorage.uk-london-1.oraclecloud.com>
- <https://objectstorage.us-ashburn-1.oraclecloud.com>
- <https://objectstorage.us-phoenix-1.oraclecloud.com>

#### Amazon S3 Compatibility API

##### [API reference](#)

- [https://<object\\_storage\\_namespace>.compat.objectstorage.ap-mumbai-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.ap-mumbai-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.ap-seoul-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.ap-seoul-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.ap-sydney-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.ap-sydney-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.ap-tokyo-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.ap-tokyo-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.ca-toronto-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.ca-toronto-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.eu-frankfurt-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.eu-frankfurt-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.eu-zurich-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.eu-zurich-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.sa-saopaulo-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.sa-saopaulo-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.uk-london-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.uk-london-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.us-ashburn-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.us-ashburn-1.oraclecloud.com)
- [https://<object\\_storage\\_namespace>.compat.objectstorage.us-phoenix-1.oraclecloud.com](https://<object_storage_namespace>.compat.objectstorage.us-phoenix-1.oraclecloud.com)



### Tip

See [Understanding Object Storage Namespaces](#) for information regarding how to find your Object Storage namespace.

### Swift API (for use with Oracle RMAN)

- <https://swiftobjectstorage.ap-mumbai-1.oraclecloud.com>
- <https://swiftobjectstorage.ap-seoul-1.oraclecloud.com>
- <https://swiftobjectstorage.ap-sydney-1.oraclecloud.com>
- <https://swiftobjectstorage.ap-tokyo-1.oraclecloud.com>
- <https://swiftobjectstorage.ca-toronto-1.oraclecloud.com>
- <https://swiftobjectstorage.eu-frankfurt-1.oraclecloud.com>
- <https://swiftobjectstorage.eu-zurich-1.oraclecloud.com>
- <https://swiftobjectstorage.sa-saopaulo-1.oraclecloud.com>
- <https://swiftobjectstorage.uk-london-1.oraclecloud.com>
- <https://swiftobjectstorage.us-ashburn-1.oraclecloud.com>
- <https://swiftobjectstorage.us-phoenix-1.oraclecloud.com>

### Oracle Cloud My Services API

#### [API reference](#)

- <https://itra.oraclecloud.com/>



#### **Important**

The My Services dashboard and APIs are deprecated.

### Registry

Oracle Cloud Infrastructure Registry fully implements a Docker protocol that enables you to use the Docker Registry HTTP API (rather than the Oracle Cloud Infrastructure API) to manage images using the following endpoints. See the [Docker documentation](#) for information about

using the Docker Registry HTTP API .

- Mumbai:
  - <https://bom.ocir.io>
  - <https://ap-mumbai-1.ocir.io>
- Seoul:
  - <https://icn.ocir.io>
  - <https://ap-seoul-1.ocir.io>
- Sydney:
  - <https://syd.ocir.io>
  - <https://ap-sydney-1.ocir.io>
- Tokyo:
  - <https://nrt.ocir.io>
  - <https://ap-tokyo-1.ocir.io>
- Toronto:
  - <https://yyz.ocir.io>
  - <https://ca-toronto-1.ocir.io>
- Frankfurt:
  - <https://fra.ocir.io>
  - <https://eu-frankfurt-1.ocir.io>
- Zurich:
  - <https://zrh.ocir.io>
  - <https://eu-zurich-1.ocir.io>
- Sao Paulo:
  - <https://gru.ocir.io>
  - <https://sa-saopaulo-1.ocir.io>

- London:
  - <https://lhr.ocir.io>
  - <https://uk-london-1.ocir.io>
- Ashburn:
  - <https://iad.ocir.io>
  - <https://us-ashburn-1.ocir.io>
- Phoenix:
  - <https://phx.ocir.io>
  - <https://us-phoenix-1.ocir.io>

### Resource Manager API

#### [API reference](#)

- <https://resourcemanager.ap-mumbai-1.oraclecloud.com>
- <https://resourcemanager.ap-seoul-1.oraclecloud.com>
- <https://resourcemanager.ap-sydney-1.oraclecloud.com>
- <https://resourcemanager.ap-tokyo-1.oraclecloud.com>
- <https://resourcemanager.ca-toronto-1.oraclecloud.com/>
- <https://resourcemanager.eu-frankfurt-1.oraclecloud.com/>
- <https://resourcemanager.eu-zurich-1.oraclecloud.com/>
- <https://resourcemanager.sa-saopaulo-1.oraclecloud.com>
- <https://resourcemanager.uk-london-1.oraclecloud.com/>
- <https://resourcemanager.us-ashburn-1.oraclecloud.com/>
- <https://resourcemanager.us-phoenix-1.oraclecloud.com/>

### Search API

#### [API reference](#)

- <https://query.ap-mumbai-1.oraclecloud.com/>
- <https://query.ap-seoul-1.oraclecloud.com/>
- <https://query.ap-sydney-1.oraclecloud.com/>
- <https://query.ap-tokyo-1.oraclecloud.com/>
- <https://query.ca-toronto-1.oraclecloud.com/>
- <https://query.eu-frankfurt-1.oraclecloud.com/>
- <https://query.eu-zurich-1.oraclecloud.com/>
- [https://query.sa-saopaulo-1.oraclecloud.com](https://query.sa-saopaulo-1.oraclecloud.com/)
- <https://query.uk-london-1.oraclecloud.com/>
- <https://query.us-ashburn-1.oraclecloud.com/>
- <https://query.us-phoenix-1.oraclecloud.com/>

### Service Limits and Quotas API

#### [API reference](#)

- <https://limits.ap-mumbai-1.oci.oraclecloud.com>
- <https://limits.ap-seoul-1.oci.oraclecloud.com>
- <https://limits.ap-sydney-1.oci.oraclecloud.com>
- <https://limits.ap-tokyo-1.oci.oraclecloud.com>
- <https://limits.ca-toronto-1.oci.oraclecloud.com>
- <https://limits.eu-frankfurt-1.oci.oraclecloud.com>
- <https://limits.uk-london-1.oci.oraclecloud.com>
- <https://limits.us-ashburn-1.oci.oraclecloud.com>
- <https://limits.us-phoenix-1.oci.oraclecloud.com>

### Streaming API

#### [API reference](#)

- <https://streaming.ap-mumbai-1.oci.oraclecloud.com>
- <https://streaming.ap-seoul-1.oci.oraclecloud.com>
- <https://streaming.ap-sydney-1.oci.oraclecloud.com>
- <https://streaming.ap-tokyo-1.oci.oraclecloud.com>
- <https://streaming.ca-toronto-1.oci.oraclecloud.com>
- <https://streaming.eu-frankfurt-1.oci.oraclecloud.com>
- <https://streaming.eu-zurich-1.oci.oraclecloud.com>
- <https://streaming.sa-saopaulo-1.oci.oraclecloud.com>
- <https://streaming.uk-london-1.oci.oraclecloud.com>
- <https://streaming.us-ashburn-1.oci.oraclecloud.com>
- <https://streaming.us-phoenix-1.oci.oraclecloud.com>

### Web Application Acceleration and Security API

#### [API reference](#)

- <https://waas.ap-mumbai-1.oraclecloud.com>
- <https://waas.ap-seoul-1.oraclecloud.com>
- <https://waas.ap-sydney-1.oraclecloud.com>
- <https://waas.ap-tokyo-1.oraclecloud.com>
- <https://waas.ca-toronto-1.oraclecloud.com>
- <https://waas.eu-frankfurt-1.oraclecloud.com>
- <https://waas.eu-zurich-1.oraclecloud.com>
- <https://waas.sa-saopaulo-1.oraclecloud.com>
- <https://waas.uk-london-1.oraclecloud.com>

- <https://waas.us-ashburn-1.oraclecloud.com>
- <https://waas.us-phoenix-1.oraclecloud.com>

### Work Requests API (for Compute and Database work requests)

#### [API reference](#)

- <https://iaas.ap-mumbai-1.oraclecloud.com>
- <https://iaas.ap-seoul-1.oraclecloud.com>
- <https://iaas.ap-sydney-1.oraclecloud.com>
- <https://iaas.ap-tokyo-1.oraclecloud.com>
- <https://iaas.ca-toronto-1.oraclecloud.com>
- <https://iaas.eu-frankfurt-1.oraclecloud.com>
- <https://iaas.eu-zurich-1.oraclecloud.com>
- <https://iaas.sa-saopaulo-1.oraclecloud.com>
- <https://iaas.uk-london-1.oraclecloud.com>
- <https://iaas.us-ashburn-1.oraclecloud.com>
- <https://iaas.us-phoenix-1.oraclecloud.com>

### API Errors

#### **Common Errors Returned by All Services**

The following table lists the common errors returned by all the services for Oracle Cloud Infrastructure.

## CHAPTER 34 Developer Tools

HTTP Status Code	Error Code	Description	Retry
400	CannotParseRequest	The request is incorrectly formatted.	No.
400	InvalidParameter	A parameter is invalid or incorrectly formatted.	No.
400	LimitExceeded	Fulfilling this request exceeds the Oracle-defined limit for this tenancy for this resource type.	No.
400	MissingParameter	A required parameter is missing.	No.
400	QuotaExceeded	Fulfilling this request exceeds the administrator-defined quota for this compartment for this resource.	No.
400	RelatedResourceNotAuthorizedOrNotFound	A resource specified in the body of the request was not found, or you do not have authorization to access that resource.	Yes, with backoff.
401	NotAuthenticated	The required authentication information was not provided or was incorrect. There are other reasons why this error code is generated. For more information, see <a href="#">HTML Status Code 401</a> .	Yes, with backoff.
402	SignUpRequired	This operation requires opt-in before it may be called.	No.
403	NotAuthorized	You do not have authorization to update one or more of the fields included in this request.	No.

## CHAPTER 34 Developer Tools

HTTP Status Code	Error Code	Description	Retry
404	NotAuthorizedOr NotFound	A resource specified via the URI (path or query parameters) of the request was not found, or you do not have authorization to access that resource. For more information, see <a href="#">HTML Status Code 404</a> .	Yes, with backoff.
404	NotFound	There is no operation supported at the URI path and HTTP method you specified in the request.	No.
405	MethodNotAllowed	The target resource does not support the HTTP method.	No.
409	IncorrectState	The requested state for the resource conflicts with its current state.	Yes, with backoff.
409	InvalidatedRetryToken	The provided retry token was used in an earlier request that resulted in a system update, but a subsequent operation invalidated the token. This can happen, for example, in cases where an entity created with the same token has since been deleted. If the system state change that is associated with this request should be performed again, retry it using a different token.	No.

HTTP Status Code	Error Code	Description	Retry
409	NotAuthorizedOrResourceAlreadyExists	You do not have authorization to perform this request, or the resource you are attempting to create already exists. This error code is returned only from <b>create</b> operations, where it is returned instead of the more general <b>NotAuthorizedOrNotFound</b> error code."	Yes, with backoff.
412	NoEtagMatch	The ETag specified in the request does not match the ETag for the resource.	No.
429	TooManyRequests	You have issued too many requests to the Oracle Cloud Infrastructure APIs in too short of an amount of time.	Yes, with backoff.
500	InternalServerError	An internal server error occurred.	Yes, with backoff.
501	MethodNotImplemented	The HTTP request target does not recognize the HTTP method.	No.

## Error Details and Troubleshooting

### HTTP STATUS CODE: 401

- **Missing or incorrect authentication information.** Verify that all the required information (tenant OCID, user OCID, fingerprint, and private key) is provided and accurate. Verify that the public key corresponding to the fingerprint has been

uploaded for the user. For more information, see [Required Keys and OCIDs](#).

- **Clock skew.** This status code is returned if the client's clock is skewed more than five (5) minutes from the server's clock. For more information, see [Maximum Allowed Client Clock Skew](#).
- **API request signature error.** This status code is returned if a required header is missing from a signing string. For more information, see [Request Signatures](#).

### **ERROR CODES: NOTAUTHORIZEDORNOTFOUND, RELATEDRESOURCENOTAUTHORIZEDORNOTFOUND , NOTAUTHORIZEDORRESOURCEALREADYEXISTS**

- **Authorization error.** Verify that the user is in a group that has the permissions to work with resources in a compartment.
- **Compartment or resource not found.** Verify that the compartment or resource exist and is referenced correctly. For example, this status code is returned for either of the following errors:
  - `CompartmentNotFound` if a compartment doesn't exist
  - `VolumeNotFound` if a volume doesn't exist

## Asynchronous Work Requests

This topic describes asynchronous work requests for long-running operations against Oracle Cloud Infrastructure services. It also provides guidance on obtaining request status, and for inspecting the request response to enable filtering for affected resources.

### **Overview**

API calls to Oracle Cloud Infrastructure services can launch long-running operations that do not complete the client's request before a response is returned. In these cases, the service spawns an *asynchronous work request* that allows for visibility into the progress of long-running, asynchronous operations. The response to the REST API call contains a work request ID in the `opc-work-request-id` header, which allows you to monitor its progress and status. The work request itself remains in a queue until the operation has completed.

You can monitor the status of the work request at any time by calling `GetWorkRequest` and passing in the work request ID.



### Note

The Compute and Database services support work requests using [Work Requests API](#), which contains the `GetWorkRequest` operation.

The Container Engine for Kubernetes, Load Balancing, Object Storage and IAM services each support work requests through the service APIs. These service APIs each include operations that work in a similar manner to the `GetWorkRequest` operation used by the Work Requests API.

Two features of the request response are of particular interest: the status of the work request, and a list of the resources that are affected by the work request. The status is important because asynchronous work requests must know when an operation has completed, is still running, or whether it has failed altogether.

To retrieve information about work request failures or errors, each service provides APIs for fetching information about errors, and logs. For links to API reference documentation for each of the service, see the section [For More Information](#).

Also important in cases where a work request operation affects several resources is having a list of the resources that a work request affects, along with each one's `entityType` and `actionType` attributes.

### Work Request Status

Asynchronous work requests allow you to monitor their progress by providing a status attribute on the `WorkRequest` object. Each of the supported services provides its own API for obtaining status, as listed in the following sections.



**Note**

There is a [ContainerEngineWaiters](#) class that allows you to create a callback using the `forWorkRequest` method. Use this API to forward a notification when an operation's status changes, for example, from `IN_PROGRESS` to `COMPLETED`.

The following table lists status attributes that are supported by the `WorkRequest` object on the respective services.

Service	Status Attributes
Compute	<ul style="list-style-type: none"> <li>• ACCEPTED</li> <li>• IN_PROGRESS</li> <li>• FAILED</li> <li>• SUCCEEDED</li> <li>• CANCELING</li> <li>• CANCELED</li> </ul>
Container Engine for Kubernetes	<ul style="list-style-type: none"> <li>• ACCEPTED</li> <li>• IN_PROGRESS</li> <li>• FAILED</li> <li>• SUCCEEDED</li> <li>• CANCELING</li> <li>• CANCELED</li> </ul>

## CHAPTER 34 Developer Tools

---

Service	Status Attributes
Database	<ul style="list-style-type: none"><li>• ACCEPTED</li><li>• IN_PROGRESS</li><li>• FAILED</li><li>• SUCCEEDED</li></ul>
IAM	<ul style="list-style-type: none"><li>• ACCEPTED</li><li>• IN_PROGRESS</li><li>• FAILED</li><li>• SUCCEEDED</li><li>• CANCELING</li><li>• CANCELED</li></ul>
Load Balancing	<ul style="list-style-type: none"><li>• CREATING</li><li>• FAILED</li><li>• ACTIVE</li><li>• DELETING</li><li>• DELETED</li></ul>
Object Storage	<ul style="list-style-type: none"><li>• ACCEPTED</li><li>• IN_PROGRESS</li><li>• FAILED</li><li>• COMPLETED</li><li>• CANCELING</li><li>• CANCELED</li></ul>

### Filtering the Request Response

You sometimes need to know which resources are affected by a given asynchronous work request. In cases where the request response includes just one or two affected resources, the body of the request response is probably sufficient. However, in cases where a request response affects a great many resources, you must filter the response to identify the resources that you're interested in.

Filtering of resources listed in a work request response relies on two attributes of the `WorkRequestResource` `type`: `entityType` and `actionType`.

- **entityType**: Represents the resource type which the work request affects. This is an optional attribute, but each resource can have only one `entityType`.
- **actionType**: Represents how the specified resource is affected by the operation associated with the work request. Each service specifies a fixed list of allowable `actionType` values (shown in the sections following).

To obtain resource information on a work request, call `GetWorkRequest` and pass in the work request ID. The call returns a response in JSON format. Following is an example from calling [GetWorkRequest](#) on the Object Storage service.

```
{
 operationType: "COPY_OBJECT",
 status: "IN_PROGRESS",
 id: "f54527d6-029b-4221-9046-a811b7686202",
 resources: [
 {
 entityType: "object",
 actionType: "READ",
 entityUri: "/n/mynamespace/b/backups/o/myobject"
 },
 {
 entityType: "object",
 actionType: "WRITTEN",
 entityUri: "/n/mynamespace/b/backups/o/copyofmyobject"
 },
],
 timeAccepted: 2017-10-13T17:23:46.000Z,
 timeStarted: 2017-10-13T17:23:52.198Z,
}
```

## CHAPTER 34 Developer Tools

```
percentComplete: 10.0
}
```



### Note

Different services provide slightly different responses. See the reference documentation for each service's work request API for details. Links to each are provided in the [For More Information](#) section.

The following table lists the entity types and action types that are supported by Oracle Cloud Infrastructure services.

Service Name	Operation	entityType	actionType
Container Engine for Kubernetes	<a href="#">CreateCluster</a>	cluster	ACCEPTED
	<a href="#">DeleteCluster</a>	nodepool	IN_PROGRESS
	<a href="#">UpdateCluster</a>		FAILED
	<a href="#">CreateNodePool</a>		SUCCEEDED
	<a href="#">DeleteNodePool</a>		CANCELING
	<a href="#">UpdateNodePool</a>		CANCELED
Load Balancing	<a href="#">CreateLoadBalancer</a>	LoadBalancer	ACCEPTED
	<a href="#">UpdateLoadBalancer</a>		IN_PROGRESS
	<a href="#">DeleteLoadBalancer</a>		FAILED
			SUCCEEDED
Object Storage	<a href="#">CopyObject</a>	object	READ
			WRITTEN

### Request/Response Sample

Following is a sequence of REST API calls to create a cluster, which is a common long-running operation. The caller retrieves the work request ID from the response to the initial `POST` call and then periodically polls the `WorkRequest` to determine the status of the operation. The request/response sequence that follows depicts this workflow:

1. The user issues a `CreateCluster` API call.
2. The service responds with status code 202, indicating that the request has been accepted and returns a work request ID in the `opc-work-request-id` header.
3. Next, the user issues a `GET` call on the work request ID to obtain the status of the work request.
4. The service responds with status code 200, indicating in the response body that the `CLUSTER_CREATE` operation has the status `ACCEPTED`.
5. With continued polling, we see another `GET` call for the work request.
6. The service responds with status code 200. The response body reports that the operation `SUCCEEDED`.

#### Step 1. Initial API call to initiate a `CLUSTER_CREATE` operation.

```
POST https://containerengine.eu-frankfurt-1.oraclecloud.com/20180222/clusters
Accept: application/json
authorization: <Redacted>
content-length: 480
Content-Type: application/json
date: Mon, 02 Jul 2018 18:20:03 GMT
host: containerengine.eu-frankfurt-1.oraclecloud.com
opc-client-info: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT
opc-request-id: D7A390ED909C47038C438BA3629FB612
User-Agent: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT (Mac OS X/10.13.5; Java/1.8.0_172; Java HotSpot (TM)
64-Bit Server VM/25.172-b11)
x-content-sha256: S8U8OKQHytLNViAzgexkxvF4ctncJJHTjuRfXn0ya4={
 "name": "JavaSDK.CRUD",
 "compartmentId": "ocid1.compartment.oc1..<unique_ID>",
 "vcnId": "ocid1.vcn.oc1.eu-frankfurt-1.<unique_ID>",
 "kubernetesVersion": "v1.10.3",
 "options": {"serviceLbSubnetIds": ["ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>",
 "ocid1.subnet.oc1.eu-frankfurt-1.<unique_ID>"]}}
```

**Step 2.** The response to the initial API call, which contains the work request ID in the `OpC-Work-Request-Id` header.

```
202
Access-Control-Allow-Methods: DELETE, GET, HEAD, OPTIONS, PATCH, POST, PUT
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: opc-work-request-id
Content-Length: 0
Date: Mon, 02 Jul 2018 18:20:04 GMT
OpC-Request-Id:
D7A390ED909C47038C438BA3629FB612/33EEDCAAB2E84508B34AA75CD0FD86F4/8261D1CC89814E9BB934440A1F43DA09
OpC-Work-Request-Id: ocid1.clustersworkrequest.oc1.eu-frankfurt-1.exampleuniqueID
Uri: /20180222/clusters
Vary: Accept-Encoding
X-Rate-Limit-Duration: 1
X-Rate-Limit-Limit: 16.70
X-Rate-Limit-Request-Forwarded-For: 10.237.10.0, 10.237.9.51
X-Rate-Limit-Request-Remote-Addr: 10.237.9.51:53077
```

**Step 3.** Because this is a long-running operation, the user periodically polls the work request using a `GET` call to determine its status.

```
GET https://containerengine.eu-frankfurt-1.oraclecloud.com/20180222/workRequests/<clusters_work_request_
OCID>
Accept: application/json
authorization: <Redacted>
date: Mon, 02 Jul 2018 18:20:04 GMT
host: containerengine.eu-frankfurt-1.oraclecloud.com
opc-client-info: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT
opc-request-id: E8F20DAC443346B3B0EA599F367EE294
User-Agent: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT (Mac OS X/10.13.5; Java/1.8.0_172; Java HotSpot(TM)
64-Bit Server VM/25.172-b11)
```

**Step 4.** The `GET` call returns the following response, which indicates in the response body that the `CLUSTER_CREATE` operation has a status of `ACCEPTED`.

```
200
Access-Control-Allow-Methods: DELETE, GET, HEAD, OPTIONS, PATCH, POST, PUT
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: opc-work-request-id
Content-Length: 717
Content-Type: application/json
Date: Mon, 02 Jul 2018 18:20:05 GMT
```

## CHAPTER 34 Developer Tools

```
Etag: 56a41efaf33d81a54933495ee910c24d7bce7a83adf18810f95e07bdd2055805
Opc-Request-Id:
E8F20DAC443346B3B0EA599F367EE294/8B19C9FC3B4442CEA14685D1973D0856/0BA60B0711764DE4A4373071632708C7
Retry-After: 30
Uri: /20180222/workRequests/_id_
Vary: Accept-Encoding
X-Rate-Limit-Duration: 1
X-Rate-Limit-Limit: 16.70
X-Rate-Limit-Request-Forwarded-For: 10.237.10.0, 10.237.9.51
X-Rate-Limit-Request-Remote-Addr: 10.237.9.51:43533
{
 "id": "ocidl.clustersworkrequest.oc1.eu-frankfurt-1.exampleuniqueID",
 "operationType": "CLUSTER_CREATE",
 "status": "ACCEPTED",
 "compartmentId": "ocidl.compartment.oc1..exampleuniqueID",
 "resources": [
 {
 "actionType": "IN_PROGRESS",
 "entityType": "cluster",
 "identifier": "ocidl.cluster.oc1.eu-frankfurt-1.exampleuniqueID",
 "entityUri": "/clusters/ocidl.cluster.oc1.eu-frankfurt-1.exampleuniqueID"
 }
],
 "timeAccepted": "2018-07-02T18:20:05Z",
 "timeStarted": null,
 "timeFinished": null
}
```

**Step 5.** The operation continues, and the user continues to poll the work request using the GET method.

```
GET https://containerengine.eu-frankfurt-1.oraclecloud.com/20180222/workRequests/<clusters_work_request_
OCID>
Accept: application/json
authorization: <Redacted>
date: Mon, 02 Jul 2018 18:24:13 GMT
host: containerengine.eu-frankfurt-1.oraclecloud.com
opc-client-info: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT
opc-request-id: 64595B97E39A471A886DA29966BB6B1D
User-Agent: Oracle-JavaSDK/1.2.42-preview1-SNAPSHOT (Mac OS X/10.13.5; Java/1.8.0_172; Java HotSpot(TM)
64-Bit Server VM/25.172-b11)
```

**Step 6.** The last GET call produced the following response, which indicates that the operation has completed. Note the `entityType` is "cluster" and the `actionType` is "CREATED".

```
200
Access-Control-Allow-Methods: DELETE, GET, HEAD, OPTIONS, PATCH, POST, PUT
Access-Control-Allow-Origin: *
Access-Control-Expose-Headers: opc-work-request-id
Content-Length: 750
Content-Type: application/json
Date: Mon, 02 Jul 2018 18:24:14 GMT
Etag: 023d2a8ccb6d893fa8c875f64652353f21d22607825f49eeeb15b5394ae24918
Opc-Request-Id:
64595B97E39A471A886DA29966BB6B1D/3A81140991C94794AF365016E31DBE82/6245FBD8C25842B6BDF15187EA6ADB21
Uri: /20180222/workRequests/_id_
Vary: Accept-Encoding
X-Rate-Limit-Duration: 1
X-Rate-Limit-Limit: 16.70
X-Rate-Limit-Request-Forwarded-For: 10.237.3.0, 10.237.40.183
X-Rate-Limit-Request-Remote-Addr: 10.237.40.183:55856

{
 "id": "ocidl.clustersworkrequest.oc1.eu-frankfurt-1.exampleuniqueID",
 "operationType": "CLUSTER_CREATE",
 "status": "SUCCEEDED",
 "compartmentId": "ocidl.compartment.oc1..exampleuniqueID",
 "resources": [
 {
 "actionType": "CREATED",
 "entityType": "cluster",
 "identifier": "ocidl.cluster.oc1.eu-frankfurt-1.exampleuniqueID",
 "entityUri": "/clusters/ocidl.cluster.oc1.eu-frankfurt-1.exampleuniqueID"
 }
],
 "timeAccepted": "2018-07-02T18:20:05Z",
 "timeStarted": "2018-07-02T18:20:10Z",
 "timeFinished": "2018-07-02T18:24:01Z"
}
```

### For More Information

- Compute and Database:
  - [Using the Console to View Work Requests](#)
  - [WorkRequest API](#)
- Container Engine for Kubernetes: [WorkRequest API](#)
- IAM:
  - [WorkRequest API \(Deleting Compartments\)](#)
  - [TaggingWorkRequest API \(Deleting Tag Key Definitions and Namespace\)](#)
- Load Balancing:
  - [Viewing the State of a Work Request](#)
  - [WorkRequest API](#)
- Object Storage: [WorkRequest API](#)

### Request Signatures

This topic describes how to sign Oracle Cloud Infrastructure API requests.

Signing samples are included for the following:

- [Bash](#)
- [PowerShell](#)
- [C#](#)
- [Java](#)
- [NodeJS](#)
- [Perl](#)
- [PHP](#)
- [Python](#)

- [Ruby](#)
- [Go](#)

### Signature Version 1

The signature described here is *version 1* of the Oracle Cloud Infrastructure API signature. In the future, if Oracle modifies the method for signing requests, the version number will be incremented and your company will be notified.

### Required Credentials and OCIDs

You need an API signing key in the correct format. See [Required Keys and OCIDs](#).



#### Warning

##### *Client Clock Skew*

If the client's clock is skewed more than 5 minutes, a 401 (NotAuthenticated) HTTP status code is returned. This will affect your API requests. For more information, see [Maximum Allowed Client Clock Skew](#).

You also need the OCIDs for your tenancy and user. See [Where to Get the Tenancy's OCID and User's OCID](#).

### Summary of Signing Steps

In general, these are the steps required to sign a request:

1. Form the HTTPS request (SSL protocol TLS 1.2 is required).
2. Create the signing string, which is based on parts of the request.
3. Create the signature from the signing string, using your private key and the RSA-SHA256 algorithm.

4. Add the resulting signature and other required information to the `Authorization` header in the request.

See the remaining sections in this topic for details about these steps.

### Specification You Need to Be Familiar With

To learn how to perform steps 2-4 in the process above, refer to [draft-cavage-http-signatures-08](#). It's a draft specification that forms the basis for how Oracle handles request signatures. It describes generally how to form the signing string, how to create the signature, and how to add the signature and required information to the request. The remaining sections in this topic assume you're familiar with it. Important details of the Oracle Cloud Infrastructure implementation of the reference are listed in the next section.

### Special Implementation Details

The following sections describe important items to note about the Oracle Cloud Infrastructure implementation of the spec.

#### AUTHORIZATION HEADER

The Oracle Cloud Infrastructure signature uses the "Signature" Authentication scheme (with an `Authorization` header), and not the Signature HTTP header.

#### REQUIRED HEADERS

This section describes the headers that must be included in the signing string.



#### Note

*Error if Required Header is Missing*

If a required header is missing, your client will receive a 401 "Unauthorized" response.

For GET and DELETE requests (when there's no content in the request body), the signing string must include at least these headers:

- `(request-target)` (as described in [draft-cavage-http-signatures-08](#))
- `host`
- `date` or `x-date` (if both are included, Oracle uses `x-date`)

For PUT and POST requests (when there's content in the request body), the signing string must include at least these headers:

- `(request-target)`
- `host`
- `date` or `x-date` (if both are included, Oracle uses `x-date`)
- `x-content-sha256` (except for Object Storage PUT requests; see the next section)
- `content-type`
- `content-length`



### Warning

For PUT and POST requests, your client must compute the `x-content-sha256` and include it in the request and signing string, even if the body is an empty string. Also, the `content-length` is always required in the request and signing string, even if the body is empty. Some HTTP clients will not send the `content-length` if the body is empty, so you must explicitly ensure your client sends it. If `date` and `x-date` are both included, Oracle uses `x-date`. The `x-date` is used to protect against the reuse of the signed portion of the request (replay attacks).

The one exception is for Object Storage PUT requests on objects (see the next section).

### SPECIAL INSTRUCTIONS FOR OBJECT STORAGE PUT

For Object Storage [PutObject](#) and [UploadPart](#) PUT requests, the signing string must include at least these headers:

- `(request-target)`
- `host`
- `date` or `x-date` (if both are included, Oracle uses `x-date`)

If the request also includes any of the other headers that are normally required for PUT requests (see the list above), then those headers must also be included in the signing string.

### CASE AND ORDER OF HEADERS

The headers must be all lowercase in the signing string.

The order of the headers in the signing string does not matter. Just make sure to specify the order in the `headers` parameter in the `Authorization` header, as described in the [draft-cavage-http-signatures-05](#).



#### Warning

The `(request-target)` includes the path and query string from the request. Oracle expects that you will create the signing string with the query parameters in the *same order* as they appear in the request. If the request query parameters change order after signing occurs, authentication will fail.

### URL ENCODING OF PATH AND QUERY STRING

When forming the signing string, you must URL encode all parameters in the path and query string (but not the headers) according to [RFC 3986](#).

## CHAPTER 34 Developer Tools

---

### KEY IDENTIFIER

You must set `keyId="<TENANCY OCID>/<USER OCID>/<KEY FINGERPRINT>"` in the `Authorization` header that you add to the request. To get those values, see [Where to Get the Tenancy's OCID and User's OCID](#). An example `keyId` looks like this (wrapped to better fit the page):

```
ocid1.tenancy.oc1..exampleujnv47knr7uuuvqar5bshnspi6xoxsfebh3vy72fi4swgrkvuvq/ocid1.user.oc1..exampleb
a3pv6wkr4jqae5f44n2b2m2yt2j6rx32uzr4h25vqstifsfdsq/40:a4:f8:a0:40:4f:a3:2f:e0:fd:4e:b9:25:72:81:5f
```

### SIGNING ALGORITHM

The signing algorithm must be RSA-SHA256, and you must set `algorithm="rsa-sha256"` in the `Authorization` header (notice the quotation marks).

### SIGNATURE VERSION

You should include `version="1"` in the `Authorization` header (notice the quotation marks). If you do not, it's assumed that you're using whatever the current version is (which is version 1 at this time).

### EXAMPLE HEADER

Here's an example of the general syntax of the `Authorization` header (for a request with content in the body):

```
Authorization: Signature version="1",keyId="<tenancy_ocid>/<user_ocid>/<key_
fingerprint>",algorithm="rsa-sha256",headers="(request-target) date x-content-sha256 content-type
content-length",signature="Base64 (RSA-SHA256 (<signing_string>)) "
```

### Test Values

Here's an example key pair, two example requests, and the resulting `Authorization` header for each.



### Warning

The example signatures use the RSA 2048-bit keys below. Use these keys only for testing your signing code, not for sending production requests.

```
-----BEGIN PUBLIC KEY-----
MIGfMA0GCsGqGSIB3DQEBAQUAA4GNADCBiQKBgQDCFENGw33yGihy92pDjZQh10C3
6rPJj+CvfSC8+q28hxAl61QFNud13wuCTUcq0Qd2qsBe/2hFyc2DCJjG0h1L78+6
Z4UMR7EOcpfdUE9Hf3m/hs+FUR45uBJeDK1HSFHD8bHKD6kv8FPGfJTotc+2xjJw
oYi+1hqplfIekaxsyQIDAQAB
-----END PUBLIC KEY-----

-----BEGIN RSA PRIVATE KEY-----
MIICXgIBAAKBgQDCFENGw33yGihy92pDjZQh10C36rPJj+CvfSC8+q28hxAl61QF
Nud13wuCTUcq0Qd2qsBe/2hFyc2DCJjG0h1L78+6Z4UMR7EOcpfdUE9Hf3m/hs+F
UR45uBJeDK1HSFHD8bHKD6kv8FPGfJTotc+2xjJwoYi+1hqplfIekaxsyQIDAQAB
AoGBAJR8ZkCUvx5kzv+utdl7T5MnordT1TvoXXJGxK7ZZ+UuvMNUcdN2QPc4sBiA
QWvLwlcSKt5DsKZ8UETpYPy8pPYnnDEz2dDYiaew9+xEpubyew2oH4Zx71wqBtOK
kqwrXa/pzdpiucRRjk6vE6YY7EBBs/g7uanVpGibOVAEsqH1AkEA7DkjVH28WDUg
f1nqvfn2Kj6CT7nIcE3jGJsZZ7z1ZmBmHFDONMLUrXR/Zm3pR5m0tCmBqa5RK95u
4l2jtlPIwJBANJT3v8pnkth48bQo/fKel6uEYyboRtA5/uHuHkZ6FQF7OUkGogc
mSJluOdc5t6hI1VsLn0QZEjQZMEOWr+wKSMCQQCC4kXJEsHAve77oP6HtG/IiEn7
kpyUXRNvFsDE0czpJJBvL/aRFUJxurK91jhjC68sA7NsKMGg5OXb5I5Jj36xAkEA
gIT7aFOYBfWgGQAQkWNkLvySgKbAZRTELBacpHMUqd11DfdntvAyqpAZ01Y0RkmW
G6aFKaqQfOXKCYWoUiVknQJAXrlgySFci/2ueK1IE1QqIiLSZ8V80lpFLRnb1pzI
7U1yQXnTAEFYM560yJlZUpOb1V4cScGd365tiSMvxLOvTA==
-----END RSA PRIVATE KEY-----
```

The public key is stored under keyId:

```
ocid1.tenancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq/ocid1.user.oc1..aaaaaaa
at5nvwcnas5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq/73:61:a2:21:67:e0:df:be:7e:4b:93:1e:15:98:a5:b7
```

For the following GET request (line breaks inserted between query parameters for easier reading; also notice the URL encoding as mentioned earlier):

## CHAPTER 34 Developer Tools

```
GET https://iaas.us-phoenix-1.oraclecloud.com/20160918/instances
?availabilityDomain=Pjwf%3A%20PHX-AD-1
&compartmentId=ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdflvjsnog7mlr6rtdb25gilchfeyjxa
&displayName=TeamXInstances
&volumeId=ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwfaux7kcvzfs3s66izvxf2h4lgvyndsdsnoiwr5q
Date: Thu, 05 Jan 2014 21:31:40 GMT
```

The signing string would be (line breaks inserted into the (request-target) header for easier reading):

```
date: Thu, 05 Jan 2014 21:31:40 GMT
(request-target): get /20160918/instances?availabilityDomain=Pjwf%3A%20PH
X-AD-1&compartmentId=ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2i
dnccdflvjsnog7mlr6rtdb25gilchfeyjxa&displayName=TeamXInstances&
volumeId=ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwfaux7kcvzfs3s66izvxf2h
4lgvyndsdsnoiwr5q
host: iaas.us-phoenix-1.oraclecloud.com
```

The Authorization header would be:

```
Signature version="1",headers="date (request-target) host",keyId="ocid1.t
enancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq/
ocid1.user.oc1..aaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3ryn
jq/73:61:a2:21:67:e0:df:be:7e:4b:93:1e:15:98:a5:b7",algorithm="rsa-sha256
",signature="GBas7grhyrhSKHP6AVIj/h5/Vp8bd/peM79H9Wv8kjoaCivujVX1pbKLjMPe
DUhxkFIWtTtLBj3sUzaFj34XE6YZAHc9r2DmE4pMwOAY/kiITcZxaloHPOeRheC0jP2dqBT1l
8fmTZVwKZOKHYPtrLJIJQHJjNvxFWeHQjMaR7M="
```

For the following POST request:

```
POST https://iaas.us-phoenix-1.oraclecloud.com/20160918/volumeAttachments
Date: Thu, 05 Jan 2014 21:31:40 GMT
{
 "compartmentId":
"ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdflvjsnog7mlr6rtdb25gilchfeyjxa",
 "instanceId": "ocid1.instance.oc1.phx.abuw4ljrlsfiqw6vzxb43vyypt4pkodawglp3wqxjqofakrwvou52gb6s5a",
 "volumeId": "ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwfaux7kcvzfs3s66izvxf2h4lgvyndsdsnoiwr5q"
}
```

The signing string would be:

```
date: Thu, 05 Jan 2014 21:31:40 GMT
(request-target): post /20160918/volumeAttachments
```

## CHAPTER 34 Developer Tools

```
host: iaas.us-phoenix-1.oraclecloud.com
content-length: 316
content-type: application/json
x-content-sha256: V9Z20UJTvkvPJ50f1BzKE32+6m2zJjweHpDMX/U4Uy0=
```

The Authorization header would be:

```
Signature version="1",headers="date (request-target) host content-length c
ontent-type x-content-sha256",keyId="ocidl.tenancy.oc1..aaaaaaaaaba3pv6wkcr
4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq/ocidl.user.oc1..aaaaaaaat5nvwc
na5j6aqzjcaty5eqbb6qt2jvvpkanghtgdaqedqw3rynjq/73:61:a2:21:67:e0:df:be:7e:4b
:93:1e:15:98:a5:b7",algorithm="rsa-sha256",signature="Mje8vIDPlwIHmD/cTDwR
xE7HaAvBgl6JnVcsuqaNRim23fFPgQfLoOOxae6WqKbluPjYE10qIdazWaBy/M18DRhqlocMwo
SXv0fbukP8J5N80LCmzT/FFBvIvTB91XuXI3hYfP9Zt117S6ieVadHUfqBedWH0itrtPJBgKmrWso="
```

### Sample Code

This section shows the basic code for signing API requests.

#### BASH

.

```
Version: 1.0.2
Usage:
oci-curl <host> <method> [file-to-send-as-body] <request-target> [extra-curl-args]
#
ex:
oci-curl iaas.us-ashburn-1.oraclecloud.com get "/20160918/instances?compartmentId=some-compartment-ocid"
oci-curl iaas.us-ashburn-1.oraclecloud.com post ./request.json "/20160918/vcns"

function oci-curl {
 # TODO: update these values to your own
 local tenancyId="ocidl.tenancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq";
 local authUserId="ocidl.user.oc1..aaaaaaaat5nvwcna5j6aqzjcaty5eqbb6qt2jvvpkanghtgdaqedqw3rynjq";
 local keyFingerprint="20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34";
 local privateKeyPath="/Users/someuser/.oci/oci_api_key.pem";

 local alg=rsa-sha256
 local sigVersion="1"
```

## CHAPTER 34 Developer Tools

---

```
local now="$(LC_ALL=C \date -u "+%a, %d %h %Y %H:%M:%S GMT")"
local host=$1
local method=$2
local extra_args
local keyId="$tenancyId/$authUserId/$keyFingerprint"

case $method in

 "get" | "GET")
 local target=$3
 extra_args=("${@: 4}")
 local curl_method="GET";
 local request_method="get";
 ;;

 "delete" | "DELETE")
 local target=$3
 extra_args=("${@: 4}")
 local curl_method="DELETE";
 local request_method="delete";
 ;;

 "head" | "HEAD")
 local target=$3
 extra_args="--head "${@: 4}")
 local curl_method="HEAD";
 local request_method="head";
 ;;

 "post" | "POST")
 local body=$3
 local target=$4
 extra_args=("${@: 5}")
 local curl_method="POST";
 local request_method="post";
 local content_sha256="$(openssl dgst -binary -sha256 < $body | openssl enc -e -base64)";
 local content_type="application/json";
 local content_length="$(wc -c < $body | xargs)";
 ;;

 "put" | "PUT")
 local body=$3
```

## CHAPTER 34 Developer Tools

```
 local target=$4
 extra_args=("${@: 5}")
 local curl_method="PUT"
 local request_method="put"
 local content_sha256="$(openssl dgst -binary -sha256 < $body | openssl enc -e -base64)";
 local content_type="application/json";
 local content_length="$(wc -c < $body | xargs)";
 ;;

 *) echo "invalid method"; return;;
esac

This line will url encode all special characters in the request target except "/", "?", "=", and "&",
since those characters are used
in the request target to indicate path and query string structure. If you need to encode any of "/",
"?", "=", or "&", such as when
used as part of a path value or query string key or value, you will need to do that yourself in the
request target you pass in.

local escaped_target="$(echo $(rawurlencode "$target"))"
local request_target="(request-target): $request_method $escaped_target"
local date_header="date: $now"
local host_header="host: $host"
local content_sha256_header="x-content-sha256: $content_sha256"
local content_type_header="content-type: $content_type"
local content_length_header="content-length: $content_length"
local signing_string="$request_target\n$date_header\n$host_header"
local headers="(request-target) date host"
local curl_header_args
curl_header_args=(-H "$date_header")
local body_arg
body_arg=()

if ["$curl_method" = "PUT" -o "$curl_method" = "POST"]; then
 signing_string="$signing_string\n$content_sha256_header\n$content_type_header\n$content_length_header"
 headers=$headers" x-content-sha256 content-type content-length"
 curl_header_args=("${curl_header_args[@]}" -H "$content_sha256_header" -H "$content_type_header" -H
"$content_length_header")
 body_arg=(--data-binary @$body)
fi

local sig=$(printf '%b' "$signing_string" | \
```

## CHAPTER 34 Developer Tools

```
openssl dgst -sha256 -sign $privateKeyPath | \
openssl enc -e -base64 | tr -d '\n')

curl "${extra_args[@]}" "${body_arg[@]}" -X $curl_method -sS https://${host}${escaped_target} "${curl_
header_args[@]}" \
 -H "Authorization: Signature
version=\"${sigVersion}\",keyId=\"${keyId}\",algorithm=\"${alg}\",headers=\"${headers}\",signature=\"${sig}\""
}
url encode all special characters except "/", "?", "=", and "&"
function rawurlencode {
 local string="${1}"
 local strlen=${#string}
 local encoded=""
 local pos c o

 for ((pos=0 ; pos<strlen ; pos++)); do
 c=${string:$pos:1}
 case "$c" in
 [_.~a-zA-Z0-9] | "/" | "?" | "=" | "&") o="$c" ;;
 *) printf -v o '%%%02x' "$c"
 esac
 encoded+="$o"
 done

 echo "${encoded}"
}
}
```

An example of a request.json file that could be used with the preceding Bash code is shown next:

```
{
 "compartmentId": "some-compartment-id",
 "displayName": "some-vcn-display-name",
 "cidrBlock": "10.0.0.0/16"
}
```

### POWERSHELL

The following is an example for creating a request signature for an Oracle Cloud Infrastructure REST API call using a PowerShell script (`oci-rest.ps1`). The example uses the Bouncy Castle library .dll file to enable crypto functionality. Download the DLL from <https://www.bouncycastle.org> and place it in the same directory as the PowerShell script file.

## CHAPTER 34 Developer Tools

```
.

param (
 [Parameter(Mandatory=$true)][string]$method,
 [Parameter(Mandatory=$true)][string]$ocihost,
 [Parameter(Mandatory=$true)][string]$target,

 [bool]$echo_debug = $false,

 # TODO: Update these defaults or override them on the command line.
 [string]$tenancyId =
'ocidl.tenancy.oc1..aaaaaaaaaba3pv6wkr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq',
 [string]$authUserId= 'ocidl.user.oc1..aaaaaaaaalj3z3isgtuqd5uqft424he7r3cuqfr3e5gpidgnmqsxwd5qevkha',
 [string]$keyFingerprint = '29:f3:01:46:07:b8:dc:8c:16:c3:2b:b3:8d:dc:26:c5',
 [string]$privateKeyPath = $PSScriptRoot + '/oci_api_key.pem',

 [Parameter(Mandatory=$false)][string]$body,
 [Parameter(Mandatory=$false)][string]$bouncycastlelib
)

#####
This is a powershell example of how to create request signatures for an
Oracle Cloud Infra REST API call. It was modeled after the bash example.
#
Note that it utilizes the Bouncy Castle library dll for crypto functionality.
It is assumed to be in the same directory as this script,
but can be changed via commandline argument.
See https://www.bouncycastle.org for more details.
#
Usage:
oci-rest.ps1 -host <host> -method <method> -body [file-to-send-as-body] -target <request-target> -
bouncycastlelib [BouncyCastle.Crypto.dll]
#
Examples:
./oci-rest.ps1 -method get -ocihost iaas.us-ashburn-1.oraclecloud.com -target
"/20160918/instances?compartmentId=ocidl.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdf1vjsnog7mlr6rt
db25gilchfeyjxa"
./oci-rest.ps1 -method post -ocihost iaas.us-ashburn-1.oraclecloud.com -target "/20160918/vcns" -body
./request.json
#
#####
```

## CHAPTER 34 Developer Tools

```
#####
Creates a message digest for the request body.
#####
function Digest($body_file_path) {
 $sha256digest = New-Object org.bouncycastle.crypto.digests.SHA256Digest
 $content = Get-Item $body_file_path
 $bytes = $null
 if ($PSVersionTable.PSVersion.Major -ge 6) {
 [byte[]]$bytes = Get-Content $body_file_path -AsByteStream
 } else {
 [byte[]]$bytes = Get-Content $body_file_path -Encoding byte
 }
 $sha256digest.BlockUpdate($bytes, 0, $bytes.Length)
 $result_size = $sha256digest.GetDigestSize()
 $result_bytes = New-Object Byte[] $result_size
 $sha256digest.DoFinal($result_bytes, 0) | Out-Null
 $content_sha256 = [Convert]::ToBase64String($result_bytes)
 return $content_sha256
}

#####
Creates the signature to be put in the Authorization request header.
#####
function Sign($signing_string, $privateKeyPath) {
 $sha256digest = New-Object org.bouncycastle.crypto.digests.SHA256Digest
 $signer = New-Object Org.BouncyCastle.Crypto.Signers.RSADigestSigner $sha256digest

 $privateKeyFile = [System.IO.File]::OpenText($privateKeyPath)
 $pemReader = New-Object Org.BouncyCastle.OpenSsl.PemReader $privateKeyFile
 $keyPair = [Org.BouncyCastle.Crypto.AsymmetricCipherKeyPair]$pemReader.ReadObject()
 # $keyParameter = [Org.BouncyCastle.Security.DotNetUtilities]::ToRSAParameters($keyPair.Private)
 $keyParameter = $keyPair.Private
 $signer.Init($true, $keyParameter)

 $encoding = [System.Text.Encoding]::UTF8
 [byte[]]$bytes = $encoding.GetBytes($signing_string)
 $signer.BlockUpdate($bytes, 0, $bytes.Length)
 $signature = $signer.GenerateSignature()
 $signedString = [Convert]::ToBase64String($signature)

 return $signedString
}
}
```

```
#####
Makes the Oracle Cloud API REST call.
#####
function RestCall($method, $ocihost, $target, $privateKeyPath, $keyId,
 $body_file_path='', $echo_debug=$false) {
 $alg = 'rsa-sha256'
 $sigVersion = '1'
 $now = Get-Date
 $now = $now.ToUniversalTime()
 $now_string = $now.ToString("ddd, dd MMM yyyy HH:mm:ss") + " GMT"

 $content_type = ''
 $content_length = 0
 $content_sha256 = ''
 $request_method = $method.ToLower()
 If ($request_method -eq "get") {
 $method = "Get"
 } ElseIf ($request_method -eq "delete") {
 $method = "Delete"
 } ElseIf ($request_method -eq "head") {
 $method = "Head"
 } ElseIf ($request_method -eq "post") {
 if ($body_file_path.Length -eq 0) {
 echo "body parameter must be specified and point to valid json body file."
 Exit 1
 }
 $method = "Post"
 $content_type = 'application/json'
 $content_length = (Get-Item $body_file_path).length
 $content_sha256 = Digest $body_file_path
 if ($echo_debug) {
 output_debug "digest=$content_sha256"
 }
 } ElseIf ($request_method -eq "put") {
 if ($body_file_path.Length -eq 0) {
 echo "body parameter must be specified and point to valid json body file."
 Exit 1
 }
 $method = "Put"
 $content_type = 'application/json'
 $content_length = (Get-Item $body_file_path).length
 }
}
```

```

$content_sha256 = Digest $body_file_path
if ($echo_debug) {
 output_debug "digest=$content_sha256"
}
} Else {
 echo "invalid method"
 Exit 1
}

$escaped_target = rawurlencode $target
$request_target = $request_method + " " + $escaped_target
if ($echo_debug) {
 output_debug "escaped target=$escaped_target"
}

$headers = @{}
$header_list = "(request-target) date host"
$headers["date"] = $now_string
$headers["host"] = $ocihost

#$nl = [Environment]::NewLine # This doesn't work in windows environments
$nl = "`n"
$signing_string = "(request-target): " + $request_target + $nl + "date: " + $now_string + $nl + "host:
" + $ocihost

if (($request_method -eq "put") -or ($request_method -eq "post")) {
 $headers["x-content-sha256"] = $content_sha256
 $headers["content-type"] = $content_type
 $headers["content-length"] = $content_length
 $header_list = $header_list + " x-content-sha256 content-type content-length"
 $signing_string = $signing_string + $nl + "x-content-sha256: " + $content_sha256 + $nl +
 "content-type: " + $content_type + $nl + "content-length: " + $content_length
}

if ($echo_debug) {
 output_debug "signing string=$signing_string"
}
$sig = Sign $signing_string $privateKeyPath
if ($echo_debug) {
 output_debug "signature=$sig"
}
$authorization = 'Signature version="' + $sigVersion + '",keyId="' + $keyId + '",algorithm="' +

```

## CHAPTER 34 Developer Tools

```
 $alg + '",headers="' + $header_list + '",signature="' + $sig + ''
$headers["Authorization"] = $authorization

$url = "https://" + $ocihost + $escaped_target
if ($echo_debug) {
 output_debug "authorization=$authorization"
 output_debug "url=$url"
 output_debug "headers=$headers"
 $headers.GetEnumerator() | Out-String
}

Without this setting, Invoke-RestMethod was failing on windows with a connection error.
[Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12

if ($body_file_path.Length -gt 0) {
 PostPutRequest $method $url $headers $body_file_path
} else {
 Invoke-RestMethod -Method $method -Uri $url -Headers $headers | ConvertTo-Json
}
}

#####
Makes a Post or Put call.
We do this b/c Invoke-RestMethod doesn't seem to give the granular control
needed, but HttpRequest works well.
#####
Function PostPutRequest($method, $url, $headers, $body_file_path) {
 $junk = [System.Reflection.Assembly]::LoadWithPartialName("System.Web")
 $request = [System.Net.HttpWebRequest]::Create($url)
 $request.Method = $method.ToUpper()
 $request.Accept = "application/json";
 $request.ProtocolVersion = "1.0"
 $request.ContentType = $headers["content-type"]
 $request.ContentLength = $headers["content-length"]
 $request.Date = $headers["date"]
 $request.Host = $headers["host"]
 $request.Headers["x-content-sha256"] = $headers["x-content-sha256"]
 $request.Headers["authorization"] = $headers["authorization"]
 # $request.Headers["(request-target)"] = $headers["(request-target)"]

 # Create the input stream to the REST API
 $requestInputStream = $request.GetRequestStream()
}
```

```

Create a stream writer to write the json
$writer = New-Object System.IO.StreamWriter($requestInputStream)
$writer.AutoFlush = $true

Write the json
Try {
 $bytes = $null
 if ($PSVersionTable.PSVersion.Major -ge 6) {
 [byte[]]$bytes = Get-Content $body_file_path -AsByteStream
 } else {
 [byte[]]$bytes = Get-Content $body_file_path -Encoding byte
 }

 $writer.Write($bytes, 0, $bytes.Length)
} Catch [System.IO.IOException] {
 Throw "Cannot write to stream. Exception $($_.Exception)"
} Catch [System.Exception] {
 Throw "Some other weird error caught...$($_.Exception)"
} Finally {
 $writer.Close()
}
Get-WebResponseOutput $request
}

#####
Gets the response output for a request.
#####
Function Get-WebResponseOutput($request) {
 $junk = [System.Reflection.Assembly]::LoadWithPartialName("System.Web")

 $response = $null
 Try {
 $response = $request.GetResponse()
 } Catch [System.Net.WebException] {
 echo "Exception from server: " $_.Exception.Message
 $ex = $_.Exception.Response.StatusCode
 if ($response -ne $null) {
 $response.Close()
 }

 Throw "Exception from server: $ex"
 } Catch [System.Exception] {
 if ($response -ne $null) {

```

## CHAPTER 34 Developer Tools

```
 $response.Close()
 }
 echo "Some other random error: " $_.Exception.Message
 Throw "Some other random error..${($_.Exception)}"
}

$readStream = $response.GetResponseStream()
$reader = New-Object System.IO.StreamReader($readStream)

Try {
 $output = $reader.readtoend()
} Catch {
 echo "Exception reading stream from server. Exception: " $_.Exception.Message
 Throw "Exception reading stream from server. Exception: ${($_.Exception)}"
} Finally {
 if ($response -ne $null) {
 $response.Close()
 }

 $reader.Close()
}
return $output
}

#####
url encode all special characters except "/", "?", "=", and "&"
This will url encode all special characters in the request target except "/", "?", "=", and "&",
since those characters are used in the request target to indicate path and query string structure.
If you need to encode any of "/", "?", "=", or "&", such as, when used
as part of a path value or query string key or value,
you will need to do that yourself in the request target you pass in.
#####
function rawurlencode($target) {
 Add-Type -AssemblyName System.Web
 $chars_to_skip = "-_~.abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789/?=&";
 $encoded = ""
 for ($i=0; $i -lt $target.Length; $i++) {
 $ch = $target[$i]
 $o = $ch
 if ($chars_to_skip.IndexOf($ch) -lt 0) {
 $o = [System.Web.HttpUtility]::UrlEncode($ch)
 }
 }
}
```

## CHAPTER 34 Developer Tools

---

```
 $encoded = $encoded + $o
 }
 return $encoded
}

#####
Trivial function to output debug messages to console.
#####
function output_debug($msg) {
 echo "[debug] $msg"
 #Write-Verbose $msg
}

#####
Main entry point logic.
#####
if ($bouncycastlelib.Length -eq 0) {
 $bouncycastlelib = $PSScriptRoot + "/BouncyCastle.Crypto.dll"
}
if ($echo_debug) {
 [Reflection.Assembly]::LoadFile($bouncycastlelib)
 output_debug "Bouncy Castle loaded and ready"
} else {
 [Reflection.Assembly]::LoadFile($bouncycastlelib) | Out-Null
}

$keyId = $tenancyId + "/" + $authUserId + "/" + $keyFingerprint
if ($echo_debug) {
 output_debug "keyId=$keyId"
}

RestCall $method $ocihost $target $privateKeyPath $keyId $body $echo_debug
```

Following is an example of a `request.json` file that you can use with the preceding PowerShell code:

```
{
 "compartmentId": "some-compartment-id",
 "displayName": "some-vcn-display-name",
```

## CHAPTER 34 Developer Tools

---

```
"cidrBlock": "10.0.0.0/16"
}
```

### C#

```
.

// Version 1.0.1
namespace Oracle.Oci
{
 using System;
 using System.Collections.Generic;
 using System.IO;
 using System.Net;
 using System.Security.Cryptography;
 using System.Text;

 //
 // Nuget Package Manager Console: Install-Package BouncyCastle
 // Nuget CLI: nuget install BouncyCastle
 //
 using Org.BouncyCastle.Crypto;
 using Org.BouncyCastle.Crypto.Parameters;
 using Org.BouncyCastle.OpenSsl;
 using Org.BouncyCastle.Security;

 public class Signing
 {
 public static void Main(string[] args)
 {
 var tenancyId =
"ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq";
 var compartmentId =
"ocidl.compartment.oc1..aaaaaaaam3we6vgnherjq5q2idnccdfvlvsnog7mlr6rtdb25gilchfeyjxa";
 var userId = "ocidl.user.oc1..aaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq";
 var fingerprint = "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34";
 var privateKeyPath = "private.pem";
 var privateKeyPassphrase = "password";

 var signer = new RequestSigner(tenancyId, userId, fingerprint, privateKeyPath,
privateKeyPassphrase);

 // Oracle Cloud Infrastructure APIs require TLS 1.2

```

## CHAPTER 34 Developer Tools

---

```
// uncomment the line below if targeting < .NET Framework 4.6 (HttpWebRequest does not
enable TLS 1.2 by default in earlier versions)
// ServicePointManager.SecurityProtocol = SecurityProtocolType.Tls12;

// GET with query parameters (gets user)
var uri = new Uri($"https://identity.us-phoenix-1.oraclecloud.com/20160918/users/{userId}");
var request = (HttpWebRequest)WebRequest.Create(uri);
request.Method = "GET";
request.Accept = "application/json";

signer.SignRequest(request);

Console.WriteLine($"Authorization header: {request.Headers["authorization"]}");

ExecuteRequest(request);

// POST with body (creates a VCN)
uri = new Uri($"https://iaas.us-phoenix-1.oraclecloud.com/20160918/vcns");
var body = string.Format(@"{{"cidrBlock" : "10.0.0.0/16","compartmentId" : ""
{0}","displayName" : "MyVcn"}}", compartmentId);
var bytes = Encoding.UTF8.GetBytes(body);

request = (HttpWebRequest)WebRequest.Create(uri);
request.Method = "POST";
request.Accept = "application/json";
request.ContentType = "application/json";
request.Headers["x-content-sha256"] = Convert.ToBase64String(SHA256.Create().ComputeHash
(bytes));

using (var stream = request.GetRequestStream())
{
 stream.Write(bytes, 0, bytes.Length);
}

signer.SignRequest(request);

Console.WriteLine($"Authorization header: {request.Headers["authorization"]}");

ExecuteRequest(request);
```

```
// GET with query parameters (gets namespace)
uri = new Uri($"https://objectstorage.us-phoenix-1.oraclecloud.com/n/");
request = (HttpWebRequest)WebRequest.Create(uri);
request.Method = "GET";
request.Accept = "application/json";

signer.SignRequest(request);

Console.WriteLine($"Authorization header: {request.Headers["authorization"]}");

string namespaceName = ExecuteRequest(request);

namespaceName = JsonConvert.DeserializeObject<string>(namespaceName);

// POST with body (creates a bucket)
uri = new Uri($"https://objectstorage.us-phoenix-1.oraclecloud.com/n/{namespaceName}/b/");
body = string.Format(@"{{"name" : "bucket01","compartmentId" : ""
{0}","publicAccessType" : "ObjectRead"}}", compartmentId);
bytes = Encoding.UTF8.GetBytes(body);

request = (HttpWebRequest)WebRequest.Create(uri);
request.Method = "POST";
request.Accept = "application/json";
request.ContentType = "application/json";
request.Headers["x-content-sha256"] = Convert.ToBase64String(SHA256.Create().ComputeHash
(bytes));

using (var stream = request.GetRequestStream())
{
 stream.Write(bytes, 0, bytes.Length);
}

signer.SignRequest(request);

Console.WriteLine($"Authorization header: {request.Headers["authorization"]}");

ExecuteRequest(request);

// PUT with body (puts a object)
uri = new Uri($"https://objectstorage.us-phoenix-1.oraclecloud.com/n/
{namespaceName}/b/bucket01/o/object01");
```

## CHAPTER 34 Developer Tools

---

```
body = "Hello Object Storage Service!!!";
bytes = Encoding.UTF8.GetBytes (body);

request = (HttpWebRequest)WebRequest.Create (uri);
request.Method = "PUT";
request.Accept = "application/json";
request.ContentType = "application/json";

using (var stream = request.GetRequestStream())
{
 stream.Write (bytes, 0, bytes.Length);
}

signer.SignRequest (request, true);

Console.WriteLine ($"Authorization header: {request.Headers["authorization"]}");

ExecuteRequest (request);

// POST with body (create multipart upload)
uri = new Uri ($"https://objectstorage.us-phoenix-1.oraclecloud.com/n/
{namespaceName}/b/bucket01/u");
body = "{\"object\" : \"object02\"}";
bytes = Encoding.UTF8.GetBytes (body);

request = (HttpWebRequest)WebRequest.Create (uri);
request.Method = "POST";
request.Accept = "application/json";
request.ContentType = "application/json";
request.Headers["x-content-sha256"] = Convert.ToBase64String (SHA256.Create ().ComputeHash
(bytes));

using (var stream = request.GetRequestStream())
{
 stream.Write (bytes, 0, bytes.Length);
}

signer.SignRequest (request);

Console.WriteLine ($"Authorization header: {request.Headers["authorization"]}");

ExecuteRequest (request);
```

```
 Console.ReadKey();
 }

 private static string ExecuteRequest(HttpWebRequest request)
 {
 try
 {
 var webResponse = (HttpWebResponse)request.GetResponse();
 var response = new StreamReader(webResponse.GetResponseStream()).ReadToEnd();
 Console.WriteLine($"Response: {response}");

 return response;
 }
 catch (WebException e)
 {
 Console.WriteLine($"Exception occurred: {e.Message}");
 Console.WriteLine($"Response: {new StreamReader(e.Response.GetResponseStream()
 ().ReadToEnd())}");

 return String.Empty;
 }
 }

 public class RequestSigner
 {
 private static readonly IDictionary<string, List<string>> RequiredHeaders = new
 Dictionary<string, List<string>>
 {
 { "GET", new List<string>{"date", "(request-target)", "host" }},
 { "HEAD", new List<string>{"date", "(request-target)", "host" }},
 { "DELETE", new List<string>{"date", "(request-target)", "host" }},
 { "PUT", new List<string>{"date", "(request-target)", "host", "content-length",
 "content-type", "x-content-sha256" }},
 { "POST", new List<string>{"date", "(request-target)", "host", "content-length",
 "content-type", "x-content-sha256" }},
 { "PUT-LESS", new List<string>{"date", "(request-target)", "host" }}
 };

 private readonly string keyId;
 private readonly ISigner signer;
 }
}
```

## CHAPTER 34 Developer Tools

```
 /// <summary>
 /// Adds the necessary authorization header for signed requests to Oracle Cloud
Infrastructure services.
 /// Documentation for request signatures can be found here:
https://docs.cloud.oracle.com/Content/API/Concepts/signingrequests.htm
 /// </summary>
 /// <param name="tenancyId">The tenancy OCID</param>
 /// <param name="userId">The user OCID</param>
 /// <param name="fingerprint">The fingerprint corresponding to the provided key</param>
 /// <param name="privateKeyPath">Path to a PEM file containing a private key</param>
 /// <param name="privateKeyPassphrase">An optional passphrase for the private key</param>
public RequestSigner(string tenancyId, string userId, string fingerprint, string
privateKeyPath, string privateKeyPassphrase="")
{
 // This is the keyId for a key uploaded through the console
 this.keyId = $"{tenancyId}/{userId}/{fingerprint}";

 AsymmetricCipherKeyPair keyPair;
 using (var fileStream = File.OpenText(privateKeyPath))
 {
 try {
 keyPair = (AsymmetricCipherKeyPair)new PemReader(fileStream, new Password
(privateKeyPassphrase.ToCharArray())).ReadObject();
 }
 catch (InvalidCipherTextException) {
 throw new ArgumentException("Incorrect passphrase for private key");
 }
 }

 RsaKeyParameters privateKeyParams = (RsaKeyParameters)keyPair.Private;
 this.signer = SignerUtilities.GetSigner("SHA-256withRSA");
 this.signer.Init(true, privateKeyParams);
}

public void SignRequest(HttpWebRequest request, bool useLessHeadersForPut = false)
{
 if (request == null) { throw new ArgumentNullException(nameof(request)); }

 // By default, request.Date is DateTime.MinValue, so override to DateTime.UtcNow, but
preserve the value if caller has already set the Date
 if (request.Date == DateTime.MinValue) { request.Date = DateTime.UtcNow; }
```

## CHAPTER 34 Developer Tools

---

```
var requestMethodUpper = request.Method.ToUpperInvariant();
var requestMethodKey = useLessHeadersForPut ? requestMethodUpper + "-LESS" :
requestMethodUpper;

List<string> headers;
if (!RequiredHeaders.TryGetValue(requestMethodKey, out headers)) {
 throw new ArgumentException($"Don't know how to sign method: {request.Method}");
}

// for PUT and POST, if the body is empty we still must explicitly set content-length =
0 and x-content-sha256
// the caller may already do this, but we shouldn't require it since we can determine it
here
if (request.ContentLength <= 0 && (string.Equals(requestMethodUpper, "POST") ||
string.Equals(requestMethodUpper, "PUT")))
{
 request.ContentLength = 0;
 request.Headers["x-content-sha256"] = Convert.ToBase64String(SHA256.Create
().ComputeHash(new byte[0]));
}

var signingStringBuilder = new StringBuilder();
var newline = string.Empty;
foreach (var headerName in headers)
{
 string value = null;
 switch (headerName)
 {
 case "(request-target)":
 value = buildRequestTarget(request);
 break;
 case "host":
 value = request.Host;
 break;
 case "content-length":
 value = request.ContentLength.ToString();
 break;
 default:
 value = request.Headers[headerName];
 break;
 }
}
```

## CHAPTER 34 Developer Tools

```
 if (value == null) { throw new ArgumentException($"Request did not contain required
header: {headerName}"); }
 signingStringBuilder.Append(newline).Append($" {headerName}: {value}");
 newline = "\n";
 }

 // generate signature using the private key
 var bytes = Encoding.UTF8.GetBytes(signingStringBuilder.ToString());
 this.signer.BlockUpdate(bytes, 0, bytes.Length);
 var signature = Convert.ToBase64String(this.signer.GenerateSignature());
 var authorization = $"Signature version=""1"",headers=""{string.Join(" ",
headers)}"",keyId=""{keyId}"",algorithm=""rsa-sha256"",signature=""{signature}""";
 request.Headers["authorization"] = authorization;
}

private static string buildRequestTarget (HttpWebRequest request)
{
 // ex. get /20160918/instances
 return $"{request.Method.ToLowerInvariant()} {request.RequestUri.PathAndQuery}";
}

}

/// <summary>
/// Implements Bouncy Castle's IPasswordFinder interface to allow opening password protected
private keys.
/// </summary>
public class Password : IPasswordFinder
{
 private readonly char[] password;

 public Password(char[] password) { this.password = password; }

 public char[] GetPassword() { return (char[])password.Clone(); }
}
}
}
```

### JAVA

This sample omits the optional `version` field in the Authorization header.

```
/*
 * @version 1.0.1
```

## CHAPTER 34 Developer Tools

---

```
*
<dependency>
 <groupId>com.google.guava</groupId>
 <artifactId>guava</artifactId>
 <version>19.0</version>
</dependency>
*/
import com.google.common.collect.ImmutableList;
import com.google.common.collect.ImmutableMap;
import com.google.common.hash.Hashing;
/*
<dependency>
 <groupId>org.apache.httpcomponents</groupId>
 <artifactId>httpcore</artifactId>
 <version>4.4.1</version>
</dependency>

(or transitively via org.apache.httpcomponents:httpclient:jar:4.5)
*/
import org.apache.http.HttpEntity;

/*
<dependency>
 <groupId>org.apache.httpcomponents</groupId>
 <artifactId>httpclient</artifactId>
 <version>4.5</version>
</dependency>
*/
import org.apache.http.client.methods.HttpEntityEnclosingRequestBase;
import org.apache.http.client.methods.HttpGet;
import org.apache.http.client.methods.HttpPost;
import org.apache.http.client.methods.HttpRequestBase;
import org.apache.http.entity.ByteArrayEntity;

/*
<dependency>
 <groupId>org.tomitribe</groupId>
 <artifactId>tomitribe-http-signatures</artifactId>
 <version>1.0</version>
</dependency>
*/
import org.apache.http.entity.StringEntity;
import org.tomitribe.auth.signatures.MissingRequiredHeaderException;
```

## CHAPTER 34 Developer Tools

---

```
import org.tomitribe.auth.signatures.PEM;
import org.tomitribe.auth.signatures.Signature;
import org.tomitribe.auth.signatures.Signer;

import java.io.ByteArrayOutputStream;
import java.io.IOException;
import java.io.InputStream;
import java.io.UnsupportedEncodingException;
import java.net.URI;
import java.nio.charset.StandardCharsets;
import java.nio.file.Files;
import java.nio.file.Paths;
import java.security.Key;
import java.security.PrivateKey;
import java.security.spec.InvalidKeySpecException;
import java.text.SimpleDateFormat;
import java.util.*;
import java.util.stream.Collectors;

/**
 * This example creates a {@link RequestSigner}, then prints out the Authorization header
 * that is inserted into the HttpGet object.
 *
 * <p>
 * apiKey is the identifier for a key uploaded through the console.
 * privateKeyFilename is the location of your private key (that matches the uploaded public key for
 * apiKey).
 * </p>
 *
 * The signed HttpGet request is not executed, since instanceId does not map to a real instance.
 */
public class Signing {
 public static void main(String[] args) throws UnsupportedEncodingException {
 HttpRequestBase request;

 // This is the keyId for a key uploaded through the console
 String apiKey =
("ocid1.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq/"
 +
"ocid1.user.oc1..aaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq/"
 + "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34");
```

## CHAPTER 34 Developer Tools

---

```
String privateKeyFilename = "../sample-private-key";
PrivateKey privateKey = loadPrivateKey(privateKeyFilename);
RequestSigner signer = new RequestSigner(apiKey, privateKey);

// GET with query parameters
String uri = "https://iaas.us-ashburn-
1.oraclecloud.com/20160918/instances?availabilityDomain=%s&compartmentId=%s&displayName=%s&volumeId=%s";
uri = String.format(uri,
 "Pjwf%3A%20PHX-AD-1",
 // Older ocid formats included ":" which must be escaped
"ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdf1vjsnog7mlr6rtdb25gilchfeyjxa".replace(":",
"%3A"),
 "TeamXInstances",
"ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwoffaux7kcvzfs3s66izvxf2h41gvyndsdsnoiwr5q".replace(":", "%3A")
);

request = new HttpGet(uri);
// Uncomment to use a fixed date
// request.setHeader("Date", "Thu, 05 Jan 2014 21:31:40 GMT");

signer.signRequest(request);
System.out.println(uri);
System.out.println(request.getFirstHeader("Authorization"));

// POST with body
uri = "https://iaas.us-ashburn-1.oraclecloud.com/20160918/volumeAttachments";
request = new HttpPost(uri);
// Uncomment to use a fixed date
// request.setHeader("Date", "Thu, 05 Jan 2014 21:31:40 GMT");
HttpEntity entity = new StringEntity("{\n" +
 " \"compartmentId\":
\n\"ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdf1vjsnog7mlr6rtdb25gilchfeyjxa\", \n" +
 " \"instanceId\":
\n\"ocid1.instance.oc1.phx.abuw4ljrlsfiqw6vzzxb43vyypt4pkodawglp3wqxjqofakrwvou52gb6s5a\", \n" +
 " \"volumeId\":
\n\"ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwoffaux7kcvzfs3s66izvxf2h41gvyndsdsnoiwr5q\" \n" +
 "}");
((HttpPost) request).setEntity(entity);
```

## CHAPTER 34 Developer Tools

---

```
signer.signRequest(request);
System.out.println("\n" + uri);
System.out.println(request.getFirstHeader("Authorization"));

}

/**
 * Load a {@link PrivateKey} from a file.
 */
private static PrivateKey loadPrivateKey(String privateKeyFilename) {
 try (InputStream privateKeyStream = Files.newInputStream(Paths.get(privateKeyFilename))) {
 return PEM.readPrivateKey(privateKeyStream);
 } catch (InvalidKeySpecException e) {
 throw new RuntimeException("Invalid format for private key");
 } catch (IOException e) {
 throw new RuntimeException("Failed to load private key");
 }
}

/**
 * A light wrapper around https://github.com/tomitribe/http-signatures-java
 */
public static class RequestSigner {
 private static final SimpleDateFormat DATE_FORMAT;
 private static final String SIGNATURE_ALGORITHM = "rsa-sha256";
 private static final Map<String, List<String>> REQUIRED_HEADERS;

 static {
 DATE_FORMAT = new SimpleDateFormat("EEE, dd MMM yyyy HH:mm:ss zzz", Locale.US);
 DATE_FORMAT.setTimeZone(TimeZone.getTimeZone("GMT"));
 REQUIRED_HEADERS = ImmutableMap.<String, List<String>>builder()
 .put("get", ImmutableList.of("date", "(request-target)", "host"))
 .put("head", ImmutableList.of("date", "(request-target)", "host"))
 .put("delete", ImmutableList.of("date", "(request-target)", "host"))
 .put("put", ImmutableList.of("date", "(request-target)", "host", "content-length",
"content-type", "x-content-sha256"))
 .put("post", ImmutableList.of("date", "(request-target)", "host", "content-length",
"content-type", "x-content-sha256"))
 .build();
 }

 private final Map<String, Signer> signers;

 /**
```

## CHAPTER 34 Developer Tools

```
* @param apiKey The identifier for a key uploaded through the console.
* @param privateKey The private key that matches the uploaded public key for the given apiKey.
*/
public RequestSigner(String apiKey, Key privateKey) {
 this.signers = REQUIRED_HEADERS
 .entrySet().stream()
 .collect(Collectors.toMap(
 entry -> entry.getKey(),
 entry -> buildSigner(apiKey, privateKey, entry.getKey())));
}

/**
 * Create a {@link Signer} that expects the headers for a given method.
 * @param apiKey The identifier for a key uploaded through the console.
 * @param privateKey The private key that matches the uploaded public key for the given apiKey.
 * @param method HTTP verb for this signer
 * @return
 */
protected Signer buildSigner(String apiKey, Key privateKey, String method) {
 final Signature signature = new Signature(
 apiKey, SIGNATURE_ALGORITHM, null, REQUIRED_HEADERS.get(method.toLowerCase()));
 return new Signer(privateKey, signature);
}

/**
 * Sign a request, optionally including additional headers in the signature.
 *
 *
 * If missing, insert the Date header (RFC 2822).
 * If PUT or POST, insert any missing content-type, content-length, x-content-sha256
 * Verify that all headers to be signed are present.
 * Set the request's Authorization header to the computed signature.
 *
 *
 * @param request The request to sign
 */
public void signRequest(HttpRequestBase request) {
 final String method = request.getMethod().toLowerCase();
 // nothing to sign for options
 if (method.equals("options")) {
 return;
 }
}
```

```

final String path = extractPath(request.getURI());

// supply date if missing
if (!request.containsHeader("date")) {
 request.addHeader("date", DATE_FORMAT.format(new Date()));
}

// supply host if missing
if (!request.containsHeader("host")) {
 request.addHeader("host", request.getURI().getHost());
}

// supply content-type, content-length, and x-content-sha256 if missing (PUT and POST only)
if (method.equals("put") || method.equals("post")) {
 if (!request.containsHeader("content-type")) {
 request.addHeader("content-type", "application/json");
 }
 if (!request.containsHeader("content-length") || !request.containsHeader("x-content-
sha256")) {
 byte[] body = getRequestBody((HttpEntityEnclosingRequestBase) request);
 if (!request.containsHeader("content-length")) {
 request.addHeader("content-length", Integer.toString(body.length));
 }
 if (!request.containsHeader("x-content-sha256")) {
 request.addHeader("x-content-sha256", calculateSHA256(body));
 }
 }
}

final Map<String, String> headers = extractHeadersToSign(request);
final String signature = this.calculateSignature(method, path, headers);
request.setHeader("Authorization", signature);
}

/**
 * Extract path and query string to build the (request-target) pseudo-header.
 * For the URI "http://www.host.com/somePath?example=path" return "/somePath?example=path"
 */
private static String extractPath(URI uri) {
 String path = uri.getRawPath();
 String query = uri.getRawQuery();

```

```

 if (query != null && !query.trim().isEmpty()) {
 path = path + "?" + query;
 }
 return path;
 }

/**
 * Extract the headers required for signing from a {@link HttpRequestBase}, into a Map
 * that can be passed to {@link RequestSigner#calculateSignature}.
 *
 * <p>
 * Throws if a required header is missing, or if there are multiple values for a single header.
 * </p>
 *
 * @param request The request to extract headers from.
 */
private static Map<String, String> extractHeadersToSign(HttpRequestBase request) {
 List<String> headersToSign = REQUIRED_HEADERS.get(request.getMethod().toLowerCase());
 if (headersToSign == null) {
 throw new RuntimeException("Don't know how to sign method " + request.getMethod());
 }
 return headersToSign.stream()
 // (request-target) is a pseudo-header
 .filter(header -> !header.toLowerCase().equals("(request-target)"))
 .collect(Collectors.toMap(
 header -> header,
 header -> {
 if (!request.containsHeader(header)) {
 throw new MissingRequiredHeaderException(header);
 }
 if (request.getHeaders(header).length > 1) {
 throw new RuntimeException(
 String.format("Expected one value for header %s", header));
 }
 return request.getFirstHeader(header).getValue();
 }
));
}

/**
 * Wrapper around {@link Signer#sign}, returns the {@link Signature} as a String.
 *
 * @param method Request method (GET, POST, ...)

```

```

 * @param path The path + query string for forming the (request-target) pseudo-header
 * @param headers Headers to include in the signature.
 */
private String calculateSignature(String method, String path, Map<String, String> headers) {
 Signer signer = this.signers.get(method);
 if (signer == null) {
 throw new RuntimeException("Don't know how to sign method " + method);
 }
 try {
 return signer.sign(method, path, headers).toString();
 } catch (IOException e) {
 throw new RuntimeException("Failed to generate signature", e);
 }
}

/**
 * Calculate the Base64-encoded string representing the SHA256 of a request body
 * @param body The request body to hash
 */
private String calculateSHA256(byte[] body) {
 byte[] hash = Hashing.sha256().hashBytes(body).asBytes();
 return Base64.getEncoder().encodeToString(hash);
}

/**
 * Helper to safely extract a request body. Because an {@link HttpEntity} may not be
repeatable,
 * this function ensures the entity is reset after reading. Null entities are treated as an
empty string.
 *
 * @param request A request with a (possibly null) {@link HttpEntity}
 */
private byte[] getRequestBody(HttpEntityEnclosingRequestBase request) {
 HttpEntity entity = request.getEntity();
 // null body is equivalent to an empty string
 if (entity == null) {
 return "".getBytes(StandardCharsets.UTF_8);
 }
 // May need to replace the request entity after consuming
 boolean consumed = !entity.isRepeatable();
 ByteArrayOutputStream content = new ByteArrayOutputStream();
 try {

```

## CHAPTER 34 Developer Tools

---

```
 entity.writeTo(content);
 } catch (IOException e) {
 throw new RuntimeException("Failed to copy request body", e);
 }
 // Replace the now-consumed body with a copy of the content stream
 byte[] body = content.toByteArray();
 if (consumed) {
 request.setEntity(new ByteArrayEntity(body));
 }
 return body;
}
}
```

### NODEJS

```
/*
 Version 1.0.1
 Before running this example, install necessary dependencies by running:
 npm install http-signature jssha
*/

var fs = require('fs');
var https = require('https');
var os = require('os');
var httpSignature = require('http-signature');
var jsSHA = require("jssha");

// TODO: update these values to your own
var tenancyId = "ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq";
var authUserId = "ocidl.user.oc1..aaaaaaaat5nvwcna5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq";
var keyFingerprint = "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34";
var privateKeyPath = "~/oci/oci_api_key.pem";

var identityDomain = "identity.us-ashburn-1.oraclecloud.com";
var coreServicesDomain = "iaas.us-ashburn-1.oraclecloud.com";

if(privateKeyPath.indexOf("~/") === 0) {
 privateKeyPath = privateKeyPath.replace("~/", os.homedir())
}
}
```

## CHAPTER 34 Developer Tools

---

```
var privateKey = fs.readFileSync(privateKeyPath, 'ascii');

// signing function as described at
https://docs.cloud.oracle.com/Content/API/Concepts/signingrequests.htm
 function sign(request, options) {

 var apiKeyId = options.tenancyId + "/" + options.userId + "/" + options.keyFingerprint;

 var headersToSign = [
 "host",
 "date",
 "(request-target)"
];

 var methodsThatRequireExtraHeaders = ["POST", "PUT"];

 if(methodsThatRequireExtraHeaders.indexOf(request.method.toUpperCase()) !== -1) {
 options.body = options.body || "";

 var shaObj = new jsSHA("SHA-256", "TEXT");
 shaObj.update(options.body);

 request.setHeader("Content-Length", options.body.length);
 request.setHeader("x-content-sha256", shaObj.getHash('B64'));

 headersToSign = headersToSign.concat([
 "content-type",
 "content-length",
 "x-content-sha256"
]);
 }

 httpSignature.sign(request, {
 key: options.privateKey,
 keyId: apiKeyId,
 headers: headersToSign
 });

 var newAuthHeaderValue = request.getHeader("Authorization").replace("Signature ", "Signature
version=\"1\",");
 request.setHeader("Authorization", newAuthHeaderValue);
 }
}
```

## CHAPTER 34 Developer Tools

---

```
}

// generates a function to handle the https.request response object
function handleRequest(callback) {

 return function(response) {
 var responseBody = "";

 response.on('data', function(chunk) {
 responseBody += chunk;
 });

 response.on('end', function() {
 callback(JSON.parse(responseBody));
 });
 }
}

// gets the user with the specified id
function getUser(userId, callback) {

 var options = {
 host: identityDomain,
 path: "/20160918/users/" + encodeURIComponent(userId),
 };

 var request = https.request(options, handleRequest(callback));

 sign(request, {
 privateKey: privateKey,
 keyFingerprint: keyFingerprint,
 tenancyId: tenancyId,
 userId: authUserId
 });

 request.end();
};

// creates a Oracle Cloud Infrastructure VCN in the specified compartment
function createVCN(compartmentId, displayName, cidrBlock, callback) {

 var body = JSON.stringify({
```

```
 compartmentId: compartmentId,
 displayName: displayName,
 cidrBlock: cidrBlock
 });

 var options = {
 host: coreServicesDomain,
 path: '/20160918/vcns',
 method: 'POST',
 headers: {
 "Content-Type": "application/json",
 }
 };

 var request = https.request(options, handleRequest(callback));

 sign(request, {
 body: body,
 privateKey: privateKey,
 keyFingerprint: keyFingerprint,
 tenancyId: tenancyId,
 userId: authUserId
 });

 request.end(body);
};

// test the above functions
console.log("GET USER:");

getUser(authUserId, function(data) {
 console.log(data);

 console.log("\nCREATING VCN:");

 // TODO: replace this with a compartment you have access to
 var compartmentIdToCreateVcnIn = tenancyId;

 createVCN(compartmentIdToCreateVcnIn, "Test-VCN", "10.0.0.0/16", function(data) {
 console.log(data);
 });
});
});
```

### PERL

This sample omits the optional `version` field in the Authorization header.

```
#!/usr/bin/perl
Version 1.0.1
Need the following:
Modules - Authen::HTTP::Signature, DateTime, DateTime::Format::HTTP, Mozilla::CA, File::Slurp,
LWP::UserAgent, LWP::Protocol::https
LWP::UserAgent and LWP::Protocol::https >= 6.06
OpenSSL >= 1.0.1

use strict;
use warnings;

{
 package OCISigner;

 use Authen::HTTP::Signature;
 use Digest::SHA qw(sha256_base64);
 use DateTime;
 use DateTime::Format::HTTP;

 my @generic_headers = (
 'date', '(request-target)', 'host'
);
 my @body_headers = (
 'content-length', 'content-type', 'x-content-sha256'
);
 my @all_headers = (@generic_headers, @body_headers);
 my %required_headers = (
 get => \@generic_headers,
 delete => \@generic_headers,
 head => \@generic_headers,
 post => \@all_headers,
 put => \@all_headers
);

 sub new {
 my ($class, $api_key, $private_key) = @_;
 my $self = {
 _api_key => $api_key,
 _private_key => $private_key
```

```
};
bless $self, $class;
return $self;
}

sub sign_request {
 my ($self, $request) = @_;
 my $verb = lc($request->method);
 my $sign_body = grep(/^$verb$/, ('post', 'put'));
 $self->inject_missing_headers($request, $sign_body);
 my $headers = $required_headers{$verb};

 my $all_auth = Authen::HTTP::Signature->new(
 key => $self->{_private_key},
 request => $request,
 key_id => $self->{_api_key},
 headers => $headers,
);
 $all_auth->sign();
}

sub inject_missing_headers {
 my ($self, $request, $sign_body) = @_;
 $request->header('content-type', 'application/json') unless $request->header('content-type');
 $request->header('accept', '*//*') unless $request->header('accept');
 my $class = 'DateTime::Format::HTTP';
 $request->header('date', $class->format_datetime(DateTime->now)) unless $request->header('date');

 $request->header('host', $request->uri->host) unless $request->header('host');
 if ($sign_body) {
 $request->content('') unless $request->content;
 $request->header('content-length', length($request->content)) unless $request->header('content-length');
 $request->header('x-content-sha256', $self->compute_sha256($request->content)) unless $request->header('x-content-sha256');
 }
}

sub compute_sha256 {
 my ($self, $content) = @_;
 my $digest = sha256_base64($content);
 while (length($digest) % 4) {
```

## CHAPTER 34 Developer Tools

---

```
 $digest .= '=';
 }
 return $digest;
}
} # OCISigner

{
package OCIClient;

use LWP::UserAgent;
use Mozilla::CA;

sub new {
 my ($class, $api_key, $private_key) = @_;
 my $ua = LWP::UserAgent->new;
 $ua->ssl_opts(
 verify_hostname => 1,
 SSL_ca_file => Mozilla::CA::SSL_ca_file()
);
 my $self = {
 _signer => OCISigner->new($api_key, $private_key),
 _ua => $ua
 };
 bless $self, $class;
 return $self;
}

sub make_request {
 my ($self, $request) = @_;
 print "Sending request\n";
 $self->{_signer}->sign_request($request);

 my $response = $self->{_ua}->request($request);
 if ($response->is_success) {
 my $message = $response->decoded_content;
 print "Received reply: $message\n";
 }
 else {
 print "HTTP GET error code: ", $response->code, "\n";
 print "HTTP GET error message: ", $response->message, "\n";
 }
}
}
```

## CHAPTER 34 Developer Tools

```
} # OCIClient

use File::Slurp qw(read_file);

my $api_key = "ocid1.tenancy.oc1..aaaaaaaaaba3pv6wkr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq/" .
 "ocid1.user.oc1..aaaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq/" .
 "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34";

my $private_key = read_file('../sample-private-key') or die $!;

my $client = OCIClient->new($api_key, $private_key);

Uncomment to use a fixed date
#my $fixed_date = 'Thu, 05 Jan 2014 21:31:40 GMT';
my $fixed_date;

GET with query parameters
Note: Older ocid formats included ":" which must be escaped
my %query_args = (
 availability_domain => "Pjwf%3A%20PHX-AD-1",
 compartment_id =>
"ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdflvjsnog7mlr6rtdb25gilchfeyjxa",
 display_name => "TeamXInstances",
 volume_id => "ocid1.volume.oc1.phx.abyhqljrgvttnlx73nmrwwfaux7kcvzfs3s66izvxf2h4lgyvndsdsnoiwr5q"
);

my $uri = "https://iaas.us-phoenix-1.oraclecloud.com/20160918/instances?availabilityDomain=" .
 $query_args{availability_domain} .
 "&compartmentId=" .
 $query_args{compartment_id} .
 "&displayName=" .
 $query_args{display_name} .
 "&volumeId=" .
 $query_args{volume_id};

my $request = HTTP::Request->new(GET => $uri);
$request->header('date', $fixed_date) if $fixed_date;
$client->make_request($request);

POST with body
$uri = "https://iaas.us-ashburn-1.oraclecloud.com/20160918/volumeAttachments";
my $body = q|{
 "compartmentId":
"ocid1.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdflvjsnog7mlr6rtdb25gilchfeyjxa",
```

## CHAPTER 34 Developer Tools

---

```
"instanceId": "ocidl.instance.oc1.phx.abuw41jr1sfiqw6vzzxb43vyypt4pkodawglp3wqxjqofakrwwou52gb6s5a",
"volumeId": "ocidl.volume.oc1.phx.abyhqljrgvttnlx73nmrfaux7kcvzfs3s66izvxf2h41lgvyndsdsnoiwr5q"
}||;
$request = HTTP::Request->new(POST => $uri);
$request->header('date', $fixed_date) if $fixed_date;
$request->content($body);
$client->make_request($request);
```

### PHP

.

```
<? php

// Version 1.0.0
//
// Dependencies:
// - PHP curl extension
// - Guzzle (composer require guzzlehttp/guzzle)
//

require 'vendor/autoload.php';

use Psr\Http\Message\RequestInterface;
use GuzzleHttp\HandlerStack;
use GuzzleHttp\Handler\CurlHandler;
use GuzzleHttp\Client;
use GuzzleHttp\Middleware;

// TODO: Update these for your tenancy
$tenancy_id = 'ocidl.tenancy.oc1..aaaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq';
$user_id = 'ocidl.user.oc1..aaaaaaat5nvwncna5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq';
$thumbprint = '20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34';
$region = 'us-phoenix-1';
$key_location = 'file://private.pem';
$key_passphrase = 'password';

$namespace = 'MyNamespace';
$bucket_name = 'MyBucket';
$file_to_upload = 'myfile.txt';

$key_id = "$tenancy_id/$user_id/$thumbprint";
```

## CHAPTER 34 Developer Tools

---

```
function sign_string($data, $key_path, $passphrase){
 $pkeyid = openssl_pkey_get_private($key_path, $passphrase);
 if (!$pkeyid) {
 exit('Error reading private key');
 }

 openssl_sign($data, $signature, $pkeyid, OPENSSL_ALGO_SHA256);

 // free the key from memory
 openssl_free_key($pkeyid);

 return base64_encode($signature);
}

function oci_signer_middleware(RequestInterface $request) {
 global $key_id;
 global $key_location;
 global $key_passphrase;

 // headers required for all HTTP verbs
 $headers = "date (request-target) host";

 // example: Thu, 05 Jan 2014 21:31:40 GMT
 $date=gmtime("D, d M Y H:i:s T", time());
 $method = strtolower($request->getMethod());
 $request_target = $request->getRequestTarget();
 $host = $request->getHeader('Host')[0];

 $request = $request->withHeader('Date', $date);

 $signing_string = "date: $date\n(request-target): $method $request_target\nhost: $host";

 // additional required headers for POST and PUT requests
 if ($method == 'post' || $method == 'put') {
 $content_length = $request->getHeader('Content-Length')[0];

 // if content length is 0 we still need to explicitly send the Content-Length header
 if (!$content_length){
 $content_length = 0;
 $request = $request->withHeader('Content-Length', 0);
 }
 }
}
```

## CHAPTER 34 Developer Tools

```
$content_type = $request->getHeader('Content-Type')[0];
$content_sha256 = base64_encode(hex2bin(hash("sha256", $request->getBody())));

$request = $request->withHeader('x-content-sha256', $content_sha256);

$headers = $headers . " content-length content-type x-content-sha256";
$signing_string = $signing_string . "\ncontent-length: $content_length\ncontent-type: $content_
type\nx-content-sha256: $content_sha256";
}

echo "Signing string:\n$signing_string".PHP_EOL;

$signature = sign_string($signing_string, $key_location, $key_passphrase);

$authorization_header = "Signature version=\"1\",keyId=\"$key_id\",algorithm=\"rsa-
sha256\",headers=\"$headers\",signature=\"$signature\"";
$request = $request->withHeader('Authorization', $authorization_header);

echo "\nRequest headers:".PHP_EOL;
foreach ($request->getHeaders() as $name => $values) {
 echo $name . ': ' . implode(' ', $values) . "\n";
}

return $request;
}

// EXAMPLE REQUESTS
$handler = new CurlHandler();
$stack = HandlerStack::create($handler);

// place signing middleware after prepare-body so it can access Content-Length header
$stack->after('prepare_body', Middleware::mapRequest('oci_signer_middleware'));

$client = new Client([
 'handler' => $stack
]);

// GET current user
echo "*****".PHP_EOL;
echo "Getting user: $user_id...".PHP_EOL;
echo "*****".PHP_EOL;
$response = $client->get("https://identity.$region.oraclecloud.com/20160918/users/$user_id");
```

## CHAPTER 34 Developer Tools

```
echo "\nResponse:\n";
echo $response->getStatusCode().PHP_EOL;
echo $response->getBody().PHP_EOL.PHP_EOL;

// Create a VCN
echo "*****".PHP_EOL;
echo "Creating VCN...".PHP_EOL;
echo "*****".PHP_EOL;
$body = "{\n \"cidrBlock\" : \"10.0.0.0/16\",\n \"compartmentId\" : \"\${tenancy_id}\",\n \"displayName\" :\n \"MyPhpVcn\"}";
$response = $client->post("https://iaas.\${region}.oraclecloud.com/20160918/vcns", ["body" => $body,
'headers' => ['Content-Type' => 'application/json']]);
echo "\nResponse:".PHP_EOL;
echo $response->getStatusCode().PHP_EOL;
echo $response->getBody().PHP_EOL.PHP_EOL;

// PUT object with no content
echo "*****".PHP_EOL;
echo "Putting object 'NewObject'...".PHP_EOL;
echo "*****".PHP_EOL;
$body = '';
$response = $client->put("https://objectstorage.\${region}.oraclecloud.com/n/\${namespace}/b/\${bucket}_
name/o/NewObject", ["body" => $body, 'headers' => ['Content-Type' => 'application/json']]);
echo "\nResponse:\n";
echo $response->getStatusCode().PHP_EOL;
echo $response->getBody().PHP_EOL;

// PUT object with content
echo "*****".PHP_EOL;
echo "Putting object 'NewObject2'...".PHP_EOL;
echo "*****".PHP_EOL;

$file_handle = fopen($file_to_upload, "rb");
$body = "";
while (!feof($file_handle)) {
 $body = $body . fgets($file_handle);
}
fclose($file_handle);

$response = $client->put("https://objectstorage.\${region}.oraclecloud.com/n/\${namespace}/b/\${bucket}_
name/o/NewObject2", ["body" => $body, 'headers' => ['Content-Type' => 'application/octet-stream']]);
echo "\nResponse:\n";
```

## CHAPTER 34 Developer Tools

---

```
echo $response->getStatusCode().PHP_EOL;
echo $response->getBody().PHP_EOL;

?>
```

### PYTHON

This sample omits the optional `version` field in the Authorization header.



#### Important

This Python sample code requires TLS 1.2, which is not included with the default Python on Mac OS X.

```
import base64
import email.utils
import hashlib

pip install httpsig_cffi requests six
import httpsig_cffi.sign
import requests
import six
Version 1.0.1

class SignedRequestAuth(requests.auth.AuthBase):
 """A requests auth instance that can be reused across requests"""
 generic_headers = [
 "date",
 "(request-target)",
 "host"
]
 body_headers = [
 "content-length",
 "content-type",
 "x-content-sha256",
]
 required_headers = {
 "get": generic_headers,
 "head": generic_headers,
 "delete": generic_headers,
```

```
"put": generic_headers + body_headers,
"post": generic_headers + body_headers
}

def __init__(self, key_id, private_key):
 # Build a httpsig_cffi.requests_auth.HTTPSignatureAuth for each
 # HTTP method's required headers
 self.signers = {}
 for method, headers in six.iteritems(self.required_headers):
 signer = httpsig_cffi.sign.HeaderSigner(
 key_id=key_id, secret=private_key,
 algorithm="rsa-sha256", headers=headers[:])
 use_host = "host" in headers
 self.signers[method] = (signer, use_host)

def inject_missing_headers(self, request, sign_body):
 # Inject date, content-type, and host if missing
 request.headers.setdefault(
 "date", email.utils.formatdate(usegmt=True))
 request.headers.setdefault("content-type", "application/json")
 request.headers.setdefault(
 "host", six.moves.urllib.parse.urlparse(request.url).netloc)

 # Requests with a body need to send content-type,
 # content-length, and x-content-sha256
 if sign_body:
 body = request.body or ""
 if "x-content-sha256" not in request.headers:
 m = hashlib.sha256(body.encode("utf-8"))
 base64digest = base64.b64encode(m.digest())
 base64string = base64digest.decode("utf-8")
 request.headers["x-content-sha256"] = base64string
 request.headers.setdefault("content-length", len(body))

def __call__(self, request):
 verb = request.method.lower()
 # nothing to sign for options
 if verb == "options":
 return request
 signer, use_host = self.signers.get(verb, (None, None))
 if signer is None:
 raise ValueError(
```

## CHAPTER 34 Developer Tools

---

```
 "Don't know how to sign request verb {}".format(verb))

 # Inject body headers for put/post requests, date for all requests
 sign_body = verb in ["put", "post"]
 self.inject_missing_headers(request, sign_body=sign_body)

 if use_host:
 host = six.moves.urllib.parse.urlparse(request.url).netloc
 else:
 host = None

 signed_headers = signer.sign(
 request.headers, host=host,
 method=request.method, path=request.path_url)
 request.headers.update(signed_headers)
 return request

-----BEGIN RSA PRIVATE KEY-----
...
-----END RSA PRIVATE KEY-----
with open("../sample-private-key") as f:
 private_key = f.read().strip()

This is the keyId for a key uploaded through the console
api_key = "/" .join([
 "ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq",
 "ocidl.user.oc1..aaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvvpkanghtgdaqedqw3rynjq",
 "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34"
])

auth = SignedRequestAuth(api_key, private_key)

headers = {
 "content-type": "application/json",
 "date": email.utils.formatdate(usegmt=True),
 # Uncomment to use a fixed date
 # "date": "Thu, 05 Jan 2014 21:31:40 GMT"
}

GET with query parameters
```

## CHAPTER 34 Developer Tools

```
uri = "https://iaas.us-ashburn-1.oraclecloud.com/20160918/instances?availabilityDomain={availability_
domain}&compartmentId={compartment_id}&displayName={display_name}&volumeId={volume_id}"
uri = uri.format(
 availability_domain="Pjwf%3A%20PHX-AD-1",
 # Older ocid formats included ":" which must be escaped
 compartment_
id="ocidl.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdf1vjsnog7mlr6rtdb25gilchfeyjxa".replace(":",
"%3A"),
 display_name="TeamXInstances",
 volume_
id="ocidl.volume.oc1.phx.abyhq1jrgvttnlx73nmrwoffaux7kcvzfs3s66izvxf2h4lgvyndsdsnoiwr5q".replace(":",
"%3A")
)
response = requests.get(uri, auth=auth, headers=headers)
print(uri)
print(response.request.headers["Authorization"])

POST with body
uri = "https://iaas.us-ashburn-1.oraclecloud.com/20160918/volumeAttachments"
body = """{
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaaam3we6vgnherjq5q2idnccdf1vjsnog7mlr6rtdb25gilchfeyjxa",
 "instanceId": "ocidl.instance.oc1.phx.abuw41jr1sfiqw6vzzxb43vyyp4pkodawglp3wqxjqofakrwwou52gb6s5a",
 "volumeId": "ocidl.volume.oc1.phx.abyhq1jrgvttnlx73nmrwoffaux7kcvzfs3s66izvxf2h4lgvyndsdsnoiwr5q"
}"""
response = requests.post(uri, auth=auth, headers=headers, data=body)
print("\n" + uri)
print(response.request.headers["Authorization"])
```

### RUBY

```
require 'base64'
require 'digest'
require 'openssl'
require 'time'
require 'uri'

gem 'httparty', '~> 0.13.0'
require 'httparty'

Version 1.0.1
class Client
```

## CHAPTER 34 Developer Tools

---

```
include HTTParty
attr_reader :signer

def initialize(key_id, private_key)
 @signer = Signer.new(key_id, private_key)
end

nothing to sign for :options

[:get, :head, :delete].each do |method|
 define_method(method) do |uri, headers: {}|
 self.signer.sign(method, uri, headers, body: nil)
 self.class.send(method, uri, :headers => headers)
 end
end

[:put, :post].each do |method|
 define_method(method) do |uri, headers: {}, body: ""|
 self.signer.sign(method, uri, headers, body)
 self.class.send(method, uri, :headers => headers, :body => body)
 end
end
end

class Signer
 class << self
 attr_reader :headers
 end

 attr_reader :key_id, :private_key

 generic_headers = [:"date", :"(request-target)", :("host")]
 body_headers = [
 :("content-length", :("content-type", :("x-content-sha256"))
 @headers = {
 get: generic_headers,
 head: generic_headers,
 delete: generic_headers,
 put: generic_headers + body_headers,
 post: generic_headers + body_headers
 }
}
```

```
def initialize(key_id, private_key)
 @key_id = key_id
 @private_key = private_key
end

def sign(method, uri, headers, body)
 uri = URI(uri)
 path = uri.query.nil? ? uri.path : "#{uri.path}?#{uri.query}"
 self.inject_missing_headers(headers, method, body, uri)
 signature = self.compute_signature(headers, method, path)
 unless signature.nil?
 self.inject_authorization_header(headers, method, signature)
 end
end

def inject_missing_headers(headers, method, body, uri)
 headers["content-type"] ||= "application/json"
 headers["date"] ||= Time.now.utc.httpdate
 headers["accept"] ||= "*/*"
 headers["host"] ||= uri.host
 if method == :put or method == :post
 body ||= ""
 headers["content-length"] ||= body.length.to_s
 headers["x-content-sha256"] ||= Digest::SHA256.base64digest(body)
 end
end

def inject_authorization_header(headers, method, signature)
 signed_headers = self.class.headers[method].map(&:to_s).join(" ")
 headers["authorization"] = [
 %(Signature version="1"),
 %(headers="#{signed_headers}"),
 %(keyId="#{self.key_id}"),
 %(algorithm="rsa-sha256"),
 %(signature="#{signature}")
].join(",")
end

def compute_signature(headers, method, path)
 return if self.class.headers[method].empty?
 signing_string = self.class.headers[method].map do |header|
```

## CHAPTER 34 Developer Tools

```
 if header == :"(request-target)"
 "#{header}: #{method.downcase} #{path}"
 else
 "#{header}: #{headers[header.to_s]}"
 end
 end.join("\n")
 signature = self.private_key.sign(
 OpenSSL::Digest::SHA256.new,
 signing_string.encode("ascii"))
 Base64.strict_encode64(signature)
end
end

api_key = [
 "ocidl.tenancy.oc1..aaaaaaaaba3pv6wkcr4jqae5f15p2b2m2yt2j6rx32uzr4h25vqstifsfdsq",
 "ocidl.user.oc1..aaaaaaaat5nvwcn5j6aqzjcaty5eqbb6qt2jvpkanghtgdaqedqw3rynjq",
 "20:3b:97:13:55:1c:5b:0d:d3:37:d8:50:4e:c5:3a:34"
].join("/")
private_key = OpenSSL::PKey::RSA.new(File.read("../sample-private-key"))
client = Client.new(api_key, private_key)

headers = {
 # Uncomment to use a fixed date
 # "date" => "Thu, 05 Jan 2014 21:31:40 GMT"
}

GET with query parameters
uri = "https://iaas.us-ashburn-1.oraclecloud.com/20160918/instances?availabilityDomain=#{availability_
domain}&compartmentId=#{compartment_id}&displayName=#{display_name}&volumeId=#{volume_id}"
uri = uri % {
 :availability_domain => "Pjwf%3A%20PHX-AD-1",
 # Older ocid formats included ":" which must be escaped
 :compartment_id =>
"ocidl.compartment.oc1..aaaaaaaam3we6vgnherjq5q2idnccdflvjsnog7mlr6rtdb25gilchfeyjxa".sub(":", "%3A"),
 :display_name => "TeamXInstances",
 :volume_id =>
"ocidl.volume.oc1.phx.abyhqljrgvttnlx73nmrfaux7kcvzfs3s66izvxf2h4lgvyndsdsnoiwr5q".sub(":", "%3A")
}
response = client.get(uri, headers: headers)
puts uri
```

## CHAPTER 34 Developer Tools

```
puts response.request.options[:headers]["authorization"]
puts response.response

POST with body
uri = "https://iaas.us-ashburn.oraclecloud.com/20160918/volumeAttachments"
body = %q({
 "compartmentId":
"ocidl.compartment.oc1..aaaaaaaam3we6vgnerjq5q2idnccdf1vjsnog7mlr6rtdb25gilchfeyjxa",
 "instanceId": "ocidl.instance.oc1.phx.abuw41jr1sfiqw6vzzxb43vyypt4pkodawglp3wqxjqofakrwwou52gb6s5a",
 "volumeId": "ocidl.volume.oc1.phx.abyhqljrgvttnlx73nmrfaux7kcvzfs3s66izvxf2h4lgyndsdsnoiwr5q"
})
response = client.post(uri, headers: headers, body: body)
puts "\n" + uri
puts response.request.options[:headers]["authorization"]
puts response.response
```

### Go

The following example shows how to create a default signer.



#### Note

The Go SDK exposes a stand-alone signer that you can use to sign custom requests. You can find related code at <http://signer.go>.

```
client := http.Client{}
var request http.Request
request = ... // some custom request

// Set the Date header
request.Header.Set("Date", time.Now().UTC().Format(http.TimeFormat))

// And a provider of cryptographic keys
provider := common.DefaultConfigProvider()

// Build the signer
signer := common.DefaultSigner(provider)

// Sign the request
```

## CHAPTER 34 Developer Tools

---

```
signer.Sign(&request)

// Execute the request
client.Do(request)
```

The following example shows how the signer can allow more granular control on the headers used for signing:

```
client := http.Client{}
var request http.Request
request = ... // some custom request

// Set the Date header
request.Header.Set("Date", time.Now().UTC().Format(http.TimeFormat))

// Mandatory headers to be used in the sign process
defaultGenericHeaders = []string{"date", "(request-target)", "host"}

// Optional headers
optionalHeaders = []string{"content-length", "content-type", "x-content-sha256"}

// A predicate that specifies when to use the optional signing headers
optionalHeadersPredicate := func (r *http.Request) bool {
 return r.Method == http.MethodPost
}

// And a provider of cryptographic keys
provider := common.DefaultConfigProvider()

// Build the signer
signer := common.RequestSigner(provider, defaultGenericHeaders, optionalHeaders,
optionalHeadersPredicate)

// Sign the request
signer.Sign(&request)

// Execute the request
c.Do(request)
```

# GLOSSARY

## A

---

### **AD-specific subnet**

A subnet that is specific to a particular availability domain (AD). Historically all subnets were AD-specific. Compare with regional subnets, which Oracle recommends over AD-specific subnets.

### **alarm**

The trigger rule and query to evaluate and related configuration, such as notification details to use when the trigger is breached. Alarms passively monitor your cloud resources using metrics in Monitoring.

### **API key**

A credential for securing requests to the Oracle Cloud Infrastructure REST API.

### **attach**

Link a volume and instance together. Allows an instance to connect and mount the volume as a hard drive.

### **auth token**

Oracle Cloud Infrastructure-generated token you use to authenticate with third-party APIs, such as a Swift client.

### **availability domain**

One or more isolated, fault-tolerant Oracle data centers that host cloud resources such as instances, volumes, and subnets. A region contains one or more availability domains.

## **B**

---

### **backend set**

A logical entity defined by a list of backend servers, a load balancing policy, and a health check policy.

### **bare metal IaaS**

A cloud infrastructure that allows you to utilize hosted physical hardware, as opposed to traditional software-based virtual machines, ensuring a high level of security and performance.

### **block storage volume**

A virtual disk that provides persistent storage space for instances in the cloud.

### **bucket**

A logical container for storing objects.

### C

---

#### **CHAP**

Stands for Challenge-Handshake-Authentication-Protocol. It is a security protocol used by iSCSI for authentication between a volume and an instance.

#### **Cloud Block Storage**

A service that allows you to add block storage volumes to an instance in order to expand the available storage on that resource.

#### **cloud network**

A virtual version of a traditional network—including CIDRs, subnets, route tables, and gateways—on which your instance runs.

#### **cluster network**

A pool of high performance computing (HPC) instances that are connected with a high-bandwidth, ultra low-latency network.

#### **compartment**

A collection of related resources that can be accessed only by certain groups that have been given permission by an administrator in your organization.

#### **Compute**

A service that lets you provision and manage compute hosts, known as instances.

#### **connect**

Make an attached volume usable by an instance's guest OS.

### **CPE**

The router at the edge of your on-premises network. The Networking service also has an object called a CPE, which is a virtual representation of your edge router. You create that object when setting up VPN Connect (an IPSec VPN) between Oracle and your on-premises network.

### **cross-connect**

Used with Oracle Cloud Infrastructure FastConnect, specifically if you're using a third-party provider or colocated with Oracle in a FastConnect location. A cross-connect is the physical cable connecting your existing network to Oracle in the FastConnect location.

### **cross-connect group**

Used with Oracle Cloud Infrastructure FastConnect, specifically if you're using a third-party provider or colocated with Oracle in a FastConnect location. A cross-connect group is a link aggregation group (LAG) that contains at least one cross-connect.

### **customer-premises equipment**

The router at the edge of your on-premises network. The Networking service also has an object called a CPE, which is a virtual representation of your edge router. You create that object when setting up VPN Connect (an IPSec VPN) between Oracle and your on-premises network.

## **D**

---

### **data point**

(Monitoring service) A timestamp-value pair for the specified metric. Example: 2018-05-10T22:19:00Z, 10.4

### **DB System**

A dedicated bare metal instance running Oracle Linux, optimized for running one or more Oracle databases. A DB System is a Database Service resource.

### **DHCP options**

Configuration information that is automatically provided to the instances when they boot up.

### **dimension**

(Monitoring service) A qualifier provided in a metric definition. Example: Resource identifier (resourceId), provided in the definitions of oci\_computeagent metrics.

### **display name**

A friendly name or description that helps you easily identify the resource.

### **DRG**

An optional virtual router that you can add to your VCN to provide a path for private network traffic between your VCN and on-premises network.

### **DRG attachment**

When you attach a dynamic routing gateway (DRG) to a virtual cloud network (VCN), the result is a DRG attachment object. To detach the DRG, you delete that attachment object.

### **dynamic group**

A special type of IAM group that contains instances that match rules that you define (thus the membership can change dynamically as matching instances are terminated or launched). These instances act as "principal" actors and can make API calls to Oracle Cloud Infrastructure services according to IAM policies that you write for the dynamic group.

### **dynamic routing gateway**

An optional virtual router that you can add to your VCN to provide a path for private network traffic between your VCN and on-premises network.

## **E**

---

### **ephemeral public IP**

A public IP address (and related properties) that is temporary and exists for the life of the instance it's assigned to. It can be assigned only to the primary private IP on a VNIC. Compare with reserved public IP.

### **Export**

Controls how file systems are accessed by NFS clients when they connect to a mount target.

### **Export Options**

A set of parameters that specify the level of access granted to NFS clients when they connect to a mount target.

## **F**

---

### **FastConnect**

FastConnect provides an easy way to create a dedicated, private connection between your data center or existing network and Oracle Cloud Infrastructure. FastConnect provides higher-bandwidth options, and a more reliable and consistent networking experience compared to internet-based connections.

### **FastConnect location**

A specific data center where you can connect to Oracle Cloud Infrastructure by using FastConnect.

### **fault domain**

A logical grouping of hardware and infrastructure within an availability domain to provide isolation of resources in case of hardware failure or unexpected software changes.

### **File System**

An organized system of directories and folders where data is stored.

### **frequency**

(Monitoring service) The time period between each posted raw data point for a given metric. (Raw data points are posted by the metric namespace to the Monitoring service.)

## **G**

---

### **group**

A collection of users who all need a particular type of access to a set of resources or compartment.

### **guest operating system**

An operating system installed on a cloud instance.

### **guest OS**

An operating system installed on a cloud instance.

### H

---

#### **health check**

A test to confirm the availability of backend servers.

### I

---

#### **IaaS**

A service that allows customers to rapidly scale up or down their computer infrastructure (computing, storage, or network).

#### **IAM**

The service for controlling authentication and authorization of users who need to use your cloud resources. Also called "IAM".

#### **Identity and Access Management Service**

The service for controlling authentication and authorization of users who need to use your cloud resources. Also called "IAM".

#### **identity provider**

A service that provides identifying credentials and authentication for federated users.

#### **IdP**

Short for "identity provider", which is a service that provides identifying credentials and authentication for federated users.

### **image**

A template of a virtual hard drive that determines the operating system and other software for an instance.

### **Infrastructure-as-a-Service**

A service that allows customers to rapidly scale up or down their computer infrastructure (computing, storage, or network).

### **instance**

A bare metal or virtual machine (VM) compute host. The image used to launch the instance determines its operating system and other software. The shape specified during the launch process determines the number of CPUs and memory allocated to the instance.

### **instance wallet**

An Autonomous Database instance wallet contains only credentials and keys for a single database instance.

### **internet gateway**

An optional virtual router that you can add to your VCN. It provides a path for network traffic between your VCN and the internet.

### **interval**

(Monitoring service) The time window used to convert the given set of raw data points. Example: 5 minutes

### **IPSec connection**

The secure connection between a dynamic routing gateway (DRG) and customer-premises equipment (CPE), consisting of multiple IPSec tunnels. The IPSec connection is one of the components forming a site-to-site VPN between a virtual cloud network (VCN) and your on-premises network.

### **IPv6**

An object that contains an IPv6 address and related properties. Currently IPv6 addressing is supported only in the US Government Cloud. Only instances in IPv6-enabled VCNs and IPv6-enabled subnets can have IPv6 addresses.

### **IQN**

A unique ID assigned to an iSCSI device. Used when connecting a volume to an instance.

### **iSCSI**

A TCP/IP based standard used for communication between a volume and attached instance.

### **iSCSI Qualified Name**

A unique ID assigned to an iSCSI device. Used when connecting a volume to an instance.

## **K**

---

### **key pair**

A security mechanism consisting of a public key and a private key. Required (for example) for Secure Shell (SSH) access to an instance.

### L

---

#### **listener**

An entity that checks for incoming traffic on the load balancer's public floating IP address.

#### **local peering gateway**

A component on a VCN for routing traffic to a locally peered VCN. "Local" peering means the two VCNs are in the same region. Compare with a remote peering connection.

#### **local VCN peering**

The process of connecting two VCNs in the same region so that their resources can communicate without routing the traffic over the internet or through your on-premises network.

### **LPG**

A component on a VCN for routing traffic to a locally peered VCN. "Local" peering means the two VCNs are in the same region. Compare with a remote peering connection.

### M

---

#### **message**

(Notifications and Monitoring services) An alert published to all subscriptions in the specified topic. Each message is delivered at least once per subscription.

#### **metric**

(Monitoring service) A measurement related to health, capacity, or performance of a given resource. Example: CpuUtilization

### **metric definition**

(Monitoring service) A set of references, qualifiers, and other information provided by a metric namespace for a given metric.

### **metric namespace**

(Monitoring service) Indicator of the resource, service, or application that emits the metric. Provided in the metric definition. Example: `oci_computeagent`

### **metric stream**

(Monitoring service) An individual set of aggregated data for a metric. Typically specific to a resource.

### **Monitoring Query Language**

(Monitoring service) The syntax used for metric and alarm queries.

### **Mount Point**

A directory from which a client may access a remote File Storage Service file system.

### **Mount Target**

An NFS endpoint that allows a file system to be accessed by clients.

### **MQL**

(Monitoring service) Monitoring Query Language. The syntax used for metric and alarm queries. In the Console, MQL syntax of queries is displayed in Advanced Mode.

### N

---

#### **NAT gateway**

An optional virtual router that you can add to your VCN to perform Network Address Translation (NAT). A NAT gateway gives cloud resources without public IP addresses access to the internet without exposing those resources to incoming internet connections.

#### **network security group**

One method for implementing security rules in a VCN. A network security group consists of a set of resources (VNICs or resources with VNICs) and security rules that apply to those resources. See also security rules and security lists.

#### **notification destination**

(Monitoring service) Protocol and other details for sending messages when the alarm transitions to another state, such as from "OK" to "FIRING."

### **NSG**

One method for implementing security rules in a VCN. A network security group consists of a set of resources (VNICs or resources with VNICs) and security rules that apply to those resources. See also security rules and security lists.

### O

---

#### **object**

Any type of data, regardless of content type, is stored as an object. The object is composed of the object itself and metadata about the object. Each object is stored in a bucket.

### **OCID**

An Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). This ID is included as part of the resource's information in both the Console and API.

### **one-time password**

A single-use Console password that Oracle assigns to a new user, or to an existing user who requested a password reset.

### **Oracle Cloud Identifier**

An Oracle-assigned unique ID called an Oracle Cloud Identifier (OCID). This ID is included as part of the resource's information in both the Console and API.

### **OTP**

A single-use Console password that Oracle assigns to a new user, or to an existing user who requested a password reset.

## **P**

---

### **policy**

An IAM document that specifies who has what type of access to your resources. It is used in different ways: to mean an individual statement written in the policy language; to mean a collection of statements in a single, named "policy" document (which has an Oracle Cloud ID (OCID) assigned to it); and to mean the overall body of policies your organization uses to control access to resources.

### **policy statement**

Policies can contain one or more individual statements. Each statement gives a group a certain type of access to certain resources in a particular compartment.

**primary IP**

The private IP that is automatically created and assigned to a VNIC during creation.

**primary VNIC**

The VNIC that is automatically created and attached to an instance during launch.

**private IP**

An object that contains a private IPv4 address and related properties such as a hostname for DNS. Each instance automatically comes with a primary private IP, and you can add secondary ones.

**private peering**

One of the ways to use FastConnect. Private peering lets you extend your existing infrastructure into a virtual cloud network (VCN) in Oracle Cloud Infrastructure (for example, to implement a hybrid cloud, or a lift and shift scenario). Communication across the connection is with IPv4 private addresses (typically RFC 1918).

**private subnet**

A subnet in which instances are not allowed to have public IP addresses

**private virtual circuit**

A FastConnect virtual circuit that supports private peering.

**public IP**

An object that contains a public IP address and related properties. You control whether each private IP on an instance has an assigned public IP. There are two types: reserved public IPs and ephemeral public IPs.

### **public peering**

One of the way to use FastConnect. Public peering lets your on-premises network access public services in Oracle Cloud Infrastructure without using the internet. For example, Object Storage, the Oracle Cloud Infrastructure Console and APIs, or public load balancers in your VCN. Communication across the connection is with IPv4 public IP addresses. Without FastConnect, the traffic destined for public IP addresses would be routed over the internet. With FastConnect, that traffic goes over your private physical connection.

### **public subnet**

A subnet in which instances are allowed to have public IP addresses. When you launch an instance in a public subnet, you specify whether the instance should have a public IP address.

### **public virtual circuit**

A FastConnect virtual circuit that supports public peering.

## **Q**

---

### **query**

(Monitoring service) The expression to evaluate for returning aggregated data. A valid query includes a metric, statistic, and interval. In the Console, you can view a query in Basic Mode or Advanced Mode. The latter displays the Monitoring Query Language (MQL) syntax.

## **R**

---

### **realm**

A logical collection of regions. Realms are isolated from each other and do not share any data. Your tenancy exists in a single realm and can access the regions that belong to that realm.

**region**

A collection of availability domains located in a single geographic location.

**regional subnet**

A subnet that spans all availability domains (ADs) in the region. Oracle recommends using regional subnets because they are more flexible and make it easier to implement failover across ADs. Compare with AD-specific subnets.

**regional wallet**

An Autonomous Database regional wallet contains credentials and keys for all Autonomous Databases in a specified region.

**remote peering connection**

A component on a dynamic routing gateway (DRG) for routing traffic to a remotely peered VCN. "Remote" peering means the two VCNs are in different regions. Compare with a local peering gateway.

**remote VCN peering**

The process of connecting two VCNs in different regions so that their resources can communicate without routing their traffic over the internet or through your on-premises network.

**reserved public IP**

A public IP address (and related properties) that you create in your tenancy and assign to your instances in a given region as you like. It persists in your tenancy until you delete it. It can be assigned to any private IP on a given VNIC, not just the primary private IP. Compare with ephemeral private IP.

### **resolution**

(Monitoring service) The period between time windows, or the regularity at which time windows shift. Example: 1 minute

### **resource**

The cloud objects that your company's employees create and use when interacting with Oracle Cloud Infrastructure.

### **route table**

Virtual route table for your VCN that provides mapping for the traffic from subnets via gateways to external destinations.

### **RPC**

A component on a dynamic routing gateway (DRG) for routing traffic to a remotely peered VCN. "Remote" peering means the two VCNs are in different regions. Compare with a local peering gateway.

## **S**

---

### **secondary IP address**

An additional private IP you've added to a VNIC on an instance. Each VNIC automatically comes with a primary private IP that cannot be removed.

### **secondary VNIC**

An additional VNIC you've added to an instance. Each instance automatically comes with a primary VNIC that cannot be removed.

### **security list**

One method for implementing security rules in a VCN. A security list consists of security rules that apply to all resources in any subnet that uses the security list. See also security rules and network security groups.

### **security rule**

Virtual firewall rules for your VCN. Each security rule specifies a type of ingress or egress traffic allowed in or out of a resource or VNIC. Also see network security groups and security lists.

### **service gateway**

An optional virtual router that you can add to your VCN. The gateway enables on-premises hosts or VCN hosts to privately access Oracle services (such as Object Storage and Autonomous Database) without exposing the resources to the public internet.

### **shape**

A template that determines the number of CPUs and the amount of memory allocated to a newly created instance.

### **statement**

Policies can contain one or more individual statements. Each statement gives a group a certain type of access to certain resources in a particular compartment.

### **statistic**

The aggregation function applied to the given set of raw data points. Example: SUM

### **subnet**

Subdivision of your VCN used to separate your network into multiple smaller, distinct networks.

### **subscription**

(Notifications service) An endpoint for a topic; typically a URL or email address. Published messages are sent to each subscription for a topic.

### **suppression**

(Monitoring service) A configuration to avoid publishing messages during the specified time range. Useful for suspending alarm notifications during system maintenance.

### **Swift password**

(Deprecated. Use an auth token to authenticate with your Swift client.) Swift is the OpenStack object store service. A Swift password enables you to use an existing Swift client with Oracle Cloud Infrastructure Object Storage.

## **T**

---

### **tenancy**

The root compartment that contains all of your organization's compartments and other Oracle Cloud Infrastructure cloud resources.

### **tenant**

The name assigned to a particular company's or organization's overall environment. Users provide their tenant when signing in to the Console.

### **topic**

(Notifications service) A communication channel for sending messages to the subscriptions in the topic.

### **transit routing**

A network setup in which your on-premises network uses a connected virtual cloud network (VCN) to reach Oracle resources or services beyond that VCN. You connect the on-premises network to the VCN with a FastConnect private virtual circuit or VPN Connect, and then configure the VCN routing so that traffic transits through the VCN to its destination beyond the VCN. You can use transit routing to access multiple VCNs from your on-premises network over a single FastConnect or VPN Connect. Or you can use it to give your on-premises network private access to Oracle services so that on-premises hosts use their private IP addresses and the traffic does not go over the internet.

### **trigger rule**

(Monitoring service) The condition that must be met for the alarm to be in the firing state. A trigger rule can be based on a threshold or absence of a metric.

## **U**

---

### **user**

An individual employee or system that needs to manage or use your company's Oracle Cloud Infrastructure resources.

## **V**

---

### **VCN**

A virtual version of a traditional network—including CIDRs, subnets, route tables, and gateways—on which your instance runs.

### **virtual circuit**

Used with Oracle Cloud Infrastructure FastConnect. An isolated network path that runs over one or more physical network connections to provide a single, logical connection between the edge of

your existing network and Oracle Cloud Infrastructure.

### **virtual cloud network**

A virtual version of a traditional network—including CIDRs, subnets, route tables, and gateways—on which your instance runs.

### **virtual machine**

A software-based emulation of a full computer that runs within a physical host computer.

### **virtual network interface card**

A VNIC enables an instance to connect to a VCN and determines how the instance connects with endpoints inside and outside the VCN. Each instance automatically comes with a primary VNIC, and you can add secondary ones. Other types of cloud resources also automatically get a VNIC upon creation (examples: load balancers, DB systems).

### **VM**

A software-based emulation of a full computer that runs within a physical host computer.

### **VNIC**

A VNIC enables an instance to connect to a VCN and determines how the instance connects with endpoints inside and outside the VCN. Each instance automatically comes with a primary VNIC, and you can add secondary ones. Other types of cloud resources also automatically get a VNIC upon creation (examples: load balancers, DB systems).

### **volume**

A detachable block storage device that allows you to dynamically expand the storage capacity of an instance.

### W

---

#### **work request**

An object that reports on the current state of an asynchronous service request.

# RELEASE NOTES

You can find the Oracle Cloud Infrastructure [Release Notes](#) online.