

Oracle® Private Cloud Appliance

User Guide for Release 3.0.1

ORACLE®

F49445-01
January 2022

Oracle Legal Notices

[Copyright](#) © 2013, 2022, Oracle and/or its affiliates.

Table of Contents

Preface	ix
1 Working in the Compute Enclave	1
1.1 Using the Compute Web UI	1
1.1.1 Logging In	1
1.1.2 Navigating the Dashboard	2
1.1.3 Using Resource Type and Resource Detail Pages	3
1.1.4 Locating Tenancy and Profile Information	14
1.2 Using the OCI CLI	14
1.2.1 Before You Begin	14
1.2.2 Installing the OCI CLI	15
1.2.3 Configuring the OCI CLI	17
1.2.4 Testing the OCI CLI Configuration	20
1.2.5 Using Multiple Profiles	22
1.2.6 Working with API Signing Keys	24
1.2.7 Understanding Command Syntax and Finding Help	27
1.2.8 Using JSON for Complex Command Input	30
1.2.9 Formatting and Filtering Command Output	33
2 Identity and Access Management	35
2.1 Creating and Managing User Accounts	35
2.1.1 Creating a User	36
2.1.2 Providing a Temporary Compute Web UI Password	38
2.1.3 Setting Your Own Compute Web UI Password	39
2.1.4 Viewing User Information and Group Membership	40
2.1.5 Adding a User to a Group by Updating the User	40
2.1.6 Removing a User from a Group by Updating the User	41
2.1.7 Modifying a User	41
2.1.8 Deleting a User	42
2.2 Creating and Managing User Groups	42
2.2.1 Creating a Group	43
2.2.2 Viewing Group Information and Group Membership	44
2.2.3 Adding a User to a Group by Updating the Group	45
2.2.4 Removing a User from a Group by Updating the Group	45
2.2.5 Modifying a Group	46
2.2.6 Deleting a Group	46
2.3 Creating and Managing Compartments	47
2.3.1 Understanding the Tenancy	47
2.3.2 Listing Compartments	48
2.3.3 Creating a Compartment	49
2.3.4 Applying Tag Defaults	50
2.3.5 Adding Policies for Access Control	50
2.3.6 Adding Resources to a Compartment	50
2.3.7 Updating a Compartment	51
2.3.8 Moving a Compartment to a Different Compartment	52
2.3.9 Deleting a Compartment	52
2.4 Managing Policies	53
2.4.1 Creating a Policy	54
2.4.2 Writing Policy Statements	55
2.4.3 Deleting a Policy	62
2.5 Federating with Microsoft Active Directory	62
2.5.1 Gathering Required Information from ADFS	63
2.5.2 Verifying Identity Provider Self-Signed Certificates	63

2.5.3	Managing Identity Providers	64
2.5.4	Working with Group Mappings for an Identity Provider	67
2.5.5	Adding Oracle Private Cloud Appliance as a Trusted Relying Party in ADFS	69
2.5.6	Setting Up Policies for the Groups	71
2.5.7	Providing Federated Users Sign In Information	71
3	Resource Tag Management	73
3.1	Creating and Managing Tag Namespaces	73
3.1.1	Creating a Tag Namespace	73
3.1.2	Updating a Tag Namespace	74
3.1.3	Retiring a Tag Namespace	75
3.1.4	Reactivating a Tag Namespace	76
3.1.5	Moving a Tag Namespace to a Different Compartment	76
3.1.6	Deleting a Tag Namespace	77
3.2	Creating and Managing Tag Key Definitions	78
3.2.1	Creating a Tag Key Definition	78
3.2.2	Updating a Tag Key Definition	80
3.2.3	Retiring a Tag Key Definition	82
3.2.4	Reactivating a Tag Key Definition	82
3.2.5	Deleting a Tag Key Definition	83
3.3	Configuring Tag Defaults	84
3.3.1	Creating a Tag Default	84
3.3.2	Updating the Value of a Tag Default	86
3.3.3	Deleting a Tag Default	86
3.4	Working with Resource Tags	87
3.4.1	Adding Tags at Resource Creation	87
3.4.2	Applying Tags to an Existing Resource	89
3.4.3	Filtering a List of Resources by Tag	90
4	Networking	93
4.1	Managing VCNs and Subnets	94
4.1.1	Creating a VCN	94
4.1.2	Creating a Subnet	95
4.1.3	Editing a Subnet	96
4.1.4	Deleting a Subnet	98
4.1.5	Terminating a VCN	98
4.1.6	Deleting a VCN	98
4.2	Configuring VCN Rules and Options	99
4.2.1	Working with DHCP Options	99
4.2.2	Working with Route Tables	104
4.2.3	Controlling Traffic with Security Lists	109
4.2.4	Controlling Traffic with Network Security Groups	116
4.3	Configuring VCN Gateways	124
4.3.1	Enabling Public Connections through a NAT Gateway	126
4.3.2	Providing Public Access through an Internet Gateway	128
4.3.3	Connecting VCNs through a Local Peering Gateway	129
4.3.4	Connecting to the On-Premises Network through a Dynamic Routing Gateway	131
4.3.5	Accessing Oracle Services through a Service Gateway	133
4.4	Configuring VNICs and IP Addressing	135
4.4.1	Public IP Addresses	136
4.4.2	Managing Public IP Addresses	136
4.4.3	Creating and Managing VNICs	141
4.4.4	Assigning IP Addresses to VNICs	151
4.5	Managing Public DNS Zones	152
4.5.1	Creating a Public DNS Zone	152
4.5.2	Working with Zone Records	154

4.5.3	Editing a Public DNS Zone	157
4.5.4	Working with Transaction Signature Keys	157
4.5.5	Deleting a Public DNS Zone	159
4.6	Managing Traffic with Steering Policies	160
4.6.1	Creating a Load Balancer Steering Policy	160
4.6.2	Creating an IP Prefix Steering Policy	163
4.6.3	Editing a Steering Policy	165
4.6.4	Moving a Steering Policy to a Different Compartment	166
4.6.5	Attaching a Domain to a Steering Policy	166
4.6.6	Editing an Attached Domain	167
4.6.7	Deleting a Steering Policy Attachment	168
4.6.8	Deleting a Steering Policy	168
4.7	Networking Scenarios	169
4.7.1	Logical Routers	169
4.7.2	Using Firewalls	170
4.7.3	Use of Network Segmentation	170
4.7.4	Use of Tunneling	170
4.7.5	Use of Virtual Cloud Networks	171
5	Compute Images	175
5.1	Accessing the Management Node Images	175
5.1.1	Initial User Account for Management Node Images	175
5.1.2	Importing Images from the Management Node (Direct Method)	176
5.1.3	Importing Images from the Management Node (Indirect Method)	176
5.1.4	Downloading an Image From the Management Node	177
5.1.5	Best Practices for Sharing an Image Across Tenancies	178
5.1.6	Importing an Image	178
5.1.7	Exporting an Image to Object Storage	180
5.2	Managing Images	182
5.2.1	Overview	182
5.2.2	Listing Images and Details	182
5.2.3	Creating an Image From an Instance	184
5.2.4	Editing the Image Name or Compatible Shapes	185
5.2.5	Moving an Image to a Different Compartment	187
5.2.6	Deleting an Image	187
5.3	Bring Your Own Image (BYOI)	188
5.3.1	Importing Custom Linux Images	188
5.3.2	Importing Custom Microsoft Windows Images	189
6	Compute Instance Deployment	197
6.1	Tutorial – Launching Your First Instance	197
6.1.1	Task Flow to Launch an Instance	198
6.1.2	Prerequisites	198
6.1.3	Log into Oracle Private Cloud Appliance	198
6.1.4	Create a Compartment	199
6.1.5	Import an Image	199
6.1.6	Create a Virtual Cloud Network (VCN)	201
6.1.7	Create a Subnet	202
6.1.8	Create an Internet Gateway and Configure Route Rules	203
6.1.9	Launch an Instance	204
6.1.10	Get the Instance IP Address	206
6.1.11	Connect to Your Instance	206
6.1.12	Add a Block Volume	208
6.1.13	Attach the Block Volume to an Instance	208
6.1.14	(Optional) Clean Up Resources	209
6.2	Working with Instances	211

6.2.1	Creating an Instance	212
6.2.2	Retrieving Instance Metadata from Within the Instance	218
6.2.3	Updating an Instance	219
6.2.4	Stopping, Starting and Resetting an Instance	220
6.2.5	Terminating an Instance	221
6.3	Working with Instance Configurations and Instance Pools	222
6.3.1	Creating an Instance Configuration	222
6.3.2	Updating an Instance Configuration	224
6.3.3	Moving an Instance Configuration to a Different Compartment	224
6.3.4	Deleting an Instance Configuration	225
6.3.5	Using an Instance Configuration to Launch an Instance	225
6.3.6	Creating an Instance Pool	226
6.3.7	Updating an Instance Pool	228
6.3.8	Stopping and Starting Instances in an Instance Pool	229
6.3.9	Deleting an Instance Pool	230
6.4	Connecting to a Compute Instance	231
6.4.1	Prerequisites	231
6.4.2	Managing Key Pairs	232
6.4.3	Connecting to a Linux or Oracle Solaris Instance	234
6.4.4	Connecting to a Windows Instance	236
6.4.5	Connecting to an Instance Using a Console Connection	238
7	Block Volume Storage	241
7.1	Block Volumes	242
7.2	Creating and Attaching Block Volumes	242
7.2.1	Task Flow to Attach Block Volumes to Instances	242
7.2.2	Creating a Block Volume	242
7.2.3	Attaching a Volume	244
7.2.4	Attaching a Volume to Multiple Instances	246
7.2.5	Find Your Volume in the Instance	247
7.2.6	Configuring Volumes to Automatically Mount (Linux Instances)	249
7.3	Managing Block Volumes	250
7.3.1	Listing Block Volumes and Block Volume Details	250
7.3.2	Listing Block Volume Attachments	252
7.3.3	Editing a Volume's Configuration	254
7.3.4	Moving a Volume to a Different Compartment	255
7.3.5	Cloning a Volume	256
7.3.6	Detaching a Block Volume	258
7.3.7	Deleting a Block Volume	259
7.4	Managing Boot Volumes	259
7.4.1	Overview	260
7.4.2	Listing Boot Volumes	260
7.4.3	Listing Boot Volume Attachments	261
7.4.4	Detaching a Boot Volume	262
7.4.5	Reattaching a Boot Volume	262
7.4.6	Deleting a Boot Volume	263
7.5	Resizing Volumes	264
7.5.1	Resizing Overview	264
7.5.2	Online Volume Resizing	265
7.5.3	Offline Volume Resizing	267
7.6	Managing Volume Groups	271
7.6.1	Overview	271
7.6.2	Viewing the Volumes in a Volume Group	271
7.6.3	Creating a Volume Group from Existing Volumes	273
7.6.4	Adding Volumes to a Group	274

7.6.5 Removing Volumes from a Group	275
7.6.6 Creating a Clone of a Volume Group	276
7.6.7 Deleting a Volume Group	277
7.7 Backing Up Block Volumes	277
7.7.1 Block Volume Backups Overview	277
7.7.2 Viewing Volume Backups	278
7.7.3 Creating a Manual Boot or Block Volume Backup	279
7.7.4 Creating a Manual Backup of a Volume Group	281
7.7.5 Restoring a Backup to a New Volume	283
7.7.6 Restoring a Volume Group from a Volume Group Backup	284
7.8 Managing Backup Policies	286
7.8.1 Overview	286
7.8.2 Creating Backup Policies and Schedules	286
7.8.3 Accessing the Backups	289
7.8.4 Viewing Backup Policies	290
7.8.5 Editing a Backup Policy Schedule	291
7.8.6 Deleting a Backup Policy Schedule	292
7.8.7 Deleting a Backup Policy	293
8 File System Storage	295
8.1 Creating a File System, Mount Target, and Export	295
8.1.1 Overview	295
8.1.2 Task Flow to Create and Export A File System	296
8.1.3 Creating a Mount Target	296
8.1.4 Creating a File System	299
8.1.5 Creating an Export for a File System	300
8.2 Controlling Access to File Storage	303
8.2.1 Overview	303
8.2.2 Configuring VCN Security Rules for File Storage	303
8.2.3 Adding File Storage to a Network Security Group	304
8.2.4 Setting NFS Export Options	305
8.3 Mounting File Systems on UNIX-Type Instances	307
8.3.1 Mounting Overview	307
8.3.2 Obtaining the Mount Target IP Address	308
8.3.3 Mounting a File System on Linux, RedHat, or CentOS	309
8.3.4 Mounting a File System on Ubuntu or Debian	311
8.3.5 Configuring a File System to Automatically Mount (Linux Instances)	312
8.4 Mounting File Systems On Windows Instances	313
8.4.1 Mounting a File System On a Windows Instance Using NFS	313
8.4.2 Mounting a File System on a Window Instance Using SMB	316
8.5 Managing Mount Targets and Exports	318
8.5.1 Overview	318
8.5.2 Listing Mount Targets and Viewing Details	319
8.5.3 Changing the Mount Target Name	320
8.5.4 Listing Exports	321
8.5.5 Listing Export Sets	322
8.5.6 Deleting an Export	323
8.5.7 Moving a Mount Target to a Different Compartment	323
8.5.8 Deleting a Mount Target	324
8.6 Managing File Systems	324
8.6.1 Listing and Viewing the Details of a File System	325
8.6.2 Changing the File System Name	326
8.6.3 Moving a File System to a Different Compartment	327
8.6.4 Deleting a File System	328
8.7 Managing Snapshots	328

8.7.1 Snapshots Overview	328
8.7.2 Listing and Getting Snapshot Details	328
8.7.3 Creating a Snapshot	330
8.7.4 Accessing a Snapshot on the Mounted File System	331
8.7.5 Restoring a Snapshot (UNIX-Type Instances)	332
8.7.6 Deleting a Snapshot	332
9 Object Storage	335
9.1 Object Storage	336
9.2 Managing Buckets	336
9.2.1 Object Storage Buckets	336
9.2.2 Obtaining the Object Storage Namespace	336
9.2.3 Listing Buckets	336
9.2.4 Viewing Bucket Details	337
9.2.5 Creating a Bucket	338
9.2.6 Moving a Bucket to a Different Compartment	340
9.2.7 Deleting a Bucket	341
9.3 Managing Storage Objects	342
9.3.1 Objects	342
9.3.2 Viewing Objects in a Bucket	342
9.3.3 Creating a Folder or SubFolder	344
9.3.4 Uploading an Object	345
9.3.5 Performing a Multi-Part Upload	346
9.3.6 Listing the Parts of an Unfinished or Failed Multi-part Upload	347
9.3.7 Canceling a Multi-Part Upload	347
9.3.8 Performing a Bulk Object Upload	348
9.3.9 Copying an Object to a Different Bucket	349
9.3.10 Downloading an Object	350
9.3.11 Performing a Multi-Part Download	351
9.3.12 Performing a Bulk Download	352
9.3.13 Deleting an Object	352
9.3.14 Performing a Bulk Delete of All Objects in a Bucket	353
9.4 Managing Object Versioning	354
9.4.1 Object Versioning	354
9.4.2 Enabling Versioning During Bucket Creation	355
9.4.3 Enabling, Disabling, or Suspending Versioning (After Bucket Creation)	356
9.4.4 Viewing Object Versions and Details	357
9.4.5 Deleting the Previous Version of an Object	358
9.4.6 Recovering a Deleted Object Version	358
9.5 Using Pre-Authenticated Requests	359
9.5.1 Pre-Authenticated Requests	359
9.5.2 Listing Pre-Authenticated Requests	359
9.5.3 Creating a Pre-Authenticated Request for All Objects in a Bucket	361
9.5.4 Creating a Pre-Authenticated Request for a Specific Object	362
9.5.5 Constructing the Pre-Authenticated Request URL	364
9.5.6 Deleting a Pre-Authenticated Request	364
9.5.7 Listing Objects for Pre-Authenticated Requests	365
9.5.8 Uploading an Object Using a Pre-Authenticated Request	365
9.5.9 Downloading an Object Using a Pre-Authenticated Request	365
9.6 Defining Retention Rules	366
9.6.1 Retention Rules	366
9.6.2 Viewing Retention Rules and Details	366
9.6.3 Creating a Retention Rule	368
9.6.4 Modifying a Retention Rule	370
9.6.5 Deleting a Retention Rule	371

Preface

This publication is part of the customer documentation set for Oracle Private Cloud Appliance Release 3.0.1. Note that the documentation follows the release numbering scheme of the appliance software, not the hardware on which it is installed. All Oracle Private Cloud Appliance product documentation is available at:

<https://docs.oracle.com/en/engineered-systems/private-cloud-appliance/index.html>.

Oracle Private Cloud Appliance Release 3.x is a flexible general purpose Infrastructure as a Service solution, engineered for optimal performance and compatibility with Oracle Cloud Infrastructure. It allows customers to consume the core cloud services from the safety of their own network, behind their own firewall.

Audience

This documentation is intended for owners, administrators and operators of Oracle Private Cloud Appliance. It provides architectural and technical background information about the engineered system components and services, as well as instructions for installation, administration, monitoring and usage.

Oracle Private Cloud Appliance has two strictly separated operating areas, known as enclaves. The Compute Enclave offers a practically identical experience to Oracle Cloud Infrastructure: it allows users to build, configure and manage cloud workloads using compute instances and their associated cloud resources. The Service Enclave is where privileged administrators configure and manage the appliance infrastructure that provides the foundation for the cloud environment. The target audiences of these enclaves are distinct groups of users and administrators. Each enclave also provides its own separate interfaces.

It is assumed that readers have experience with system administration, network and storage configuration, and are familiar with virtualization technologies. Depending on the types of workloads deployed on the system, it is advisable to have a general understanding of container orchestration, and UNIX and Windows operating systems.

Feedback

Provide feedback about this documentation at:

<http://www.oracle.com/goto/docfeedback>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, code in examples, text that appears on the screen, or text that you enter.
<code>\$ prompt</code>	The dollar sign (\$) prompt indicates a command run as a non-root user.
<code># prompt</code>	The pound sign (#) prompt indicates a command run as the <code>root</code> user.

Document Revision

Document generated on: 2022-01-31 (revision: 1388)

Documentation Accessibility

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

For information about the accessibility of the Oracle Help Center, see the Oracle Accessibility Conformance Report at <https://www.oracle.com/corporate/accessibility/templates/t2-11535.html>.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

Chapter 1 Working in the Compute Enclave

Table of Contents

1.1 Using the Compute Web UI	1
1.1.1 Logging In	1
1.1.2 Navigating the Dashboard	2
1.1.3 Using Resource Type and Resource Detail Pages	3
1.1.4 Locating Tenancy and Profile Information	14
1.2 Using the OCI CLI	14
1.2.1 Before You Begin	14
1.2.2 Installing the OCI CLI	15
1.2.3 Configuring the OCI CLI	17
1.2.4 Testing the OCI CLI Configuration	20
1.2.5 Using Multiple Profiles	22
1.2.6 Working with API Signing Keys	24
1.2.7 Understanding Command Syntax and Finding Help	27
1.2.8 Using JSON for Complex Command Input	30
1.2.9 Formatting and Filtering Command Output	33

The Compute Enclave is the part of the Oracle Private Cloud Appliance where you work with and manage cloud resources. This section describes the general usage principles of the graphical user interface and command line interface to the Compute Enclave.

1.1 Using the Compute Web UI

The Compute Web UI is the graphical interface to the Compute Enclave. You can use the Compute Web UI on its own or with the OCI CLI to complete tasks. The Compute Web UI provides the same core functionality as the OCI CLI, however, the OCI CLI does have some additional functionality.

This section provides instructions for logging into the Compute Web UI, navigating the dashboard, and working with resources using resource type and resource detail pages. Within the rest of the Oracle Private Cloud Appliance User Guide you learn how to use the Compute Web UI to complete tasks within the context of the step-by-step procedures.

1.1.1 Logging In

Before you log into the Compute Web UI, make sure you have the Oracle Private Cloud Appliance system and domain names, the tenancy name, and your user name and password. If you do not have these details, ask your administrator. If you have access to the Service Web UI, you can locate the tenancy name and the system and domain names for your Oracle Private Cloud Appliance.

To log into the Compute Web UI, complete the following steps.

1. From a browser, enter the URL for your Oracle Private Cloud Appliance.

For example, <https://console.pcasys1.mycompany.com> where *pcasys1* is the name of your Oracle Private Cloud Appliance and *mycompany.com* is your domain.

The Compute Enclave Select Tenancy page is displayed.

2. Enter your tenancy name and click Continue.

The Sign In page is displayed.

3. Enter your Username and Password, and then click Sign In.

The Oracle Private Cloud Appliance dashboard displays with quick action tiles.

Note

If you are prompted to change a temporary password, see [Section 2.1.3.1, "Setting Your Password"](#).

1.1.2 Navigating the Dashboard

When you log into the Compute Enclave, the dashboard is displayed with quick action tiles for common tasks, such as viewing compute instances, block and file storage, and VCNs. There is also a quick action tile to create a virtual machine instance.

Note

The dashboard is static and not configurable.

The navigation menu, which you can click on or tab to, is organized by services. When you click on a service, the sub-menu expands and displays the service's resource types. When you click on a resource type, a page is displayed that contains a tabular list of resources related to the resource type. The following table provides the Oracle Private Cloud Appliance services and their respective resource types as they are displayed in the navigation menu.

Service	Resource Types in Sub-Menu
Compute	<ul style="list-style-type: none"> • Instances • Instance Configurations • Instance Pools • Custom Images <p>For more information, see Chapter 6, Compute Instance Deployment.</p>
Block Storage	<ul style="list-style-type: none"> • Block Volumes • Block Volume Backups • Boot Volumes • Boot Volume Backups • Backup Policies • Volume Group • Volume Group Backups <p>For more information, see Chapter 7, Block Volume Storage.</p>
File Storage	<ul style="list-style-type: none"> • File Systems • Mount Targets

Service	Resource Types in Sub-Menu
	For more information, see Chapter 8, File System Storage .
Object Storage	<ul style="list-style-type: none"> Object Storage <p>For more information, see Chapter 9, Object Storage.</p>
Networking	<ul style="list-style-type: none"> Virtual Cloud Networks Dynamic Routing Gateways <p>For more information, see Chapter 4, Networking.</p>
DNS	<ul style="list-style-type: none"> Zones TSIG Keys Steering Policies <p>For more information, see Chapter 4, Networking.</p>
Identity	<ul style="list-style-type: none"> Users Federation Groups Policies Compartments <p>For more information, see Chapter 2, Identity and Access Management.</p>
Governance	<ul style="list-style-type: none"> Tag Namespaces <p>For more information, see Section 3.1, "Creating and Managing Tag Namespaces".</p>

1.1.3 Using Resource Type and Resource Detail Pages

Resource type and resource detail pages are what you use to work with a compartment's or tenancy's resources. A resource type page displays a list of all resources of that type and also contains the service's sub-menu. When you click on a resource in the list, its own detail page is displayed. Every resource detail page has some general information about the resource, such as its OCID, when it was created, the compartment it is in, and any tags associated with it.

1.1.3.1 About Resource Type Pages

All resource type pages contain a list of resources, provide similar information about the resources, and give you the ability to perform some actions on the resources. The list of resources is displayed in table format and all resource lists contain some common elements:

- Resource Name, Created, and Actions columns.

- Auto Reload, Refresh, and Filter by Tag actions.

Note

The compute service's instance resource type page also has a Filter by Status option.

For each resource in the list, there is an Actions menu (three dots) that allows you to, for example, view details, edit, and delete the resource provided you have the appropriate permissions.

For example, if you want to view all users:

1. Open the navigation menu and click Identity.

The Identity sub-menu opens displaying its resource types: Users, Groups, Policies, Compartments, and Federation.

2. Click Users.

The Users detail page opens and the Identity sub-menu is displayed.

3. From the Users detail page, you can:

- View a list of users in table format.
- View user status.
- Act upon a user to view details, edit user information, or delete a user.
- Click Create User to create a new user.
- Click another resource type in the Identity sub-menu.

If you click Groups, Policies, or Compartments, you can see that the resource type pages give you the option of creating a new resource and contain tabular list of resources.

On some resource type pages, you are able to see resources in other compartments by using the compartment filter in the page title. For example, if you are viewing the Boot Volumes resource type page and you have more than one compartment, the page title displays the compartment name and a down arrow:

Boot Volumes in the [MyCompartmentName](#) (down arrow) compartment

To view boot volumes in another compartment within the tenancy, you simply click the down arrow and select a different compartment.

Resource lists can also contain additional information that is specific to the resource type. For example, a list of block volumes also shows the status and size of the volumes and name of the backup policies if they exist. The remainder of this section shows you the additional information provided in resource lists based on the Oracle Private Cloud Appliance services.

Compute

The following compute resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Instances	Status or Filter by Status - Allows you to see the status of an instance or filter by the instance state:

Resource Type	Resource-specific Elements
	<ul style="list-style-type: none">• Creating Image• Provisioning• Running• Starting• Stopped• Stopping• Terminated• Terminating <p>Shape - The shape of the instance, which determines the number of CPUs and the amount of memory allocated to the instance.</p> <p>Fault Domain - The name of the fault domain (a grouping of hardware and infrastructure) the instance is running in. Fault domains let you distribute your instances so that they are not on the same physical hardware.</p> <p>For more information, see Section 6.2, “Working with Instances”.</p>
Instance Pools	<p>Lifecycle State - The current state of the instance pool:</p> <ul style="list-style-type: none">• Provisioning• Scaling• Starting• Running• Stopping• Stopped• Terminating• Terminated <p>Target Instance Count - Number of instances in a pool.</p> <p>Instance Configuration - The name of the instance configuration associated with the instance pool.</p> <p>For more information, see Section 6.3, “Working with Instance Configurations and Instance Pools”.</p>
Custom Images	<p>Status - The current state of a custom image:</p> <ul style="list-style-type: none">• Provisioning• Importing

Resource Type	Resource-specific Elements
	<ul style="list-style-type: none"> • Available • Exporting • Stopping • Disabled • Deleted <p>For more information, see Section 5.2, “Managing Images”.</p>

Block Storage

The following block storage resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Block Volumes	<p>Status - The current state of a volume:</p> <ul style="list-style-type: none"> • Provisioning • Restoring • Available • Terminating • Terminated • Faulty <p>Size - The size of the volume in GBs.</p> <p>Backup Policy - The name of the backup policy.</p> <p>For more information, see Section 7.3, “Managing Block Volumes”.</p>
Block Volume Backups	<p>Status - The current state of a volume backup:</p> <ul style="list-style-type: none"> • Creating • Available • Terminating • Terminated • Faulty • Request Received <p>Total Size - The size used by the backup, in GBs, which is typically smaller than size of the block volume depending on the space consumed on the boot volume and whether the backup is full or incremental.</p> <p>For more information, see Section 7.7, “Backing Up Block Volumes”.</p>
Boot Volumes	<p>State - The current state of a boot volume:</p>

Resource Type	Resource-specific Elements
	<ul style="list-style-type: none"> • Provisioning • Restoring • Available • Terminating • Terminated • Faulty <p>Attached to Instance - Displays Yes if the boot volume is attached to an instance and No if it is not.</p> <p>Size in GB - The size of the boot volume in GBs.</p> <p>For more information, see Section 7.4, “Managing Boot Volumes”.</p>
<p>Boot Volume Backups</p>	<p>Status - The current state of a boot volume backup:</p> <ul style="list-style-type: none"> • Creating • Available • Terminating • Terminated • Faulty • Request Received <p>Total Size - The size used by the backup, in GBs, which is typically smaller than size of the boot volume depending on the space consumed on the boot volume.</p> <p>For more information, see Section 7.7, “Backing Up Block Volumes”.</p>
<p>Volume Groups</p>	<p>Status - The current state of a volume group:</p> <ul style="list-style-type: none"> • Provisioning • Available • Terminating • Terminated • Faulty <p>Total Size - The aggregate size of the volume group in GBs.</p> <p>Source Volume Group - Specifies the source for a volume group which can be a volume group backup ID, a volume group ID, or a volume ID.</p> <p>For more information, see Section 7.6, “Managing Volume Groups”.</p>

Resource Type	Resource-specific Elements
Volume Group Backups	<p>Status - The current state of a volume group backup:</p> <ul style="list-style-type: none"> • Creating • Committed • Available • Terminating • Terminated • Faulty • Request Received <p>Backup Size (in GB) - The aggregate size of the volume group backup, in GBs, which is typically smaller than the size of a volume group depending on the space consumed on the volume group.</p> <p>For more information, see Section 7.7, "Backing Up Block Volumes".</p>

File Storage

The following file storage resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
File Systems	<p>State - The current state of the file system:</p> <ul style="list-style-type: none"> • Creating • Active • Deleting • Deleted <p>Utilization - The number of bytes consumed by the file system, including any snapshots. This number reflects the metered size of the file system and is updated asynchronously with respect to updates to the file system.</p> <p>For more information, see Section 8.6, "Managing File Systems".</p>
Mount Targets	<p>State - The current state of the mount target:</p> <ul style="list-style-type: none"> • Creating • Active • Deleting • Deleted • Failed <p>For more information, see Section 8.5, "Managing Mount Targets and Exports".</p>

Object Storage

The following object storage resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Object Storage	<p>Default Storage Tier - The storage tier type of Standard or Archive assigned to the bucket. A bucket is set to Standard tier by default, which means objects uploaded or copied to the bucket will be in the standard storage tier. When the Archive tier type is set explicitly for a bucket, objects uploaded or copied to the bucket will be stored in archive storage.</p> <p>Visibility - Whether or not this bucket is read only. By default, a bucket is not read-only. A bucket is set to read-only when it is configured as a destination in a replication policy.</p> <p>For more information, see Section 9.2, “Managing Buckets”.</p>

Networking

The following networking resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Virtual Cloud Networks	<p>Status - The current state of the virtual cloud network:</p> <ul style="list-style-type: none">• Provisioning• Available• Terminating• Terminated• Updating <p>CIDR Block - The list of IPv4 CIDR blocks the VCN uses.</p> <p>DNS Domain Name - Name of the associated DNS domain.</p> <p>For more information, see Section 4.1, “Managing VCNs and Subnets”.</p>
Dynamic Routing Gateways	<p>Status - The current state of the dynamic routing gateway:</p> <ul style="list-style-type: none">• Provisioning• Available• Terminating• Terminated <p>For more information, see Section 4.3.4, “Connecting to the On-Premises Network through a Dynamic Routing Gateway”.</p>

DNS

The following DNS resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Zones	<p>Status - The current state of the zone resource:</p> <ul style="list-style-type: none">• Creating• Active• Deleting• Deleted• Failed• Updating <p>Zone Type - The type of the zone which must be either Primary or Secondary.</p> <p>For more information, see Section 4.5, "Managing Public DNS Zones".</p>
Steering Policies	<p>Status - The current state of the steering policy:</p> <ul style="list-style-type: none">• Creating• Active• Deleting• Deleted <p>Policy Type - The type of steering policy which is either Load Balancer or IP Prefix Steering. Load Balancer policies allow distribution of traffic across multiple endpoints. Endpoints can be assigned equal weights to distribute traffic evenly across the endpoints or custom weights may be assigned for ratio load balancing. IP Prefix steering policies enable customers to steer DNS traffic based on the IP Prefix of the originating query.</p> <p>For more information, see Section 4.6, "Managing Traffic with Steering Policies".</p>
TSIG Keys	<p>Status - The current state of the tag signature key:</p> <ul style="list-style-type: none">• Creating• Active• Deleting• Deleted• Failed• Updating <p>Algorithm - The type of algorithm used for the tag signature key:</p> <ul style="list-style-type: none">• hmac-md5• hmac-sha1

Resource Type	Resource-specific Elements
	<ul style="list-style-type: none">• hmac-sha224• hmac-sha256• hmac-h384• hmac-sha512 <p>For more information, see Section 4.5.4, “Working with Transaction Signature Keys”.</p>

Identity

The following identity resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Users	<p>Status - The current state of the user:</p> <ul style="list-style-type: none">• Creating• Active• Inactive• Deleting• Deleted <p>Email - The email address assigned to the user which does not have to be unique across all users in the tenancy; multiple user accounts can have the same email address.</p> <p>For more information, see Section 2.1, “Creating and Managing User Accounts”.</p>
Federation	<p>Status - The current state of an identity provider:</p> <ul style="list-style-type: none">• Creating• Active• Inactive• Deleting• Deleted <p>Type - The type identity provider service or product which is either Security Assertion Markup Language (SAML) 2.0 protocol or Microsoft Active Directory Federation Services (ADFS).</p> <p>Redirect URL - The identity provider-provided URL that enables a service provider to get required information to federate with that identity provider.</p> <p>For more information, see Section 2.5, “Federating with Microsoft Active Directory”.</p>
Groups	<p>Status - The current state of a group:</p> <ul style="list-style-type: none">• Creating

Resource Type	Resource-specific Elements
	<ul style="list-style-type: none"> • Active • Inactive • Deleting • Deleted <p>For more information, see Section 2.2, “Creating and Managing User Groups”.</p>
Policies	<p>Status - The current state of a policy:</p> <ul style="list-style-type: none"> • Creating • Active • Inactive • Deleting • Deleted <p>Statements - The number of statements attached to a policy.</p> <p>For more information, see Section 2.4, “Managing Policies”.</p>
Compartments	<p>Status - The current state of a compartment:</p> <ul style="list-style-type: none"> • Creating • Active • Inactive • Deleting • Deleted <p>For more information, see Section 2.3, “Creating and Managing Compartments”.</p>

Governance

The following governance resource types have additional information in their resource list.

Resource Type	Resource-specific Elements
Tag Namespaces	<p>Status - The current state of a tag namespace:</p> <ul style="list-style-type: none"> • Active • Inactive • Deleting • Deleted <p>For more information, see Section 3.1, “Creating and Managing Tag Namespaces”.</p>

1.1.3.2 About Resource Detail Pages

Each resource listed on a resource type page has its own page with additional details about it. Every resource detail page has some general information about the resource, such as its OCID, when it was created, the compartment it is in, and any tags associated with it.

Additionally, there is a Resources section where you can take additional actions upon the resource or view more information. All the available properties of the current item can be viewed in this area. Controls to update or add to these properties are also displayed within the Resources section. Most often this is the only place you can accomplish these actions on the resource or see this information. For example, the only way to attach an instance to a boot volume and to see the compatible shapes on a custom image is through their respective resource detail pages.

The following table shows you some of the tasks you can only do from a resource's detail page.

Task	Resource Detail Page
Create a volume group clone	Volume group
Create a schedule for a backup policy	Backup policy
Create file system export	Mount target
Create file system snapshot or create file system export	File system
Upload an API key	User
View or configure policy statements	Policy
View or create a tag key definition	Tag namespace
Attach a DRG to a VCN	Dynamic routing gateway
View or add the following for VCNs:	Virtual cloud network
<ul style="list-style-type: none"> • Subnets • Route Tables <ul style="list-style-type: none"> View or add route rules from a route table's detail page • Internet Gateways • Local Peering Gateway • DHCP Options • Security Lists <ul style="list-style-type: none"> View or create ingress or egress rules from a security list's detail page • NAT Gateways • Network Security Groups • Service Gateways • Dynamic Routing Gateway 	
View or create DNS zone records	Zone
View or add attached domains page	Steering policy

1.1.4 Locating Tenancy and Profile Information

For many tasks in Oracle Private Cloud Appliance OCI CLI you need the tenancy OCID, which you can find on a tenancy's detail page in the Compute Web UI. You can find this page by clicking your user name in the top menu bar and selecting Tenancy or using the OCI CLI after you have installed and configured it.

A tenancy detail page provides you with some general information (which includes the OCID), object storage settings, and any associated tags. Within the Compute Web UI you cannot make any changes to a tenancy; rather, this is done by an administrator of the Service Web UI. For more information about tenancies, see [Section 2.3.1, "Understanding the Tenancy"](#).

Every user in the system has an associated profile. The information in a user's profile can be found in the user detail pages. Users with administrative privileges (or through group membership or policies) have access to all user profiles.

You can find your profile page by logging into a tenancy for which you have access, clicking your user name in the top menu bar and selecting Profile. From your profile page, you can view general information and your OCID, view any tags associated with your profile, and view, add or delete API keys. You can also see which groups you belong to; however, you cannot change any of your group assignments unless you have administrative privileges.

For more information about user profiles, see [Section 2.1, "Creating and Managing User Accounts"](#).

1.2 Using the OCI CLI

This section provides instructions for installing and configuring the OCI CLI as well as some general information to help you use it. The rest of the Oracle Private Cloud Appliance User Guide shows you how to use the OCI CLI to complete tasks within the context of the step-by-step procedures.

The OCI CLI is the command line interface to the Compute Enclave. You can use the OCI CLI on its own or with the Compute Web UI to complete tasks. The OCI CLI provides the same core functionality as the Compute Web UI, plus additional commands, such as the ability to run scripts that extend the functionality. The OCI CLI's functionality is based on REST APIs which you can access from a browser with this URL:

```
https://console.pcasysname.mycompanydomain/api-reference
```

where *pcasysname* is the name of your Oracle Private Cloud Appliance and *mycompanydomain* is your domain. You can find the system and domain names on the dashboard of the Service Web UI or you can ask an administrator for this information.

1.2.1 Before You Begin

To install and use the OCI CLI, you must have:

- A user account for the Compute Web UI.
- An RSA public/private keypair used for signing API requests, with the public key uploaded for the user through the Compute Web UI.

Important

In the configuration steps, you have the option of using existing API public and private keys or creating new keys. If you do not already have an existing key pair, we recommend that you create them as part of the manual or automated OCI CLI configuration.

- An Oracle Private Cloud Appliance self-signed certificate.

This requirement is satisfied during the configuration steps.

You can install the OCI CLI on Mac OS, Windows, or any supported Linux/Unix operating system:

- Oracle Linux 7 and Oracle Linux 8
- CentOS 7.0 and CentOS 8.x
- Ubuntu 16.04, Ubuntu 18.04, and Ubuntu 20.04

1.2.2 Installing the OCI CLI

You can install the OCI CLI on Oracle Linux or Mac OS operating systems using a package manager. If you are using Windows or some other operating system, you use the install script.

Important

If you already have the OCI CLI installed and configured, you can skip to [Section 1.2.3, “Configuring the OCI CLI”](#) to learn how to further configure the CLI for Oracle Private Cloud Appliance.

To install the OCI CLI, its dependencies, and Python, follow the steps for your operating system.

Oracle Linux 8

Run the following commands to install the OCI CLI:

```
$ sudo dnf -y install oraclelinux-developer-release-el8
$ sudo dnf install python36-oci-cli
```

To uninstall the OCI CLI, run:

```
$ sudo dnf remove python36-oci-cli
```

Oracle Linux 7

Run the following command to install the OCI CLI:

```
$ sudo yum install python36-oci-cli
```

To uninstall the OCI CLI, run:

```
$ sudo yum remove python36-oci-cli
```

Mac OS

You can use [Homebrew](#) to install, upgrade, and uninstall the OCI CLI on Mac OS.

Note

Optionally, you can install the OCI CLI using the install script. See *Using the Install Script for Other Operating Systems* in this section for details.

- To install the OCI CLI, run:

```
$ brew update && brew install oci-cli
```

- To upgrade the OCI CLI, run:

```
$ brew update && brew upgrade oci-cli
```

- To uninstall the OCI CLI, run:

```
$ brew uninstall oci-cli
```

Windows

You can use Windows PowerShell to install the OCI CLI.

1. Open the PowerShell console using the Run as Administrator option.
2. The installer enables auto-complete by installing and running a script. To allow this script to run, you must enable the RemoteSigned execution policy.

To configure the remote execution policy for PowerShell, run the following command:

```
$ Set-ExecutionPolicy RemoteSigned
```

3. Force PowerShell to use TLS 1.2 for Windows 2012 and Windows 2016:

```
$ [Net.ServicePointManager]::SecurityProtocol = [Net.SecurityProtocolType]::Tls12
```

4. Download the installer script:

```
$ Invoke-WebRequest https://raw.githubusercontent.com/oracle/oci-cli/master/scripts/ ^
install/install.ps1 -OutFile install.ps1
```

5. Run the installer script with or without prompts:
 - a. To run the installer script with prompts, run the following command and respond to the installation script prompts:

```
$ iex ((New-Object System.Net.WebClient).DownloadString ^
('https://raw.githubusercontent.com/oracle/oci-cli/master/scripts/install ^
/install.ps1'))
```

- b. To run the installer script without prompts and accept the default settings:

```
$ install.ps1 -AcceptAllDefaults
```

Using the Install Script for Other Operating Systems

For any other operating system, you can use the install script to install the OCI CLI, its dependencies, and Python.

1. From a terminal, enter the following command to run the installer script:

```
$ bash -c "$(curl -L https://raw.githubusercontent.com/oracle/oci-cli/master/scripts/ ^
install/install.sh)"
```

Note

To run a silent install that accepts all default values with no prompts, use the `--accept-all-defaults` option.

2. Respond to the installation script prompts.
 - If you do not have a compatible version of Python installed in Linux, you are prompted to provide a location for installing the binaries and executables. The script installs Python for you.
 - If you do not have a compatible version of Python installed in MacOS, you are notified that your version of Python is incompatible. You must upgrade before you can proceed with the installation. The script does not install Python for you.

- When prompted to upgrade the OCI CLI to the newest version, respond with **Y** to overwrite an existing installation.
- When prompted to update your PATH, respond with **Y** to be able to invoke the OCI CLI without providing the full path to the executable.

1.2.3 Configuring the OCI CLI

Before using the OCI CLI, you must configure it for working with Oracle Private Cloud Appliance and obtain the system's certificate authority (CA) chain. You can complete the configuration manually or use the config setup tool to help you.

Important

If you are already using the OCI CLI and have it configured for other purposes, read this section entirely before proceeding with any of the configuration steps.

1.2.3.1 Obtain the Required Information

Whether you manually configure the OCI CLI or use the setup config tool, there is required information you must provide for the configuration file. Before you begin the configuration process, ensure you have the following:

- User OCID

The user's OCID is in `ocid1.user...unique_ID` format and you can copy it from the user details page in the Compute Web UI. To navigate to your user details page, click your user name in the Compute Web UI dashboard and then click My Profile.

- Tenancy OCID

The tenancy OCID is in `ocid1.tenancy...unique_ID` format and you can copy it from the tenancy details page in the Compute Web UI. To navigate to the tenancy details page, click your user name in the Compute Web UI dashboard and then click Tenancy.

- Region name

The region name is in `pcasys1.example.com` format, where `pcasys1` is the name of your Oracle Private Cloud Appliance system and `example.com` is your domain.

If you have access to the Service Web UI, you can find the system and domain names on the dashboard. Otherwise, ask a Service Web UI administrator for the information.

If you do not already have an existing API public and private key pair, we recommend that you create them as part of the manual or automated OCI CLI configuration. For more information, see the [Section 1.2.3.2, "Manual Configuration"](#) or [Section 1.2.3.3, "Automated Configuration"](#) section.

If you have an existing API public and private key pair that you want to use, make sure:

- They are in PEM format.
- Your public key is added to your user profile.
- You know the full path and file name of the private key, for example, `~/oci/oci_api_key.pem`.
- You have your public key fingerprint, which is in `xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx` format. You can find this fingerprint on your profile page in the Compute Web UI or through a terminal using a command, for example:

```
openssl rsa -pubout -outform DER -in ~/.oci/oci_api_key.pem | openssl md5 -c
```

1.2.3.2 Manual Configuration

Complete the following steps to manually configure the OCI CLI for Oracle Private Cloud Appliance. Ensure you have gathered all the required information.

The steps below assume you are on a Linux system and that you have already created a user through the Compute Web UI. However, the basic procedure is the same for other system types.

1. From a terminal, log into the system where you installed the OCI CLI and create an API key pair, for example:

```
$ oci setup keys

Enter a passphrase for your private key (empty for no passphrase):
Public key written to: /home/myuserdir/.oci/oci_api_key_public.pem
Private key written to: /home/myuserdir/.oci/oci_api_key.pem
Public key fingerprint: xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
```

2. From a browser, log into the Compute Web UI.
3. Navigate to your user details page, click your user name in the dashboard, and then click My Profile. Your user details page is displayed.
4. In the Resources section, click API Keys and then click Add API Key.
5. Navigate to the location of your public key or paste the public key contents and then click Upload Key.
6. In the `/home/myuserdir/.oci` directory, create a new file named `config` and add a new profile section and the required information, for example:

```
$ vi config

[DEFAULT]
user=ocidl.user...unique_id
key_file=/home/myuserdir/.oci/oci_api_key.pem
tenancy=ocidl.tenancy...unique_id
region=pcasys1.example.com
fingerprint=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
```

where `pcasys1` is the name of your Oracle Private Cloud Appliance and `example.com` is your domain.

If you have access to the Service Web UI, you can find the system and domain names on the dashboard. Otherwise, ask a Service Web UI administrator for the information.

1.2.3.3 Automated Configuration

If this is the first time you are using the OCI CLI, the setup config tool helps you walk you through setup process. When you enter the `oci setup config` command it prompts you for the information required for the config file and the API public/private keys and then generates an API key pair and creates the config file.

To configure the OCI CLI using the setup config tool:

1. From a command window, enter `oci setup config` and follow the prompts, for example:

```
$ oci setup config
```

```
This command provides a walkthrough of creating a valid CLI config file.
Enter a location for your config [/home/myuserdir/.oci/config]:
Enter a user OCID: ocidl.user...uniqueid
Enter a tenancy OCID: ocidl.tenancy...uniqueid
```

Important

For the step *Enter a region by index or name*, you cannot enter the region in the required `system.domain` format. Rather, enter any value from the list as the value is meaningless to Oracle Private Cloud Appliance. In Step 2 you will modify the config file to provide the information needed by Oracle Private Cloud Appliance.

```
Enter a region by index or name(e.g.
1: ap-chiyoda-1, 2: ap-chuncheon-1, 3: ap-hyderabad-1, 4: ap-melbourne-1, 5: ap-mumbai-1,
6: ap-osaka-1, 7: ap-seoul-1, 8: ap-sydney-1, 9: ap-tokyo-1, 10: ca-montreal-1,
11: ca-toronto-1, 12: eu-amsterdam-1, 13: eu-frankfurt-1, 14: eu-zurich-1, 15: me-dubai-1,
16: me-jeddah-1, 17: sa-santiago-1, 18: sa-saopaulo-1, 19: sa-vinhedo-1, 20: uk-cardiff-1,
21: uk-gov-cardiff-1, 22: uk-gov-london-1, 23: uk-london-1, 24: us-ashburn-1,
25: us-gov-ashburn-1, 26: us-gov-chicago-1, 27: us-gov-phoenix-1, 28: us-langley-1,
29: us-luke-1, 30: us-phoenix-1, 31: us-sanjose-1): 24
```

```
Do you want to generate a new API Signing RSA key pair?
(If you decline you will be asked to supply the path to an existing key.) [Y/n]: Y
Enter a directory for your keys to be created [/home/myuserdir/.oci]:
Enter a name for your key [oci_api_key]:
Public key written to: /home/myuserdir/.oci/oci_api_key_public.pem
Enter a passphrase for your private key (empty for no passphrase):
Private key written to: /home/myuserdir/.oci/oci_api_key.pem

Fingerprint: xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
Config written to /home/myuserdir/.oci/config
```

2. Navigate to the `~/home/myuserdir/.oci` directory and modify the `config` file to use the correct region, for example:

```
$ vi config

[DEFAULT]
user=ocidl.user...unique_id
fingerprint=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
key_file=/home/myuserdir/.oci/oci_api_key.pem
tenancy=ocidl.tenancy...unique_id
region=pcasys1.example.com
```

where `pcasys1` is the name of your Oracle Private Cloud Appliance and `example.com` is your domain.

If you have access to the Service Web UI, you can find the system and domain names on the dashboard. Otherwise, ask a Service Web UI administrator for the information.

3. If you haven't already, upload your API Signing public key through the Compute Web UI. For more information, see [Section 1.2.6.2, "Adding an API Public Key to a User Profile"](#).

1.2.3.4 Obtaining the Certificate Authority Bundle

Whether you configured the OCI CLI manually or used the automated tool, you must obtain the Oracle Private Cloud Appliance external silo CA chain before you can run any commands.

The external silo CA chain must be copied to the system where you are installing the OCI CLI and referenced in the `oci_cli_rc` file.

1. Navigate to the `~/home/myuserdir/.oci` directory.

Important

As a best practice, consider creating a subdirectory within the `.oci` directory for each profile you have in the `.oci/config` file. For example, if you have sections `[PCA1]` and `[PCA2]` for two different Oracle Private Cloud Appliance systems, create two subdirectories called `PCA1` and `PCA2` to store the API keys and external silo CA chain.

2. Copy the external silo CA chain:

```
curl -ko ca.crt https://iaas.system-name.domain-name/cachain
```

3. In the `~/home/myuserdir/.oci` directory, create a new file named `oci_cli_rc`. Using the same profile name, add the path to the external silo CA chain, for example:

```
$ vi oci_cli_rc
[DEFAULT]
cert-bundle=/home/myuserdir/.oci/ca.crt
```

1.2.4 Testing the OCI CLI Configuration

After you have installed and configured the OCI CLI, you can enter a few list commands to verify it is working correctly. For example:

```
$ oci iam compartment list --include-root
{
  "data": [
    {
      "compartment-id": null,
      "defined-tags": {},
      "description": "Root compartment",
      "freeform-tags": {},
      "id": "ocidl.tenancy...unique-id",
      "inactive-status": null,
      "is-accessible": null,
      "lifecycle-state": "ACTIVE",
      "name": "pca3tenantdefault",
      "time-created": "2021-10-13T00:39:50.916561+00:00"
    },
    {
      "compartment-id": "ocidl.tenancy...unique-id",
      "defined-tags": {},
      "description": "Compartment 1",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique-id",
      "inactive-status": null,
      "is-accessible": null,
      "lifecycle-state": "ACTIVE",
      "name": "pca3compartment1",
      "time-created": "2021-10-13T15:51:28.245810+00:00"
    },
    {
      "compartment-id": "ocidl.tenancy...unique-id",
      "defined-tags": {},
      "description": "Compartment 2",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique-id",
      "inactive-status": null,
      "is-accessible": null,
      "lifecycle-state": "ACTIVE",
      "name": "pca3compartment2",
```

```

    "time-created": "2021-10-13T17:16:53.413803+00:00"
  },
]
}

```

```

$ oci iam user list
{
  "data": [
    {
      "capabilities": null,
      "compartment-id": "ocidl.tenancy...unique-id",
      "defined-tags": {},
      "description": "Default Administrator User",
      "email": null,
      "email-verified": null,
      "external-identifier": null,
      "freeform-tags": {},
      "id": "ocidl.user...unique-id",
      "identity-provider-id": null,
      "inactive-status": null,
      "is-mfa-activated": null,
      "last-successful-login-time": null,
      "lifecycle-state": "ACTIVE",
      "name": "admin",
      "previous-successful-login-time": null,
      "time-created": "2021-10-13T00:39:50.956416+00:00"
    },
    {
      "capabilities": null,
      "compartment-id": "ocidl.tenancy...unique-id",
      "defined-tags": {},
      "description": "Compute Admin User",
      "email": "null",
      "email-verified": null,
      "external-identifier": null,
      "freeform-tags": {},
      "id": "ocidl.user...unique-id",
      "identity-provider-id": null,
      "inactive-status": null,
      "is-mfa-activated": null,
      "last-successful-login-time": null,
      "lifecycle-state": "ACTIVE",
      "name": "comp-admin1",
      "previous-successful-login-time": null,
      "time-created": "2021-10-13T22:05:33.993284+00:00"
    },
    {
      "capabilities": null,
      "compartment-id": "ocidl.tenancy...unique-id",
      "defined-tags": {},
      "description": "Network Admin User",
      "email": null,
      "email-verified": null,
      "external-identifier": null,
      "freeform-tags": {},
      "id": "ocidl.user...unique-id",
      "identity-provider-id": null,
      "inactive-status": null,
      "is-mfa-activated": null,
      "last-successful-login-time": null,
      "lifecycle-state": "ACTIVE",
      "name": "nw-admin1",
      "previous-successful-login-time": null,
      "time-created": "2021-10-13T22:31:58.616036+00:00"
    }
  ]
}

```

1.2.5 Using Multiple Profiles

You can use the OCI CLI configuration files (`config` and `oci_cli_rc`) to store more than one profile. Each profile section in the `config` file references a tenancy within Oracle Private Cloud Appliance; however, these tenancies do not have to exist within the same Oracle Private Cloud Appliance, for example:

```
[DEFAULT]
user=ocidl.user...unique_id
fingerprint=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
key_file=/home/myuserdir/.oci/oci_api_key.pem
tenancy=ocidl.tenancy...unique_id
region=mypcaname.mydomainname

[PCA2]
user=ocidl.user...unique_id
fingerprint=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
key_file=/root/.oci/oci_api_key.pem
tenancy=ocidl.tenancy...unique_id
region=mypca2name.mydomainname

[PCA3]
user=ocidl.user...unique_id
fingerprint=xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx
key_file=/root/.oci/oci_api_key.pem
tenancy=ocidl.tenancy...unique_id
region=mypca3name.mydomainname
```

Important

If you do not specify a profile in your commands, the default profile is used. If you do not have a default profile, you must pass `--profile` in your commands.

If you use multiple profiles in your `config` file, then you must setup your `oci_cli_rc` file for multiple profiles as well, for example:

```
[DEFAULT]
cert-bundle=/home/myuserdir/.oci/default-ca/ca-crt

[PCA2]
cert-bundle=/home/myuserdir/.oci/pca2-ca/ca-crt

[PCA3]
cert-bundle=/home/myuserdir/.oci/pca3-ca/ca-crt
```

Important

As a best practice, consider creating a subdirectory within the `.oci` directory for each profile you have in the `.oci/config` file. For example, if you have sections `[DEFAULT]`, `[PCA2]` and `[PCA3]` for three different Oracle Private Cloud Appliance systems, create three subdirectories called `default-ca`, `pca2-ca` and `pca3-ca` to store the API keys and external silo CA chain.

By having multiple profiles in the same configuration files, you can easily specify which tenancy profile you want to pass a command against. If you have a `[DEFAULT]` profile, you can run a command without using the `--profile` option. For any other profile, you must use the `--profile PROFILE-NAME` to pass a command against a specific profile, for example:

```
$ oci --profile PCA2 iam compartment list
```



```

{
  "data": [
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Network compartment",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "PCA2-Network"
    },
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Instance compartment",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "PCA2-Instance"
    },
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Test compartment",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "PCA2-Test"
    }
  ]
}

```

```
$ oci --profile PCA3 iam compartment list
```

```

{
  "data": [
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Compartment for block storage",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "Block-storage"
    },
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Compartment for networking",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "Networks"
    },
    {
      "compartment-id": "ocidl.tenancy...unique_id",
      "defined-tags": {},
      "description": "Compartment for instances",
      "freeform-tags": {},
      "id": "ocidl.compartment...unique_id",
      "lifecycle-state": "ACTIVE",
      "name": "Instances"
    }
  ]
}

```

1.2.6 Working with API Signing Keys

If you need to use the OCI CLI or make REST API requests, you must have an API signing public and private key pair. API requests are signed with the private key and the public key is used to verify the authenticity of the request. The private key is stored locally and the public key is uploaded to a user account. You can have a maximum of three (3) public keys per user account.

Important

The API signing key pair is **not** the SSH key that you use to access compute instances. Both the private key and public key (minimum 2048 bits) must be in PEM format (not SSH-RSA format).

1.2.6.1 Generating an API Key Pair

If you do not already have an existing API signing public and private key pair, we recommend that you create the key pair as part of the manual or automated configuration. To do so, use the `oci setup keys` command as shown in the [Section 1.2.3.2, “Manual Configuration”](#) section or follow the prompts in the [Section 1.2.3.3, “Automated Configuration”](#) section.

If you want to create a key pair independent of the OCI CLI configuration, the following sections show you how to do this on Linux, Mac, and Windows operating systems. You can then use these keys when you configure the OCI CLI.

Using Linux or Mac OS X

1. Generate the private key.
 - Generate the key encrypted with a passphrase:

```
$ openssl genrsa -out ~/.oci/oci_api_key.pem -aes128 2048
```

Note

Use of a passphrase is strongly recommended.

- Generate the key with no passphrase:

```
$ openssl genrsa -out ~/.oci/oci_api_key.pem 2048
```

2. Check the permission on the private key file and change if necessary.

The file permission should be 600 or 400 to ensure that only you can read the private key file.

3. Generate the public key from your new private key:

```
$ openssl rsa -pubout -in ~/.oci/oci_api_key.pem -out ~/.oci/oci_api_key_public.pem
```

This public key file can have the same permissions as the private key file or can be readable by everyone.

Using Windows

1. Install Git Bash for Windows.

See <https://git-scm.com/download/win>.

2. Include the OpenSSL binary in your Windows path.

On default installations, the `openssl.exe` binary is in the following directory:

```
C:\Program Files\Git\mingw64\bin
```

3. Generate the private key.

- Generate the key encrypted with a passphrase:

```
$ openssl genrsa -out %HOMEDRIVE%%HOMEPATH%\oci\oci_api_key.pem -aes128 -passout ^
stdin 2048
```

Note

Use of a passphrase is strongly recommended.

- Generate the key with no passphrase:

```
$ openssl genrsa -out %HOMEDRIVE%%HOMEPATH%\oci\oci_api_key.pem 2048
```

4. Check the permission on the private key file and change if necessary.

The file permission should be set so that only you can read the private key file.

5. Generate the public key from your new private key:

```
$ openssl rsa -pubout -in %HOMEDRIVE%%HOMEPATH%\oci\oci_api_key.pem -out ^
%HOMEDRIVE%%HOMEPATH%\oci\oci_api_key_public.pem
```

This public key file can have the same permissions as the private key file or can be readable by everyone.

1.2.6.2 Adding an API Public Key to a User Profile

You add your own API public key to your profile using the Compute Web UI. If you don't have a login and password for the Compute Web UI, contact an administrator.

Using the Compute Web UI

1. From a browser, log into the Compute Web UI.
2. Navigate to your user details page, click your user name in the dashboard, and then click My Profile.
Your user details page is displayed.
3. In the Resources section, click API Keys and then click Add API Key.
4. In the Add Public Key dialog, navigate to the location of your public key or paste the public key contents and then click Upload Key.

Users can have a maximum of three (3) public keys added to their profile.

Using the OCI CLI

After you've installed and configured the OCI CLI, you can also use the `api-key upload` command to upload additional keys. If you have more than one API public key, you must specify the key's fingerprint to indicate which key you're using to sign the request.

1. Get the OCID of the user that needs an API signing key (`oci iam user list`).
2. Use the user API key list command to ensure that your account does not already have the maximum three API signing keys.

Syntax

```
$ oci iam user api-key list --user-id user_OCID
```

3. Run the API key upload command.

Syntax

```
$ oci iam user api-key upload --user-id ocid1.user.unique_ID { --key key | --key-file \
file://keyfile.pem }
```

- *key* – an RSA public key in PEM format
- *keyfile.pem* – a file that contains an RSA public key in PEM format

1.2.6.3 Finding an API Public Key Fingerprint

Your public key fingerprint, which is in `xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx` format, can be found on your profile page in the Compute Web UI or through a terminal using the following OpenSSL commands:

For Linux and Mac OS X:

```
$ openssl rsa -pubout -outform DER -in ~/.oci/oci_api_key.pem | openssl md5 -c
```

For Windows:

```
$ openssl rsa -pubout -outform DER -in \.oci\oci_api_key.pem | openssl md5 -c
```

1.2.6.4 Deleting an API Signing Key from a User Profile

You can delete your own API signing keys and tenancy administrators can delete API signing keys for any user in their tenancy.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.
2. Click the name of the user account for which you want to delete an API signing key.
3. Scroll to the Resources section of the user details page.
4. For the API key that you want to delete, click the Actions icon (three dots) and then click Delete.

Using the OCI CLI

1. Get the following information:
 - The OCID of the user account from which you want to delete an API signing key.

```
$ oci iam user list
```

- The fingerprint of the API signing key that you want to delete.

```
$ oci iam user api-key list --user-id user_OCID
```

2. Run the user API key delete command.

Syntax

```
$ oci iam user api-key delete --user-id user_OCID --fingerprint fingerprint
```

1.2.7 Understanding Command Syntax and Finding Help

This section provides some basic information to help you along as you begin using the OCI CLI, such as command syntax, how to find OCIDs, and where to get help with commands.

1.2.7.1 Command Syntax

In general, commands entered in the OCI CLI have the following syntax:

```
$ oci <service> <type> <action> <required-parameters> <optional-parameters>
```

For example, in the `oci iam user create --name joeb --description "user for tenancy admin" --email joeb@example.com` command:

- `iam` is the service
- `user` is the resource type
- `create` is the action
- `name` and `description` are the required parameters
- `email` is an optional parameter

1.2.7.2 Obtaining OCIDs

When you use the OCI CLI, the majority of commands require an OCID. In general,

- - `list` commands require the OCID of the compartment where you are looking for the resource.
- - `create` commands require the OCID of the compartment where you want to create the resource.
- - `get`, - `update`, - `delete` commands require the OCID of the resource.
- - `move` commands require the OCID of the resource and the OCID of the destination compartment.

Some commands require the OCID of a different resource. For example, creating a DRG route table requires the OCID of the DRG, and adding or removing a route rule requires the OCID of the route table.

You can find OCIDs using the OCI CLI or the Compute Web UI. The following lists show you how to find the most commonly needed OCIDs using the OCI CLI.

Block volume service OCIDs

- Boot volume

```
$ oci bv boot-volume list --availability-domain availability_domain_name --compartment-id \
compartment_OCID
```

- Volume

```
$ oci bv volume list --compartment-id compartment_OCID
```

- Volume backup policy

```
$ oci bv volume-backup-policy list --compartment-id compartment_OCID
```

- Volume group

```
$ oci bv volume-group list --compartment-id compartment_OCID
```

Compute service OCIDs

- Instance

```
$ oci compute instance list --compartment-id compartment_OCID
```

- Instance VNIC

```
$ oci compute instance list-vnics --compartment-id compartment_OCID
```

- Volume attachment

```
$ oci compute volume-attachment list --compartment-id compartment_OCID
```

Identity and access management service OCIDs

- Availability domain name

```
$ oci iam availability-domain list
```

- Compartments within a tenancy

```
$ oci iam compartment list
```

- Compartments within a tenancy and including the tenancy

```
$ oci iam compartment list --include-root
```

- Compartments and all sub-compartments in a tenancy

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- Compartment including its sub-compartments

```
$ oci iam compartment list --compartment-id compartment_OCID --compartment-id-in-subtree
```

- Group

```
$ oci iam group list
```

- Policy

```
$ oci iam policy list --compartment-id compartment_OCID
```

- Tag namespace

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

- User

```
$ oci iam user list
```

Network service OCIDs

- DHCP options

```
$ oci network dhcp-options list --compartment-id compartment_OCID
```

- Route table

```
$ oci network route-table list --compartment-id compartment_OCID
```

- Subnet

```
$ oci network subnet list --compartment-id compartment_OCID
```

- VCN

```
$ oci network vcn list --compartment-id compartment_OCID
```

1.2.7.3 Getting Help with Commands

You can get inline help by appending `--help`, `-h` or `-?` to a command:

- `oci --help` returns a list of supported commands and general command options.
- `oci service --help` returns the complete command reference for a service. For example, `oci iam --help` returns:

```
NAME
    iam -

DESCRIPTION
    OCI CLI for managing users, groups, compartments, and policies.

AVAILABLE COMMANDS
    o auth-token
    o create
    o delete
    o list
    o update
    o authentication-policy
    o get
    o update
    o availability-domain
    o list
    o bulk-action-resource-type-collection
    o list
    o compartment
    o bulk-delete-resources
    o bulk-move-resources
    ...
```

- `oci service type --help` returns the complete command reference for a service type. For example `oci iam user --help` returns:

```
NAME
    iam_user -

DESCRIPTION
    An individual employee or system that needs to manage or use your
    company's Oracle Cloud Infrastructure resources. Users might need to
    launch instances, manage remote disks, work with your cloud network,
    etc.
    ...

    These users are created directly within the Oracle Cloud Infrastructure
    system, via the IAM service. They are different from federated users,
    who authenticate themselves to the Oracle Cloud Infrastructure Console
    via an identity provider.
    ...

AVAILABLE COMMANDS
    o api-key
```

```

    o delete
    o list
    o upload
o create
o delete
o get
o list
o list-groups
...

```

- `oci service type action --help` returns the complete command reference for a service type action. For example, `oci iam user list --help` returns:

```

NAME
    iam_user_list -
    o Description
    o Usage
    o Optional Parameters
    o Global Parameters
    o Examples

DESCRIPTION
    Lists the users in your tenancy. You must specify your tenancy's OCID
    as the value for the compartment ID (remember that the tenancy is
    simply the root compartment).

USAGE
    oci iam user list [OPTIONS]

OPTIONAL PARAMETERS
    --all
    Fetches all pages of results. If you provide this option, then you
    cannot provide the --limit option.

    --compartment-id, -c [text]
    The OCID of the compartment (remember that the tenancy is simply the
    root compartment). If not provided, this parameter will use the
    tenancy's OCID (root compartment's OCID) from the config file.

    --external-identifier [text]
    The id of a user in the identity provider.
    ...

```

For more information, see the [Oracle Cloud Infrastructure CLI Command Reference](#).

1.2.8 Using JSON for Complex Command Input

Complex input, such as arrays and objects with more than one value, are passed in JSON format and can be provided as a string at the command line, as a file, or as a command line string and as a file. You can combine arguments on the command line with file input. However, if the same values are provided in a file and on the command line, the command line takes precedence.

1.2.8.1 Using a JSON String

To pass complex input to the OCI CLI as a JSON string on MacOS, Linux or Unix systems, you enclose the JSON string with single quotes. On Windows systems, you must enclose the JSON string in double quotes. And, within the JSON string, each double quote for the key and value pairs must be escaped with a backslash (\) character.

For example:

- MacOS, Linux or Unix


```
$ oci os bucket create -ns mynamespace --name mybucket --metadata \
'{"key1":"value1","key2":"value2"}' --compartment-id <compartment_OCID>
```

- Windows

```
$ oci os bucket create -ns mynamespace --name mybucket --metadata ^
"{\"key1\": \"value1\", \"key2\": \"value2\"}" --compartment-id <compartment_OCID>
```

Note

If you receive a JSON error message "Parameter '<PARAMETER NAME>' must be in JSON format.", this indicates that the value you passed for the parameter was not valid JSON. Typically, this type of error message is the result of the JSON string not being escaped correctly.

1.2.8.2 Using a JSON File

To pass complex input using a JSON file, you reference it from the command line by providing a path to the file using the `file://` prefix.

The following types of file paths are supported:

- Relative paths from the same directory
- Absolute paths on Linux, MacOS or Unix
- Full file paths on Windows

The following are examples of file locations:

- Your home directory

```
$ oci os bucket create -ns mynamespace --name mybucket --compartment-id \
ocidl.compartment.ocl...uniqueid --metadata file://~/testfile.json
```

- The current directory

```
$ oci os bucket create -ns mynamespace --name mybucket --compartment-id \
ocidl.compartment.ocl..uniqueid --metadata file://testfile.json
```

- The /tmp directory (Linux, Unix, or MacOS)

```
$ oci os bucket create -ns mynamespace --name mybucket --compartment-id \
ocidl.compartment.ocl..uniqueid --metadata file:///tmp/testfile.json
```

- The C:\temp directory (Windows)

```
$ oci os bucket create -ns mynamespace --name mybucket --compartment-id ^
ocidl.compartment.ocl..uniqueid --metadata file://C:\temp\testfile.json
```

1.2.8.3 Generating JSON Format

From the OCI CLI, you can generate the correct JSON format for an entire command or a complex type command option and save the resulting content to a file.

The `--generate-full-command-json-input` option can be used to generate the correct JSON to be used with a command. The key names are pre-populated and match the command option names. You must add the key values and save the JSON to a file before you can use it as an input to the command. For any command option that accepts multiple values, the value of the key can be a JSON array.

To generate the entire correct JSON format for a command, append `--generate-full-command-json-input` to the command.

Syntax

```
$ oci <service> <type> <action> --generate-full-command-json-input
```

Example

```
$ oci network vcn create --generate-full-command-json-input
{
  "cidrBlock": "string",
  "cidrBlocks": [
    "string",
    "string"
  ],
  "compartmentId": "string",
  "definedTags": {
    "tagNamespacel": {
      "tagKey1": "tagValue1",
      "tagKey2": "tagValue2"
    },
    "tagNamespace2": {
      "tagKey1": "tagValue1",
      "tagKey2": "tagValue2"
    }
  },
  "displayName": "string",
  "dnsLabel": "string",
  "freeformTags": {
    "tagKey1": "tagValue1",
    "tagKey2": "tagValue2"
  },
  "isIpv6Enabled": true,
  "maxWaitSeconds": 0,
  "waitForState": [
    "PROVISIONING|AVAILABLE|TERMINATING|TERMINATED|UPDATING"
  ],
  "waitIntervalSeconds": 0
}
```

A command option that is a complex type must use valid JSON for its value, which can be entered on the command line or passed in as a file using the `file://path/to/file` syntax. To generate the correct JSON format for a complex type command option, use `--generate-param-json-input` with the command.

Syntax

```
$ oci <service> <type> <action> --generate-param-json-input <command-option>
```

Example

```
$ oci network route-table create --generate-param-json-input route-rules
[
  {
    "cidrBlock": "string",
    "description": "string",
    "destination": "string",
    "destinationType": "string",
    "networkEntityId": "string"
  },
  {
    "cidrBlock": "string",
    "description": "string",

```

```

        "destination": "string",
        "destinationType": "string",
        "networkEntityId": "string"
    }
]

```

1.2.9 Formatting and Filtering Command Output

By default, all responses to a command are returned in JSON format. For example, a response like the following is returned when you issue the `oci iam region list` command:

```

{
  "data": [
    {
      "key": "pcasysname",
      "name": "pcasysname"
    }
  ]
}

```

For readability, command output can be formatted as a table. For example, issuing `oci iam region list --output table` returns a two-column table:

```

+-----+-----+
| key   | name   |
+-----+-----+
| pcasysname | pcasysname |
+-----+-----+

```

You can filter output using the JMESPath query option for JSON. Filtering is useful when dealing with large amounts of output. For example, if you issue the following command with the `--output table` option, too much data is returned and it overflows the width of the terminal. In addition, you might not need all the information that's returned.

```
$ oci compute image list -c <compartment_OCID> --output table
```

```

| base-image-id | compartment-id | create-image-allowed | display-name
+-----+-----+-----+-----+
| None          | None          | True                 | Windows-Server-2012-R2-Standard-Edition-VM-2017.
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Windows             | Serv
er 2012 R2 Standard | 2017-07-25T23:59:59.311000+00:00 |
| None          | None          | True                 | Windows-Server-2012-R2-Standard-Edition-VM-2017.
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Windows             | Serv
er 2012 R2 Standard | 2017-04-03T19:42:22.938000+00:00 |
| None          | None          | True                 | Windows-Server-2012-R2-Standard-Edition-BM-2017.
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Windows             | Serv
er 2012 R2 Standard | 2017-07-25T20:55:37.937000+00:00 |
| None          | None          | True                 | Windows-Server-2012-R2-Standard-Edition-BM-2017.
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Windows             | Serv
er 2012 R2 Standard | 2017-04-13T17:36:50.840000+00:00 |
| None          | None          | True                 | Oracle-Linux-7.4-2017.09.29-0
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Oracle Linux        | 7.4
| 2017-10-05T22:36:17.246000+00:00 |
| None          | None          | True                 | Oracle-Linux-7.4-2017.08.25-1
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Oracle Linux        | 7.4
| 2017-09-11T23:12:18.644000+00:00 |
| None          | None          | True                 | Oracle-Linux-7.4-2017.08.25-0
| 1.image.oc1.phx...uniqueid | AVAILABLE    | Oracle Linux        | 7.4| 2017-08-25T01:21:37.176000+00:00 |

```

You can limit the amount of data returned by combining the `--query` option with `--output table` to filter the information within the table format. For example:

```
$ oci compute image list -c <compartment_OCID> --output table --query "data [*].{ImageName: \"display-name\", OCID:id}"
```

ImageName	OCID
Windows-Server-2012-R2-Standard-Edition-VM-2017.07.25-0	ocidl.image.oc1.phx...uniqueid
Windows-Server-2012-R2-Standard-Edition-VM-2017.04.03-0	ocidl.image.oc1.phx...uniqueid
Windows-Server-2012-R2-Standard-Edition-BM-2017.07.25-0	ocidl.image.oc1.phx...uniqueid
Windows-Server-2012-R2-Standard-Edition-BM-2017.04.13-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-7.4-2017.09.29-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-7.4-2017.08.25-1	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-7.4-2017.08.25-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-7.3-2017.07.17-1	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-7.3-2017.07.17-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-6.9-2017.09.29-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-6.9-2017.08.25-0	ocidl.image.oc1.phx...uniqueid
Oracle-Linux-6.9-2017.07.17-0	ocidl.image.oc1.phx...uniqueid
CentOS-7-2017.09.14-0	ocidl.image.oc1.phx...uniqueid
CentOS-7-2017.07.17-0	ocidl.image.oc1.phx...uniqueid
CentOS-7-2017.04.18-0	ocidl.image.oc1.phx...uniqueid

For more information about the JMESPath query language for JSON, see [JMESPath](#).

Chapter 2 Identity and Access Management

Table of Contents

2.1 Creating and Managing User Accounts	35
2.1.1 Creating a User	36
2.1.2 Providing a Temporary Compute Web UI Password	38
2.1.3 Setting Your Own Compute Web UI Password	39
2.1.4 Viewing User Information and Group Membership	40
2.1.5 Adding a User to a Group by Updating the User	40
2.1.6 Removing a User from a Group by Updating the User	41
2.1.7 Modifying a User	41
2.1.8 Deleting a User	42
2.2 Creating and Managing User Groups	42
2.2.1 Creating a Group	43
2.2.2 Viewing Group Information and Group Membership	44
2.2.3 Adding a User to a Group by Updating the Group	45
2.2.4 Removing a User from a Group by Updating the Group	45
2.2.5 Modifying a Group	46
2.2.6 Deleting a Group	46
2.3 Creating and Managing Compartments	47
2.3.1 Understanding the Tenancy	47
2.3.2 Listing Compartments	48
2.3.3 Creating a Compartment	49
2.3.4 Applying Tag Defaults	50
2.3.5 Adding Policies for Access Control	50
2.3.6 Adding Resources to a Compartment	50
2.3.7 Updating a Compartment	51
2.3.8 Moving a Compartment to a Different Compartment	52
2.3.9 Deleting a Compartment	52
2.4 Managing Policies	53
2.4.1 Creating a Policy	54
2.4.2 Writing Policy Statements	55
2.4.3 Deleting a Policy	62
2.5 Federating with Microsoft Active Directory	62
2.5.1 Gathering Required Information from ADFS	63
2.5.2 Verifying Identity Provider Self-Signed Certificates	63
2.5.3 Managing Identity Providers	64
2.5.4 Working with Group Mappings for an Identity Provider	67
2.5.5 Adding Oracle Private Cloud Appliance as a Trusted Relying Party in ADFS	69
2.5.6 Setting Up Policies for the Groups	71
2.5.7 Providing Federated Users Sign In Information	71

Oracle Private Cloud Appliance Identity and Access Management (IAM) enables you to control which users have what access to which cloud resources in your tenancy.

For conceptual information, see the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

2.1 Creating and Managing User Accounts

By default, the tenancy has an administrative user in an administrators group, and a policy enables the administrators group to manage the tenancy. To limit a user to managing only a subset of resources in the

tenancy or another compartment, or to have less than full management access to some resources, create a user account, add the user account to one or more groups, and create one or more policies for those groups.

A user account is not automatically a member of any group. A user that is not a member of any group is visible in the tenancy but does not have access to any resources.

For conceptual information about user accounts and groups, see the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

2.1.1 Creating a User

When you create a user, the user is automatically created in the tenancy. You cannot specify a different compartment for the user.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.
2. Click the Create User button.
3. In the Create User dialog, enter the following information:
 - **Name:** A name for this user account. User names have the following characteristics:
 - Must be unique within the tenancy. You can create a user with the same name as a user that has been deleted.
 - Are case insensitive.
 - Cannot be changed later.
 - Can be no more than 100 characters.
 - Can contain only alphanumeric characters, period (.), hyphen (-), underscore (_), plus sign (+), and at sign (@).
 - **Description:** A description for this user, such as the full name of the person or a brief description of the account. The description has the following characteristics:
 - Must be 1-400 characters.
 - Does not need to be unique.
 - Can be changed later.
 - **Email Address:** (Optional) The email address for the user. Can be updated later.
 - **Password:** (Optional) To enable this user to log in to the Compute Web UI, check the box labeled "Generate a temporary password for this user."

You can provide a password later. See [Section 2.1.2, "Providing a Temporary Compute Web UI Password"](#).

Note

Passwords for federated users are not managed through this service. See information from your federated identity provider.

- *Tagging*: (Optional) Add defined or free-form tags for this user account as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.
4. Click the Create User button on the Create User dialog.

If you checked the box labeled "Generate a temporary password for this user," a Temporary Password for New User dialog pops up, showing the temporary password. You cannot retrieve this password again after you close this dialog. Copy the temporary password, save the password to a safe place for delivery to the user, and click the "I have made a note of the password" button.

The details page of the new user is displayed.

Next steps:

- Provide the user with a temporary password so that the user can set their own permanent Compute Web UI password.
 - If you checked the box labeled "Generate a temporary password for this user," provide the temporary password that you copied from the Temporary Password for New User dialog.
 - If you did not check the box labeled "Generate a temporary password for this user," or did not save that password, follow the instructions in [Section 2.1.2, “Providing a Temporary Compute Web UI Password”](#) to generate a temporary password for the user.
- Add this user to at least one group. See [Section 2.1.5, “Adding a User to a Group by Updating the User”](#).
- If the user wants to use the OCI CLI, see [Section 1.2.2, “Installing the OCI CLI”](#).

Using the OCI CLI

1. Get the following information:
 - A name and description for the user. See the Compute Web UI procedure for parameters. In the OCI CLI, a description must be provided but its value can be an empty string.
 - (Optional) The OCID of the tenancy for the user. By default, the root compartment OCID from the config file is used.
2. Run the user create command.

Syntax

```
oci iam user create --name text --description text
```

See the Compute Web UI procedure for characteristics of the name and description values. See [Section 3.4.1, “Adding Tags at Resource Creation”](#) to add defined and free-form tags.

Example:

```
$ oci iam user create --name flast --description "First Last" --email first.last@example.com
```

The output of this command is the same as the output of the `user get` command.

Next steps:

- Provide the user with a temporary password so that the user can set their own permanent Compute Web UI password. See [Section 2.1.2, “Providing a Temporary Compute Web UI Password”](#).
- Add this user to at least one group. See [Section 2.1.5, “Adding a User to a Group by Updating the User”](#).
- If the user wants to use the OCI CLI, see [Section 1.2.2, “Installing the OCI CLI”](#).

2.1.2 Providing a Temporary Compute Web UI Password

Perform this procedure for new users and for users who forget their password. This procedure generates a temporary one-time password. When the user signs in using this password, the user is required to change the password before proceeding. The generated temporary password expires after seven (7) days.

A tenancy administrator can provide a temporary password for any user. Users must set their own permanent passwords by following the instructions in [Section 2.1.3, “Setting Your Own Compute Web UI Password”](#).

Note

Passwords for federated users are not managed through the IAM service. See information from your federated identity provider.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.
2. For the user that needs a new password, click the Actions menu, and click the Change Password option.
3. In the Change Password dialog, click the Create Temporary Password button.

A Password Changed dialog pops up. The New Password field contains the temporary password.

4. Copy and save this temporary password.

You cannot retrieve this password again after you close this dialog. Copy the temporary password, and save the password to a safe place for delivery to the user.

5. Click the Close button on the dialog.
6. Deliver this temporary one-time password to the user. The user must follow the rules stated in [Section 2.1.3, “Setting Your Own Compute Web UI Password”](#) when setting their new password.

Using the OCI CLI

1. Get the OCID of the user that needs a password (`oci iam user list`).
2. Run the user Compute Web UI password create or reset command.

Example:

```
$ oci iam user ui-password create-or-reset --user-id ocid1.user.unique_ID
{
  "data": {
    "inactive-status": null,
    "lifecycle-state": "ACTIVE",
```



```
"password": "N59%fP9uTq6\\",  
"time-created": "2021-10-13T22:10:49.290000+00:00",  
"user-id": "ocidl.user.unique_ID"  
}  
}
```

3. Copy the `password` value from the command output and deliver this temporary one-time password to the user. The user must follow the rules stated in [Section 2.1.3, "Setting Your Own Compute Web UI Password"](#) when setting their new password.

2.1.3 Setting Your Own Compute Web UI Password

Users do not require an access policy to set or change their own Compute Web UI password.

2.1.3.1 Setting Your Password

Use this procedure to set your Compute Web UI password initially, or to reset your password if you forgot your password.

Using the Compute Web UI

1. Get the temporary password that was generated for you.
2. On the login screen for the Compute Web UI, enter your user name.
3. Enter the temporary password.

A dialog pops up that says your password has expired and you need to create a new password.

4. Click the Change my password button.
5. On the Change My Password screen, enter the temporary password in the Current Password field.
6. Enter a new password in the New Password field and again in the Confirm New Password field.

Passwords must be at least 12 characters in length and contain at least one of each of the following: uppercase character, lowercase character, number, and symbol.

7. Click the Save Changes button.

A dialog pops up that says your password has been successfully updated.

8. Click the Continue button.
9. Log in using your new password.

2.1.3.2 Changing Your Password

Use this procedure to change your Compute Web UI password while your current password still works.

Using the Compute Web UI

1. In the top right corner of the Compute Web UI, click your user menu.
2. Click Change My Password.
3. On the Change My Password screen, enter your current password in the Current Password field.
4. Enter a new password in the New Password field and again in the Confirm New Password field.

Passwords must be at least 12 characters in length and contain at least one of each of the following: uppercase character, lowercase character, number, and symbol.

5. Click the Save Changes button.

A dialog pops up that says your password has been successfully updated.

6. Click the Continue button.

2.1.4 Viewing User Information and Group Membership

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.

The Users page shows all users of the tenancy because user accounts cannot be in different compartments. All users are in the tenancy.

2. Click the name of the user for which you want more information.
3. On the details page for that user account, scroll down to the Resources section.
4. Click the Groups resource.

The list of groups where this user is a member is shown.

5. To see the full list of members of a group, click the name of the group in the Groups list.

Scroll down to the Resources section for that group and click Group Members.

Using the OCI CLI

1. Get the OCID of the user account for which you want the list of groups (`oci iam user list`).
2. Run the list groups command.

Syntax

```
oci iam user list-groups --user-id user_OCID
```

The output of the `list-groups` command is the same as the output of the `group get` command for each group where this user is a member.

The `user get` command does not show group membership.

2.1.5 Adding a User to a Group by Updating the User

A user must be a member of at least one group in order to have access to any resources.

Using the Compute Web UI

As an alternative to using the Users Compute Web UI page, you can use the Groups page as described in [Section 2.2.3, "Adding a User to a Group by Updating the Group"](#).

1. In the navigation menu, click Identity, and then click Users.
2. Click the name of the user that you want to add to a group.

3. On the details page, scroll down to the Resources section and click Groups.
4. At the top of the Groups list, click the Add User to Group button.
5. In the Add User to Group dialog, select a group from the drop-down list, and then click the OK button.

The selected group is added to the user's Groups list.

Using the OCI CLI

1. For the OCI CLI procedure, see [Section 2.2.3, “Adding a User to a Group by Updating the Group”](#).
2. Use the `user list-groups` command to show the groups where this user is a member. The output of the `user list-groups` command is the same as the output of the `group get` command for each group where this user is a member.

2.1.6 Removing a User from a Group by Updating the User

If you remove a user from all groups, the user will not have access to any resources.

Using the Compute Web UI

As an alternative to using the Users Compute Web UI page, you can use the Groups page as described in [Section 2.2.4, “Removing a User from a Group by Updating the Group”](#).

1. In the navigation menu, click Identity, and then click Users.
2. Click the name of the user that you want to remove from a group.
3. Scroll to the Resources section and click Groups.
4. For the group from which you want to remove the user, click the Actions menu, and click the Remove from Group option.

The selected group is removed from the user's Groups list.

Using the OCI CLI

1. For the OCI CLI procedure, see [Section 2.2.4, “Removing a User from a Group by Updating the Group”](#).
2. Use the `user list-groups` command to show the groups where this user is a member. The output of the `user list-groups` command is the same as the output of the `group get` command for each group where this user is a member.

2.1.7 Modifying a User

You can change a user account's description and email address. You can add, change, or remove tags as described in [Section 3.4.2, “Applying Tags to an Existing Resource”](#).

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.
2. For the user account that you want to modify, click the Actions menu, and click the Edit option.
3. In the Edit `username` dialog, modify the account's description, email address, or tags.
4. Click Save Changes.

Using the OCI CLI

1. Get the OCID of the user account that you want to modify (`oci iam user list`).
2. Run the user update command.

Syntax

```
oci iam user update --user-id user_OCID [ --description desc ] \  
[ --email email ] [ --defined-tags tags ] [ --freeform-tags tags ]
```

The output of this command is the same as the output of the `user get` command.

2.1.8 Deleting a User

You cannot delete a user if the user is a member of any group. You cannot delete your own user.

When you delete a user, all API keys associated with that user account are also deleted.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Users.
2. Click the name of the user that you want to delete.
3. Ensure that the user is not a member of any group.

On the user details page, scroll down to the Resources section and click Groups. To remove this user from a group, click the Actions menu for the group in the Groups list, and click the Remove from Group option.

4. At the top of the user details page, click the Delete button.
5. On the Delete User confirmation dialog, click the Confirm button.

Using the OCI CLI

1. Get the OCID of the user account that you want to delete (`oci iam user list`).
2. Use the `user list-groups` command to ensure that the user is not a member of any group.
3. Run the user delete command.

Syntax

```
oci iam user delete --user-id user_OCID
```

Example:

```
$ oci iam user delete --user-id ocidl.user.unique_ID  
Are you sure you want to delete this resource? [y/N]: y
```

To delete a user without confirmation, use the `--force` option.

2.2 Creating and Managing User Groups

Access to cloud resources is granted to groups, not directly to users. A user account is not automatically a member of any group. To enable a user to do any work with cloud resources, you must add the user to a group and then create an access policy for that group. A group is therefore a set of users who have the

same type of access to the same set of cloud resources. Organize users into groups according to which compartments and resources they need to access and how they need to work with those resources. A user can be a member of more than one group.

For conceptual information about user accounts and groups, see the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

2.2.1 Creating a Group

When you create a group, the group is automatically created in the tenancy. You cannot specify a different compartment for the group.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.
2. Click the Create Group button.
3. In the Create Group dialog, enter the following information:
 - **Name:** A name for this group. Group names have the following characteristics:
 - Must be unique within the tenancy. You can create a group with the same name as a group that has been deleted.
 - Are case insensitive.
 - Cannot be changed later.
 - Can be no more than 100 characters.
 - Can contain only alphanumeric characters, period (.), hyphen (-), and underscore (_).
 - **Description:** A description for this group. The description has the following characteristics:
 - Must be 1-400 characters.
 - Does not need to be unique.
 - Can be changed later.
 - **Tagging:** (Optional) Add defined or free-form tags for this group as described in [Section 3.4.1, "Adding Tags at Resource Creation"](#). Tags can also be applied later.
4. Click the Create Group button on the Create Group dialog.

The details page of the new group is displayed.

Next steps:

- Create an access policy for this group or add this group to an existing policy. A group has no permissions unless it is the subject of at least one policy. See [Section 2.4, "Managing Policies"](#).
- Add users to this group. See [Section 2.2.3, "Adding a User to a Group by Updating the Group"](#).

Using the OCI CLI

1. Get the following information:

- A name and description for the group. See the Compute Web UI procedure for limitations. In the OCI CLI, a description must be provided but its value can be an empty string.
 - (Optional) The OCID of the tenancy for the group. By default, the root compartment OCID from the config file is used.
2. Run the group create command.

Syntax:

```
oci iam group create --name text --description "text"
```

See the Compute Web UI procedure for characteristics of the name and description values. See [Section 3.4.1, "Adding Tags at Resource Creation"](#) to add defined and free-form tags.

Example:

```
$ oci iam group create --name Product-A --description "Resource management for Product A."
```

The output of this command is the same as the output of the `group get` command.

Next steps:

- Create an access policy for this group or add this group to an existing policy. A group has no permissions unless it is the subject of at least one policy. See [Section 2.4, "Managing Policies"](#).
- Add users to this group. See [Section 2.2.3, "Adding a User to a Group by Updating the Group"](#).

2.2.2 Viewing Group Information and Group Membership

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.

The Groups page shows all groups in the tenancy because group definitions cannot be in different compartments. All groups are in the tenancy.

2. Click the name of the group about which you want more information.
3. On the details page for that group, scroll down to the Resources section.
4. Click the Group Members resource.

The list of users that belong to this group is shown.

5. To see the full list of groups where a user is a member, click the name of the user in the Group Members list.

Scroll down to the Resources section for that user and click Groups.

Using the OCI CLI

1. Get the OCID of the group for which you want the list of users (`oci iam group list`).
2. Run the list users command.

Syntax:

```
oci iam group list-users --group-id group_OCID
```

The output of the `list-users` command is the same as the output of the `user get` command for each user that is a member of this group.

The `group get` command does not show member users.

2.2.3 Adding a User to a Group by Updating the Group

Users must be members of groups in order to have access to resources.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.
2. Click the name of the group where you want to add users.
3. On the details page, scroll down to the Resources section and click Group Members.
4. At the top of the Group Members list, click the Add User to Group button.
5. In the Add User to Group dialog, select a user from the drop-down list, and then click the OK button.

The selected user is added to the Group Members list.

Using the OCI CLI

1. Get the following information:
 - The OCID of the group where you want to add a user (`oci iam group list`).
 - The OCID of the user that you want to add to this group (`oci iam user list`).
2. Run the group add user command.

Syntax:

```
oci iam group add-user --group-id group_OCID --user-id user_OCID
```

Example:

```
$ oci iam group add-user --group-id ocid1.group.unique_ID --user-id ocid1.user.unique_ID
{
  "data": {
    "compartment-id": "ocid1.tenancy.unique_ID",
    "group-id": "ocid1.group.unique_ID",
    "id": "ocid1.user_group_membership.unique_ID",
    "inactive-status": null,
    "lifecycle-state": "ACTIVE",
    "time-created": null,
    "user-id": "ocid1.user.unique_ID"
  }
}
```

2.2.4 Removing a User from a Group by Updating the Group

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.
2. Click the name of the group where you want to remove a user.

3. On the details page, scroll down to the Resources section and click Group Members.
4. In the Group Members list, click the Actions menu for the user that you want to remove from the group, and click the Remove from Group option.
5. At the confirmation prompt, click OK.

The user is removed from the group.

Using the OCI CLI

1. Get the following information:
 - The OCID of the group where you want to remove a user (`oci iam group list`).
 - The OCID of the user that you want to remove from the group (`oci iam user list`).
2. Run the group remove user command.

Syntax:

```
oci iam group remove-user --group-id group_OCID --user-id user_OCID
```

2.2.5 Modifying a Group

You can change the description for a group. You can add, change, or remove tags as described in [Section 3.4.2, “Applying Tags to an Existing Resource”](#).

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.
2. For the group that you want to modify, click the Actions menu, and click the Edit option.
3. In the Edit *groupname* dialog, modify the group's description or tags.
4. Click Save Changes.

Using the OCI CLI

1. Get the OCID of the group that you want to modify (`oci iam group list`).
2. Run the group update command.

Syntax:

```
oci iam group update --group-id group_OCID [ --description desc ] \  
[ --defined-tags tags ] [ --freeform-tags tags ]
```

The output of this command is the same as the output of the `group get` command.

2.2.6 Deleting a Group

You cannot delete a group if the group has any members.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Groups.
2. Click the name of the group that you want to delete.

3. Ensure that the group does not have any members.

On the group details page, scroll down to the Resources section and click Group Members. To remove a user from the group, click the Actions menu for the user in the Group Members list, and click the Remove from Group option.

4. At the top of the group details page, click the Delete button.
5. On the Delete Group confirmation dialog, click the Confirm button.

Using the OCI CLI

1. Get the OCID of the group that you want to delete (`oci iam group list`).
2. Use the `group list-users` command to ensure that the group has no members.
3. Run the group delete command.

Syntax:

```
oci iam group delete --group-id group_OCID
```

Example:

```
$ oci iam group delete --group-id ocid1.group.unique_ID
Are you sure you want to delete this resource? [y/N]: y
```

To delete a group without confirmation, use the `--force` option.

2.3 Creating and Managing Compartments

Compartments contain resources such as cloud instances, virtual cloud networks, and block volumes. Your tenancy is the root compartment where you can create cloud resources and other compartments. You can create hierarchies of compartments that are up to six levels deep. You can limit access to compartment resources to specified user groups. Most resources can be moved between compartments later if your business needs change.

The compartments that you create in your tenancy are your primary building blocks for organizing and controlling access to your cloud resources. Before you create compartments and resources, see *Organizing Resources in Compartments* in the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

2.3.1 Understanding the Tenancy

A tenancy is a special compartment. The tenancy is the root compartment where you create and administer all of your cloud resources, including other compartments.

Users, groups, and identity providers are always attached directly to the tenancy, not to any compartment of the tenancy. You cannot specify a different compartment when you create a user, group, or identity provider. When you use the OCI CLI to operate on a user, group, or identity provider, the OCID of the tenancy from the `config` file is used by default.

Other resources can reside in the tenancy or in any other compartment. Operating on these resources often requires you to select the correct compartment in the Compute Web UI or specify the compartment OCID in the OCI CLI.

Use the following procedures to get the OCID of the tenancy.

Using the Compute Web UI

1. Click your user profile menu in the top right of the page.
2. Click the Tenancy option.
3. On the tenancy details page, use the Show or Copy button under OCID.

Using the OCI CLI

1. Use the `compartment list` command.

```
$ oci iam compartment list
```

Look for the `ocid1.tenancy.unique_ID` OCID.

- With no options, the `compartment list` command lists all compartments that are direct child compartments of the tenancy. The tenancy is the value of the first property listed (`compartment-id`) for every compartment in the list.
- If you specify the `--include-root` option, the tenancy is listed first, and the tenancy OCID is the value of the `id` property (the value of the `compartment-id` property is `null`).

As is true for other resources, in a `compartment list` or `get`, the `compartment-id` compartment is the parent compartment of the `id` compartment.

2.3.2 Listing Compartments

Using the Compute Web UI

1. In the navigation menu, click Identity and then click Compartments.

The list shows all compartments that are direct child compartments of the tenancy.

2. To view a compartment that is a subcompartment of a listed compartment, click the name of the listed compartment. On the details page for that compartment, scroll to the Resources section, and click Child Compartments.

You might need to click the name of a compartment in the Child Compartments list, and repeat this step.

To find the compartment where a particular resource is located, navigate to a list of those resources. Above the resource list, use the Compartment drop-down menu to select the compartment.

Using the OCI CLI

Use the `--help` option to learn about the `--access-level` option and about options that are common to `list` commands such as `--lifecycle-state` and `--sort-by`.

1. To list all compartments and subcompartments in the tenancy, specify the `--compartment-id-in-subtree` option with a value of `true`.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

Specifying the `--compartment-id` option, described in the next step, does not change this output: You cannot list just the compartment tree of a particular compartment other than the tenancy.

2. To list all compartments that are direct child compartments of another compartment, specify the OCID of that parent compartment:

```
$ oci iam compartment list --compartment-id ocid1.compartment.unique_ID
```

This output does not list the specified parent compartment and does not list compartments that are deeper in the hierarchy of this parent compartment. This output only lists the direct child compartments of the specified parent compartment.

If you do not specify a parent compartment, all compartments that are direct child compartments of the tenancy are listed. To list the tenancy in addition to the direct child compartments of the tenancy, specify the `--include-root` option.

3. To list just one particular compartment, you can specify the compartment name.

```
$ oci iam compartment list --name Acompartment
```

The output is the same as a `get` of that compartment.

```
$ oci iam compartment get --compartment-id OCID_of_Acompartment
```

2.3.3 Creating a Compartment

You can create a compartment in your tenancy or in another compartment. You can create hierarchies of compartments that are up to six levels deep.

Using the Compute Web UI

1. In the navigation menu, click Identity and then click Compartments.
2. Click the Create Compartment button above the list of compartments.
3. In the Create Compartment dialog box, enter the following information:
 - *Name*: A name for this compartment. Compartment names have the following characteristics:
 - Must be unique within the tenancy.
 - Are case insensitive.
 - Can be changed later.
 - Can be no more than 100 characters.
 - Can contain only alphanumeric characters, period (.), hyphen (-), and underscore (_).
 - *Description*: A description for this compartment. This description can be no more than 400 characters and can be changed later.
 - *Create in Compartment*: The compartment in which you want to create the new compartment. The new compartment will be a sub-compartment of the selected compartment.
 - *Tagging*: (Optional) Add defined or free-form tags for this compartment as described in [Section 3.4.1, "Adding Tags at Resource Creation"](#). Tags can also be applied later.
4. Click the Create Compartment button on the dialog box.

Click the name of the new compartment to view the compartment details, including tags.

Using the OCI CLI

1. Get the OCID of the compartment in which you want to create the new compartment. The new compartment will be a sub-compartment of the specified compartment.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Run the compartment create command.

Syntax:

```
oci iam compartment create --compartment-id compartment_OCID \  
--name text --description "text"
```

See the Compute Web UI procedure for characteristics of the name and description values. See [Section 3.4.1, “Adding Tags at Resource Creation”](#) to add defined and free-form tags.

Example:

```
$ oci iam compartment create -c ocidl.compartment.parent_compartment_unique_ID \  
--name ProductX --description "A child compartment of compartment Products"  
{  
  "data": {  
    "compartment-id": "ocidl.compartment.parent_compartment_unique_ID",  
    "defined-tags": {},  
    "description": "A child compartment of compartment Products",  
    "freeform-tags": {},  
    "id": "ocidl.compartment.new_compartment_unique_ID",  
    "inactive-status": null,  
    "is-accessible": null,  
    "lifecycle-state": "ACTIVE",  
    "name": "ProductX",  
    "time-created": "2021-10-05T22:58:23.216657+00:00"  
  },  
  "etag": "b212700d-45fa-46a9-90da-bcc016c587bc"  
}
```

To view this output later, use the `compartment get` command.

2.3.4 Applying Tag Defaults

Compartments can have resources called tag defaults. Tag defaults are defined tags that are inherited by all resources and child compartments that are created after the tag default is added to the parent compartment. To add a tag default to a compartment, see [Section 3.3, “Configuring Tag Defaults”](#).

2.3.5 Adding Policies for Access Control

Child compartments inherit access permissions from their parent compartments. If you want the access to a new compartment to be different from the access to the parent compartment, create an access policy for the new compartment. For example, grant group DevX permission to read all resources in compartment Products and permission to manage all resources in subcompartment ProductX. Grant group DevY permission to read all resources in compartment Products and permission to manage all resources in subcompartment ProductY. Because of inheritance, group DevX will be able to read all resources in compartment ProductY, and group DevY will be able to read all resources in compartment ProductX.

For information about creating and attaching policies, see [Section 2.4, “Managing Policies”](#).

2.3.6 Adding Resources to a Compartment

Use either of the following methods to add resources to a compartment:

- Specify the compartment when you create the resource.
- Move the resource from a different compartment.

See the documentation for the particular resource for information such as whether attached resources move with the moved resource.

Check whether the moved resources have the correct tags and policies applied. You might need to manually delete and add tags and policies.

The Resources box on a compartment details page in the Compute Web UI, and the compartment `list` and `get` commands in the OCI CLI, do not show all of the resources that belong to a compartment. For resources that are not listed, go to the Compute Web UI page for that resource, such as instances, and select the compartment from the Compartment drop-down menu above the resource list. In the OCI CLI, specify the compartment OCID when you list the resources. See also [Section 1.1, “Using the Compute Web UI”](#) and [Section 1.2, “Using the OCI CLI”](#).

2.3.7 Updating a Compartment

You can change the name and description of a compartment. You can add, change, or remove tags as described in [Section 3.4.2, “Applying Tags to an Existing Resource”](#). You cannot change the parent compartment. To change the parent compartment, see [Section 2.3.8, “Moving a Compartment to a Different Compartment”](#).

Using the Compute Web UI

1. In the navigation menu, click Identity and then click Compartments.
2. If the compartment that you want to update is not listed, navigate to the compartment that you want to update, as described in [Section 2.3.2, “Listing Compartments”](#).
3. For the compartment that you want to update, click the Actions menu, and click the Edit option.
4. In the Editing `compartment_name` Compartment dialog, make the changes.
5. Click the Save Changes button.

Click the compartment name to view the compartment details, including tags.

Using the OCI CLI

1. Get the OCID of the compartment that you want to update.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Run the compartment update command.

Syntax:

```
oci iam oci iam compartment update --compartment-id compartment_OCID \  
options_with_values_to_update
```

Example:

If you specify the `--defined-tags` or `--freeform-tags` options, then you must fully specify all defined and free-form tags that you want on this compartment, including tags that already exist on the compartment that you want to keep. Values that you provide for these tag options replace any existing values. See [Section 3.4, “Working with Resource Tags”](#). You will be prompted to confirm unless you specify the `--force` option.

```
$ oci iam compartment update --compartment-id ocid1.compartment.unique_ID \  
--defined-tags '{\"tag_key\": \"tag_value\"}' --force
```

```
--defined-tags '{"Product":{"LMN":"Development"}}' --freeform-tags '{"MyTag":"val-u}'  
WARNING: Updates to freeform-tags and defined-tags will replace any existing values.  
Are you sure you want to continue? [y/N]: y
```

The output of this command is the same as the output of the `compartment get` command.

2.3.8 Moving a Compartment to a Different Compartment

You can move a compartment to a different parent compartment in the same tenancy. When you move a compartment, all subcompartments of the compartment are moved. Some resources of the moved compartment are moved. You can separately move other resources as needed. See the documentation for the particular resource type for more information.

After you move a compartment to a new parent compartment, the access policies of the new parent take effect and the policies of the previous parent no longer apply. Groups who had access to the compartment and its resources in the previous parent compartment lose their access when the compartment is moved. Groups who have access in the new parent compartment gain access to the moved compartment and its resources.

Tag defaults that are automatically applied to all resources created in the new parent are not automatically applied to the newly moved compartment and its resources. You might need to separately delete and add tag defaults to the moved compartment and delete and add tags to moved resources.

See also *Moving a Compartment to a Different Parent Compartment* in the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

You must belong to a group that has `manage all-resources` permissions on the lowest shared parent compartment of the current compartment and the destination compartment.

To move a compartment, you must use the OCI CLI.

Using the OCI CLI

1. Get the OCID of the compartment that you want to move, and the OCID of the destination compartment.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Run the compartment move command.

Syntax:

```
oci iam compartment move --compartment-id compartment_to_move_OCID \  
--target-compartment-id destination_compartment_OCID
```

Use the `iam work-request get` command to check the status of the compartment move, or view the work request details in the Compute Web UI. Some resources might take longer to move than the compartment.

2.3.9 Deleting a Compartment

To delete a compartment, you must first move, delete, or terminate all of the resources in the compartment, including any policies that are attached to the compartment. Before you begin, check the move and delete capabilities for all resources in the compartment.

Using the Compute Web UI

1. In the navigation menu, click Identity and then click Compartments.

The compartments in the tenancy are listed.

2. If the compartment that you want to delete is not listed, navigate to the compartment that you want to update, as described in [Section 2.3.2, “Listing Compartments”](#).
3. For the compartment that you want to delete, click the Actions menu, and click the Delete Compartment option.

If the Delete Compartment option is not selectable, then you might not have permission to delete this compartment.

4. In the Delete Compartment confirmation dialog, click Delete.

The compartment status changes to Deleting.

In the Resources box on the compartment details page, click Work Requests and view the details of the compartment delete. When the work request is completed, the compartment is removed from the compartments list. If the work request fails, the compartment status returns to Active.

Using the OCI CLI

1. Get the OCID of the compartment that you want to delete.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Run the compartment delete command.

Syntax:

```
oci iam compartment delete --compartment-id compartment_OCID
```

Use the `iam work-request get` command to check the status of the compartment delete.

2.4 Managing Policies

A policy is a named set of one or more policy statements. Policy statements grant permissions to users to access resources.

When designing access policies, remember the following policy characteristics:

- The policy will apply to the compartment where you attach the policy and to all subcompartments of that compartment. Permissions granted in a particular compartment, including the tenancy, are inherited by all subcompartments of that compartment.
- A user can be a member of more than one group. A group can be the subject of more than one policy. A policy can have up to 50 policy statements.
- If some users need full access to the named resources and other users only need to use the resources, you need to create multiple groups and multiple policies. A tenancy can have up to 100 policies.
- Users who have full access to resources in a subcompartment probably also need view or use access to related resources in that compartment and in parent compartments. For example, users who have access to create instances in a compartment might also need access to use tag namespaces to apply defined tags to the instances, or access to read images in a different compartment.

For general information about policies, see *How Policies Work* in the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*. For specific information about policy statements, see [Section 2.4.2, “Writing Policy Statements”](#).

2.4.1 Creating a Policy

Before You Begin

A policy must have at least one policy statement. You cannot create an empty policy and add statements later. Decide what you want your policy to allow, and see [Section 2.4.2, “Writing Policy Statements”](#) to design the necessary statements.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Policies.
2. Click Create Policy.
3. In the Create Policy dialog, enter the following information:
 - *Name*: The policy name. Policy names have the following characteristics:
 - Must be unique within the tenancy.
 - Are case insensitive.
 - Cannot be changed later.
 - Can be no more than 100 characters.
 - *Description*: A description for the policy. This description can be no more than 400 characters.
 - *Create in Compartment*: Select the compartment where you want to attach this policy. The policy will apply to this compartment and all child compartments of this compartment.
 - *Statements*: Enter a policy statement. For information about how to write policy statements, see [Section 2.4.2, “Writing Policy Statements”](#).

To add a second policy statement, click the Another Statement button. You can enter up to 50 statements. If you create more than one policy statement, you can click the X button next to a statement to delete that statement.
 - *Tagging*: (Optional) Add defined or free-form tags for this policy as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.
4. Click the Submit button.

The details page for the new policy is displayed. The Resources section of the page shows the policy statements.

Using the OCI CLI

1. Get the OCID of the compartment where you want to attach the policy. The policy will apply to this compartment and all child compartments of this compartment.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Construct an argument for the `--statements` option.

The value of the `--statements` option argument is an array of policy statements in JSON format. This argument can be provided as a string on the command line or in a file. For information about how to write policy statements, see [Section 2.4.2, “Writing Policy Statements”](#).

- (Optional) Construct arguments for defined or free-form tags for this policy as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.
- Run the policy create command.

Syntax:

```
oci iam policy create -c compartment_OCID --name text --description "text" \
{ --statements '["statement", "statement"]' | --statements file://policy.json }
```

The *compartment_OCID* is the compartment where you want to attach this policy. See the Compute Web UI procedure for characteristics of the name and description values. See [Section 3.4.1, “Adding Tags at Resource Creation”](#) to add defined and free-form tags.

This command returns the same output as the `policy get` command.

2.4.2 Writing Policy Statements

A policy can have up to 50 policy statements. A tenancy can have up to 100 policies. Decide what you want your policy to permit, and use the information in this section to write the necessary statements.

Ensure that the groups and compartments that you plan to name in the policy statements exist. Note the name or OCID of each group and compartment that you want to use.

Note

If you use names in your policy statements instead of OCIDs, the policy will still be valid if the name of the group or compartment is subsequently changed. Internally, the OCID, not the name, is used. However, the policy could be more difficult for administrators to understand if the name of the group or compartment changes.

If you plan to use a tag to apply the policy to more than one group or more than one compartment, ensure that the tag exists. Note the name of the tag namespace, the name of the key, and the value that you want to use in the policy statement.

2.4.2.1 Policy Statement Syntax

Policy statements grant permissions to users to access resources. The users are also called the subject of the policy, and the permissions are also called the verb. The resource type and compartment define the set of possible resources to which the users are granted access. This set of resources is also called the target. Conditions can be used to narrow the set of users, the set of resources, and the operations that can be performed on the resources.

The following is the policy statement syntax:

```
allow users to permissions [resource_type] in compartment [ where conditions ]
```

Keywords `allow`, `to`, and `in` are required and are case insensitive.

users

One or more user groups or `any-user`. To specify more than one group, use a comma between each two group names or OCIDs in a list. You cannot specify both group names and group OCIDs. You can specify the keyword `any-user` to grant permission to all users.

- `group group_name [, group_name ...]`
- `group id group_ocid [, group_ocid ...]`
- `any-user`

permissions • One of *inspect*, *read*, *use*, or *manage*. For descriptions of the access that these permission aggregators grant, see *Verb* in *Policy Syntax* in the *Identity and Access Management Overview* in the *Oracle Private Cloud Appliance Concepts Guide*.

- One or more specific permissions, such as `INSTANCE_UPDATE`, in the following form:

```
{ PERMISSION_1[, PERMISSION_2]... }
```

If you use this option, do not specify a *resource_type*. The resource type is included in the permission.

To grant both a permission aggregator and one or more specific permissions for the same resource, use two policy statements.

resource_type A single keyword that represents one of the following:

- A single resource type, such as `instances` or `volumes`
- A family of resource types. For example, the `instance-family` resource type family includes the following resource types:
 - `app-catalog-listing`
 - `console-histories`
 - `instances`
 - `instance-console-connection`
 - `instance-images`
- `all-resources`

If you list specific permissions instead of one of the four permissions aggregator keywords, do not specify a *resource_type*. The resource type is included in the permission.

For a list of resource types, see the table in *Resource Types* in *Policy Syntax* in the *Identity and Access Management Overview* in the *Oracle Private Cloud Appliance Concepts Guide*.

compartment A single compartment name or OCID or *tenancy*.

- `compartment compartment_name`
- `compartment id compartment_OCID`
- `tenancy`

To grant access in multiple compartments, use multiple statements.

condition A predefined variable followed by an operator and a value. See [Section 2.4.2.2, “Using Conditions”](#).

2.4.2.2 Using Conditions

Conditions can be specified in a policy statement to narrow the set of users who are granted access, the set of resources to which the users are granted access, and the operations that can be performed on the resources. A condition is a predefined variable with a value that you specify. You can specify a list of conditions with AND and OR relationships. The entire condition clause must evaluate to true in order for access to be granted. For information about conditions that might unexpectedly evaluate to false, see [Section 2.4.2.3, “Conditions that Are Not Applicable”](#).

The following is the syntax of the condition clause:

```
where condition
where all|any {condition[, condition]...}
```

The syntax of *condition* is:

```
variable op 'value'
```

variable See [Table 2.1, “Supported Predefined Variables for Conditions”](#).

op

- = (equal) or != (not equal) – Applies to all variables.
- in or not in – See [Section 2.4.2.4, “Using Defined Tags in Conditions”](#).

value The *value* can be a fully specified string or can use the * wildcard. If the *value* is a fully specified string, enclose the value in single quotation marks. If you use *, enclose the value in forward slashes (/):

```
'BU1-ProdX'
/*Prod*/
/*ProdX/
/BU1-Prod*/
```

Condition values are case insensitive. For example, a condition with a value of BucketA also applies to bucket bucketA in the same compartment if such a bucket exists.

In the following table, variables that begin with *request* refer to the request that is being made: A user has clicked a Compute Web UI option or entered a OCI CLI command. Variables that begin with *target* refer to the resource that the user clicked or named in the command.

Table 2.1 Supported Predefined Variables for Conditions

Variable	Description
<code>request.groups.id</code>	The list of groups that the requesting user belongs to.
<code>request.operation</code>	The name of the operation that is being attempted.
<code>request.permission</code>	The names of the permissions that are required to perform the operation.
<code>target.compartment.id</code>	The OCID of the compartment that contains the target resource. The compartment that contains the target resource could be a child compartment of the compartment specified in the <i>in</i> clause in the policy statement.
<code>target.compartment.name</code>	The name of the compartment that contains the target resource. The compartment that contains the target resource could be a child compartment of the

Variable	Description
	compartment specified in the <code>in</code> clause in the policy statement.
<code>target.user.id</code>	The OCID of the target user. This OCID is not available when the requested permission is to create the user.
<code>target.user.name</code>	The name of the target user.
<code>target.group.id</code>	The OCID of the target group. This OCID is not available when the requested permission is to create the group.
<code>target.group.name</code>	The name of the target group.
<code>target.group.member</code>	True if the requesting user belongs to the target group.
<code>target.policy.id</code>	The OCID of the target policy. This OCID is not available when the requested permission is to create the policy.
<code>target.policy.name</code>	The name of the target policy.
<code>target.tag-namespace.id</code>	The OCID of the tag namespace that the user is requesting to list, update, or delete.
<code>target.tag-namespace.name</code>	The name of the tag namespace that the user is requesting to create or update. Use commas to separate multiple names.
<code>request.principal.group.tag</code>	See Section 2.4.2.4, “Using Defined Tags in Conditions” .
<code>target.resource.tag</code>	See Section 2.4.2.4, “Using Defined Tags in Conditions” .
<code>target.resource.compartment.tag</code>	See Section 2.4.2.4, “Using Defined Tags in Conditions” .

Example: Specify Permissions Using `request.permission`

To grant users the ability to create objects but not the ability to delete objects, you can grant `manage` access and then specify a condition that says only create and inspect access are granted:

```
allow group ObjectWriters to manage objects in compartment ABC
where any {request.permission='OBJECT_CREATE', request.permission='OBJECT_INSPECT'}
```

Example: Specify Compartments Using `target.compartment.name` and Wildcards

The following example grants users the ability to manage all resources in `virtual-network-family` in any compartment that has a name that begins with X except for compartment XYZ:

```
allow group NetworkAdmins to manage virtual-network-family in tenancy
where all {target.compartment.name=/X*/,target.compartment.name!='XYZ'}
```

Example: Nested Conditions

The following policy enables users in group `BucketAdmins` to either read, update, or manage retention rules for `BucketA` in compartment ABC:

```
allow group BucketAdmins to manage buckets in compartment ABC
where all {target.bucket.name='BucketA',
any {request.permission='BUCKET_UPDATE', request.permission='BUCKET_READ',
RETENTION_RULE_MANAGE}}
```

Because the policy is for a specific named bucket, this policy does not permit users to retrieve a list of buckets. To permit users to retrieve a list of buckets, add the following separate statement:

```
allow group BucketAdmins to inspect buckets in compartment ABC
```

See [Section 2.4.2.3, “Conditions that Are Not Applicable”](#).

Example: Apply Defined Tags

The following example enables users in groups `BucketAdmins` and `ObjectWriters` to apply tags in the `StorageTags` tag namespace:

```
allow group BucketAdmins,ObjectWriters to use tag-namespaces in tenancy
where target.tag-namespace.name='StorageTags'
```

Example: Edit Any Group Where You Are a Member

The following example enables all users to edit any group where they are members:

```
allow any-user to use groups in tenancy where target.group.member='true'
```

2.4.2.3 Conditions that Are Not Applicable

If a condition is not applicable to the rest of the policy statement, then that condition evaluates to false and access is not granted.

A condition is not applicable if it is testing information that is not available in the request. For example, the following policy statement grants `use` access to the resource `users`, but does not allow the requesting users to list or update users, even though those permissions are included in the `use` permission:

```
allow group GroupAdmins to use users in tenancy where target.group.name != 'Administrators'
```

The request to list users or update a user does not include information about groups. The list users and update user requests have no value for `target.group.name`. The test fails, and the request to list or update a user is denied.

To fix this example, you could remove the `where` clause and allow only `inspect` or `read` access.

2.4.2.4 Using Defined Tags in Conditions

Certain conditions evaluate the value of a defined tag that has been applied to a user, compartment, or resource. In these conditions, the predefined variable can be called a tag variable.

Using conditions with tag variables enables you to do the following:

- Write a single policy statement that applies to multiple user groups, compartments, or resources.
- Change the permissions that are granted without changing the policy statement. Instead, to allow or revoke access, apply tags to different resources or remove tags from resources.

See [Chapter 3, Resource Tag Management](#) for information about how to create and apply defined tags.

The general syntax of a condition that uses tag variables is the same as the syntax of a condition that uses other condition variables:

```
variable op 'value'
```

The value of each of these three parts is specialized for tags.

variable Tag condition variables include the name of the tag namespace and the name of the key in the variable name:

```
base_variable_name.tag_namespace_name.tag_key_name
```

op One of =, !=, in, or not in –.

The *in* and *not in* operations refer to members of the set of possible values for the tag.

value The *value* is a value of the defined tag. The value can be a single value or a list of values.

The following tag variables are supported:

`request.principal.group.tag`

This variable potentially grants access to multiple groups in one statement. The following statement allows any user that is a member of a group that has been tagged with tag Operations>Project>ABC to manage instance resources in compartment ProdX:

```
allow any-user to manage instance-family in compartment ProdX
where request.principal.group.tag.Operations.Project='ABC'
```

If you replace 'ABC' in the preceding statement with '*' or /*/, a user that is a member of a group that has been tagged with any value of Operations>Project could manage instance resources in compartment ProdX.

`target.resource.compartment.tag`

This variable potentially grants access to multiple compartments in one statement. The following statement allows users in group NetAdmins to use network resources in any compartment that has been tagged with either tag Operations>Project>ABC or tag Operations>Personnel>Test:

```
allow group NetAdmins to use virtual-network-family in tenancy where
any { target.resource.compartment.tag.Operations.Project='ABC' ,
target.resource.compartment.tag.Operations.Personnel='Test' }
```

If you replace *any* with *all* in the preceding statement, the statement allows users in group NetAdmins to use network resources in any compartment that has been tagged with both tag Operations>Project>ABC and tag Operations>Personnel>Test.

The following statement allows users in group NetAdmins to use network resources in any compartment that has been tagged with either tag Operations>Personnel>Development or tag Operations>Personnel>Test:

```
allow group NetAdmins to use virtual-network-family in tenancy where
target.resource.compartment.tag.Operations.Personnel in ('Development', 'Test')
```

`target.resource.tag`

This variable grants access to one or more resources of the specified type. The following statement allows group Xadmins to use any instance in compartment ProdX that is tagged with tag Operations>Project>XYZ.

```
allow group Xadmins to use instances in compartment ProdX
where target.resource.tag.Operations.Project = 'XYZ'
```

2.4.2.5 Updating a Policy

Using the Compute Web UI To Modify the Policy Description or Tags

1. In the navigation menu, click Identity, and then click Policies.
2. If the policy that you want to modify is not listed, select the correct compartment from the Compartment drop-down menu above the policies list.

3. For the policy that you want to modify, click the Actions menu for the policy, and click the Edit option.
4. Update the description or tags.
To modify tags, see [Section 3.4.2, “Applying Tags to an Existing Resource”](#).
5. Click the Save Changes button.

Using the Compute Web UI To Modify the Policy Statements

1. In the navigation menu, click Identity, and then click Policies.
2. If the policy that you want to modify is not listed, select the correct compartment from the Compartment drop-down menu above the policies list.
3. Click the name of the policy that you want to modify.
4. On the policy details page, scroll to the Resources section.
5. In the Statements list, click the Configure Policy Statements button.
6. In the Edit Statements in the `policy_name` Policy dialog, change or add policy statements.

To modify policy statements, see [Section 2.4.2, “Writing Policy Statements”](#).

To add a policy statement, click the Another Statement button. You can enter up to 50 statements.

If more than one policy statement exists, you can click the X button next to a statement to delete that statement.

7. Click the Submit button.

Using the OCI CLI

1. Get the policy OCID.

```
$ oci iam policy list --compartment-id compartment_OCID
```

2. (Optional) To change or add policy statements, construct an argument for the `--statements` option.

The value of the `--statements` option argument is an array of policy statements in JSON format. This argument can be provided as a string on the command line or in a file. For information about how to write policy statements, see [Section 2.4.2, “Writing Policy Statements”](#).

The argument that you provide for the `--statements` option replaces the existing statements in the policy. Be sure to include any statements that you want to keep from the existing policy. Use the `policy get` command to view and copy current policy statements.

If you do not specify the `--force` option, the system will display the existing statements in the policy and request that you confirm that you want to replace them.

3. (Optional) Construct arguments for defined or free-form tags for this policy as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#).
4. Run the policy update command.

Syntax:

```
oci iam policy update --policy-id policy_OCID [ --description desc ] \
```

```
[ --defined-tags tags ] [ --freeform-tags tags ] \  
[ --statements policy_statements --version-date "" ]
```

If you specify `--statements`, then you must include `--version-date ""`.

This command returns the same output as the `policy get` command.

2.4.3 Deleting a Policy

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Policies.
2. If the policy that you want to delete is not listed, select the correct compartment from the Compartment drop-down menu above the policies list.
3. For the policy that you want to delete, click the Actions menu, and click the Delete option.
4. In the confirmation dialog, click the Delete button.

Using the OCI CLI

1. Get the policy OCID.

```
$ oci iam policy list --compartment-id compartment_OCID
```

2. Run the policy delete command.

Syntax:

```
oci iam policy delete --policy-id policy_OCID
```

This command returns the same output as the `policy get` command.

2.5 Federating with Microsoft Active Directory

Many companies use an identity provider to manage user logins and passwords and to authenticate users for access to secure websites, services, and resources. To access the Oracle Private Cloud Appliance Compute Web UI, users must also sign in with a user name and password. An administrator can *federate* with a supported identity provider so that each user can use their existing login and password, rather than having to create new credentials to access and use cloud resources.

Federation involves setting up a trust relationship between the identity provider and Oracle Private Cloud Appliance. When an administrator has established this relationship, federated users are prompted with a *single sign-on* when accessing the Compute Web UI.

For more information, see *Federating with Identity Providers* in the [Oracle Private Cloud Appliance Concepts Guide](#).

You can federate multiple Active Directory (AD) accounts with Oracle Private Cloud Appliance (for example, one for each division of the organization), but each federation trust that you set up must be for a *single* AD account. To set up a trust, you perform some tasks in the Oracle Private Cloud Appliance Compute Web UI and some tasks in Active Directory Federation Services (ADFS).

Before you begin federating, make sure you already have:

- Installed and configured Microsoft Active Directory Federation Services for your organization.
- Set up groups in Active Directory that will map to groups in Oracle Private Cloud Appliance.

- Created users in Active Directory who will sign into the Oracle Private Cloud Appliance Compute Web UI.

Note

Consider naming Active Directory groups that you intend to map to Oracle Private Cloud Appliance groups with a common prefix to make it easy to apply a filter rule, for example, PCA_Administrators, PCA_NetworkAdmins, PCA_InstanceLaunchers.

2.5.1 Gathering Required Information from ADFS

To federate with Oracle Private Cloud Appliance you need to have the SAML metadata document and the names of the Active Directory (AD) groups that you want to map to Oracle Private Cloud Appliance groups.

1. Locate and download the SAML metadata document for your ADFS, which is by default at:

```
https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml
```

This is the document you will upload when you create the identity provider.

2. Make a note of all the AD groups that you want to map to Oracle Private Cloud Appliance groups.

Important

Ensure that you have all the Oracle Private Cloud Appliance groups configured before you add AD as an identity provider.

2.5.2 Verifying Identity Provider Self-Signed Certificates

Important

You can skip this section if your ADFS certificate is signed by a known certificate authority because they should already exist in the Oracle Private Cloud Appliance certificate bundle.

The Oracle Private Cloud Appliance Certificate Authority (CA), is self signed openssl generated root and intermediate x.509 certificate. These CA certificates are used to issue x.509 server/client certificates allowing you to add outside Certificate Authority (CA) trust information to the rack. If you use a self-signed certificate for ADFS, you will need to add outside CA trust information from ADFS to the management nodes on the rack.

Note

If you are using the metadataUrl property to create or update an identity provider, you will need to add the identity provider's web server's certificate chain to the Oracle Private Cloud Appliance outside CA bundle. See your identity provider's documentation on how to find the web server's certificate chain and then follow steps 3-8.

To add outside CA trust information, complete the following steps:

1. From a browser, enter the following URL and download the SAML metadata document for your ADFS, which is by default at:

```
https://<yourservname>/FederationMetadata/2007-06/FederationMetadata.xml
```

2. Open the file in a text or XML editor and locate the signing certificate section, for example:

```
<KeyDescriptor use="signing">
<KeyInfo>
<X509Data>
<X509Certificate>
<!--CERTIFICATE IS HERE-->
</X509Certificate>
</X509Data>
</KeyInfo>
</KeyDescriptor>
```

3. Log on to management node 1 whose default name is `pcamn01`.
4. Navigate to `/etc/pca3.0/vault` and create a new directory named `customer_ca`.

Note

You can use this directory for multiple files. For example you can create a file for the identity provider certificate and one for the web server's certificate chain.

5. In the `customer_ca` directory, create a new file in PEM format.
6. Copy the certificate from the `FederationMetadata.xml` file, which is located inbetween the `<X509Certificate>` and `</X509Certificate>` tag set and paste into the new PEM file. Be sure to include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----`, for example:

```
-----BEGIN CERTIFICATE-----
<CERTIFICATE CONTENT>
-----END CERTIFICATE-----
```

7. Save the file and close.
8. Run the following command to update the `ca_outside_bundle.crt` on all management nodes:

```
python3 /usr/lib/python3.6/site-packages/pca_foundation/secret_service/cert_generator/cert_generator_app.py
```

2.5.3 Managing Identity Providers

To federate with an identity provider in Oracle Private Cloud Appliance you create it in the Compute Web UI and map account groups.

After you create your identity provider, you might have the need to make an update. For example, you will need to update your metadata XML file when it expires. You can also view all identity providers, view details of or delete an identity provider.

2.5.3.1 Adding Active Directory as an Identity Provider

To federate with Active Directory in Oracle Private Cloud Appliance you must add it as an identity provider. At the same time, you can set up the group mappings or you can set them up later.

To add AD as an identity provider, follow the procedure for the Compute Web UI.

Using the Compute Web UI

1. Sign in with your Oracle Private Cloud Appliance login and password.
2. Open the navigation menu, click Identity, and then click Federation.
3. On the Federation page, click Create Identity Provider.
4. On the Create an Identity Provider page, provide the following information:

a. *Display Name*

The name that the federated users see when choosing which identity provider to use for signing in to the Compute Web UI. This name must be unique across all identity providers you add to the tenancy and cannot be changed.

b. *Description*

A friendly description of the identity provider.

c. *Authentication Contexts*

Click Add Class Reference and select an authentication context from the list.

When one or more values are specified, Oracle Private Cloud Appliance (the relying party), expects the identity provider to use one of the specified authentication mechanisms when authenticating the user. The returned SAML response from the identity provider must contain an authentication statement with that authentication context class reference. If the SAML response authentication context does not match what is specified here, the Oracle Private Cloud Appliance authentication service rejects the SAML response with a 400.

d. *Encrypt Assertion (Optional)*

When enabled, the authorization service expects encrypted assertions from the identity provider. Only the authorization service can decrypt the assertion. When not enabled, the authorization service expects SAML tokens to be unencrypted, but protected, by SSL.

e. *Force Authentication (Optional)*

When enabled, users are always asked to authenticate at their identity provider when redirected by the authorization service. When not enabled, users are not asked to re-authenticate if they already have an active login session with the identity provider.

f. *Metadata*

Upload the FederationMetadata.xml document from your SAML 2.0 compliant identity provider. You can drag and drop the file or you can paste the XML content.

g. *Tagging (Optional)*

Add any free-form or defined tags.

5. Click Create Identity Provider.

Your new identity provider is assigned an OCID and is displayed on the Federations page

After the identity provider is added to your tenancy, you must set up the group mappings between Oracle Private Cloud Appliance and Active Directory.

To set up group mappings, see [Section 2.5.4.1, "Creating Group Mappings"](#).

2.5.3.2 Updating an Identity Provider

To update an identity provider, follow the procedure for the Compute Web UI.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers in your tenancy is displayed.

2. For the identity provider you want to update, click the Actions icon (three dots) and then click Edit.
3. Change any of the following information; however, be aware that changing this information can affect the federation:

- *Description*

- *Authentication Contexts*

Add or delete a class reference.

- *Encrypt Assertion*

Enable or disable encrypted assertions from the identity provider.

- *Force Authentication*

Enable or disable redirect authentication from the identity provider.

- *Metadata*

Upload a new FederationMetadata.xml document from the identity provider.

- *Tagging*

Add or delete any free-form or defined tags.

For more information, see step 4 in [Section 2.5.3.1, “Adding Active Directory as an Identity Provider”](#)

4. Click Update Identity Provider.

2.5.3.3 Viewing Identity Provider Details

The identity provider details page displays general information such as OCID and authentication contexts. It also provides the identity provider's settings, which include the redirect URL.

From this page, you can also edit the identity provider and manage the group mappings.

To view details for an identity provider, follow the procedure for either the Compute Web UI or the OCI CLI.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers in your tenancy is displayed.

2. For the identity provider whose details you want to view, click the Actions icon (three dots) and then click View Details.

The identity provider details page is displayed.

Using the OCI CLI

1. Find the required OCID:

- `(oci iam identity-provider list)`

2. Enter the command followed by the required parameters.

Syntax

```
oci iam identity-provider get --identity-provider-id identity-provider-ocid
```

Example

```
# oci iam identity-provider get --identity-provider-id ocid1...unique-id
```

2.5.3.4 Listing Identity Providers

To list the identity providers for a tenancy, follow the procedure for either the Compute Web UI or the OCI CLI.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers in your tenancy is displayed.

Using the OCI CLI

1. Find the required OCID:
 - (`oci iam compartment list -include-root`)
2. Enter the command followed by the required parameters.

Syntax

```
oci iam identity-provider list --compartment-id tenancy-ocid
```

Example

```
# oci iam identity-provider list --compartment-id ocid1.tenancy...unique-id
```

2.5.3.5 Deleting an Identity Provider

If you want to remove the option for federated users to log into Oracle Private Cloud Appliance you must delete the identity provider, which also deletes all of the associated group mappings.

To delete an identity provider, follow the procedure for the Compute Web UI.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers in your tenancy is displayed.
2. For the identity provider you want to delete, click the Actions icon (three dots) and then click Delete.
3. At the Delete Identity Provider prompt, click Confirm.

You will see a brief Success pop-up and the identity provider is no longer in the Federation list.

2.5.4 Working with Group Mappings for an Identity Provider

When working with group mappings, keep in mind the following:

- A given Active Directory group is mapped to a single Oracle Private Cloud Appliance group.
- Oracle Private Cloud Appliance group names cannot contain spaces and cannot be changed later. Allowed characters are letters, numerals, hyphens, periods, underscores, and plus signs (+).
- You can't update a group mapping, but you can delete the mapping and add a new one.

2.5.4.1 Creating Group Mappings

After you have created an identity provider, you must create mappings from ADFS groups to Oracle Private Cloud Appliance groups.

To create a group mapping, follow the procedure for the Compute Web UI. Repeat the steps for each identity provider group you want to map.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.
A list of the identity providers in your tenancy is displayed.
2. Click the identity provider whose group mappings you want to create.
The identity provider's details page is displayed.
3. In the Resources section, click Add Mappings.
The IDP Group Mapping Form is displayed
4. In the Name field, enter the *exact* name of the identity provider group.
5. From the Group list, select the Oracle Private Cloud Appliance group you want to map to the identity provider group.
6. Click Create IDP Group Mapping.
The new group mapping is displayed in the list.

2.5.4.2 Viewing Group Mappings

To view group mapping details, follow the procedure for the Compute Web UI.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.
A list of the identity providers in your tenancy is displayed.
2. Click the name of the identity provider.
The identity provider's details page is displayed.
3. In the Resources section, click Group Mappings.
A list of identity provider groups is displayed.

2.5.4.3 Deleting a Group Mapping

To delete a group mapping, follow the procedure for the Compute Web UI. Repeat the steps for each identity provider group you want to delete.

Using the Compute Web UI

1. Open the navigation menu, click Identity and then click Federation.

A list of the identity providers in your tenancy is displayed.

2. Click the name of the identity provider that contains a group mapping you want to delete.

The identity provider's details page is displayed.

3. In the Resources section, click Group Mappings.

For the group mapping you want to delete, click the Actions icon (three dots) and then click Delete.

4. Confirm when prompted.

You will see a brief Success pop-up and the identity provider group is no longer in the Group Mappings list.

2.5.5 Adding Oracle Private Cloud Appliance as a Trusted Relying Party in ADFS

Important

The Oracle Private Cloud Appliance certificate bundle must be added to Active Directory, so that ADFS can trust the Oracle Private Cloud Appliance certificate. If you do not do this, user logins will fail. For more information about the Oracle Private Cloud Appliance certificate bundle, see [Section 1.2.3.4, "Obtaining the Certificate Authority Bundle"](#).

To complete the federation process, you must add Oracle Private Cloud Appliance as a trusted relying party in ADFS and then add associated relying party claim rules.

1. In the Compute Web UI on the Federation page, view the following text block:

You need the Private Cloud Appliance Federation Metadata document when setting up a trust with Microsoft Active Directory Federation Services or with other SAML 2.0-compliant identity providers. This is an XML document that describes the Private Cloud Appliance endpoint and certificate information. [Click Here](#)

2. Click "Click Here".

A metadata XML file opens in the browser with a URL similar to:

```
https://console.system-name.domain-name/wsapi/rest/saml/metadata/ocid1.tenancy...unique-id
```

3. Copy the metadata XML file URL.
4. From the system installed with ADFS, open a browser window and paste the URL.
5. Save the file making sure to use the .xml extension, for example, *my-sp-metadata.xml*.
6. Go to the AD FS Management Console and sign in to the account you want to federate.
7. Add Oracle Private Cloud Appliance as a trusted relying party.
 - a. Under AD FS, right-click Relying Party Trusts and then select Add Relying Party Trust.

- b. In the Add Relying Party Trust Wizard Welcome page, select Claims Aware and then click Start.
- c. On the Select Data Source page, select "Import data about the relying party from a file".
- d. Click Browse and navigate to your *my-sp-metadata.xml* and then click Open.
- e. On the Specify Display Name page, enter a display name, add any optional notes for the relying party, and then click Next.
- f. On the Choose Access Control Policy page, select the type of access you want to grant and then click Next.
- g. On the Ready to Add Trust page, review the settings, and then click Next to save your relying party trust information.
- h. On the Finish page, check "Configure claims issuance policy for this application" and then click Close.

The Edit Claim Issuance Policy dialog appears, which you can leave open for the next section.

Adding Relying Party Claim Rules

After you add Oracle Private Cloud Appliance as a trusted relying party, you must add the claim rules so that the elements required (Name ID and groups) are added to the SAML authentication response.

To add a Name ID rule:

1. In the Edit Claim Issuance Policy dialog, click Add Rule.

The Select Rule Template dialog is displayed.

2. For Claim rule template, select Transform an Incoming Claim and then click Next.

3. Enter the following:

- *Claim rule name*: Enter a name for this rule, for example, `nameid`.
- *Incoming claim type*: Select Windows account name.
- *Outgoing claim type*: Select a claim type, for example, Name ID.
- *Outgoing name ID format*: Select Persistent Identifier.
- Select Pass through all claim values and then click Finish.

The rule is displayed in the rules list.

The Issuance Transform Rules dialog displays the new rule.

If your Active Directory users are in no more than 100 groups, you simply add the groups rule. However, if your Active Directory users are in more than 100 groups, those users cannot be authenticated to use the Oracle Private Cloud Appliance Compute Web UI. For these groups, you must apply a filter to the groups rule.

To add the groups rule:

1. In the Issuance Transform Rules dialog, click Add Rule.

The Select Rule Template dialog is displayed.

2. For Claim rule template, select Send Claims Using a Custom Rule and then click Next.
3. In the Add Transform Claim Rule Wizard, enter the following:
 - a. *Claim rule name*: Enter groups.
 - b. *Custom rule*: Enter the custom rule.

```
c:[Type == "http://schemas.microsoft.com/ws/2008/06/identity/claims/windowsaccountname", Issuer == "
```
 - c. Click Finish.

The Issuance Transform Rules dialog displays the new rule.

2.5.6 Setting Up Policies for the Groups

If you haven't already done so, set up IAM policies to control the access the federated users have to your organization's Oracle Private Cloud Appliance resources. For more information, see [Identity and Access Management Overview](#)

2.5.7 Providing Federated Users Sign In Information

Before federated users can log into Oracle Private Cloud Appliance Compute Web UI, you must provide them with the URL and the name of the tenant to which they have access. You must also ensure that you have configured the groups mappings otherwise a federated user cannot do any work in Oracle Private Cloud Appliance.

Important

A federated user cannot log into Oracle Private Cloud Appliance using the OCI CLI.

Chapter 3 Resource Tag Management

Table of Contents

3.1 Creating and Managing Tag Namespaces	73
3.1.1 Creating a Tag Namespace	73
3.1.2 Updating a Tag Namespace	74
3.1.3 Retiring a Tag Namespace	75
3.1.4 Reactivating a Tag Namespace	76
3.1.5 Moving a Tag Namespace to a Different Compartment	76
3.1.6 Deleting a Tag Namespace	77
3.2 Creating and Managing Tag Key Definitions	78
3.2.1 Creating a Tag Key Definition	78
3.2.2 Updating a Tag Key Definition	80
3.2.3 Retiring a Tag Key Definition	82
3.2.4 Reactivating a Tag Key Definition	82
3.2.5 Deleting a Tag Key Definition	83
3.3 Configuring Tag Defaults	84
3.3.1 Creating a Tag Default	84
3.3.2 Updating the Value of a Tag Default	86
3.3.3 Deleting a Tag Default	86
3.4 Working with Resource Tags	87
3.4.1 Adding Tags at Resource Creation	87
3.4.2 Applying Tags to an Existing Resource	89
3.4.3 Filtering a List of Resources by Tag	90

Oracle Private Cloud Appliance Tagging enables you to add metadata to resources by applying key/value pairs called defined tags or free-form tags. You can create tag defaults on compartments, which are tags that are automatically applied to all newly created resources in the tagged compartment. Uses for tags include:

- Applying access policies to resources. For example, you can change ownership of a resource to a different product group by changing a tag value on the resource rather than changing the resource access policy directly. See [Section 2.4.2.4, “Using Defined Tags in Conditions”](#)).
- Filtering resource lists in the Compute Web UI.

For conceptual information, see the [Tagging Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide* and *How Policies Work* in the [Identity and Access Management Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

3.1 Creating and Managing Tag Namespaces

Tag namespaces enable you to create collections of related tags. After you create a tag namespace, create tag key definitions within that tag namespace. See [Section 3.2, “Creating and Managing Tag Key Definitions”](#). Tag namespaces with tag key definitions must exist in the tenancy before users can apply a defined tag to a resource.

3.1.1 Creating a Tag Namespace

A tenancy can have at most 100 tag namespaces.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. Above the list of tag namespaces, click the Create Namespace Definitions button.
3. In the Create Namespace Definition window, enter the following information:
 - *Create in Compartment*: The compartment in which you want to create the namespace definition.
 - *Namespace Definition Name*: A name for this tag namespace. Tag namespace names have the following characteristics:
 - Must be unique within the tenancy.
 - Are case insensitive.
 - Cannot be changed later.
 - Can be no more than 100 characters.
 - Cannot contain period (.) or space characters.
 - *Description*: A description for this set of tags. This description can be no more than 256 characters.
 - *Tagging*: (Optional) Add defined or free-form tags for this tag namespace as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.
4. Click Create Namespace Definition.

The new tag namespace definition is displayed on the Tag Namespaces page.

Using the OCI CLI

1. Get the OCID of the compartment where you want to create the tag namespace.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

2. Run the tag namespace create command.

Syntax:

```
oci iam tag-namespace create --compartment-id compartment_OCID --name tag_namespace_name \
--description "text"
```

You can tag the new tag namespace during creation by adding the options described in [Section 3.4.1, “Adding Tags at Resource Creation”](#).

Example:

```
$ oci iam tag-namespace create --compartment-id ocid1.compartment.unique_ID --name Products \
--description "Identify resources used in product development."
```

This command returns the same output as the `tag-namespace get` command.

3.1.2 Updating a Tag Namespace

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.

2. If the tag namespace that you want to modify is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. For the namespace that you want to modify, click the Actions menu, and click the Edit option.
The Edit dialog is displayed.
4. Update the tag namespace.
You can modify the description of a tag namespace and add or modify tags on the tag namespace.
5. Click Update Tag Namespace.

Using the OCI CLI

1. Get the tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Run the tag namespace update command.

Syntax:

```
oci iam tag-namespace update --tag-namespace-id tag_namespace_OCID --description "text"
```

To add or modify a tag on the tag namespace, add the options described in [Section 3.4.2, “Applying Tags to an Existing Resource”](#).

Example:

```
$ oci iam tag-namespace update --tag-namespace-id ocid1.tagnamespace.unique_ID \  
--description "Identify resources used to develop different products."
```

This command returns the same output as the `tag-namespace get` command.

3.1.3 Retiring a Tag Namespace

When you retire a tag namespace, all tag key definitions and tags in that tag namespace are retired, and you cannot create new tag key definitions in that tag namespace. Retired tags cannot be applied to resources. However, retired tags remain applied to any resources where they were already applied and can still be used in operations such as listing, sorting, and filtering.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace that you want to retire is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. For the tag namespace that you want to retire, click the Actions menu, and click the Retire option.
4. At the Retire Tag Namespace confirmation prompt, click Confirm.

The state of the tag namespace changes to Inactive. On the details page for the tag namespace, the tag key definitions are also in state Inactive.

Using the OCI CLI

1. Get the tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Run the tag namespace retire command.

Syntax:

```
oci iam tag-namespace retire --tag-namespace-id tag_namespace_OCID
```

This command returns the same output as the `tag-namespace get` command.

3.1.4 Reactivating a Tag Namespace

You can reactivate a tag namespace that is retired. When you reactivate a tag namespace, you can create new tag key definitions in that tag namespace.

When you reactivate a retired tag namespace, the tag key definitions and tags are not reactivated. To use tag key definitions that were retired with the namespace, you must explicitly reactivate each tag key definition.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace that you want to reactivate is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. For the tag namespace that you want to reactivate, click the Actions menu, and click the Reactivate option.
4. At the Reactivate Tag Namespace confirmation prompt, click Confirm.

The state of the tag namespace changes to Active. On the details page for the tag namespace, the tag key definitions are still shown as Inactive.

Using the OCI CLI

1. Get the tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Run the tag namespace reactivate command.

Syntax:

```
oci iam tag-namespace reactivate --tag-namespace-id tag_namespace_OCID
```

This command returns the same output as the `tag-namespace get` command.

3.1.5 Moving a Tag Namespace to a Different Compartment

You can move an active or retired tag namespace and its tag key definitions to a different compartment within the same tenancy.

To move a tag namespace, you must be granted `manage tag-namespaces` access in both compartments.

To move a tag namespace, you must use the OCI CLI.

Using the OCI CLI

1. Get the following information:

- The OCID of the compartment where you want to move the tag namespace.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- The tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Run the tag namespace move command.

Syntax:

```
oci iam tag-namespace change-compartment --compartment-id destination_compartment_OCID \  
--tag-namespace-id tag_namespace_OCID
```

Use the `tag-namespace get` command to verify the new `compartment-id`.

3.1.6 Deleting a Tag Namespace

A tag namespace must be retired before it can be deleted. See [Section 3.1.3, “Retiring a Tag Namespace”](#).

To delete a tag namespace, all tag key definitions in that namespace must be deleted. See [Section 3.2.5, “Deleting a Tag Key Definition”](#).

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace that you want to delete is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Ensure that the tag namespace that you want to delete is retired (in state Inactive).
4. Click the name of the tag namespace that you want to delete.
5. In the Resources box on the tag namespace details page, click Tag Key Definitions.
6. Ensure that all tag key definitions are deleted.
7. On the Controls menu at the top of the details page, click the Delete option.
8. At the Delete Tag Namespace confirmation prompt, click Confirm.

The tag namespace is removed from the Tag Namespace list, and tags in that tag namespace are removed from resources.

Using the OCI CLI

1. Get the tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Ensure that the tag namespace that you want to delete is retired (in state Inactive).

```
$ oci iam tag-namespace get --tag-namespace-id tag_namespace_OCID
```

3. Ensure that all tag key definitions for this tag namespace are deleted.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

4. Run the tag namespace delete command.

Syntax:

```
oci iam tag-namespace delete --tag-namespace-id tag_namespace_OCID
```

Example:

```
$ oci iam tag-namespace delete --tag-namespace-id ocidl.tagnamespace.unique_ID
Are you sure you want to delete this resource? [y/N]: y
{
  "opc-work-request-id": "ocidl.workrequest.unique_ID"
}
```

Use the following command to check the status of the tag namespace delete:

```
$ oci iam tagging-work-request get --work-request-id ocidl.workrequest.unique_ID
```

To delete a tag namespace without confirmation, use the `--force` option.

3.2 Creating and Managing Tag Key Definitions

A tag namespace contains tag key definitions. Instances of tag key definitions that are applied to resources are called *defined tags*.

A tag key definition includes a tag key name and tag value type. A tag key definition might include a value, depending on the tag value type. The tag value type specifies whether the tag user enters a value or selects a predefined value when applying a tag to a resource.

For more information about tag key definitions, including required permissions, see the [Tagging Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

3.2.1 Creating a Tag Key Definition

Create a tag key definition within a tag namespace. A tag namespace can have at most 100 tag key definitions.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace where you want to add the tag key definition is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Click the name of the tag namespace where you want to add a tag key definition.
4. In the Resources box on the tag namespace details page, click Tag Key Definitions.
5. In the Tag Key Definitions area, click the Create Tag Key Definition button.
6. In the Create Tag Key Definition window, enter the following information:
 - *Name*: The key name. Tag key names have the following characteristics:
 - Must be unique within the namespace.
The same tag key name can be used in different tag namespaces.
 - Are case insensitive.

- Cannot be changed later.
- Can be no more than 100 characters.
- Cannot contain period (.) or space characters.
- *Description*: A description for the tag key definition. This description can be no more than 256 characters.

7. Select the Tag Value Type.

- *Static Value*: Not populated. The user must enter a value when the tag is applied to a resource. This is the default selection.
- *A List of Values*: A predefined list of values. Enter at least one value. The user must select one of these predefined values to apply to a resource.

Separate multiple values with new lines. Duplicate values and blank lines are invalid. Values are case sensitive and can be no more than 256 characters.

A value can include one of the following variables:

`${iam.principal.name}` The name of the user that tagged the resource.

`${iam.principal.type}` The type of principal that tagged the resource. One of: root user, IAM user, or Instance principal.

`${oci.datetime}` The date and time that the tag was created.

8. Click Create Tag Key Definition.

The new tag key definition is displayed on the details page for the tag namespace.

Using the OCI CLI

1. Get the tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

2. Run the tag key definition create command.

The following command creates a tag key definition in which the user must enter a tag value when the tag is applied to a resource. Omitting the `--validator` option is comparable to selecting the *Static Value* value type in the Compute Web UI.

Syntax:

```
oci iam tag create --tag-namespace-id tag_namespace_OCID --name text --description "text"
```

Example:

```
$ oci iam tag create --tag-namespace-id ocid1.tagnamespace.unique_ID --name VolumeType \
--description "Identify volumes by type"
```

The following command creates a tag key definition with a list of values from which the user must select when applying the tag.

Syntax:

```
oci iam tag create --tag-namespace-id tag_namespace_OCID --name text --description "text" \
--validator values
```

The value of the `--validator` option argument is a JSON definition of the tag values. This JSON definition can be provided as a string on the command line or in a file.

You can generate a template of the correct JSON to provide by using the `--generate-param-json-input` option with the base command that you will use to tag the resource. The argument for the `--generate-param-json-input` option is the name of the option that you will use to specify the tag values (`--validator`) without the option indicator (`--`), as shown in the following example:

```
$ oci iam tag create --generate-param-json-input validator > volume_types.json
```

The following is the content of the output `volume_types.json` file:

```
{
  "validatorType": "ENUM",
  "values": [
    "string",
    "string"
  ]
}
```

Edit this template to provide new tag values. The value of `values` is a single string or an array of strings. The `validator-type` must be `ENUM`. Specify the result to the `--validator` option in the final command.

In the following example, the tag values are provided in an inline JSON string.

```
$ oci iam tag create --tag-namespace-id ocidl.tagnamespace.unique_ID --name VolumeType \
--description "Identify volumes by type" \
--validator '{"validator-type": "ENUM", "values": ["typeA", "typeB", "typeC"]}'
```

The following example includes a variable value:

```
$ oci iam tag create --tag-namespace-id ocidl.tagnamespace.unique_ID --name Product-XYZ \
--description "Identify resources assigned to XYZ development." \
--validator '{"validator-type": "ENUM", "values": ["Assigned by: ${iam.principal.name}"]}'
```

In the following example, the tag values are provided in a JSON file. Use the `file://` syntax to specify a file as the option argument.

```
$ oci iam tag create --tag-namespace-id ocidl.tagnamespace.unique_ID --name VolumeType \
--validator file://volume_types.json
```

This command returns the same output as the `tag get` command.

3.2.2 Updating a Tag Key Definition

You can update the description of a tag key definition and change the tag value type and value. You cannot update a tag key definition that is retired.

If the value of the tag key definition that you are updating is a predefined list, you cannot remove or change any value that is the value of a tag default. To remove or change a tag key definition value that is the value of a tag default, first update the tag default to use a different value. See [Section 3.3, “Configuring Tag Defaults”](#).

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.

2. If the tag namespace where you want to update a tag key definition is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Click the name of the tag namespace where you want to update a tag key definition.
4. In the Resources box on the tag namespace details page, click Tag Key Definitions.
5. For the tag key definition that you want to update, click the Actions menu, and click the Edit option.
6. In the Edit Tag Key Definition dialog, you can modify the description or change the tag value type. If you choose A List of Values for the type, you must add at least one value in the Value box. Separate multiple values with new lines. Duplicate values and blank lines are invalid. Values are case sensitive and can be no more than 256 characters.
7. Click Save Changes.

Using the OCI CLI

1. Get the following information:

- The tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

- The name of the tag key.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

2. Run the tag key definition update command.

The following command updates only the tag key description. Values settings remain the same.

Syntax:

```
oci iam tag update --tag-namespace-id tag_namespace_OCID --tag-name text --description "text"
```

Example:

```
$ oci iam tag update --tag-namespace-id ocid1.tagnamespace.unique_ID --tag-name VolumeType \
--description "Identify the type of the volume."
```

The following command updates only the values.

```
$ oci iam tag update --tag-namespace-id ocid1.tagnamespace.unique_ID --tag-name VolumeType \
--validator file://volume_types.json
WARNING: Updates to freeform-tags and defined-tags and validator will replace any existing
values.
Are you sure you want to continue? [y/N]: y
```

See [Section 3.2.1, “Creating a Tag Key Definition”](#) for the content of the `volume_types.json` file. Values that you provide to the `--validator` option replace any existing values. To add values or to change only some of the values, provide the complete list in this update.

The following command updates the value type to allow the user to enter a value rather than select from a predefined list of values.

```
$ oci iam tag update --tag-namespace-id ocid1.tagnamespace.unique_ID --tag-name VolumeType \
--validator '{}'
```

This command returns the same output as the `tag get` command.

3.2.3 Retiring a Tag Key Definition

When you retire a tag key definition, you cannot apply tags that are based on this tag key definition to resources. Existing tag defaults that are based on this tag key definition will not be automatically applied to newly created resources. However, the tag is not removed from resources where it was already applied. The tag still exists as metadata on those resources and you can still use the retired tag in operations such as listing, sorting, and reporting.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace where you want to retire a tag key definition is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Click the name of the tag namespace where you want to retire a tag key definition.
4. In the Resources box on the tag namespace details page, click Tag Key Definitions.
5. For the tag key definition that you want to retire, click the Actions menu, and click the Retire option.
6. At the Retire Tag Key Definition confirmation prompt, click Confirm.

Using the OCI CLI

1. Get the following information:

- The tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

- The name of the tag key.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

2. Run the tag key definition retire command.

Syntax:

```
oci iam tag retire --tag-name text --tag-namespace-id tag_namespace_OCID
```

3.2.4 Reactivating a Tag Key Definition

When you reactivate a tag key definition, it is again available for you to apply to resources. You cannot reactivate a tag key definition if the parent tag namespace is retired.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace where you want to reactivate a tag key definition is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Click the name of the tag namespace where you want to reactivate a tag key definition.
4. In the Resources box on the tag namespace details page, click Tag Key Definitions.
5. For the tag key definition that you want to reactivate, click the Actions menu, and click the Reactivate option.
6. At the Reactivate Tag Namespace confirmation prompt, click Confirm.

Using the OCI CLI

1. Get the following information:

- The tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

- The name of the tag key.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

2. Run the tag key definition reactivate command.

Syntax:

```
oci iam tag reactivate --tag-name text --tag-namespace-id tag_namespace_OCID
```

3.2.5 Deleting a Tag Key Definition

A tag key definition must be retired before it can be deleted. See [Section 3.2.3, “Retiring a Tag Key Definition”](#).

When you delete a tag key definition, tags that are based on this tag key definition are removed from all resources. Tag defaults that are based on this tag key definition are not removed from compartments.

Using the Compute Web UI

1. In the navigation menu, click Governance, and then click Tag Namespaces.
2. If the tag namespace where you want to delete a tag key definition is not listed, use the Compartment drop-down menu above the tag namespaces list to select the correct compartment.
3. Click the name of the tag namespace where you want to delete a tag key definition.
4. In the Resources box on the tag namespace details page, click Tag Key Definitions.
5. Ensure that the tag key definition that you want to delete is retired (in state Inactive).
6. For the tag key definition that you want to delete, click the Actions menu, and click the Delete option.

You cannot restore a deleted tag key definition.

7. At the Delete Tag Key Definition prompt, click Confirm.

The tag key definition status changes to Deleting, and all tags that are based on this tag key definition are removed from resources.

When the tag removal process is finished, the tag key definition status changes to Deleted. You can create a new tag key definition with the same name as the deleted tag key definition.

Using the OCI CLI

1. Get the following information:

- The tag namespace OCID.

```
$ oci iam tag-namespace list --compartment-id compartment_OCID
```

- The name of the tag key.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

2. Ensure that the tag key definition that you want to delete is retired (in state Inactive).

```
$ oci iam tag get --tag-namespace-id tag_namespace_OCID
```

3. Run the tag key definition delete command.

Syntax:

```
oci iam tag delete --tag-name text --tag-namespace-id tag_namespace_OCID
```

Example:

```
$ oci iam tag delete --tag-name volume-types --tag-namespace-id ocidl.tagnamespace.unique_ID
Are you sure you want to delete this resource? [y/N]: y
{
  "opc-work-request-id": "ocidl.workrequest.unique_ID"
}
```

Use the following command to check the status of the tag delete:

```
$ oci iam tagging-work-request get --work-request-id ocidl.workrequest.unique_ID
```

To delete a tag key definition without confirmation, use the `--force` option.

3.3 Configuring Tag Defaults

A tag default is a defined tag that is automatically applied to resources that are created in the specified compartment.

Tag defaults have the following characteristics:

- The tag default is applied to all new resources that are created in that compartment, including in child compartments.
- The tag default is not applied to resources that already existed before the tag default was created.
- Tag defaults cannot be changed by creating or editing resources. With permission to use the tag namespace, you can change the value of the tag when you create or modify a resource. To change the tag default that will be applied to all new resources in the compartment, you must update the tag default on the compartment.
- If you change the default value of the tag default, existing occurrences of that tag default on resources are not updated.
- If you change a value of a tag key definition whose value is a defined list, or if you delete a tag key definition, existing occurrences of a tag default that is based on that tag key definition are not updated. Tag default values must be separately updated.

See [Section 3.2, “Creating and Managing Tag Key Definitions”](#) for information about the effect on tag defaults of retiring or deleting tag key definitions. For more information about tag defaults, see *Tag Defaults* in the [Tagging Overview](#) in the *Oracle Private Cloud Appliance Concepts Guide*.

3.3.1 Creating a Tag Default

To create a tag default, specify a compartment, a tag key definition, and a value. If the value of the selected tag key definition is *Static Value*, then you can select *User-Defined Value* for the value of the tag default.

A compartment can have at most five tag defaults.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Compartments.
2. If the compartment where you want to add a tag default is not listed, navigate to the correct compartment.

Click the name of the top-level parent compartment and, on the compartment details page, scroll to the Child Compartments box. If necessary, click on the name of another compartment to view those child compartments.

3. Click the name of the compartment to which you want to add a tag default.
4. In the Resources section on the compartment details page, click Tag Defaults.

Ensure that no more than four tag defaults already exist in this compartment.

5. Click the Create Tag Default button.
6. In the Tag Default dialog, select the Tag Namespace and the Tag Key.
7. For Required Tag Value Options, choose one of the following value types:
 - *Default Value*: Enter the value for this tag default. If the selected tag key has a defined list of values, then this *Default Value* must be a member of that list.
 - *User-Defined Value*: Users are required to enter the value when a resource is created. Selecting *User-Defined Value* is invalid if the selected tag key definition has a predefined list of values.
8. Click Submit.

The new tag default is displayed on the compartment details page.

Using the OCI CLI

1. Get the following information:
 - The OCID of the compartment on which you want to create the tag default.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- The OCID of the tag key definition.

```
$ oci iam tag list --tag-namespace-id tag_namespace_OCID
```

2. Ensure that no more than four tag defaults already exist in this compartment.

```
oci iam tag-default list --compartment-id compartment_OCID
```

3. Run the tag default create command.

Syntax:

```
oci iam tag-default create --compartment-id compartment_OCID \  
--tag-definition-id tag_definition_OCID --value text
```

Example:

```
$ oci iam tag-default create --compartment-id ocid1.compartment.unique_ID \  

```

```
--tag-definition-id ocidl.tag.unique_ID --value 789
```

Depending on your shell, you might need to escape the dollar symbol to specify a variable value:

```
$ oci iam tag-default create --compartment-id ocidl.compartment.unique_ID \  
--tag-definition-id ocidl.tag.unique_ID --value "Assigned by: \${iam.principal.name}"
```

This command returns the same output as the `tag-default get` command.

3.3.2 Updating the Value of a Tag Default

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Compartments.
2. If the compartment where you want to update a tag default is not listed, navigate to the correct compartment.

Click the name of the top-level parent compartment and, on the compartment details page, scroll to the Child Compartments box. If necessary, click on the name of another compartment to view those child compartments.

3. Click the name of the compartment that has the tag default whose value you want to change.
4. In the Resources section of the compartment details page, click Tag Defaults.
5. For the tag default that you want to change, click the Actions menu, and click the Edit option.
6. In the Tag Defaults dialog, specify the type of value you want the tag default to have:
 - *Default Value*: Enter the value for this tag default. If the selected tag key has a defined list of values, then this *Default Value* must be a member of that list.
 - *User-Defined Value*: Users are required to enter the value when a resource is created. Selecting *User-Defined Value* is invalid if the selected tag key definition has a predefined list of values.
7. Click Submit.

The updated tag default is displayed on the compartment's details page.

Using the OCI CLI

1. Get the OCID of the tag default that you want to modify.

```
$ oci iam tag-default list --compartment-id compartment_OCID
```

2. Run the tag default update command.

Syntax:

```
oci iam tag-default update --tag-default-id tag_default_OCID --value text
```

This command returns the same output as the `tag-default get` command.

3.3.3 Deleting a Tag Default

When you delete a tag default from a compartment, existing occurrences of the tag are not removed from resources.

Using the Compute Web UI

1. In the navigation menu, click Identity, and then click Compartments.
2. If the compartment where you want to delete a tag default is not listed, navigate to the correct compartment.

Click the name of the top-level parent compartment and, on the compartment details page, scroll to the Child Compartments box. If necessary, click on the name of another compartment to view those child compartments.

3. Click the name of the compartment that has the tag default that you want to delete.
4. In the Resources section of the compartment details page, click Tag Defaults.
5. For the tag default that you want to delete, click the Actions menu, and click the Delete option.
6. At the Delete Default Tag confirmation prompt, click Confirm.

On the compartment details page, the state of the tag default is Deleting.

Using the OCI CLI

1. Get the OCID of the tag default that you want to delete.

```
$ oci iam tag-default list --compartment-id compartment_OCID
```

2. Run the tag default delete command.

Syntax:

```
oci iam tag-default delete --tag-default-id tag_default_OCID
```

Example:

```
$ oci iam tag-default delete --tag-default-id ocid1.tag-default.unique_ID
Are you sure you want to delete this resource? [y/N]: y
```

To delete a tag default without confirmation, use the `--force` option.

3.4 Working with Resource Tags

Tagging resources enables you to identify characteristics of resources and apply the same policies to a group of resources. For more information about policies, see [Section 2.4.2.4, “Using Defined Tags in Conditions”](#).

Note the following requirements:

- To apply, modify, or delete a tag on a resource, you must have permission to update the resource.
- To apply, modify, or delete a defined tag on a resource, you must also have `use` access on the tag namespace.
- A resource can have at most 64 defined tags and ten free-form tags.

3.4.1 Adding Tags at Resource Creation

Any tag defaults that are defined on a compartment are automatically added to all resources that are created in that compartment, or any child compartment of that compartment, after the tag default was

defined. A tag default might require you to enter a value for the tag in order to create the resource. See [Section 3.3, “Configuring Tag Defaults”](#).

Using the Compute Web UI

1. In the Create dialog for the resource, scroll to the Tagging section.
2. Select the Tag Namespace or select None (apply a free-form tag).
 - If you selected a Tag Namespace, then select the Tag Key, and enter a value or select a value from the list.
 - If you selected None (apply a free-form tag), then enter a Tag Key and enter a value.
3. To apply another tag, click the Additional Tag button.

You cannot specify more than one tag with the same tag namespace and the same tag key for a defined tag, or more than one tag with the same tag key for a free-form tag.

4. To review the tags on the resource, to go the details page for the new resource.

On the resource details page, click the Tags tab to display the tags that are applied to this resource.

Using the OCI CLI

To add a tag to a resource when you create the resource, use the resource `create` or `launch` command.

1. Get the information for each tag that you want to add to the resource.
 - Get the namespace, key, and value for each defined tag that you want to add to the resource.

Construct an argument for the `--defined-tags` option. Specify each tag namespace and tag key pair only one time.

- Get the key and value for each free-form tag that you want to add to the resource.

Construct an argument for the `--freeform-tags` option. Specify each tag key only one time.

The value of the `--defined-tags` option argument and the `--freeform-tags` option argument is a JSON definition of the tags. This JSON definition can be provided as a string on the command line or in a file.

You can generate a template of the correct JSON to provide by using the `--generate-param-json-input` option with the base command that you will use to tag the resource. The argument for the `--generate-param-json-input` option is the name of the option that you will use to specify the tags (`--defined-tags` in this example) without the option indicator (`--`), as shown in the following example:

```
$ oci service resource create \
--generate-param-json-input defined-tags > defined_tags.json
```

The content of the output `defined_tags.json` file is:

```
{
  "tagNamespace1": {
    "tagKey1": "tagValue1",
    "tagKey2": "tagValue2"
  },
  "tagNamespace2": {
    "tagKey1": "tagValue1",
```

```
  "tagKey2": "tagValue2"
}
```

If you specify `freeform-tags` instead of `defined-tags` in the preceding command, you get the following output:

```
{
  "tagKey1": "tagValue1",
  "tagKey2": "tagValue2"
}
```

Edit these templates to provide the desired tags. Specify the result in the final command as shown in the following step.

2. Run the resource create or launch command.

If you want to add one or more defined tags, use the `--defined-tags` option. If you want to add one or more free-form tags, use the `--freeform-tags` option.

Syntax:

```
oci service resource create --compartment-id compartment_OCID \
--defined-tags defined_tags_json --freeform-tags freeform_tags_json \
other_resource_create_options
```

Example:

In the following example, one or more defined tags is added using a file, and a free-form tag is added using a string argument. Use the `file://` syntax to specify a file as the option argument.

```
$ oci service resource create --compartment-id ocidl.compartment.unique_ID \
--defined-tags file://defined_tags.json --freeform-tags '{ "MyTag": "val-u" }' \
other_resource_create_options
```

The output of the resource `create` or `launch` command is the same as the output of the resource `get` command. The output shows the defined and free-form tags.

3.4.2 Applying Tags to an Existing Resource

Using the Compute Web UI

1. In the resource Edit dialog, scroll to the Tagging section.

You can add tags, and you can modify or delete any tags that already exist.

2. To add tags, click the Additional Tag button if necessary, and select the Tag Namespace or select None (apply a free-form tag).
 - If you selected a Tag Namespace, then select a Tag Key, and enter a value or select a value from the list.
 - If you selected None (apply a free-form tag), then enter a Tag Key and enter a value.

You cannot specify more than one tag with the same tag namespace and the same tag key for a defined tag, or more than one tag with the same tag key for a free-form tag.

3. To modify existing tags, change the selections or enter different values.
4. To delete a tag, click the trash can.

5. When you are finished adding and modifying tags, click Save Changes.
6. To review the tags on the resource, go to the details page for the resource.

On the resource details page, click the Tags tab to display the list of tags that are applied to this resource.

Using the OCI CLI

To add tags to an existing resource, and modify or delete any tags that already exist, use the resource `update` command.

1. Get the namespace, key, and value information for each tag that you want to add to the resource.
2. Create a JSON definition of the tags that you want to apply. See [Section 3.4.1, “Adding Tags at Resource Creation”](#) for information about how to create correct JSON.

Note

Any defined tags that already exist on this resource will be replaced by the `--defined-tags` argument. Any free-form tags that already exist on this resource will be replaced by the `--freeform-tags` argument. Be sure to include any existing tags that you want to keep in the new arguments for these options.

Use the resource `get` command to show the current defined and free-form tags.

3. Run the resource update command.

If you want to apply one or more defined tags, use the `--defined-tags` option. If you want to apply one or more free-form tags, use the `--freeform-tags` option.

The output of the resource `update` command is the same as the output of the resource `get` command. The output shows the defined and free-form tags.

3.4.3 Filtering a List of Resources by Tag

Using the Compute Web UI

1. Display a list of resources.
2. Under Filter by Tag(s), click Select Tag(s).
3. Do the following in the Filter by Tag dialog:
 - a. Select either Defined Tag or Free-Form Tag.
 - If you select Defined Tag, then select a Tag Namespace and a Tag Key.
 - If you selected Free-Form Tag, then enter a Tag Key.
 - b. Optionally enter values in the Select Values (optional) field.
 - Leave the Select Values (optional) field blank. This option returns all resources that are tagged with the selected namespace and key (for defined tags) or returns all resources that are tagged with the specified key (for free-form tags), regardless of the tag value.
 - Enter values in the Select Values (optional) field. This option returns all resources that are tagged with any of the tag value(s) that you enter. Select a value or enter a single value in the text box.

To specify multiple values for the same namespace and key, enter Enter or Return, and then enter a new value. Each value is displayed below the text box.

- c. Click Filter by Tag.

The filter that is currently applied is displayed on the Select Tag(s) button.

4. To filter by multiple tags, click the + on the Select Tag(s) button, and repeat the previous step.
5. To remove a filter, click the filter definition on the Select Tag(s) button.

Using the OCI CLI

1. Use the `list` command for the resource to show each resource in the specified compartment.

The information for each resource shows the existing defined tags and free-form tags.

2. Use tools for your operating system to filter the list.

Chapter 4 Networking

Table of Contents

4.1 Managing VCNs and Subnets	94
4.1.1 Creating a VCN	94
4.1.2 Creating a Subnet	95
4.1.3 Editing a Subnet	96
4.1.4 Deleting a Subnet	98
4.1.5 Terminating a VCN	98
4.1.6 Deleting a VCN	98
4.2 Configuring VCN Rules and Options	99
4.2.1 Working with DHCP Options	99
4.2.2 Working with Route Tables	104
4.2.3 Controlling Traffic with Security Lists	109
4.2.4 Controlling Traffic with Network Security Groups	116
4.3 Configuring VCN Gateways	124
4.3.1 Enabling Public Connections through a NAT Gateway	126
4.3.2 Providing Public Access through an Internet Gateway	128
4.3.3 Connecting VCNs through a Local Peering Gateway	129
4.3.4 Connecting to the On-Premises Network through a Dynamic Routing Gateway	131
4.3.5 Accessing Oracle Services through a Service Gateway	133
4.4 Configuring VNICs and IP Addressing	135
4.4.1 Public IP Addresses	136
4.4.2 Managing Public IP Addresses	136
4.4.3 Creating and Managing VNICs	141
4.4.4 Assigning IP Addresses to VNICs	151
4.5 Managing Public DNS Zones	152
4.5.1 Creating a Public DNS Zone	152
4.5.2 Working with Zone Records	154
4.5.3 Editing a Public DNS Zone	157
4.5.4 Working with Transaction Signature Keys	157
4.5.5 Deleting a Public DNS Zone	159
4.6 Managing Traffic with Steering Policies	160
4.6.1 Creating a Load Balancer Steering Policy	160
4.6.2 Creating an IP Prefix Steering Policy	163
4.6.3 Editing a Steering Policy	165
4.6.4 Moving a Steering Policy to a Different Compartment	166
4.6.5 Attaching a Domain to a Steering Policy	166
4.6.6 Editing an Attached Domain	167
4.6.7 Deleting a Steering Policy Attachment	168
4.6.8 Deleting a Steering Policy	168
4.7 Networking Scenarios	169
4.7.1 Logical Routers	169
4.7.2 Using Firewalls	170
4.7.3 Use of Network Segmentation	170
4.7.4 Use of Tunneling	170
4.7.5 Use of Virtual Cloud Networks	171

4.1 Managing VCNs and Subnets

The VCN is the basic networking unit of the Oracle Private Cloud Appliance product. VCNs can be further divided into IP subnets, and individual VCNs can communicate with each other through various types of gateways, each type intended for a particular purpose.

4.1.1 Creating a VCN

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click Create VCN to open the Create a Virtual Cloud Network window.
3. Enter the following details:
 - **Name:** Enter a descriptive name for the VCN.
 - **Compartment:** Select the compartment in which to create the VCN.
 - **CIDR Block:** Specify which CIDR range can be used within the VCN.
 - **DNS:** Indicate whether you want to use DNS host names in the VCN. If so, enter a DNS label or leave the field blank to let the system generate a name for you. The first character must be a letter. Only use letters and numbers. Up to 15 characters are allowed.
4. Optionally, add one or more tags to this VCN resource.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

5. Click Create Virtual Cloud Network. The new VCN is displayed in the list of Virtual Cloud Networks for the compartment where you created it.

Using the CLI

1. Copy the OCID of the compartment in which to create the VCN.
2. Enter the `vcn create` command using at least the compartment ID and CIDR block options.

If you want to use DNS host names in the VCN, include the DNS label in the create command. It cannot be added later.

In this example we also set a user-friendly display name for the VCN.

```
# oci network vcn create --compartment-id "<target-compartment-ocid>" \
--cidr-block "10.1.0.0/16" --dns-label "vcn1" --display-name "myvcn1"
{
  "data": {
    "cidr-block": "10.1.0.0/16",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "default-dhcp-options-id": "ocidl.dhcptoptions.oc1.pca.....uniqueID",
    "default-route-table-id": "ocidl.routetable.oc1.pca.....uniqueID",
    "default-security-list-id": "ocidl.security_list.oc1.pca.....uniqueID",
    "defined-tags": null,
    "display-name": "myvcn1",
    "dns-label": "vcn1",
    "freeform-tags": null,
```



```

    "id": "ocidl.vcn.oc1.pca.....uniqueID",
    "ipv6-cidr-block": null,
    "ipv6-public-cidr-block": null,
    "lifecycle-state": "AVAILABLE",
    "time-created": "2021-01-11T14:08:15+00:00",
    "vcn-domain-name": "vcn1.oraclevcn.com"
  },
  "etag": "1"
}

```

4.1.2 Creating a Subnet

VCNs can be divided into subnets. Although it is possible to have an enormous VCN with a thousand IP addresses, it often makes sense from a performance and fault isolation standpoint to create multiple subnets within a VCN. The subnets can still communicate if configured properly.

IP subnet calculation can be a difficult task, especially when figuring out which IP addresses in the range are reserved. The wide range of allowable CIDR block addresses complicates the issue. There are free subnet calculation tools to help available online, such as <https://www.calculator.net/ip-subnet-calculator.html>.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN in which you want to create a new subnet. The VCN Detail page is displayed.
3. Click Create Subnet to open the Create Subnet window.
4. Enter the following details:
 - **Name:** Enter a descriptive name for the subnet.
 - **CIDR Block:** Specify which CIDR range can be used within the subnet. It must be within the VCN CIDR block and must not overlap with other subnets.
 - **Route Table (Optional):** Select the route table to associate with this subnet. If you enabled compartment selection, specify the compartment that contains the route table. If you do not select a route table, a default route table is used.
 - **Subnet Access:** Select whether or not instances in this subnet are allowed to obtain a public IP address.
 - **DNS Label (Optional):** This option is available only if you provided a DNS label for the VCN during creation. Indicate whether you want to use DNS host names in the subnet. If so, enter a DNS label that is unique across the entire system.
 - **DHCP Options (Optional):** Select the set of DHCP options to associate with the subnet. If you enabled compartment selection, specify the compartment that contains the set of DHCP options. If you do not set these values, a default set of values is used.
 - **Security Lists (Optional):** Select one or more security lists to associate with the subnet. If you do not select a security list, a default security list is used.
5. Optionally, add one or more tags to this subnet resource.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

6. Click Create Subnet. The new subnet is displayed in the list of Subnets for the VCN where you created it.

Using the CLI

1. Gather the information required for these resources:
 - Compartment OCID (`oci iam compartment list --all`) From the display of all compartments, select the OCID of the compartment you are creating the subnet in.
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Enter the `subnet create` command using at least the compartment ID, VCN ID and CIDR block options.

If you want to use DNS host names in the subnet, include the DNS label in the create command. It cannot be added later. This option is available only if you provided a DNS label for the VCN during creation.

In this example we also set a user-friendly display name for the subnet. No set of DHCP options is specified, so the subnet will use the VCN default set.

```
# oci network subnet create --compartment-id "<target-compartment-ocid>" \
--vcn-id "<parent-vcn-ocid>" --cidr-block "10.1.1.0/24" --dns-label "sn1" --display-name "mysn1"
{
  "data": {
    "availability-domain": "PCANETWORK",
    "cidr-block": "10.1.1.0/24",
    "compartment-id": "ocidl.tenancy.oc1.pca.....uniqueID",
    "defined-tags": null,
    "dhcp-options-id": "ocidl.dhcpoptions.oc1.pca.....uniqueID",
    "display-name": "mysn1",
    "dns-label": "sn1",
    "freeform-tags": null,
    "id": "ocidl.subnet.oc1.pca.....uniqueID",
    "ipv6-cidr-block": null,
    "ipv6-public-cidr-block": null,
    "ipv6-virtual-router-ip": null,
    "lifecycle-state": "AVAILABLE",
    "prohibit-public-ip-on-vnic": false,
    "route-table-id": "ocidl.routetable.oc1.pca.....uniqueID",
    "security-list-ids": [
      "ocidl.security_list.oc1.pca.....uniqueID"
    ],
    "subnet-domain-name": "sn1.testvcn1.example.com",
    "time-created": "2021-01-11T17:51:54+00:00",
    "vcn-id": "ocidl.vcn.oc1.pca.....uniqueID",
    "virtual-router-ip": "10.1.1.1",
    "virtual-router-mac": "00:13:97:0e:8f:ff"
  },
  "etag": "1"
}
```

4.1.3 Editing a Subnet

Once a VCN subnet has been established, properties such as the name of the subnet, the route tables and security lists used by the subnet, and DHCP options can be changed.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN that contains the subnet you want to edit. The VCN Detail page is displayed.

3. In the Subnets list, locate the item to edit. In the Actions menu, click Edit to open the Edit Subnet window.
4. Make the necessary changes to the subnet. The following properties can be edited:
 - **Name:** Change the name of the subnet.
 - **Route Table:** Select a different route table to associate with this subnet. If you enabled compartment selection, specify the compartment that contains the route table.
 - **DHCP Options:** Select a different set of DHCP options to associate with the subnet. If you enabled compartment selection, specify the compartment that contains the set of DHCP options.
 - **Security Lists:** Select a different security list to associate with this subnet. If you enabled compartment selection, specify the compartment that contains the route table.
5. Optionally, add or delete tags for this subnet resource.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

6. Click Save Changes. The subnet properties are updated.

Using the CLI

1. Gather the information you need.
 - Compartment OCID (`oci iam compartment list --all`)
 - Subnet OCID (`oci network subnet list --compartment-id <compartment_OCID>`)
2. Enter the `subnet update` command using the subnet ID and the parameters you want to change.

In this example we change the DHCP options and route table for the subnet.

```
# oci network subnet update --subnet-id "<subnet-ocid>" \
--dhcp-options-id "<dhcp-options-ocid>" \
--route-table-id "<route-table-ocid>"
{
  "data": {
    "availability-domain": "PCANETWORK",
    "cidr-block": "10.1.1.0/24",
    "compartment-id": "ocidl.tenancy.oc1.pca.....uniqueID",
    "defined-tags": null,
    "dhcp-options-id": "ocidl.dhcpoptions.oc1.pca.....uniqueID",
    "display-name": "mysn1",
    "dns-label": "sn1",
    "freeform-tags": null,
    "id": "ocidl.subnet.oc1.pca.....uniqueID",
    "ipv6-cidr-block": null,
    "ipv6-public-cidr-block": null,
    "ipv6-virtual-router-ip": null,
    "lifecycle-state": "AVAILABLE",
    "prohibit-public-ip-on-vnic": false,
    "route-table-id": "ocidl.routetable.oc1.pca.....uniqueID",
    "security-list-ids": [
      "ocidl.security_list.oc1.pca.....uniqueID"
    ],
    "subnet-domain-name": "sn1.testvcn1.oraclevcn.com",
    "time-created": "2021-01-11T17:51:54+00:00",
    "vcn-id": "ocidl.vcn.oc1.pca.....uniqueID",
```

```
"virtual-router-ip": "10.1.1.1",
"virtual-router-mac": "00:13:97:0e:8f:ff"
},
"etag": "4"
}
```

4.1.4 Deleting a Subnet

A subnet can only be deleted if it is empty. Before deleting a subnet, make sure that all compute instances and other resources have been removed.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN that contains the subnet you want to delete. The VCN Detail page is displayed.
3. In the Subnets list, locate the item to delete. In the Actions menu, click Delete. Confirm the operation when prompted.

Using the CLI

1. Copy the OCID of the subnet you want to delete.
2. Enter the `subnet delete` command using the subnet ID.

```
# oci network subnet delete --subnet-id <subnet-ocid>
Are you sure you want to delete this resource? [y/N]: y
```

4.1.5 Terminating a VCN

To delete a VCN, it must first be empty and have no related resources or attached gateways (for example: no internet gateway, dynamic routing gateway, or so on). To delete a VCN's subnets, they must first be empty.

If you terminate a VCN, there should be no errors when you delete the VCN.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Select the VCN you want to terminate from the dropdown list. The VCN Detail page is displayed, along with any allocate gateways in the Resources list.
3. Click Terminate. Confirm the operation when prompted.

4.1.6 Deleting a VCN

You can delete a VCN if it is empty (terminated and detached). Make sure you are aware of all compute instances and other resources that are impacted before deleting a VCN.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you want to delete. The VCN Detail page is displayed.
3. Click Terminate. Confirm the operation when prompted.

Using the CLI

1. Gather the information for these resources:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2.

```
# oci network vcn delete --vcn-id <vcn-ocid>
Are you sure you want to delete this resource? [y/N]: y
```

4.2 Configuring VCN Rules and Options

VCNs and their subnets have various rules and options associated with them. The main categories are the use of DHCP, route tables, and security. These areas are so important that even if you do not configure them explicitly, the system uses a set of default values for each category.

The default values can always be changed or added to, but it is important to know what parameters are available for DHCP options, route tables, and security lists.

4.2.1 Working with DHCP Options

When you create a subnet, you specify which set of DHCP options to associate with the subnet. If you don't, the default set of DHCP options for the VCN is used. You can change DHCP options for a subnet at any time, but a subnet can only be assigned one set of DHCP options at a time. The set of DHCP options is a resource and has its own OCID. The assigned DHCP option set applies to all of the instances in that subnet.

When creating a new set of DHCP options, you may optionally assign it a friendly name. It doesn't have to be unique, and you can change it later.

To delete a set of DHCP options, it must not be associated with a subnet. If the set of DHCP options is associated with a subnet, you must detach the DHCP option set before deleting it. You can't delete a VCN's default set of DHCP options.

For more information, see the *DHCP Options* section in the [Virtual Networking Overview](#).

4.2.1.1 Viewing a VCN's Default Set of DHCP Options

Every VCN has a distinct set of default VCN options. If you have configured VCN A and VCN B, then each VCN has a default set of DHCP options associated with it: "Default DHCP Options for VCN A" and "Default DHCP Options for VCN B." Make sure you select the correct DHCP option set.

Note

It is a requirement that the target primary private IP cannot already be assigned a public IP.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click DHCP Options.

The default set and its details are displayed in the list.

Using the CLI

1. Gather the information you need to run the command:

- **Compartment OCID** (`oci iam compartment list --all`). This must be the OCID of the primary private IP address (`private-ip-id`) for the instance to which you want to assign this public IP address. The value of `is-primary` must be "true." The compartment OCID must be for the compartment that the private IP is in, which is not necessarily the compartment of the instance.

2. Run this command.

This example shows the default set of DHCP options, however, this command will list all of the DHCP option sets available in the compartment, if more are configured.

```
oci network dhcp-options list --compartment-id \
ocidl.tenancy.....uniqueID
{
  "data": [
    {
      "compartment-id": "ocidl.tenancy.....uniqueID",
      "defined-tags": {},
      "display-name": "Default DHCP Options for VCN-10-0-0-0-1r",
      "freeform-tags": {},
      "id": "ocidl.dhcpoptions.....uniqueID",
      "lifecycle-state": "AVAILABLE",
      "options": [
        {
          "custom-dns-servers": null,
          "server-type": "VcnLocalPlusInternet",
          "type": "DomainNameServer"
        },
        {
          "search-domain-names": [
            "vcn10.oraclevcn.com"
          ],
          "type": "SearchDomain"
        }
      ],
      "time-created": "2021-05-17T20:05:57+00:00",
      "vcn-id": "ocidl.vcn.....uniqueID"
    }
  ]
}
```

4.2.1.2 Updating an Existing Set of DHCP Options

Changes you make to DHCP options will take effect after your instance reboots. For more information, refer to <https://docs.oracle.com/en-us/iaas/Content/Network/Tasks/managingDHCP.htm#Importan>

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.

2. Click the VCN you are interested in.

3. Under Resources, click DHCP Options.

4. For the set you're interested in, click the Actions icon (three dots), and then click Edit:

- **DNS Type:** If you want instances in the subnet to resolve internet hostnames and hostnames of instances in the VCN, select Internet and VCN Resolver. Or to use a DNS server of your choice,

select Custom Resolver and then enter the server's IP address (three servers maximum). For more information, see the *Name Resolution* section in the [Virtual Networking Overview](#).

- **Search Domain:** If you want instances in the subnet to append a particular search domain when resolving DNS queries, enter it here. For more information, see the *DHCP Options* section in the [Virtual Networking Overview](#).
- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#).

5. Click Save Changes.
6. If you have any existing instances in a subnet that uses this set of DHCP options, make sure to restart the DHCP client on each affected instance, or reboot the instance itself so that it picks up the new setting.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - DHCP OCID (`oci network dhcp-options list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network dhcp-options update \
--display-name <new_name>
--dhcp-id <dhcp_OCID>
```

Example

```
oci network dhcp-options update --display-name sample-vcn \
--dhcp-id ocid1.dhcpoptions.....uniqueID
{
  "data": {
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "sample-vcn",
    "freeform-tags": {},
    "id": "ocid1.dhcpoptions.....uniqueID",
    "lifecycle-state": "AVAILABLE",
    "options": [
      {
        "custom-dns-servers": null,
        "server-type": "VcnLocalPlusInternet",
        "type": "DomainNameServer"
      }
    ],
    "time-created": "2021-05-21T14:36:30+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID"
  },
  "etag": "3"
}
```

4.2.1.3 Creating a New Set of DHCP Options

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click DHCP Options.
4. Click Create DHCP Options
5. Enter the following:
 - **Name:** A friendly name for the set of options. It doesn't have to be unique, and you can change it later.
 - **Create in Compartment:** The compartment where you want to create the set of DHCP options.
 - **DNS Type:** If you want instances in the subnet to resolve internet hostnames and hostnames of instances in the VCN, select Internet and VCN Resolver. Or to use a DNS server of your choice, select Custom Resolver and then enter the server's IP address (three servers maximum). For more information, see the *Name Resolution* section in the [Virtual Networking Overview](#).
 - **Search Domain:** If you want instances in the subnet to append a particular search domain when resolving DNS queries, enter it here. Be aware that the Networking service automatically sets the Search Domain option in certain situations. For more information, see the *DHCP Options* section in the [Virtual Networking Overview](#).
 - **Tagging:** For more information about tagging, see [Section 3.4.1, "Adding Tags at Resource Creation"](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.
6. Click Create DHCP Options.

The set of options is created and then displayed on the DHCP Options page of the compartment you chose. You can now specify this set of options when creating or updating a subnet.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run this command. This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network dhcp-options create --compartment-id <compartment-id-ocid>
--vcn-id <vcn-ocid>
```

Example using the minimum required parameters:

```
oci network dhcp-options create \
--compartment-id ocid1.tenancy.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID \
```



```
--options '{"type": "DomainNameServer", "customDnsServers": "202.44.61.9", \
  "serverType": "CustomDnsServer"}'
{
  "data": {
    "compartment-id": "ocidl.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "dhcp_options20210521161734",
    "freeform-tags": {},
    "id": "ocidl.dhcpoptions.....uniqueID",
    "lifecycle-state": "PROVISIONING",
    "options": null,
    "time-created": "2021-05-21T16:17:34+00:00",
    "vcn-id": "ocidl.vcn.....uniqueID"
  },
  "etag": "1"
}
```

4.2.1.4 Changing DHCP Options for a Subnet

Updates the specified set of DHCP options. You can update the display name or other options. You can also update the freeform or other DHCP options tags.

The changes take effect within a few seconds. Note that the entire new set of DHCP options takes effect.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Click Subnets.
4. Click the Subnet you are interested in.
5. Click Edit.
6. In the DHCP Options section, select the new set of DHCP options you want the subnet to use.
7. Click Save Changes.

Using the CLI

1. Gather this information to run the `oci network dhcp-options update` command:
 - Compartment OCID (`oci iam compartment list --all`)
 - DHCP OCID (`oci network dhcp-options list --compartment-id <compartment_OCID>`)
2. Run this command.

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

```
oci network dhcp-options update --dhcp-id ocidl.dhcpoptions.....uniqueID
{
  "data": {
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "test_DHCP_options",
    "freeform-tags": "",
    "id": "ocidl.dhcpoptions.defaultrealm.....uniqueID",
    "lifecycle-state": "PROVISIONING",
```

```

"options": [
  {
    "custom-dns-servers": "",
    "server-type": "VcnLocalPlusInternet",
    "type": "DomainNameServer"
  }
],
"time-created": "2021-10-08T20:50:11.414117+00:00",
"vcn-id": "ocidl.vcn.defaultrealm.....uniqueID"
},
"etag": "83139c07-160c-4948-9bc1-866230c3ca83"
}

```

4.2.1.5 Deleting a Set of DHCP Options

To delete a set of DHCP options, it must not be associated with a subnet yet. You can't delete the default set of DHCP options in a VCN.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click DHCP Options.
4. For the set you want to delete, click the Actions icon (three dots), and then click Delete.
5. Confirm when prompted.

Using the CLI

1. Gather this information to run the `oci network dhcp-options delete` command:

- Compartment OCID (`oci iam compartment list --all`)
- DHCP OCID (`oci network dhcp-options list --compartment-id <compartment_OCID>`)

2. Run this command.

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

```

oci network dhcp-options delete --dhcp-id ocidl.vcn.....uniqueID

Are you sure you want to delete this resource? [y/N]: y

```

4.2.2 Working with Route Tables

When you create a subnet, you specify a route table to associate with the subnet. If you don't, a default route table is used, which was created when you created the VCN. You can change route table entries for a subnet at any time, but a subnet can only be assigned one route table at a time. The assigned route table applies to all of the instances in that subnet.

When creating a new route table, you may optionally assign it a friendly name. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the route table a unique identifier called an Oracle Cloud ID (OCID).

To delete a route table, it must not be associated with a subnet yet. You can't delete a VCN's default route table.

For more information, see the *Route Tables* section in the [Virtual Networking Overview](#).

4.2.2.1 Viewing a VCN's Default Route Table

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Route Tables.

The default route table is displayed in the list of tables.

4. Click the default route table to view its details.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network route-table list \
--compartment-id <compartment-id-ocid> \
--vcn-id <vcn-ocid>
```

Example using the minimum required parameters:

```
oci network route-table list \
--compartment-id ocid1.tenancy.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID
{
  "data": {
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "Default Route Table for VCN-10-0-0-0-1r",
    "freeform-tags": {},
    "id": "ocid1.routetable.....uniqueID",
    "lifecycle-state": "AVAILABLE",
    "route-rules": [],
    "time-created": "2021-05-17T20:05:57+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID"
  },
}
```

4.2.2.2 Updating Rules in an Existing Route Table

You can add, edit, or delete rules in an existing route table using the UI. However, if you update the rules in a route table using the CLI, the update replaces all the current rules with the rules specified in the update. In other words, updating is more of a replace operation than an amend operation.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.

3. Under Resources, click Route Tables.
4. Click the route table you're interested in.
5. If you want to create a route rule, click Add Route Rule and enter the following:
 - **Target Type:** Select the target type from the list of allowed target types.
 - **Destination CIDR Block:** The value is the destination CIDR block for the traffic. A value of 0.0.0.0/0 means that all non-intra-VCN traffic that is not already covered by other rules in the route table goes to the target specified in this rule.
 - **Compartment:** The compartment where the target is located.
 - **Target or Gateway:** Select the name of the gateway you want to use from the list.
 - **Description:** An optional description of the rule.
6. If you want to delete an existing rule, click the Actions icon (three dots), and then click **Delete**.
7. If you want to edit an existing rule, click the Actions icon (three dots), and then click **Edit**.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Route Table OCID (`oci network route-table list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network route-table update --rt-id <vcn-ocid> --route-rules [complex type]
```

Complex data types are usually handled by using the `--generate-full-command-json-input` option, or, in this case, `oci network route-table update --generate-param-json-input route-rules`. This generates a sample json file to be used with this command option. The key names are pre-populated and match the command option names (converted to camelCase format, e.g. compartment-id becomes compartmentId).

The values of the keys are edited by the user before the sample file can be used as an input to this command.

For any command option that accepts multiple values, the value of the key can be a JSON array.

Options can still be provided on the command line. If an option exists in both the JSON document and the command line then the command line specified value will be used.

Example using the route rules input file, a new display name, and minimum required parameters:

```
oci network route-table update --rt-id ocid1.routetable....uniqueID \
--display-name new-route-table-for-vcn-a \
--route-rules file:///root/users-stuff/updated-route-rules.json

{
  "data": [
    {
```

```

"compartment-id": "ocidl.tenancy.....uniqueID",
"defined-tags": {},
"display-name": "new-route-table-for-vcn-a",
"freeform-tags": {},
"id": "ocidl.routetable.....uniqueID",
"lifecycle-state": "AVAILABLE",
"route-rules": [
  {
    "cidr-block": null,
    "description": "new route rules for VCN A",
    "destination": "0.0.0.0/0",
    "destination-type": "CIDR_BLOCK",
    "network-entity-id": "ocidl.internetgateway.....uniqueID"
  }
],
"time-created": "2021-11-08T19:21:37.429310+00:00",
"vcn-id": "ocidl.vcn.....uniqueID"
}
]
}

```

4.2.2.3 Create a Route Table

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Route Tables.
4. Click Create Route Table.
5. Enter the following:
 - **Name:** A friendly name for the route table. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
 - **Create in Compartment:** The compartment where you want to create the route table, if different from the compartment you're currently working in.
6. Optionally, click +Additional Route Rule to add one or more route rules, each with the following information (remember, a route table can exist with no rules until you're ready to add them):
 - **Target Type:** See the list of target types in [Overview of Routing for Your VCN](#).
 - **Destination CIDR Block:** The value is the destination CIDR block for the traffic. A value of 0.0.0.0/0 means that all non-intra-VCN traffic that is not already covered by other rules in the route table goes to the target specified in this rule.
 - **Compartment:** The compartment where the target is located.
 - **Target:** Select the target for the rule.
 - **Description:** An optional description of the rule.
7. Tagging: If you have permissions to create a resource, then you also have permissions to apply free-form tags to that resource. To apply a defined tag, you must have permissions to use the tag namespace. For more information about tagging, see [Resource Tags](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click Create Route Table.

The route table is created and then displayed on the Route Tables page in the compartment you chose. You can now specify this route table when creating or updating a subnet.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID (`oci iam compartment list --all`)
- VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci network route-table create \
--compartment-id <compartment-id-ocid> \
--route-rules [complex-type] \
--vcn-id <vcn-ocid>
```

Example using the minimum required parameters:

```
oci network route-table create --compartment-id \
ocidl.compartment....uniqueID \
--route-rules '[{"cidrBlock":"0.0.0.0/0","networkEntityId": \
"ocidl.internetgateway.ocl.....uniqueID"}]' \
--vcn-id ocidl.vcn....uniqueID

oci network route-table create
--compartment-id ocidl.compartment....uniqueID \
--route-rules file:///root/users-stuff/updated-route-rules.json
--vcn-id ocidl.vcn....uniqueID

{
  "data": [
    {
      "compartment-id": "ocidl.tenancy.....uniqueID",
      "defined-tags": {},
      "display-name": "Route Table for VCN A",
      "freeform-tags": {},
      "id": "ocidl.routetable.....uniqueID",
      "lifecycle-state": "AVAILABLE",
      "route-rules": [
        {
          "cidr-block": null,
          "description": "Route Rules for VCN A",
          "destination": "0.0.0.0/0",
          "destination-type": "CIDR_BLOCK",
          "network-entity-id": "ocidl.internetgateway.....uniqueID"
        }
      ],
      "time-created": "2021-11-08T19:21:37.429310+00:00",
      "vcn-id": "ocidl.vcn.....uniqueID"
    }
  ]
}
```

4.2.2.4 Delete a Route Table

Prerequisite: To delete a route table, it must not be associated with a subnet yet. You can't delete the default route table in a VCN.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Route Tables.
4. Click the route table you are interested in.
5. Click Delete.
6. Confirm when prompted.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Route table OCID (`oci network route-table list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network route-table delete \  
--rt-id <rt-ocid>
```

Example using the minimum required parameters:

```
oci network route-table delete --rt-id \  
ocid1.routetable....uniqueID  
Are you sure you want to delete this resource? [y/N]: y
```

To suppress this prompt, use the `--force` option.

4.2.2.5 Manage Tags for a Route Table

Using the Compute Web UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Route Tables.
4. Click the route table you are interested in.
5. Click the Tags tab to view or edit the existing tags. Or click Apply tag(s) to add new ones.

For more information, see [Resource Tags](#).

4.2.3 Controlling Traffic with Security Lists

You use a security list to define the rules that apply to all inbound (ingress) and outbound (egress) traffic of a subnet. You can associate up to five security lists per subnet. In the same way as route tables, there are default security lists and dedicated security lists. For better control and management, you should always use dedicated security lists for each subnet.

4.2.3.1 Viewing a VCN's Security Lists

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Security Lists.

The security lists are displayed in the list of tables.

4. Click the security list to view its detailed ingress and egress rules.

Using the CLI

1. Gather the information you need to run the command for the VCN you are interested in:
 - Compartment OCID (`oci iam compartment list <options>`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network security-list list \
--compartment-id <compartment-id-ocid> \
--vcn-id <vcn-ocid>
```

Example using the minimum required parameters:

```
oci network security-list list \
--compartment-id ocid1.tenancy.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID

{
  "data": [
    {
      "compartment-id": "ocid1.tenancy.....uniqueID",
      "defined-tags": {},
      "display-name": "Default Security List for VCN-10.25",
      "egress-security-rules": [
        {
          "description": null,
          "destination": "0.0.0.0/0",
          "destination-type": "CIDR_BLOCK",
          "icmp-options": null,
          "is-stateless": null,
          "protocol": "all",
          "tcp-options": null,
          "udp-options": null
        }
      ],
      "freeform-tags": {},
      "id": "ocid1.security_list.....uniqueID",
      "ingress-security-rules": [
        {
          "description": null,
          "icmp-options": null,
          "is-stateless": true,
          "protocol": "all",
          "source": "0.0.0.0/0",
          "source-type": "CIDR_BLOCK",
```



```

        "tcp-options": null,
        "udp-options": null
    }
},
"lifecycle-state": "AVAILABLE",
"time-created": "2021-08-18T18:52:02.378108+00:00",
"vcn-id": "ocidl.vcn.....uniqueID"
}
}

```

4.2.3.2 Creating a New Security List

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Security Lists.
4. Click Create Security List.
5. Enter the following:
 - **Name:** A friendly name for the security list. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
 - **Create in Compartment:** The compartment where you want to create the route table, if different from the compartment you're currently working in.
6. Optionally, in the Allow Rules for Ingress box, click +New Rule to add one or more new ingress rules. Each rule needs the following information (remember, a default security list exists a basic set of rules that can be added to):
 - **Stateless:** If you want the new rule to be stateless, check this box. By default, security list rules are stateful and apply to both a request and its coordinated response. For more information on stateless and stateful rules, see the *Security Lists* section of the Virtual Networking Overview in the [Oracle Private Cloud Appliance Concepts Guide](#).
 - **Ingress Type:** Set the type of security list rule from a drop down list of types allowed. The value is the destination CIDR block for the traffic.
 - **Ingress CIDR:** If you select CIDR as the Ingress Type, then enter the CIDR block covered by the rule.
 - **IP Protocol:** The rule can apply to all IP protocols, or choices such as ICMP, TCP, or UDP. Select the protocol covered from the drop-down list.
 - **Port Range:** For some protocols, such as TCP or UDP, you can supply a port range.
 - **Parameter Type and Code:** For ICMP, you can supply a particular Parameter Type and Parameter Code to which the rules applies.
 - **Description:** An optional description of the rule.
7. Optionally, in the Allow Rules for Egress box, click +New Rule to add one or more new egress rules. Each rule needs the following information (remember, a default security list exists a basic set of rules that can be added to):

- **Stateless:** If you want the new rule to be stateless, check this box. By default, security list rules are stateful and apply to both a request and its coordinated response. For more information on stateless and stateful rules, see the *Security Lists* section of the Virtual Networking Overview in the [Oracle Private Cloud Appliance Concepts Guide](#).
 - **Egress Type:** Set the type of security list rule from a drop down list of types allowed. The value is the destination CIDR block for the traffic.
 - **Egress CIDR:** If you select CIDR as the Egress Type, then enter the CIDR block covered by the rule.
 - **IP Protocol:** The rule can apply to all IP protocols, or choices such as ICMP, TCP, or UDP. Select the protocol covered from the drop-down list.
 - **Port Range:** For some protocols, such as TCP or UDP, you can supply a port range.
 - **Parameter Type and Code:** For ICMP, you can supply a particular Parameter Type and Parameter Code to which the rules applies.
 - **Description:** An optional description of the rule.
8. Tagging: For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.
9. Click Create Security List.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list <options>`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network security-list create \
--compartment-id <compartment-id-ocid> \
--vcn-id <vcn-ocid> \
--ingress-security-rules <complex-data-type> \
--egress-security-rules <complex-data-type>
```

Example using the minimum required parameters:

```
oci network security-list create --compartment-id ocid1.tenancy.....uniqueID
--vcn-id ocid1.vcn.....uniqueID
--egress-security-rules '[{"destination": "10.0.2.0/24", "protocol": "6", "isStateless":
true, "tcpOptions": {"destinationPortRange": {"max": 1521, "min": 1521},
"sourcePortRange": {"max": 1521, "min": 1521}}}]'
--ingress-security-rules '[{"source": "10.0.2.0/24", "protocol": "6", "isStateless":
true, "tcpOptions": {"destinationPortRange": {"max": 1521, "min": 1521},
"sourcePortRange": {"max": 1521, "min": 1521}}}]'
{
  "data": {
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "defined-tags": null
  },
}
```

```

"display-name": "securitylist20210916174938",
"egress-security-rules": [
  {
    "description": null,
    "destination": "10.0.2.0/24",
    "destination-type": "CIDR_BLOCK",
    "icmp-options": null,
    "is-stateless": true,
    "protocol": "6",
    "tcp-options": {
      "destination-port-range": {
        "max": 1521,
        "min": 1521
      },
      "source-port-range": {
        "max": 1521,
        "min": 1521
      }
    },
    "udp-options": null
  }
],
"freeform-tags": {},
"id": "ocidl.security_list.....uniqueID",
"ingress-security-rules": [
  {
    "description": null,
    "icmp-options": null,
    "is-stateless": true,
    "protocol": "6",
    "source": "10.0.2.0/24",
    "source-type": "CIDR_BLOCK",
    "tcp-options": {
      "destination-port-range": {
        "max": 1521,
        "min": 1521
      },
      "source-port-range": {
        "max": 1521,
        "min": 1521
      }
    },
    "udp-options": null
  }
],
"lifecycle-state": "PROVISIONING",
"time-created": "2021-09-16T17:49:38.243708+00:00",
"vcn-id": "ocidl.vcn.....uniqueID"
},
"etag": "508fb93a-c68b-4539-817b-bd7bdbe9ec34"
}

```

4.2.3.3 Updating Rules in an Existing Security List

You can add, edit, or delete rules in an existing security list, even in the default security list, using UI. However, in the CLI, if you update the rules in a security list, the update replaces all the current rules with the rules specified in the update. In other words, in the CLI, updating is more of a replace operation than an amend operation.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Security Lists.

The default security list is displayed in the list of tables.

4. Click the default security list to view its detailed ingress and egress rules.
5. To update an Ingress rule, click on the Actions Menu icon (three dots) next to the Ingress rule you want to update.
 - To delete the rule, click on Delete.
 - To update the rule with new information, click on Edit.
6. To update an Egress rule, click on the Actions Menu icon (three dots) next to the Egress rule you want to update.
 - To delete the rule, click on Delete.
 - To update the rule with new information, click on Edit.
7. When all changes have been made, click on Save Changes. (Cancel discards the changes.)

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list <options>`)
 - Security list OCID (`oci network security-list list --compartment-id <compartment-id-ocid>`)

2. Run this command.

Syntax (entered on a single line):

```
oci network security-list update \
--compartment-id <compartment-id-ocid> \
--security-list-id <security-list-id-ocid>
--display-name <text> (this parameter is optional, but you must update something)
```

Technically, the only required parameter to update a security list is the security list OCID. However, this does not change anything because all the other parameters are optional. When you change any parameters in an egress or ingress rule, you must confirm or deny the change unless the

```
--force
```

parameter is used to remove the confirmation prompt.

This example updates an egress rule to change the security list port ranges from 1521 to 1500. The change is not forced:

```
oci network security-list update --security-list-id ocid1.security_list.....uniqueID
--egress-security-rules '[{"destination": "10.0.2.0/24", "protocol": "6", "isStateless":
true, "tcpOptions": {"destinationPortRange": {"max": 1500, "min": 1500},
"sourcePortRange": {"max": 1500, "min": 1500}}]'
```

```
WARNING: Updates to defined-tags and egress-security-rules and freeform-tags and
ingress-security-rules will replace any existing values.
Are you sure you want to continue? [y/N]: y
```

```
{
  "data": {
```

```

"compartment-id": "ocidl.tenancy.....uniqueID",
"defined-tags": {
  "Finance": {
    "CostCenter": "Assigned by: admin"
  }
},
"display-name": "securitylist20210916174938",
"egress-security-rules": [
  {
    "description": null,
    "destination": "10.0.2.0/24",
    "destination-type": "CIDR_BLOCK",
    "icmp-options": null,
    "is-stateless": true,
    "protocol": "6",
    "tcp-options": {
      "destination-port-range": {
        "max": 1500,
        "min": 1500
      },
      "source-port-range": {
        "max": 1500,
        "min": 1500
      }
    },
    "udp-options": null
  }
],
"freeform-tags": {},
"id": "ocidl.security_list.....uniqueID",
"ingress-security-rules": [
  {
    "description": null,
    "icmp-options": null,
    "is-stateless": true,
    "protocol": "6",
    "source": "10.0.2.0/24",
    "source-type": "CIDR_BLOCK",
    "tcp-options": {
      "destination-port-range": {
        "max": 1521,
        "min": 1521
      },
      "source-port-range": {
        "max": 1521,
        "min": 1521
      }
    },
    "udp-options": null
  }
],
"lifecycle-state": "AVAILABLE",
"time-created": "2021-09-16T17:49:38.243708+00:00",
"vcn-id": "ocidl.vcn.....uniqueID"
},
"etag": "1b3875c4-0bb8-4af7-9584-a5209c8d23ac"
}

```

4.2.3.4 Delete a Security List

There are some prerequisites. To delete a security list, it must not be associated with a subnet. You can't delete the default security list in a VCN.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.

2. Click the VCN you are interested in.
3. Under Resources, click Security Lists.
4. Click the security list you are interested in.
5. Click on the Actions menu icon (three dots) and select Delete.
6. Confirm the deletion when prompted.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list <options>`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
 - Security list OCID (`oci network route-table list --compartment-id <compartment_OCID> --vcn-id <vcn_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network security-list delete \
--security-list-id <security-list-ocid>
```

Example using the minimum required parameters:

```
oci network security-list delete --security-list-id \
ocid1.securitylist....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

4.2.3.5 Manage Tags for a Security List

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Security Lists.
4. Click the security list you are interested in.
5. Click the Tags tab to view or edit the existing tags. Or click Apply tag(s) to add new ones.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

4.2.4 Controlling Traffic with Network Security Groups

Security lists provide virtual firewall rules to all the VNICs in a subnet. In order to provide a set of firewall rules for a specific set of VNICs in a VCN (not a particular subnet) you can create a network security group (NSG). You can also apply an NSG to a Mount Target.

NSGs let you build rules for groups of instances, even if the instances are in different subnets. For example, the same NSG can apply to all the database servers, or all the application servers running a

certain application. Instead of applying security to a particular subnet, you create an NSG and then add the appropriate instances to the NSG.

There is no requirement to use either security lists or NSGs. You can use security lists without establishing NSGs, or create NSGs without creating any security lists. However, if you use both security lists and NSGs, the rules that apply to a VNIC are the union (both sets) of the rules that are in the security list for the VNIC and the rules specific to that VNIC from the NSG.

Both security lists and NSGs use security rules to implement their functions. However, unlike security lists, an NSG does not contain any default rules when initially created. You must formulate and configure the rules for the NSG after its creation and before it is useful.

NSG creation involves two specific configuration areas: VNICs and security rules. The VNICs referenced by the NSG must be all in the same VCN. NSG security rules determine which types of traffic are allowed in and out of the VNICs in the NSG.

Before deleting an NSG, you must remove all of the VNICs from the NSG.

4.2.4.1 Creating a VCN's Network Security Groups

Before an NSG can be used, it must be created.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Network Security Groups.
4. Click Create Network Security Group.
5. Enter the following:
 - **Name:** A friendly name for the network security group. The name doesn't have to be unique, and it cannot be changed later in the Console (but you can change it with the API). Avoid entering confidential information.
 - **Create in Compartment:** The compartment where you want to create the NSG, if different from the compartment you're currently working in.
6. Tagging: For more information about tagging, see [Section 3.4.1, "Adding Tags at Resource Creation"](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.
7. Click Create Network Security Group.

The NSG is created and you are positioned in the empty NSG to create security rules and apply the NSG to VNICs. The NSG rules are created with the same procedure as for security lists. That is, you configure egress and ingress rules that prevent or allow traffic into and out of the VNICs.

You do not have to create NSG rules at this time. If you want to create only an empty NSG, then you are done.

8. If you want to create security rules for the NSG at this time, click Create Security Rules.
9. In the Allow Rules for Ingress box, click +New Rule to add one or more new ingress rules. Each rule needs the following information:

- **Stateless:** If you want the new rule to be stateless, check this box. By default, security list rules are stateful and apply to both a request and its coordinated response.
 - **Ingress Type:** Set the type of security list rule from a drop down list of types allowed. The value is the destination CIDR block for the traffic.
 - **Ingress CIDR:** If you select CIDR as the Ingress Type, then enter the CIDR block covered by the rule.
 - **IP Protocol:** The rule can apply to all IP protocols, or choices such as ICMP, TCP, or UDP. Select the protocol covered from the drop-down list.
 - **Port Range:** For some protocols, such as TCP or UDP, you can supply a port range.
 - **Parameter Type and Code:** For ICMP, you can supply a particular Parameter Type and Parameter Code to which the rules applies.
 - **Description:** An optional description of the rule.
10. In the Allow Rules for Egress box, click +New Rule to add one or more new egress rules. Each rule needs the following information:
- **Stateless:** If you want the new rule to be stateless, check this box. By default, security list rules are stateful and apply to both a request and its coordinated response.
 - **Egress Type:** Set the type of security list rule from a drop down list of types allowed. The value is the destination CIDR block for the traffic, or a service.
 - **Egress CIDR:** If you select CIDR as the Egress Type, then enter the CIDR block covered by the rule.
 - **IP Protocol:** The rule can apply to all IP protocols, or choices such as ICMP, TCP, or UDP. Select the protocol covered from the drop-down list.
 - **Port Range:** For some protocols, such as TCP or UDP, you can supply a port range.
 - **Parameter Type and Code:** For ICMP, you can supply a particular Parameter Type and Parameter Code to which the rules applies.
 - **Description:** An optional description of the rule.
11. The NSG is now ready to attach to VNICs.

Using the CLI

Creating an NSG with useful rules is a two-step process in the CLI. First, you create the NSG, then you add security rules.

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):


```
oci network nsg create \
--compartment-id <compartment_OCID> \
--vcn-id <vcn_OCID>
```

Example using the minimum required parameters:

```
oci network nsg create
--compartment-id ocid1.compartment.....uniqueID
--vcn-id ocid1.vcn.....uniqueID
{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "securitygroup20210921221126",
    "freeform-tags": {},
    "id": "ocid1.networksecuritygroup.....uniqueID",
    "lifecycle-state": "PROVISIONING",
    "time-created": "2021-09-21T22:11:26.538684+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID"
  },
  "etag": "ab7ecc36-06a0-46e3-ac97-c794cbe8c8d8"
}
```

4.2.4.2 Attach a Network Security Group to a VNIC

In order for an NSG to be functional, you must link the NSG to an attached VNIC in a VCN. You do not specify which VNICs or instances are associated with the NSG at the NSG level. The instance must be running for the attachment to be successful. Similarly, to remove a VNIC from an NSG, you perform that action by updating the instance, not the NSG. A related topic is [Section 4.4.3.5, “Add or Remove a VNIC from a Network Security Group”](#).

Using the UI

1. Open the Navigation Menu. Under Compute, click View Instances.
2. In the list, click the instance you wish to associate with the NSG.
3. Under Resources, click Attached VNICs.

The list of VNICs attached to this instance is displayed.

4. From the list of Attached VNICs, click the name of the attached VNIC that you want to associate with the NSG.
5. Under Resources for the attached VNICs, click on Network Security Groups.
6. Click on Update Network Security Groups.
7. Check the box next to Enable Network Security Groups.
8. Select the NSG you want to attach to the VNIC from the dropdown list. You can add multiple NSGs all at once.
9. Click Update Network Security Groups for VNIC. If you have reached the maximum number of VNIC attachments for the instance, you get an error message and the operation cannot be completed. The instance must be running for the attachment to be successful.
10. The VNIC attachment is created with a generated name and appears in the list of Network Security Groups for the instance. You can examine the NSG rules by selecting the NSG for the VCN and then, under Resources, inspect the security rules and VNICs for the NSG.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Network Security Group OCID (`oci network nsg list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network vnic update \
--vnic-id <vnic-ocid> \
--nsg-id <nsg-ocid>
```

Example using the minimum required parameters:

```
oci network vnic update \
--vnic-id ocid1.vnic.....uniqueID \
--nsg-ids ["ocid1.networksecuritygroup.....uniqueID"]
,
WARNING: Updates to defined-tags and freeform-tags and nsg-ids will replace any existing values.
Are you sure you want to continue? [y/N]: y
{
  "data": {
    "availability-domain": "ad1",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "example-instance",
    "freeform-tags": {},
    "hostname-label": "example-instance",
    "id": "ocid1.vnic.1742XC3024.broom14.....uniqueID",
    "is-primary": false,
    "lifecycle-state": "PROVISIONING",
    "mac-address": "00:13:97:5c:7a:58",
    "nsg-ids": [
      "ocid1.networksecuritygroup.1.....uniqueID"
    ],
    "private-ip": "10.27.100.2",
    "public-ip": null,
    "skip-source-dest-check": true,
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-12-01T17:14:46.084875+00:00",
    "vlan-id": null
  },
  "etag": "17e8e038-3051-4ef8-96a4-21d7107d37e6"
}
```

4.2.4.3 Add or Remove Resources from a Network Security Group

Network Security Groups (NSGs) are applied to instances created in a VCN. The VNICs of an instance are a form of resource from the NSG perspective. You can add or remove a resource like VNICs from an NSG.

4.2.4.4 Manage Rules for a Network Security Group

You can change the rules that are applied by a Network Security Group (NSG).

Using the UI

1. Open the Navigation Menu. Under Networking, click View Virtual Cloud Networks.

2. In the list, click the VCN you wish to manage rules in the NSG.
3. Under Resources, click Network Security Groups.

The list of NSGs for this VCN is displayed. If there are none, you need to create the NSG first.

4. From the list of NSGs, click the name of the NSG for which you want to manage rules. You can add new rules (see [Section 4.2.4.1, "Creating a VCN's Network Security Groups"](#)) or manage existing rules created earlier.
5. In the Actions Menu (three dots) for the Egress or Ingress rule, click Edit.
6. Change the fields you wish to modify, then click Update.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Network Security Group OCID (`oci network nsg list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network nsg rules update \
--nsg-id <nsg-ocid> \
--security-rules [complex type]
```

Complex data types are usually handled by using the `--generate-full-command-json-input` option, or, in this case, `oci network nsg rules update security-rules --generate-param-json-input route-rules`. This generates a sample json file to be used with this command option. The key names are pre-populated and match the command option names (converted to camelCase format, e.g. compartment-id becomes compartmentId).

The values of the keys are edited by the user before the sample file can be used as an input to this command.

For any command option that accepts multiple values, the value of the key can be a JSON array.

Options can still be provided on the command line. If an option exists in both the JSON document and the command line then the command line specified value will be used.

Example using the minimum required parameters:

```
oci network nsg rules update
--nsg-id ocidl.networksecuritygroup.....uniqueID
--security-rules '[{"description": "TEST-RULE","destination": "0.0.0.0/0",
  "destinationType": "CIDR", "direction": "INGRESS","isStateless": true,
  "protocol": "6","source": "0.0.0.0/0","sourceType": "CIDR_BLOCK",
  "tcpOptions": {"destinationPortRange": {"max": 100,"min": 10},
  "sourcePortRange": {"max": 500,"min": 50}}]'
```

```
{
  "data": {
    "security-rules": [
      {
        "description": "TEST-RULE",
```

```

"destination": null,
"destination-type": null,
"direction": "INGRESS",
"icmp-options": null,
"id": "ocid1.security_rule.....uniqueID",
"is-stateless": true,
"is-valid": true,
"protocol": "6",
"source": "0.0.0.0/0",
"source-type": "CIDR_BLOCK",
"tcp-options": {
  "destination-port-range": {
    "max": 100,
    "min": 10
  },
  "source-port-range": {
    "max": 500,
    "min": 50
  }
},
"time-created": null,
"udp-options": null
}
]
},
"etag": "90000"
}

```

4.2.4.5 Viewing a VCN's Network Security Groups

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Network Security Groups. The number in parentheses gives the number of NSGs configured for this VCN. If the number is zero (0), there are no NSGs to view.
4. Click on the Network Security Group you are interested in.
5. If there are security rules for the NSG, you can view them. If there are no rules, you can create them.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Network Security Group OCID (`oci network nsg list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network nsg rules list \
--nsg-id <nsg-ocid>
```

Example using the minimum required parameters:

```
oci network nsg rules list
--nsg-id ocid1.networksecuritygroup.....uniqueID
```

```

{
  "data": [
    {
      "description": "test-rule-1",
      "destination": null,
      "destination-type": null,
      "direction": "INGRESS",
      "icmp-options": null,
      "id": "ocidl.security_rule.....uniqueID",
      "is-stateless": false,
      "is-valid": true,
      "protocol": "6",
      "source": "0.0.0.0/0",
      "source-type": "CIDR_BLOCK",
      "tcp-options": {
        "destination-port-range": {
          "max": 3000,
          "min": 3000
        },
        "source-port-range": {
          "max": 3000,
          "min": 3000
        }
      },
      "time-created": "2021-10-21T21:09:39.941613+00:00",
      "udp-options": null
    },
    {
      "description": "test-rule-2",
      "destination": null,
      "destination-type": null,
      "direction": "INGRESS",
      "icmp-options": null,
      "id": "ocidl.security_rule.....uniqueID",
      "is-stateless": true,
      "is-valid": true,
      "protocol": "6",
      "source": "0.0.0.0/0",
      "source-type": "CIDR_BLOCK",
      "tcp-options": {
        "destination-port-range": {
          "max": 100,
          "min": 10
        },
        "source-port-range": {
          "max": 500,
          "min": 50
        }
      },
      "time-created": "2021-10-21T23:18:17.145627+00:00",
      "udp-options": null
    }
  ]
}

```

4.2.4.6 Manage Tags for a Network Security Group

You can add tags to an NSG during its creation, or afterwards.

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.

3. Under Resources, click Network Security Groups.
4. From the dropdown list of NSGs, click the NSG to add tags to.
5. Click the Actions icon (three dots), and then click View Tags. From there you can view the existing tags, edit them, and apply new ones.
6. Alternatively, click Edit Network Security Group. From there you can view the existing tags, edit them, and apply new ones.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

4.2.4.7 Deleting a VCN's Network Security Groups

Using the UI

1. Open the Navigation Menu. Under Networking, click Virtual Cloud Networks.
2. Click the VCN you are interested in.
3. Under Resources, click Network Security Groups. The number in parentheses gives the number of NSGs configured for this VCN. If the number is zero (0), there are no NSGs to delete.
4. Click on the Actions icon (three dots) associated with the NSG and click Delete.
5. Confirm when prompted.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - Network Security Group OCID (`oci network nsg list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci network nsg delete \
--nsg-id <nsg-ocid>
```

Example using the minimum required parameters:

```
oci network nsg delete
--nsg-id ocid1.networksecuritygroup....uniqueID \
Are you sure you want to delete this resource? [y/N]: y
```

To suppress this prompt, use the `--force` option. If the NSG is not detached from the VNIC, you get a message that says the NSG cannot be deleted because it is in use.

4.3 Configuring VCN Gateways

Virtual processes communicate with other processes in a variety of ways. If two VMs are in the same subnet, meaning the network portions of their IP addresses match, there is no special configuration needed to allow them to communicate. A logical switch connects source and destination at the MAC address

level. Also, communication between VMs in the same VCN but different subnets requires no routing configuration. Routing is only needed for traffic that is going to a destination or coming from a source external to a VCN.

When communication between two virtual processes is needed and the source and destination are in two different VCNs, then configuration of one of five different types of gateway is necessary in the source VCN. In this context, a gateway is a special type of router, connecting two different IP networks by following rules set up in a route table. (A router can be thought of as a multiport gateway, and a gateway can be thought of as a two-port router.)

When you first create a VCN, various resources are listed in the UI and available for listing with a CLI command. Some of the resources are listed automatically when you create a subnet, and others must be configured explicitly.

- **Subnets.** This resource gives the number of subnets created under the VCN. All other resources also display counts for the VCN.
- **Route Tables.** This resource gives the number of route tables. Subnets can share route tables, especially default route tables, so this count is not necessarily the same as the count of subnets, especially if there is more than one subnet for the VCN.
- **Internet Gateways.** This resource gives the number of internet gateways configured. Initially, there are none.
- **Local Peering Gateway.** This resource gives the number of local peering gateways configured. Initially, there are none.
- **DHCP Options.** This resource gives the number of DHCP option lists. There is at least one for the VCN by default, but more can be created.
- **Security Lists.** This resource gives the number of Security Lists. There is at least one set of ingress and egress rules for the VCN by default, but more can be created.
- **NAT Gateways.** This resource gives the number of NAT gateways configured. Initially, there are none.
- **Network Security Groups.** This resource gives the number of Network Security Groups configured. Initially, there are none, but you can gather existing Security Lists into Network Security Groups, where all security rules are applied at once, as needed.
- **Service Gateways.** This resource gives the number of service gateways configured. Initially, there are none.
- **Dynamic Routing Gateways.** This resource gives the number of dynamic routing gateways (DRGs) configured. Initially, there are none. Note that these gateways are not configured without the VCN, but attached to the VCN.
- **Dynamic Routing Gateway Attachments.** This resource gives the number of dynamic routing gateway attachments that have been configured. You must have a DRG configured to have attachments listed.

The various types of gateways are configured for very specific reasons.

- **NAT Gateway.** A NAT gateway is used to translate IP addresses as traffic passes from one part of an IP network to another. When used between a VCN and the on-premises data center network, the NAT address becomes the source address for traffic sent on to the data center network. A NAT gateway allows egress to the on-premises network from a VCN. It does not allow connections to be initiated to the VMs in the VCN. Although essentially one-way, return traffic is allowed for connections initiated in the VCN. Contrast NAT Gateway with the Internet Gateway, which allows connections into and out of

the VCN, the NAT Gateway allows VMs with public IP addresses to be reachable from outside the PCA network.

Note

A VCN connected to the on-premises network with a Dynamic Routing Gateway cannot overlap with any on-premises CIDR, or other VCN CIDRs connected with a Dynamic Routing Gateway. In other words, the IP addresses used must be exclusive to the VCN.

- **Internet Gateway (IGW).** An IGW provides the VCN with access to the global public internet, but only through the on-premises data center network. The source and destination must have routable, public IP addresses, and a VCN can have only one IGW.
- **Local Peering Gateway (LPG).** A Local Peering Gateway (LPG) is a way to connect VCNs so that elements in each VCN can communicate, even using private IP address. Peered VCNs can be in different tenancies.
- **Dynamic Routing Gateway (DRG).** A DRG is used to connect a VCN to the data center's IP address space. That is, outside the Oracle Private Cloud Appliance rack in the data center. The data center network can, if configured that way, pass Oracle Private Cloud Appliance traffic on to other destinations.
- **Service Gateway (SG).** Some services are isolated on their own network for security and performance reasons. The service gateway (SG) allows a VCN with no external access to privately access Service Network services (such as object storage) in a private subnet.

4.3.1 Enabling Public Connections through a NAT Gateway

A NAT gateway is used to translate IP addresses as traffic passes from one part of an IP network to another. This prevents sources and destinations from having identical IP addresses, and allows RFC 1918 private addresses used in Oracle Private Cloud Appliance traffic to communicate with on-premises data center networks. A NAT gateway is attached to a VCN at the subnet level, allowing finer control of the address translations. The NAT gateway is configured separately from the VCNs, and is not required to be in the same compartment as the VCN (but can be). However, the NAT gateway is within the VCN, and only one NAT per VCN is allowed. The NAT address becomes the source address for traffic sent on to the data center network.

Using the UI

1. In the navigation menu, under Networking, click Virtual Cloud Networks. A list of previously configured VCNs in compartments appears. If the compartment you are creating the NAT gateway in is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on the VCN that you are creating the NAT gateway in.
3. In the Resources menu for that VCN, click on NAT Gateways (the number of configured NAT gateways in parentheses does not matter).
4. Click on Create NAT Gateway
5. Fill in the required NAT gateway information:
 - **Name:** Provide a name or description for the NAT gateway. Avoid using any of the organization's confidential information.
 - **Create in Compartment:** Select the compartment in which to create the NAT Gateway.
 - **Block Traffic** Choose whether to block traffic to this NAT Gateway.

- **(Yes: Traffic Not Blocked):** By default, the VCN uses the NAT gateway even if it is not completely configured.
- **(No: Traffic Blocked):** You can set the NAT gateway not see traffic until it is explicitly enabled to do so.

For more information on NAT gateways, refer to the *NAT Gateways* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

- **Tagging:** Optionally, add one or more tags to this resource.
If you are not sure whether to apply tags, skip this option (you can apply tags later).
For more information about tagging resources, see [Tagging Overview](#).

6. Click Create NAT Gateway.

The NAT Gateway is now ready for the addition of route rules or security settings.

Using the CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run the `oci network nat-gateway create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network nat-gateway create \
--compartment-id <compartment_OCID> \
--vcn-id <vcn_OCID>
```

Example:

```
oci network nat-gateway create \
--compartment-id ocid1.compartment.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID

{
  "data": {
    "block-traffic": true,
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "natgateway20210827215953",
    "freeform-tags": {},
    "id": "ocid1.vcn.....uniqueID",
    "lifecycle-state": "PROVISIONING",
    "nat-ip": "10.133.80.3",
    "public-ip-id": "ocid1.publicip.AK00661530.scasg01.....uniqueID",
    "time-created": "2021-08-27T21:59:53.858329+00:00",
    "vcn-id": "ocid1.vcn.AK00661530.scasg01.....uniqueID"
  },
}
```

```
"etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}
```

The NAT Gateway is now ready for the addition of route rules or security settings. Note that the name of the gateway (natgateway20210827215953) is assigned automatically and not by a parameter, and that the IP address of the device (10.133.80.3) is also assigned automatically.

4.3.2 Providing Public Access through an Internet Gateway

Internet Gateway (IGW). An IGW provides the VCN with access to the global public internet, but only through the on-premises data center network. The IGW is configured within the VCN, so the IGW is automatically attached to the VCN in which it is configured. The source and destination must have routable, public IP addresses, and a VCN can have only one IGW. Any traffic using public IP addresses goes through the IGW. The IGW is not required to be in the same compartment as the VCN. A subnet's route table determines which public subnets can use the IGW, and the subnet security list defines the types of traffic that can use the IGW. Like a physical router, the IGW can be disabled through the CLI, severing internet access, no matter what permissions are established.

Using the Compute Web UI

1. In the navigation menu, under Networking, click Virtual Cloud Networks. A list of previously configured VCNs in compartments appears. If the compartment you are creating the internet gateway in is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on the VCN that you are creating the internet gateway in.
3. In the Resources menu for that VCN, click on Internet Gateways (the number of configured internet gateways in parentheses does not matter).
4. Click on Create Internet Gateway
5. Fill in the required internet gateway information:

- **Name:** Provide a name or description for the internet gateway. Avoid using any of the organization's confidential information.
- **Create in Compartment:** Select the compartment in which to create the Internet Gateway.

For more information on NAT gateways, refer to the *Internet Gateways* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

- **Tagging:** Optionally, add one or more tags to this resource.

For more information about tagging, see [Section 3.4.1, "Adding Tags at Resource Creation"](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

6. Click Create Internet Gateway.

The Internet Gateway is now ready for the addition of route rules or security settings.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)

- Run the `oci network internet-gateway create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network internet-gateway create
--compartment-id <compartment_OCID>
--is-enabled <boolean: true | false>
--vcn-id <vcn_OCID>
```

Example:

```
oci network internet-gateway create \
--compartment-id ocid1.compartment.....uniqueID
--is-enabled true
--vcn-id ocid1.vcn.....uniqueID

{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "internetgateway20210830165014",
    "freeform-tags": {},
    "id": "ocid1.internetgateway.AK00661530.scasg01.....uniqueID",
    "is-enabled": true,
    "lifecycle-state": "PROVISIONING",
    "time-created": "2021-08-30T16:50:14.634466+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID",
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}
```

The Internet Gateway is now ready for the addition of route rules or security settings. Note that the name of the gateway is assigned automatically (internetgateway20210830165014) and not by a parameter.

4.3.3 Connecting VCNs through a Local Peering Gateway

A Local Peering Gateway (LPG) is a way to connect VCNs so that elements in each VCN can communicate, even using private IP address. Peered VCNs cannot be in different tenancies. There are several other requirements for LPG configuration:

- The CIDRs for the VCNs linked by the LPG cannot overlap.
- Each peered VCN must have an LPG configured correctly, and the LPGs must be connected.
- VCN route rules must be properly configured to steer VCN subnet traffic to and from the LPGs.
- Security rules must be properly configured to allow or deny certain types VCN subnet traffic use the LPGs

Using the UI

- In the navigation menu, under Networking, click Virtual Cloud Networks. A list of previously configured VCNs in compartments appears. If the compartment you are creating the local peering gateway in is not in the title bar, then use the drop-down tab to select the correct compartment.

2. Click on the VCN that you are creating the local peering gateway in.
3. In the Resources menu for that VCN, click on Local Peering Gateways (the number of configured local peering gateways in parentheses does not matter).
4. Click on Create Local Peering Gateway
5. Fill in the required Local Peering gateway information:

- **Name:** Provide a name or description for the local peering gateway. Avoid using any of the organization's confidential information.
- **Create in Compartment:** Select the compartment in which to create the Local Peering Gateway.
- **Route Table Association (Optional)** Optionally, you can associate a route table with the Local Peering Gateway. A list of configured route tables for the selected compartment is in a drop-down list. You can change the compartment by clicking (change) next to the compartment name.

For more information on local peering gateways, refer to the *Local Peering Gateways* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

- **Tagging:** Optionally, add one or more tags to this resource.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

6. Click Create Local Peering Gateway.

The Local Peering Gateway is now ready for connecting VCNs with Establish Peering Connection, and the addition of route rules or security settings.

Using the CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
2. Run the `oci network local-peering-gateway create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network local-peering-gateway create \
--compartment-id <compartment_OCID> \
--vcn-id <vcn_OCID>
```

Example:

```
oci network local-peering-gateway create \
--compartment-id ocid1.compartment.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID
```

```

{
  "data": {
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "localpeeringgateway20210830174050",
    "freeform-tags": {},
    "id": "ocidl.lpg.AK00661530.scasg01.....uniqueID",
    "is-cross-tenancy-peering": false,
    "lifecycle-state": "AVAILABLE",
    "peer-advertised-cidr": null,
    "peer-advertised-cidr-details": null,
    "peering-status": "NEW",
    "peering-status-details": null,
    "route-table-id": null,
    "time-created": "2021-08-30T17:40:50.876023+00:00",
    "vcn-id": "ocidl.vcn.....uniqueID"
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}

```

4.3.4 Connecting to the On-Premises Network through a Dynamic Routing Gateway

Dynamic Routing Gateway (DRG). A DRG is the Oracle Private Cloud Appliance equivalent of a general purpose router. A DRG is used to connect a VCN to the data center's IP address space. The router is configured separately from the VCNs, at the compartment level and is not required to be in the same compartment as the VCN (but it typically is). Once configured, the DRG can be attached to more than one VCN and, like a physical router, can be attached and detached at any time, although perhaps with traffic loss. Also like a physical router, even when attached to a VCN, the DRG must have route table rules to steer traffic to the on-premises data center network's IP address space.

4.3.4.1 Create a Dynamic Routing Gateway

Using the Compute Web UI

1. In the navigation menu, under Networking, click Dynamic Routing Gateways (DRGs). A list of previously configured DRGs in compartments appears. If the compartment you are creating the dynamic routing gateway in is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on Create Dynamic Routing Gateway
3. Fill in the required dynamic routing gateway information:
 - **Name:** Provide a name or description for the dynamic routing gateway. Avoid using any of the organization's confidential information.
 - **Create in Compartment:** Select the compartment in which to create the dynamic routing Gateway.

For more information on dynamic routing gateways, refer to the *Dynamic Routing Gateways* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

 - **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).
4. Click Create Dynamic Routing Gateway.

The Dynamic Routing Gateway is now ready for the addition of DRG attachments.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
2. Run the `oci network drg create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network drg create
--compartment-id <compartment_OCID>
```

Example:

```
oci network drg create \
--compartment-id ocid1.compartment.....uniqueID
{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "drg20210830204524",
    "freeform-tags": {},
    "id": "ocid1.drg.....uniqueID",
    "lifecycle-state": "AVAILABLE",
    "time-created": "2021-08-30T20:45:24.236954+00:00"
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}
```

4.3.4.2 Attach VCNs to a Dynamic Routing Gateway

You can connect many VCNs to a DRG, but each VCN can have only one DRG attached. You must still make sure the route tables and security lists allow communication.

Using the UI

1. In the navigation menu, under Networking, click on Dynamic Routing Gateways. A list of previously configured DRGs in compartments appears. If the compartment you are attaching the dynamic routing gateway to is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on Dynamic Routing Gateway name in the list of DRGs for that compartment.
3. Click on Attach to Virtual Cloud Network.
4. Click on the VCN to attach the DRG to from the list of VCNs in the drop down list. If the correct compartment you are is not in the title bar, then use the drop-down tab to select the correct compartment.
5. Click on Attach to DRG.

- Repeat the process to attach the other VCNs to the DRG and connect the VCNs.

The Dynamic Routing Gateway is attached to the selected VCN.

You can connect up to 10 VCNs to a DRG, but each VCN can have only one DRG attached. You must still make sure the route tables and security lists allow communication.

Using the CLI

- Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID>`)
 - Dynamic Routing Gateway OCID (`oci network drg-attachment --compartment-id <compartment_OCID>`)
- Run the `oci network drg-attachment create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network drg-attachment create \
--drg-id <drg_OCID> \
--vcn-id <vcn_OCID>
```

Example:

```
oci network drg-attachment create \
--drg-id ocid1.drg.....uniqueID \
--vcn-id ocid1.vcn.....uniqueID

{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "display-name": "drgattachment20210902221928",
    "drg-id": "ocid1.drg.....uniqueID",
    "id": "ocid1.drgattachment.AK00661530.scasg01.....uniqueID",
    "lifecycle-state": "ATTACHING",
    "route-table-id": null,
    "time-created": "2021-09-02T22:19:28.642402+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID"
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}
```

4.3.5 Accessing Oracle Services through a Service Gateway

Some services are isolated on their own network for security and performance reasons. The service gateway (SG) allows a VCN with no external access to privately access Service Network services (such as object storage) in a private subnet. These services are reached at the infrastructure level through the management node cluster.

The feature is non-functional and implemented for compatibility purposes.

A VCN can have only one service gateway. The service gateway is automatically attached to the VCN it is created in. Services use CIDR labels, and are allowed by default.

For each enabled Service, you need a route rule with the Service object's *cidrBlock* as the rule destination and the service gateway as the rule target.

Using the UI

1. In the navigation menu, under Networking, click Virtual Cloud Networks. A list of previously configured VCNs in compartments appears. If the compartment you are creating the service gateway in is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on the VCN that you are creating the service gateway in.
3. In the Resources menu for that VCN, click on Service Gateways (If you are creating a service gateway for a particular VCN< the number of configured service gateways in parentheses should be zero (0)).
4. Click on Create Service Gateway
5. Fill in the required service gateway information:

- **Name:** Provide a name or description for the service gateway. Avoid using any of the organization's confidential information.
- **Create in Compartment:** Select the compartment in which to create the service Gateway.
- **Services:** Select the service from the list.
- **Route Table Association (Optional):** Optionally, you can associate a route table with the Service Gateway. A list of configured route tables for the selected compartment is in a drop-down list. You can change the compartment by clicking (change) next to the compartment name.

For more information on service gateways, refer to the *Service Gateways* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

- **Tagging:** Optionally, add one or more tags to this resource.

For more information about tagging, see [Section 3.4.1, "Adding Tags at Resource Creation"](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

6. Click Create Service Gateway.

The Service Gateway is now ready for the addition of route rules or security settings.

Using the CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VCN OCID (`oci network vcn list --compartment-id <compartment_OCID`)
2. Run the `oci network service-gateway create` command.

Complex data types are usually handled by using the `--generate-full-command-json-input` option, or, in this case, `oci network service-gateway create --generate-param-json-input services`. This generates a sample json file to be used with this command option. The key

names are pre-populated and match the command option names (converted to camelCase format, e.g. compartment-id becomes compartmentId).

The values of the keys are edited by the user before the sample file can be used as an input to this command.

For any command option that accepts multiple values, the value of the key can be a JSON array.

Options can still be provided on the command line. If an option exists in both the JSON document and the command line then the command line specified value will be used.

```
oci network service-gateway create
--compartment-id ocid1.compartment.....uniqueID
--vcn-id ocid1.vcn.....uniqueID
--services ' [{"serviceId": "grafana"} ]'

{
  "data": {
    "displayName": "servicegateway20210830204524",
    "freeform-tags": {},
    "id": "ocid1.servicegateway.....uniqueID",
    "maxWaitSeconds": 0,
    "routeTableId": NULL,
    "services": [
      {
        "serviceId": "grafana"
      }
    ],
    "vcnId": "ocid1.vcn.....uniqueID",
    "waitForState": "PROVISIONING",
    "waitIntervalSeconds": 0
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}
```

4.4 Configuring VNICs and IP Addressing

The compute nodes in the Oracle Private Cloud Appliance have physical network interface cards (NICs). When you create and launch a compute instance on one of the servers, the Networking service ensures that a VNIC is created on top of a physical interface, so that the instance can communicate over the network. Each instance has a primary VNIC that is automatically created and attached during launch. The primary VNIC resides in the subnet you specify when creating the instance. It cannot be removed from the instance.

You can add secondary VNICs to an instance after it is launched. To be able to use a secondary VNIC, you must also configure the instance OS for it. Secondary VNICs must always be attached to an instance and cannot be moved. The process of creating a secondary VNIC automatically attaches it to the instance. The process of detaching a secondary VNIC automatically deletes it. They are automatically detached and deleted when you terminate the instance.

Instances use IP addresses for communication. Each instance has at least one private IP address and optionally a public IP address. A private IP address enables the instance to communicate with other instances inside the VCN, or with hosts in your on-premises network. A public IP address enables the instance to communicate with hosts outside of the Oracle Private Cloud Appliance network environment.

When you create a subnet, it is considered public by default, which means instances in that subnet are allowed to have public addresses and internet communication is permitted. Whoever launches the instance chooses whether it has a public address. You can override that behavior when creating the subnet and request that it be private, which disallows the use of public addresses and internet communication.

Network administrators can therefore ensure that instances in the subnet have no internet access, even if the VCN has a working internet gateway, and security rules and firewall rules allow the traffic.

4.4.1 Public IP Addresses

A public IP address is an IPv4 address that is reachable from the internet. If a resource in your tenancy needs to be directly reachable from the internet, it must have a public IP address. Depending on the type of resource, there might be other requirements.

Certain types of resources in your tenancy are designed to be directly reachable from the internet and therefore automatically come with a public IP address. For example: a NAT gateway or a public load balancer. Other types of resources are directly reachable only if you configure them to be. For example, instances in your VCN.

This section focuses on:

- The types of public IP addresses and their characteristics
- How to control whether an instance has a public IP address
- Creating and managing IP addresses

You can assign a public IP address to an instance to enable communication with the internet. The instance is assigned a public IP address from the address pool.

The assignment is actually to a private IP address object on the instance. The VNIC that the private IP is assigned to must be in a public subnet. A given instance can have multiple secondary VNICs, and a given VNIC can have multiple secondary private IPs. So you can assign a given instance multiple public IPs across one or more VNICs if you like.

There are two types of public IPs:

- **Ephemeral:** Think of this type as temporary and existing for the lifetime of the instance.
- **Reserved:** Think of this type as persistent and existing beyond the lifetime of the instance it's assigned to. You can unassign it and then reassign it to another instance whenever you like.

When you launch an instance in a public subnet, by default, the instance gets a public IP unless you say otherwise. After you create a given public IP, you can't change which type it is. For example, if you launch an instance that is assigned an ephemeral public IP with one IP address, you can't convert the ephemeral public IP to a reserved public IP with the same address.

If you try to perform any operation that assigns or moves a public IP to a VNIC or instance that has already reached its public IP limit, an error is returned.

The operations include:

- Assigning a public IP
- Creating a new secondary VNIC with a public IP
- Moving a private IP with a public IP to another VNIC
- Moving a public IP to another private IP

4.4.2 Managing Public IP Addresses

You can assign a public IP address to an instance to enable communication with the internet. The instance is assigned a public IP address from the Oracle Cloud Infrastructure address pool. The assignment is to the private IP address object already present on the instance.

All of the following are required for an instance to communicate directly with the internet:

- The instance must be in a public subnet, which is configured when the subnet is created. Private subnets cannot have a public IP address assigned to instances in the subnet.
- The instance must have a public IP address.
- The instance's VCN must have an internet gateway configured.
- The public subnet must have route table and security list entries that enable the internet communications.

This section covers public IP address creation and management. This section does not cover instance, internet gateway, route table, or security list configuration. For more information, see the sections of those component configurations and the *Network Scenarios* section of the *Networking* chapter in the [Oracle Private Cloud Appliance User Guide](#).

4.4.2.1 Viewing Public IP Addresses

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID: (`oci iam compartment list <options>`)
2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network public-ip list
--compartment-id <compartment_OCID>
--scope <REGION | AVAILABILITY_DOMAIN>
```

Note

If the public IP exists within a region and is assigned to a regional entity (such as a NAT gateway), it can be assigned to a private IP in any availability domain in the region. Reserved public IPs have scope = REGION, as do ephemeral public IPs assigned to a regional entity. Ephemeral public IPs, on the other hand, are assigned to private IPs and can have scope = AVAILABILITY_DOMAIN.

Example:

```
oci network public-ip list
--compartment-id ocid1.compartment. .... .uniqueID
--scope REGION
{
  "data": [
    {
      "assigned-entity-id": null,
      "assigned-entity-type": "PRIVATE_IP",
      "availability-domain": null,
      "compartment-id": "ocid1.compartment. .... .uniqueID",
```

```

    "defined-tags": {},
    "display-name": "publicip20211110220031",
    "freeform-tags": {},
    "id": "ocidl.publicip. .... .uniqueID",
    "ip-address": "10.80.79.158",
    "lifecycle-state": "AVAILABLE",
    "lifetime": "RESERVED",
    "private-ip-id": null,
    "public-ip-pool-id": null,
    "scope": "REGION",
    "time-created": "2021-11-10T22:00:31.040800+00:00"
  }
]
}

```

4.4.2.2 Reserving a Public IP Address

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID: (`oci iam compartment list <options>`)
2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```

oci network public-ip create \
--compartment-id <compartment_OCID> \
--lifetime <RESERVED>

```

Example

```

oci network public-ip create
--compartment-id ocidl.compartment. .... .uniqueID
--lifetime RESERVED
{
  "data": {
    "assigned-entity-id": null,
    "assigned-entity-type": "PRIVATE_IP",
    "availability-domain": null,
    "compartment-id": "ocidl.compartment. .... .uniqueID",
    "defined-tags": {},
    "display-name": "publicip20211110220031",
    "freeform-tags": {},
    "id": "ocidl.publicip. .... .uniqueID",
    "ip-address": "10.80.79.158",
    "lifecycle-state": "PROVISIONING",
    "lifetime": "RESERVED",
    "private-ip-id": null,
    "public-ip-pool-id": null,
    "scope": "REGION",
    "time-created": "2021-11-10T22:00:31.040800+00:00"
  },
  "etag": "aad5bbcc-9d89-40cd-ab10-03dcc2e4ee0a"
}

```

4.4.2.3 Assigning a Public IP Address to an Instance

The `private-ip-id` parameter option is used to assign the public IP address to a private IP address. You must create an `EPHEMERAL` public IP address to do this. The ephemeral address lifetime is the lifetime of the instance. An ephemeral IP address is deleted when its private IP is deleted, its VNIC is detached or terminated, or its instance is terminated.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID: (`oci iam compartment list --all`)
- Private IP OCID (`oci network private-ip list`)

You can also specify the `--ip-address`, which lists just one IP address. In addition, the `--vlan-id`, `oci network private-ip get --private-ip-id OCID`), and `--vnic-id` options list one IP address.

2. Run this command.

Note

For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network public-ip create \
--compartment-id <compartment_OCID> \
--lifetime <EPHEMERAL> \
--private-ip-id <private_ip_OCID> (required for EPHEMERAL lifetime)
```

Example:

```
# oci network public-ip create \
--compartment ocid1.compartment.ocid1.compartment. .... .uniqueID \
--lifetime EPHEMERAL \
--private-ip-id ocid1.privateip.. .... .uniqueID
{
  "data": {
    "assigned-entity-id": "ocidl.privateip.. .... .uniqueID",
    "assigned-entity-type": "PRIVATE_IP",
    "availability-domain": "ad1",
    "compartment-id": "ocidl.compartment.. .... .uniqueID",
    "defined-tags": {},
    "display-name": "publicip20211115191656",
    "freeform-tags": {},
    "id": "ocidl.publicip.. .... .uniqueID",
    "ip-address": "10.80.79.159",
    "lifecycle-state": "ASSIGNING",
    "lifetime": "EPHEMERAL",
    "private-ip-id": "ocidl.privateip.. .... .uniqueID",
    "public-ip-pool-id": null,
    "scope": "AVAILABILITY_DOMAIN",
    "time-created": "2021-11-15T19:16:56.063971+00:00"
  },
  "etag": "aad5bbcc-9d89-40cd-ab10-03dcc2e4ee0a"
}
```

4.4.2.4 Editing Public IP Addresses

You can update some of the parameters associated with a public IP address.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID: `(oci iam compartment list --all)`
 - Public IP address OCID: `(oci network public-ip list --compartment-id <compartment_OCID> --scope AVAILABILITY_DOMAIN)`
2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network public-ip update
--public-ip-id <public_ip_OCID>
--display-name <text> (this parameter is optional)
```

Note

The `display-name` parameter is technically optional, but used here as an example of things that can be edited or updated.

Example

```
oci network public-ip update
--public-ip-id ocid1.compartment. .... .uniqueID
--display-name public-ip-test-address
{
  "data": {
    "assigned-entity-id": null,
    "assigned-entity-type": "PRIVATE_IP",
    "availability-domain": null,
    "compartment-id": "ocid1.compartment. .... .uniqueID"
  },
  "defined-tags": {},
  "display-name": "public-ip-test-address",
  "freeform-tags": {},
  "id": "ocid1.publicip. .... .uniqueID",
  "ip-address": "10.80.79.158",
  "lifecycle-state": "AVAILABLE",
  "lifetime": "RESERVED",
  "private-ip-id": null,
  "public-ip-pool-id": null,
  "scope": "REGION",
  "time-created": "2021-11-10T22:00:31.040800+00:00"
},
"etag": "2f878ae9-c038-42f9-a410-4574810e5a25"
}
```

4.4.2.5 Deleting Public IP Addresses

You can delete a public IP address.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID: (`oci iam compartment list --all`)
- Public IP address OCID: (`oci network public-ip list --compartment-id <compartment_OCID> --all`)

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network public-ip delete
--public-ip-id <public_ip_OCID>
```

Example

```
oci network public-ip delete
--public-ip-id ocid1.publicip. .... .uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

You can suppress the confirmation by using the `--force` option.

4.4.3 Creating and Managing VNICs

The compute nodes in the Oracle Private Cloud Appliance have physical network interface cards (NICs). When you create and launch a compute instance on one of the servers, the Networking service ensures that a VNIC is created on top of a physical interface, so that the instance can communicate over the network. Each instance has a primary VNIC that is automatically created and attached during launch. The primary VNIC resides in the subnet you specify when creating the instance. It cannot be removed from the instance.

You can add secondary VNICs to an instance after it is launched. To be able to use a secondary VNIC, you must also configure the instance OS for it. The maximum number of VNICs for an instance varies by shape. Each secondary VNIC can be in a subnet in the same VCN as the primary VNIC, or in a different subnet that is either in the same VCN or a different one. Note that attaching multiple VNICs from the same subnet CIDR block to an instance can introduce asymmetric routing, especially on Linux instances. For more information, refer to the *Virtual Network Interface Cards (VNICs)* section in the *Instance Connectivity* chapter of the [Oracle Private Cloud Appliance Concepts Guide](#).

4.4.3.1 View an Instance VNIC

Using the UI

1. Open the navigation menu and click Compute. Under Compute, click Instances.
2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

Using the CLI

1. Gather the information you need to run the command:
 - Instance OCID: (`oci compute instance list --all`)
2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci compute instance list-vnics \
--instance-id <instance_OCID>
```

Example

```
oci compute instance list-vnics --instance-id ocid1.....uniqueID

{
  "data": [
    {
      "availability-domain": "PCANETWORK",
      "compartment-id": "ocid1.tenancy.....uniqueID",
      "defined-tags": {},
      "display-name": "vnic20210528203222",
      "freeform-tags": {},
      "hostname-label": null,
      "id": "ocid1.vnic.....uniqueID",
      "is-primary": true,
      "lifecycle-state": "AVAILABLE",
      "mac-address": "00:13:97:99:cc:4b",
      "nsg-ids": null,
      "private-ip": "10.0.5.3",
      "public-ip": null,
      "skip-source-dest-check": false,
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-05-28T20:32:22+00:00",
      "vlan-id": null
    }
  ]
}
```

4.4.3.2 View VNIC Attachments in a Compartment

Using the UI

1. Open the Navigation Menu. Under Networking, click Instances.
2. Make sure you are in the compartment you are interested in. If not, you can select the correct compartment from the dropdown list of compartments.
3. Click on the instance that you are interested in.
4. In the Resources list for the instance, click on Attached VNICs.

The list of attached VNICs for that instance is displayed.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci compute vnic-attachment list \
--compartment-id <compartment_OCID>
```

Example

```
oci compute vnic-attachment list --compartment-id ocid1.tenancy.....uniqueID
{
  "data": [
    {
      "availability-domain": "PCANETWORK",
      "compartment-id": "ocid1.tenancy.....uniqueID",
      "display-name": "vnicattachment20210528203222",
      "id": "ocid1.vnicattachment.....uniqueID",
      "instance-id": "ocid1.instance.....uniqueID",
      "lifecycle-state": "ATTACHED",
      "nic-index": 0,
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-05-28T20:32:22+00:00",
      "vlan-id": null,
      "vlan-tag": 0,
      "vnic-id": "ocid1.vnic.....uniqueID"
    },
    {
      "availability-domain": "PCANETWORK",
      "compartment-id": "ocid1.tenancy.....uniqueID",
      "display-name": "vnicattachment20210601150704",
      "id": "ocid1.vnicattachment.....uniqueID",
      "instance-id": "ocid1.instance.....uniqueID",
      "lifecycle-state": "ATTACHED",
      "nic-index": 0,
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-06-01T15:07:04+00:00",
      "vlan-id": null,
      "vlan-tag": 0,
      "vnic-id": "ocid1.vnic.....uniqueID"
    }
  ]
}
```

4.4.3.3 Create and Attach a Secondary VNIC

You can always configure a secondary VNIC for an instance. However, to enable the instance OS to attach to and use the secondary VNIC, you must also configure the instance to use secondary VNICs. To configure the instance OS to use the VNIC, see either [Linux: Configuring the OS for Secondary VNICs](#) or [Windows: Configuring the OS for Secondary VNICs](#), depending on the instance OS.

Using the UI

1. Open the navigation menu and click Compute. Under Compute, click Instances.
2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. Click Create VNIC.
6. In the Create VNIC Attachment dialog box, you specify which VCN and subnet to put the VNIC in. By default, the VNIC will be created in the current compartment. You can choose a VCN and subnet from the same compartment or a different compartment.

Enter the following:

- **Compartment:** The current compartment is displayed. Click (change) to select a different compartment from the dropdown list.
 - **VCN:** Select the VCN that contains the subnet of interest from the dropdown list.
 - **Subnet:** Select the subnet of interest from the dropdown list. The secondary VNIC must be in the same availability domain as the instance's primary VNIC, so the subnet list includes any [regional subnets](#) or [AD-specific subnets](#) in the primary VNIC's availability domain.
 - **Skip source/destination check:** By default, this check box is NOT selected, which means the VNIC performs the source/destination check. Only select this check box if you want the VNIC to be able to forward traffic. See [Overview of VNICs and Physical NICs](#).
 - **Assign a Public IP Address:** Whether to assign a public IPv4 address to the VNIC's primary private IP. This option is available only if the subnet is public. Choose this option to specify an existing [reserved public IP address](#) by name, or to create a new reserved IP address by assigning a name and selecting a source [IP pool](#) for the address. If you don't select an IP pool you've created, the default Oracle IP pool is used.
 - **Primary IP Information:**
 - **Private IP Address:** Optional. The address entered must be within the CIDR block range assigned to the subnet and must not already be in use.
 - **Hostname:** Optional. A hostname to be used for DNS within the cloud network. Only letters, numbers, and hyphens with no spaces are allowed, up to 63 characters maximum. Option available only if the VCN and subnet both have DNS labels.
 - **Enable Network Security Groups:** By default, you cannot attach Network Security Groups to the VNIC. Select this to allow Network Security Groups and display a dropdown list of configured NSGs to choose from.
7. Click Create Attachment. The secondary VNIC is created and then displayed on the Attached VNICs list for the instance. It can take several seconds before the secondary VNIC appears on the page.
 8. You still must configure the instance OS to use the secondary VNIC. See [Linux: Configuring the OS for Secondary VNICs](#) or [Windows: Configuring the OS for Secondary VNICs](#).

Using the CLI

1. Gather the information you need to run the command:

- Instance OCID: (`oci compute instance list --compartment-id <compartment_OCID> --all`)
- Subnet OCID: (`oci network subnet list --compartment-id <compartment_OCID> --all`)

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci compute instance attach-vnic \
--instance-id <instance_OCID> \
--subnet-id <subnet_OCID>
```

Example

```
oci compute instance attach-vnic \
--instance-id ocid1.instance.....uniqueID \
--subnet-id ocid1.subnet.....uniqueID
```

There is no output from this command if it succeeds. To see the attached VNICs, use the `oci compute instance list-vnics --instance-id <instance_OCID>` command. The new attached secondary VNIC is the non-primary entry (`is-primary = false`).

```
oci compute instance list-vnics \
--instance-id ocid1.instance.....uniqueID
{
  "data": [
    {
      "availability-domain": "ad1",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "example-instance",
      "freeform-tags": {},
      "hostname-label": "example",
      "id": "ocid1.vnic.....uniqueID",
      "is-primary": true,
      "lifecycle-state": "AVAILABLE",
      "mac-address": "00:13:97:18:e4:ec",
      "nsg-ids": null,
      "private-ip": "192.168.0.2",
      "public-ip": "10.80.79.145",
      "skip-source-dest-check": false,
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-12-01T16:31:05.681547+00:00",
      "vlan-id": null
    },
    {
      "availability-domain": "ad1",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "example-instance",
      "freeform-tags": {},
```

```

    "hostname-label": "example-instance",
    "id": "ocid1.vnic.....uniqueID",
    "is-primary": false,
    "lifecycle-state": "AVAILABLE",
    "mac-address": "00:13:97:30:9d:db",
    "nsg-ids": null,
    "private-ip": "10.27.100.2",
    "public-ip": "10.80.79.141",
    "skip-source-dest-check": false,
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-12-01T18:21:48.442239+00:00",
    "vlan-id": null
  }
]
}

```

You still must configure the instance OS to use the secondary VNIC. See [Linux: Configuring the OS for Secondary VNICs](#) or [Windows: Configuring the OS for Secondary VNICs](#).

4.4.3.4 Update an Existing VNIC

You can update the VNIC's private IP address or hostname, or whether or not to use Network Security Groups.

Using the UI

1. You update the attached and existing VNICs through the instance. Select the instance that you are interested in.
2. Select Attached VNICs from the Resources list for that instance.
3. Select the attached VNIC that you are updating.
4. Select IP Addresses or Network Security Groups from the Resources list for the instance.
5. For IP Addresses, click Assign Private IP address to change the IP address or hostname for the VNIC attachment.
 - a. **IP Address:** Enter the new IP address to use for the VNIC. The address entered must be within the CIDR block range assigned to the subnet and must not already be in use.
 - b. **Hostname:** A hostname to be used for DNS within the cloud network. Only letters, numbers, and hyphens with no spaces are allowed, up to 63 characters maximum. Option available only if the VCN and subnet both have DNS labels.
 - c. Click Attach IP Address for VNIC to save your changes.
6. For Network Security Groups, click Update Network Security Groups.
 - a. **Enable Network Security Groups:** By default, you cannot attach Network Security Groups to the VNIC. Select this to allow Network Security Groups and display a dropdown list of configured NSGs to choose from. If NSGs are enabled, clear this option to remove the NSGs from the VNIC.
 - b. Click Update Network Security Groups for VNIC to save your changes.
7. The updated information is listed for the IP address or NSG resource for the instance VNIC.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID (`oci iam compartment list --all`)
- VNIC OCID: (`oci compute vnic-attachment list --compartment-id <compartment_OCID> --all`)

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci network vnic update \
--vnic-id <vnic_OCID> \
--display-name <name> (optional parameter, but used for example)
```

Example

```
oci network vnic update --vnic-id ocid1.vnic.....uniqueID
--display-name easy-to-recall-display-name
{
  "data": [
    {
      "availability-domain": "ad1",
      "compartment-id": "ocid1.tenancy.....uniqueID",
      "display-name": "easy-to-recall-display-name",
      "id": "ocid1.vnicattachment.....uniqueID",
      "instance-id": "ocid1.instance.....uniqueID",
      "lifecycle-state": "ATTACHING",
      "nic-index": 0,
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-10-22T17:06:14.398445+00:00",
      "vlan-id": null,
      "vlan-tag": 0,
      "vnic-id": "ocid1.....uniqueID"
    }
  ]
}
```

4.4.3.5 Add or Remove a VNIC from a Network Security Group

You can change which [network security groups](#) (NSGs) a VNIC belongs to, or remove a VNIC from all NSGs. A related topic is [Section 4.2.4.2, “Attach a Network Security Group to a VNIC”](#).

Using the UI

1. Open the navigation menu and click Compute. Under Compute, click Instances.
2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. Click the VNIC you're interested in.

Each VNIC's details page includes a list of the NSGs that the VNIC belongs to (if any).

6. Click the NSG that you are interested in.
7. Click Update Network Security Group.
8. Select the NSG to add to the VNIC to.
9. Click Update Network Security Group for VNIC.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --options`)
 - VNIC OCID: (`oci compute vnic-attachment list --compartment-id <compartment_OCID> --all`)
 - NSG OCID (`oci network nsg list --compartment-id <compartment_OCID> --options`)
2. Run this command. At the WARNING, you can choose to update existing values or not.

Syntax (entered on a single line):

```
oci network vnic update \
--vnic-id <vnic_OCID> \
--nsg-ids '["<nsg_OCID>"]'
```

Example

```
oci network vnic update \
--vnic-id ocid1.vnic.....uniqueID \
--nsg-ids '["ocid1.networksecuritygroup.....uniqueID"]
,
WARNING: Updates to defined-tags and freeform-tags and nsg-ids will replace any existing values.
Are you sure you want to continue? [y/N]: y
{
  "data": {
    "availability-domain": "ad1",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "example-instance",
    "freeform-tags": {},
    "hostname-label": "example-instance",
    "id": "ocid1.vnic.1742XC3024.broom14.....uniqueID",
    "is-primary": false,
    "lifecycle-state": "PROVISIONING",
    "mac-address": "00:13:97:5c:7a:58",
    "nsg-ids": [
      "ocid1.networksecuritygroup.1.....uniqueID"
    ],
    "private-ip": "10.27.100.2",
    "public-ip": null,
    "skip-source-dest-check": true,
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-12-01T17:14:46.084875+00:00",
    "vlan-id": null
  },
  "etag": "17e8e038-3051-4ef8-96a4-21d7107d37e6"
}
```

4.4.3.6 Delete a Secondary VNIC

This operation detaches and deletes the specified secondary VNIC. This operation cannot be used on the instance's primary VNIC. When you terminate an instance, all attached VNICs (primary and secondary) are automatically detached and deleted.

Caution

If the VNIC has a private IP that is the [target of a route rule](#), deleting the VNIC causes the route rule to blackhole and traffic will be dropped.

Using the UI

1. Open the navigation menu and click Compute. Under Compute, click Instances.
2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. For the VNIC you want to delete, click the Actions icon (three dots), and then click Delete VNIC.
6. Confirm when prompted.

It takes typically a few seconds before the VNIC is deleted.

If the secondary VNIC is on a Linux instance: If you then run the provided script in [Linux: Configuring the OS for Secondary VNICs](#), it removes the secondary VNIC from the OS configuration.

Using the CLI

1. Gather the information you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - VNIC OCID: (`oci compute vnic-attachment list --compartment-id <compartment_OCID> --all`)
2. Run this command.

Syntax (entered on a single line):

```
oci compute instance detach-vnic \
--compartment-id <compartment_OCID> \
--vnic-id <vnic_OCID>
```

Example

```
oci compute instance detach-vnic \
--compartment-id ocid1.compartment.....uniqueID \
--vnic-id ocid1.vnic.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

4.4.3.7 Manage Tags for a VNIC

Using the UI

1. Open the navigation menu and click Compute. Under Compute, click Instances.

2. Confirm you're viewing the compartment that contains the instance you're interested in.
3. Click the instance to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. Click the VNIC that you're interested in.
6. Click the Tags tab to view or edit the existing tags. Or click Add Tags to add new ones.

For more information about tagging, see [Section 3.4.1, “Adding Tags at Resource Creation”](#) or [Resource Tag Management](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID (`oci iam compartment list --all`)
- VNIC OCID (`oci compute vnic-attachment list --compartment-id <compartment_OCID> --all`)
- Defined-tags (specified in a JSON string or file that lists the tags you want to modify.)

Complex data types are usually handled by using the `--generate-full-command-json-input` option, or, in this case, `oci network vnic update --generate-param-json-input defined-tags`. This generates a sample json file to be used with this command option. The key names are pre-populated and match the command option names (converted to camelCase format, e.g. compartment-id becomes compartmentId).

The values of the keys are edited by the user before the sample file can be used as an input to this command.

For any command option that accepts multiple values, the value of the key can be a JSON array.

Options can still be provided on the command line. If an option exists in both the JSON document and the command line then the command line specified value will be used.

2. Run this command. At the WARNING, you can choose to update existing values or not.

Syntax (entered on a single line):

```
oci network vnic update \
--vnic-id <vnic_OCID> \
--defined-tags <complex-type>
```

Example

```
oci network vnic update \
--vnic-id ocid1.vnic.....uniqueID \
--defined-tags '{"tagNamespace1":{"tagKey1": "tagValue1","tagKey2": "tagValue2"}, \
"tagNamespace2":{"tagKey1": "tagValue1","tagKey2": "tagValue2"}}'
```

WARNING: Updates to defined-tags and freeform-tags and nsg-ids will replace any existing values.
Are you sure you want to continue? [y/N]: y

```
{
  "data": {
    "availability-domain": "ad1",
    "compartment-id": "ocid1.compartment.....uniqueID",
```



```

"defined-tags": {
  "tagNamespace1": {
    "tagKey1": "tagValue1",
    "tagKey2": "tagValue2"
  },
  "tagNamespace2": {
    "tagKey1": "tagValue1",
    "tagKey2": "tagValue2"
  }
},
"display-name": "example-instance",
"freeform-tags": {},
"hostname-label": "example",
"id": "ocidl.vnic.....uniqueID",
"is-primary": true,
"lifecycle-state": "PROVISIONING",
"mac-address": "00:13:97:18:e4:ec",
"nsg-ids": null,
"private-ip": "192.168.0.2",
"public-ip": "10.80.79.145",
"skip-source-dest-check": false,
"subnet-id": "ocidl.subnet.....uniqueID",
"time-created": "2021-12-01T16:31:05.681547+00:00",
"vlan-id": null
},
"etag": "17e8e038-3051-4ef8-96a4-21d7107d37e6"
}

```

4.4.4 Assigning IP Addresses to VNICs

You can assign a private IP address to an instance and VNIC that does not have one.

Using the UI

1. From the Dashboard, click on Compute-View Instances.
2. From the Resources navigation menu, click Instances.
3. Click the Instance you are interested in from the dropdown list to view its details.
4. Under Resources, click Attached VNICs.

The primary VNIC and any secondary VNICs attached to the instance are displayed.

5. Click on the Instance with Attached VNICs that you are interested in.

Each instance's details page includes a list of the IP addresses associated with instance hostnames (if any).

6. Click Assign Private IP Address.

Enter the following:

- **IP Address:** Enter the private IP address to assign to the instance. The IP address must be within the CIDR block for the instance.
- **Hostname:** Enter the hostname to assign to the instance.

7. Click Attach IP Address.

Using the CLI

1. Gather the information you need to run the command:

- Compartment OCID (`oci iam compartment list --all`)
- VNIC OCID: (`oci compute vnic-attachment list --compartment-id <compartment_OCID> --all`)

2. Run this command.

Syntax (entered on a single line):

```
oci network vnic assign-private-ip --vnic-id <vnic_OCID>
```

Example

```
oci network vnic assign-private-ip \
--vnic-id ocid1.vnic.....uniqueID
{
  "data": {
    "assigned-entity-id": null,
    "assigned-entity-type": "PRIVATE_IP",
    "availability-domain": null,
    "compartment-id": "ocid1.compartment. .... .uniqueID"
  },
  "defined-tags": {},
  "display-name": "public-ip-test-address",
  "freeform-tags": {},
  "id": "ocid1.publicip. .... .uniqueID",
  "ip-address": "10.27.100.3",
  "lifecycle-state": "AVAILABLE",
  "lifetime": "RESERVED",
  "private-ip-id": null,
  "public-ip-pool-id": null,
  "scope": "REGION",
  "time-created": "2021-11-10T22:00:31.040800+00:00"
},
  "etag": "2f878ae9-c038-42f9-a410-4574810e5a25"
}
```

4.5 Managing Public DNS Zones

In its most basic form, DNS returns an IP address (if known) when given a string in the DNS name space for that zone. However, DNS is also the way that an IP host client application knows where to get its own configuration information using DHCP (DHCID records), go to send or receive email (MX records), and more. Without DNS, client devices would have to know the proper IP addresses not only for local servers, but for every server or application they interacted with, no matter where in the world they were located. With DNS, clients can always find the correct location of `www.oracle.com` or any other application.

Once you create a DNS zone inside a compartment, you cannot move the zone to another compartment.

4.5.1 Creating a Public DNS Zone

DNS zones are created in a compartment to associate IP addresses with portions of the DNS name space. Zones are created in a compartment using the DNS service.

Using the UI

1. In the navigation menu, under DNS, click Zones. A list of previously configured zones in compartments appears. If the compartment you are creating the DNS zone in is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click Create Zone.

3. Fill in the required zone information:

- **Zone Name:** Provide a name or description for the DNS zone. Avoid using any of the organization's confidential information.
- **Compartment:** Select the compartment in which to create the DNS zone.
- **Zone Type:** Choose the type of DNS zone you are creating.
 - **Primary:** A primary DNS zone is the original authoritative DNS zone of a portion of the DNS name space. When a DNS server hosts a primary zone, that DNS server is the Authoritative DNS Server and is considered the primary source of information in that zone.
 - **Secondary:** A secondary DNS zone is a read-only copy of a primary DNS zone or another secondary DNS zone. A secondary DNS zone is kept on a Secondary DNS Server and reduces the load on the primary DNS zone and eliminated a single point of failure risk to name resolution inside the zone.

For more information on DNS Zones, refer to the *Name Resolution* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

4. Click Create Zone.

The zone is now ready for the addition of zone records or for the configuration of TSIG Keys or Steering Policies.

Using the CLI

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list --all`)

2. Run the `oci dns zone create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci dns zone create \
--compartment-id <compartment_OCID> \
--name <dns_zone_name> \
--zone-type <PRIMARY | SECONDARY>
```

Example:

```
oci dns zone create \
--compartment-id ocid1.compartment.....uniqueID \
--name test-dns-zone \
--zone-type PRIMARY
```

```

{
  "data": {
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "external-masters": null,
    "freeform-tags": {},
    "id": "ocidl.dns-zone.....uniqueID",
    "is-protected": null,
    "lifecycle-state": "ACTIVE",
    "name": "test_dns_zone",
    "nameservers": [
      {
        "hostname": "ns1.example.com"
      }
    ],
    "scope": null,
    "self-uri": "https://20180115/zones/test_dns_zone",
    "serial": 1,
    "time-created": "2021-08-17T18:08:00.059867+00:00",
    "version": 1,
    "view-id": null,
    "zone-type": "PRIMARY"
  },
  "etag": "3e389cab-b3fd-4783-91c1-ede81bc132d5"
}

```

4.5.2 Working with Zone Records

Creating a DNS zone is only the beginning of working with DNS. The zone is essentially empty when created, except for a basic Start of Authority (SOA) and Name Server (NS) record. The SOA record provides a kind of history of this DNS zone and holds information such as when it was last updated and things like that. The NS record contains the fully-qualified name of the DNS server for the zone. The NS record is very important and therefore has a high TTL, usually 24 hours (86400 seconds).

To make the name server truly useful, the zone must be rounded out and filled with the DNS records that form the basis of responses to the kinds of queries that clients make. These queries include IP addresses for parts of the domain name space, email server details, and so on.

4.5.2.1 Creating a Zone Record

The RDATA field is where the content of the zone record is entered. The format of the information varies according to the type of record you are creating. However, the data must be in one of the formats that DNS understands. For example, an A-type zone record RDATA is an IP address, and an MX record contains information on how to route email. Because of the authoritative nature of the zone records within a zone, RDATA is not editable. If DNS information in a zone changes, then the old record must be deleted and a new record created.

Using the UI

1. In the navigation menu, under DNS, click Zones. A list of previously configured zones in compartments appears. If the compartment you are adding zone records to is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on the name of the zone. The information screen contains general zone information such as type and compartment, OCID (which you can show in full or copy to the clipboard), and the date and time that the zone was created. The zone records that exist are also displayed, and initially there are only SOA and NS records.
3. Fill in the required zone record information:
 - **Zone Record:** Select the type of zone record you are creating from the drop-down list.

- **A - IPv4 Address:** A host record, which is used to point a hostname to an IPv4 address. This is the most basic DNS record type.
 - You can add many other types of zone records: any types in the drop down list. For more information on the DNS zone record types and the data they require, see the *Name Resolution* section in the Virtual Networking Overview chapter of the Oracle Private Cloud Appliance Concepts Guide.
 - **Domain (Optional):** Type the name of the zone subdomain if used (this value is already filled in based on the zone itself: the initial dot (".") is used for adding the zone subdomain).
 - **TTL:** Check this box to set your own value for the TTL of that particular record type. If you do not check this box, the default TTL value for that record type is used (for example, 300 for SOA, 86400 for NS). The valid range is from 1 to 129540 seconds (from 1 second to about a day and a half).
 - **Edit RDATA:** Check this box if you wish to edit the RDATA information, such as the IP address or Target established by the zone record type. This box is only displayed for some zone record types.
 - **(RDATA):** This unlabeled field varies based on the type of zone record created. For example, you enter the 32-bit IP address that corresponds to the A-type DNS record, or Flags for a DNSKEY zone record, if that is what you are creating.
 - **A - IPv4 Address:** If you are creating an A type zone record, then the data is a properly formatted IPv4 address. This is the most basic DNS record, but there are many others.
 - The RDATA field reflects the correct information for the type of zone record selected.
4. Click Create Record.

The zone record is now added to the zone. If you click the optional box to **Add another record**, then the screen stays at the **Create DNS Zone Record** state to make record entry more efficient.

Using the CLI:

There is no "create dns zone record" command in the CLI. Instead, the command "`oci dns zone record update`" command replaces records in the specified zone with the records specified in the request body of the command. If a specified record does not exist, then it will be created. Also, if a current record is not in the records list, it will be deleted. Care is needed, because if the record exists, then it will be updated with the record information in the body of the request. The command in this section adds an A resource record (IPv4 address and domain name) to a DNS zone named "dns-test-zone."

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - DNS zone name (`oci dns zone list --compartment_OCID <compartment_OCID> --all`)
2. Run the `oci dns record zone update` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci dns record zone update \
  --zone-name-or-id <zone_name> or <compartment_OCID> \
  --items <complex type>
```

Note

DNS resource record types are provided as "items" in a complex JSON format. This is a JSON list with items of type RecordDetails. For documentation on RecordDetails please see <https://docs.cloud.oracle.com/api/#/en/dns/20180115/datatypes/RecordDetails>. This is a complex type whose value must be valid JSON. The value can be provided as a string on the command line or passed in as a file using the `file://path/to/file` syntax.

The `--generate-param-json-input` option can be used to generate an example of the JSON which must be provided. We recommend storing this example in a file, modifying it as needed and then passing it back in via the `file://` syntax.

Example:

```
oci dns record zone update
  --zone-name-or-id <zone_name> or <compartment_OCID>
  --items
  {
    "domain": "test-dns-zone.test-pca-comparment.example.com",
    "isProtected": true,
    "rdata": "10.225.15.10",
    "recordHash": "fkT4md",
    "rrsetVersion": "1",
    "rtype": "1",
    "ttl": 3600
  }
```

4.5.2.2 Editing a Zone Record

There is no `edit record` command. You can update a group of records, and if one of the records in the list is the same except for the `rdata` for example, in effect you have updated the record.

4.5.2.3 Deleting a Zone Record

You can delete many, but not all, DNS zone records. The initial SOA and NS records, created by default when the zone is created, cannot be deleted. To delete a zone record:

Using the Compute Web UI

1. In the navigation menu, under DNS, click Zones. A list of previously configured zones in compartments appears. If the compartment you are adding zone records to is not in the title bar, then use the drop-down tab to select the correct compartment.
2. Click on the name of the zone. The information screen contains general zone information such as type and compartment, OCID (which you can show in full or copy to the clipboard), and the date and time that the zone was created. The zone records that exist are also displayed.
3. Click on the Action square with the three dots on the right side of the zone record that you are deleting.
4. Click Delete.

The zone record is deleted and removed from the list for that DNS zone.

Using the CLI

To delete resource records in a zone with the CLI, use the `oci dns record rrset delete` command to delete an entire resource record set (for example, all A-type IPv4 address records for a given host name). The resource records are identified by their DNS record type (A, MX, and so on). The command in this section deletes the A resource record (IPv4 address and domain name) in a DNS zone named "dns-test-zone" for a device named "test-device-1."

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
 - DNS zone name (`oci dns zone list --compartment_OCID <compartment_OCID> --all`)
2. Run the `oci dns record rrset delete` command.

Syntax: (entered on a single line):

```
oci dns record rrset delete \
  --domain <domain-name> \
  --rtype <resource-record-type> \
  --zone-name-or-id <zone_name> or <compartment_OCID>
```

Example:

```
oci dns record rrset delete \
  --domain "test-device-1.dns-test-zone.example.com" \
  --rtype "A" \
  --zone-name-or-id "dns-test-zone"
```

The record is deleted from the zone.

4.5.3 Editing a Public DNS Zone

You can add resource tags to an existing zone. You can also edit the externalMasters field of a SECONDARY zone.

4.5.4 Working with Transaction Signature Keys

A DNS transaction signature (TSIG) is a network protocol defined in RFC 2845. The main purpose of the TSIG is to allow DNS to authenticate updates to a DNS database, so that malicious users cannot change name resolution records to point to a bogus IP address instead of (for example) the IP address on a bank. TSIG uses one-way hashing and shared secret keys to provide a secure means to authenticate the endpoints of a connection for processing (or responding to) DNS update requests.

The TSIG protocol uses timestamps to prevent replay of recorded responses. Therefore, DNS servers and TSIG clients need accurate clocks to provide the timestamps. A number of extensions to the basic TSIG protocol have been made to extend the types of cryptography and hashing methods that are supported by TSIG.

To use TSIG for a DNS zone, add TSIG keys to the DNS zone. The TSIG key must be base 64 encoded.

Using the UI

1. In the navigation menu, under DNS Zones, click TSIG Keys.
2. Click Create Key.
3. Fill in the required TSIG Key information:

- **Name:** Provide a name or description for the TSIG key. Avoid using any of the organization's confidential information.
- **Compartment:** Select the compartment in which to create the TSIG key.
- **Algorithm:** Choose the security algorithm for the TSIG Key you are creating, such as **hmac-sha256**.
- **Secret Key:** Provide the base64 string encoding the binary shared secret that corresponds to the key. The maximum is 255 characters. An example key in base64 encoding is shown in RFC3874. You can provide the key in one of two ways:
 - **Select the key file:** If you provide the TSIG shared secret key this way, you can drag and drop the key file into the space provided.
 - **Paste the key:** If you provide the TSIG shared secret key this way, you can copy and paste the contents of the key file into the space provided.
- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

4. Click Create TSIG Key.

The TSIG key now available for use in the DNS zone between TSIG client and DNS server.

Using the CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list --all`)
2. Run the `oci dns tsig-key create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci dns tsig-key create \
--algorithm <hmac-algorithm> \
--name <tsig-key-name> \
--compartment-id <compartment_OCID> \
--secret <secret-string>
```

Example:

```
oci dns tsig-key create --algorithm hmac-sha256 --name new-tsig-key \
--compartment-id ocid1.compartment.....uniqueID \
--secret 2o8goaon2168n(secret key string)e6um81vd2lwoouq461sygak0009014
{
  "data": {
    "-self": "https://20180115/tsigKeys/new-tsig-key",
    "algorithm": "hmac-sha256",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
```



```

"freeform-tags": {},
"id": "ocidl.dns-tsig-key.....uniqueID",
"lifecycle-state": "ACTIVE",
"name": "new-tsig-key",
"secret": "2o8goaon2l68n(secret key string)e6um8lvd2lwdoouq46lsygak0009014",
"time-created": "2021-10-29T17:50:31.219934+00:00",
"time-updated": null
},
"etag": "81eb0e02-e09c-4b25-9d21-ee9a9ab2aacc"
}

```

The new key appears in the list of TSIG keys for the DNS for this compartment.

4.5.4.1 Adding a TSIG Key

To add a TSIG key to an existing list of TSIG keys, simply create another key with a unique TSIG key name and a new algorithm or a new key value. To modify fields in an existing TSIG key, use the [update](#) command.

A TSIG key is a separate object from a DNS zone. You can have a SECONDARY DNS zone reference a TSIG key as part of its ExternalMaster definition. But creating a new key doesn't do anything for a PRIMARY zone.

4.5.4.2 Removing a TSIG Key

Using the UI

1. In the navigation menu, under DNS Zones, click TSIG Keys.
2. Click on the TSIG key that you want to remove from the dropdown list of TSIG keys.
3. Click Delete from the list of actions under the Action Menu icon (three bars), or click the Delete button at the top of the display window.

The TSIG key is removed from the list.

Using the CLI

1. Gather the information for these resources:
 - Compartment OCID (`oci iam compartment list --all`)
 - TSIG OCID (`oci dns tsig-key list --compartment-id <compartment_OCID>`)
2. Syntax (entered on a single line):

```
oci dns tsig-key delete --tsig-key-id <tsig-key_OCID>
```

Example:

```
oci dns tsig-key delete --tsig-key-id ocidl.dns.tsig.key.....uniqueID \
Are you sure you want to delete this resource? [y/N]: y
```

The TSIG key is deleted from the list of TSIG keys for DNS in this compartment. Use the `--force` option to suppress the "Are you sure...?" message.

4.5.5 Deleting a Public DNS Zone

Using the UI

1. In the navigation menu, under DNS Zones, click Zones.

2. Click on the zone name that you want to remove from the dropdown list of zones.
3. Click the Delete button at the top of the display window.

The DNS zone is removed from the list.

Using the CLI

1. Gather the information for these resources:
 - Compartment OCID (`oci iam compartment list --all`)
 - Zone OCID (`oci dns zone list --compartment-id <compartment_OCID>`)
2. Syntax (entered on a single line):

```
oci dns zone delete --zone-name-or-id zone_OCID-or-zone-name>
```

Example:

```
oci dns zone delete --zone-id ocid1.dns.zone.....uniqueID \  
Are you sure you want to delete this resource? [y/N]: y
```

The DNS zone is deleted from the compartment. Use the `--force` option to suppress the "Are you sure...?" message.

4.6 Managing Traffic with Steering Policies

DNS can do more than return an IP address (if known) when given a string in the DNS name space for that zone. DNS is also a part of a system of traffic management, where traffic is distributed among multiple servers depending on some criterion, such as location. Steering policies are a way to distribute access to a single full-qualified name cross multiple servers.

For example, the same content could be available from multiple source servers, whether it is a streaming video or records from a product database. One server might be in the United States, and the other in Europe. A traffic steering policy could distribute traffic based on IP address or CIDR. Other criteria can be used for this traffic distribution, such as load balancing, which strives to keep the load on multiple servers roughly equal.

Oracle Private Cloud Appliance offers two major types of traffic steering policies based on load balancing and some value of the IP address prefix (network portion of the IP address, such as 192.168.100.0/24).

4.6.1 Creating a Load Balancer Steering Policy

If you have more than one DNS server, you can distribute traffic in a load balancing fashion, based on the weight you assign to each of them.

Using the Compute Web UI

1. Open the Navigation Menu. Under DNS Zones, click Steering Policies.
2. Click Create Steering Policy.
3. Click the Load Balancer button to create a load balancer steering policy.
4. Enter the required information:
 - **Name:** Enter a name to display for the load balancer steering policy. Do not use confidential information.

- **Policy TTL:** Enter a TTL in seconds for responses to steering policy requests. The maximum is 604800 seconds (equal to 168 hours or 7 days).
- **Answer(s):** Supply the answer or answers to the DNS request for FILTER, WEIGHED, and LIMIT rules. You do not have to specify which condition the answers is for: that is all done by the load balancer template.
 - **Name:** Enter a name for the RData returned, such as Server1.
 - **Type:** Choose the type of resource record to return for the request from the dropdown list. Choices are items such as A (IPv4 address) or CNAME (canonical name).
 - **RData:** Enter the resource record RData that is returned that corresponds to the Type selected. For example, for Type = A, the RData would be an IPv4 address.
 - **Weight:** Enter a weight for this policy to use for load balancing. Values up to 256 are supported. The default is 10. Higher weights mean that policy answer is used more often. For example, if `dns-server1` and `dns-server2` have equal weights, DNS requests are split evenly between them. If `dns-server1` has a weight twice that of `dns-server2`, then `dns-server1` is used twice as often as `dns-server2`.
- 5. **Disabled:** The steering policy answer is enabled at creation by default. To disable this steering policy answer, click this toggle to change the Disabled value to TRUE.
- 6. Optionally, add or delete tags for this subnet resource.

For more information about tagging, see [Tagging Overview](#). If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click Save Changes. The load balancing steering policy is created.

Using the OCI CLI

1. Gather the information you need.
 - Compartment OCID (`oci iam compartment list --all`)
2. Run the `oci dns steering-policy create` command with the LOAD_BALANCE parameter.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option. Complex types are long json strings.

Syntax (entered on a single line):

```
oci dns steering-policy create
--compartment-id <compartment_OCID>
--display-name <dns_steering_policy_name>
--template <LOAD_BALANCE>
--answers <complex-type>
--rules <complex-type>
```

Example:

```
oci dns steering-policy create
--compartment-id ocid1.compartment.....uniqueID
--display-name test-lb-policy-1
```

```

--template LOAD_BALANCE
--answers '[{"name": "server", "pool": "server", "rdata": "10.25.11.10", "rtype": "A"},
{"name": "trial", "pool": "trial", "rdata": "10.25.11.10", "rtype": "A"}]'
--rules '[{"ruleType": "FILTER", "defaultAnswerData": [{"answerCondition":
"answer.isDisabled != true", "shouldKeep": true}], {"ruleType": "WEIGHTED", "defaultAnswerData":
[{"answerCondition": "answer.name == 'server'", "value": 90}, {"answerCondition":
"answer.name == 'trial'", "value": 10}], {"defaultCount": 1, "ruleType": "LIMIT"}]'

{
  "data": {
    "-self": "https://20180115/steeringPolicies/ocid1.dnspolicy.....uniqueID",
    "answers": [
      {
        "is-disabled": true,
        "name": "server",
        "pool": "server",
        "rdata": "10.25.11.10",
        "rtype": "A"
      },
      {
        "is-disabled": true,
        "name": "trial",
        "pool": "trial",
        "rdata": "10.25.11.10",
        "rtype": "A"
      }
    ],
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "lr-policy",
    "freeform-tags": {},
    "health-check-monitor-id": null,
    "id": "ocid1.dnspolicy.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "rules": [
      {
        "cases": null,
        "default-answer-data": [
          {
            "answer-condition": "answer.isDisabled != true",
            "should-keep": true
          }
        ],
        "description": null,
        "rule-type": "FILTER"
      },
      {
        "cases": null,
        "default-answer-data": [
          {
            "answer-condition": "answer.name == 'server'",
            "value": 90
          },
          {
            "answer-condition": "answer.name == 'trial'",
            "value": 10
          }
        ],
        "description": null,
        "rule-type": "WEIGHTED"
      },
      {
        "cases": null,
        "default-count": 1,
        "description": null,
        "rule-type": "LIMIT"
      }
    ]
  }
}

```

```

    ],
    "template": "LOAD_BALANCE",
    "time-created": "2021-11-03T23:36:25.392833+00:00",
    "ttl": 30
  },
  "etag": "2c63fca5-f747-487e-b2f3-0ae5d6fe939c"
}

```

The load balancer steering policy is created and available for attaching to a DNS domain.

4.6.2 Creating an IP Prefix Steering Policy

An IP prefix steering policy dynamically routes DNS request traffic to different servers based on the originating IP prefix (for example, 172.16.1.0/24).

Using the UI

1. Open the Navigation Menu. Under DNS Zones, click Manage DNS.
2. From the list of DNS resources, click Steering Policies. The steering policies for that compartment are displayed.
3. Click Create Steering Policy.
4. Select IP Prefix Steering and supply the following properties:
 - **Name:** The name for the new steering policy.
 - **Policy TTL:** The Time To Live (TTL) for responses from the steering policy, in seconds. The maximum allowed value is 604800 (equal to 168 hours or 7 days).
5. In the Answer(s) box, supply the following properties:
 - **Name:** The name for response to requests sent to the new steering policy.
 - **Type:** The type of request and response. The choices are A, AAA, or CNAME.
 - **RData:** The zone record data to return for the query. It must match the type expected by the type chosen.
 - **Pool:** Select the IP address pool to use of the policy from the dropdown list.
 - **+Add Answer:** Click this box to add more answers to the requests received by the steering policy.
 - **Disabled:** This toggle determines if the IP prefix answer is enabled at creation or not. The default is enabled.
6. In the IP Prefix Steering Rules box, supply the following properties:
 - **+Add Rule:** Click this box to add rules to the IP prefix steering policy.
 - **Order:** Use the directional arrows to order the rule in the sequence of configured rules.
 - **Subnet Address:** Enter the IP subnet prefix to apply to this steering policy.
 - You can add more rules to this steering policy by clicking **+Add Rule**.
7. **Tagging:** Optionally, you can add tags to the steering policy.

For more information about tagging, see [Tagging Overview](#).

If you are not sure whether to apply tags, skip this option (you can apply tags later) or ask your administrator.

8. Click Save Changes. The IP prefix steering policy is created.

Using the CLI

1. Gather the information you need.
 - Compartment OCID (`oci iam compartment list --all`)
2. Run the `oci dns steering-policy create` command with the `ROUTE_BY_IP` parameter.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option. Complex types are long json strings.

Syntax (entered on a single line):

```
oci dns steering-policy create
--compartment-id <compartment_OCID>
--display-name <dns_steering_policy_name>
--template <ROUTE_BY_IP>
--answers <complex-type>
--rules <complex-type>
```

Example:

```
oci dns steering-policy create --compartment-id ocid1.compartment.....uniqueID
--display-name test-ip-steering-1
--template ROUTE_BY_IP
--answers file:///root/users-stuff/ip-steering-answers.json
--rules file:///root/users-stuff/ip-steering-rules-2.json
{
  "data": {
    "-self": "https://20180115/steeringPolicies/ocid1.dnspolicy.....uniqueID",
    "answers": [
      {
        "is-disabled": null,
        "name": "server",
        "pool": "server",
        "rdata": "10.20.10.10",
        "rtype": "A"
      },
      {
        "is-disabled": null,
        "name": "trial",
        "pool": "trial",
        "rdata": "10.20.10.10",
        "rtype": "A"
      }
    ],
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "test-ip-steering-1",
    "freeform-tags": {},
    "health-check-monitor-id": null,
    "id": "ocid1.dnspolicy.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "rules": [
      {
```

```

"cases": null,
"default-answer-data": [
  {
    "answer-condition": "answer.isDisabled != true",
    "should-keep": true
  }
],
"description": null,
"rule-type": "FILTER"
},
{
  "cases": [
    {
      "answer-data": [
        {
          "answer-condition": "answer.pool == 'internal'",
          "value": 1
        }
      ],
      "case-condition": "query.client.address in (subnet '10.0.3.0/24')"
    },
    {
      "answer-data": [
        {
          "answer-condition": "answer.pool == 'external'",
          "value": 1
        }
      ],
      "case-condition": null
    }
  ],
  "default-answer-data": null,
  "description": null,
  "rule-type": "PRIORITY"
},
{
  "cases": null,
  "default-count": 1,
  "description": null,
  "rule-type": "LIMIT"
}
],
"template": "ROUTE_BY_IP",
"time-created": "2021-11-09T16:53:34.963177+00:00",
"ttl": 30
},
"etag": "aad5bbcc-9d89-40cd-ab10-03dcc2e4ee0a"
}
    
```

The IP steering policy is created and available to attach to a DNS domain.

4.6.3 Editing a Steering Policy

Using the UI

1. Open the Navigation Menu. Click DNS, and then click Steering Policies.
2. For the policy that you want to update, click the Actions menu, and click the Edit option.
3. Make the necessary changes on the Edit Steering Policy dialog.
4. When you have finished making changes, click the Save Changes button on the dialog.

The details page for this steering policy is displayed with the updated information.

Using the CLI

1. Get the steering policy OCID.

Use the following command to list all of the steering policies in the specified compartment to get the OCID of the steering policy that you want to update:

```
# oci dns steering-policy list --compartment-id compartment_OCID
```

2. Run the steering policy update command.

Syntax:

```
oci dns steering-policy update --steering-policy-id steering_policy_OCID \  
options_with_values_to_update
```

Example:

This example shows replacing the `answers` block of the steering policy. You can also change the display name, health check monitor, rules or rules template, TTL, and scope.

```
# oci dns steering-policy update --steering-policy-id ocid1.dnspolicy.unique_ID \  
--answers file://answers.json
```

This command returns the same output as the `steering-policy get` command.

4.6.4 Moving a Steering Policy to a Different Compartment

Using the CLI

1. Get the following information:

- The OCID of the compartment where the steering policy is currently located, and the OCID of the compartment where you want to move the steering policy.

```
# oci iam compartment list [options]
```

- The steering policy OCID.

```
# oci dns steering-policy list --compartment-id current_compartment_OCID
```

2. Run the steering policy update command.

Syntax:

```
oci dns steering-policy change-compartment -c destination_compartment_OCID \  
--steering-policy-id steering_policy_OCID
```

This command returns the same output as the `steering-policy get` command. Verify the new `compartment-id`.

4.6.5 Attaching a Domain to a Steering Policy

A steering policy must be attached to a domain for the policy to answer DNS queries for that domain. The attachment is automatically placed into the same compartment as the domain's zone.

Using the UI

1. Open the Navigation Menu. Click DNS, and then click Steering Policies.

2. Click the name of the policy to which you want attach a domain.
3. Scroll to the Resources section and click Attached Domains.
4. In the list of attached domains, click the Add Attached Domain button.
5. In the Add Attached Domain dialog, enter the domain name and select a zone.
6. Click the Submit button.

The new domain is added to the Attached Domains list for this steering policy.

Using the CLI

1. Get the following information:
 - The steering policy OCID. Use the following command to list all of the steering policies in the specified compartment to get the OCID of the steering policy to which you want to attach a domain:

```
# oci dns steering-policy list --compartment-id current_compartment_OCID
```

- The name of the domain that you want to attach to the steering policy.
- The OCID of the attached zone. Use the following command to list all of the zones in the specified compartment to get the OCID of the zone where the domain that you want to attach is located:

```
# oci dns zone list compartment_OCID
```

2. Run the steering policy attachment create command.

Syntax:

```
oci dns steering-policy-attachment create --steering-policy-id steering_policy_OCID \  
--domain-name domain_name --zone-id zone_OCID
```

The value of the `--domain-name` argument is the attached domain within the attached zone specified in the `--zone-id` argument.

This command returns the same output as the `steering-policy-attachment get` command.

4.6.6 Editing an Attached Domain

Using the UI

1. Open the Navigation Menu. Click DNS, and then click Steering Policies.
2. Click the name of the policy for which you want to edit an attached domain.
3. Scroll to the Resources section and click Attached Domains.
4. Click the name of the attached domain that you want to edit.
5. On the top of the details page for the attached domain, click the Edit button.
6. Make the necessary changes on the Edit Steering Policy Attachment dialog.
7. When you have finished making changes, click the Save Changes button on the dialog.

The details page for this steering policy attachment is displayed with the updated information.

Using the CLI

1. Get the steering policy attachment OCID.

Use the following command to list all of the steering policy attachments in the specified compartment to get the OCID of the steering policy attachment that you want to update:

```
# oci dns steering-policy-attachment list --compartment-id compartment_OCID
```

2. Run the steering policy attachment update command.

Syntax:

```
oci dns steering-policy-attachment update \  
--steering-policy-attachment-id steering_policy_attachment_OCID
```

This command returns the same output as the `steering-policy-attachment get` command.

4.6.7 Deleting a Steering Policy Attachment

Using the UI

1. Open the Navigation Menu. Click DNS, and then click Steering Policies.
2. Click the name of the policy for which you want to delete an attachment.
3. Scroll to the Resources section and click Attached Domains.
4. For the attached domain that you want to delete, click the Actions menu, click the Delete option, and confirm the deletion.

The steering policy attachment is removed from the Attached Domains list.

Using the CLI

1. Get the steering policy attachment OCID.

Use the following command to list all of the steering policy attachments in the specified compartment to get the OCID of the steering policy attachment that you want to delete:

```
# oci dns steering-policy-attachment list --compartment-id compartment_OCID
```

2. Run the steering policy attachment delete command.

Syntax:

```
oci dns steering-policy-attachment delete \  
--steering-policy-attachment-id steering_policy_attachment_OCID
```

4.6.8 Deleting a Steering Policy

A policy that is attached to any zones cannot be deleted. To detach a policy from a zone, see [Section 4.6.7, “Deleting a Steering Policy Attachment”](#).

Using the UI

1. Open the Navigation Menu. Click DNS, and then click Steering Policies.
2. Click the name of the policy that you want to delete.

3. Scroll to the Resources section, click Attached Domains, and ensure that this policy has no attached domains.
4. Click the Delete button at the top of the steering policy details page, and confirm that you want to delete this steering policy.

The steering policies list page is displayed.

Using the CLI

1. Get the steering policy OCID.

Use the following command to list all of the steering policies in the specified compartment to get the OCID of the steering policy that you want to delete:

```
# oci dns steering-policy list --compartment-id compartment_OCID
```

2. Run the steering policy delete command.

Syntax:

```
oci dns steering-policy delete --steering-policy-id steering_policy_OCID
```

4.7 Networking Scenarios

All networking scenarios for a virtualized cloud environment are similar to scenarios for individual IP subnets connected by physical switches, routers, and gateways. In other words, virtual devices still have MAC (hardware) addresses as source and destination addresses in frames and IP addresses as source and destination addresses in packets.

Content delivery works essentially the same way as well. If the network portion of the source and destination IP addresses are in the same defined VCN subnet, delivery is through a logical switch (bridge) based on source and destination MAC frame addresses. If the network portion of the source and destination IP addresses are in different VCN subnets, then delivery is based on source and destination IP packet addresses.

You do not have to configure a logical switch for Oracle Private Cloud Appliance. The MAC addresses are known to all the other entities connected to the logical switch. It is assumed that if you place two or more VMs in the same VCN subnet, it is okay if they communicate. If VM isolation is the goal, then establish separate subnets or VCNs for them.

4.7.1 Logical Routers

Traffic between different VCN subnets is handled by at least one logical router, by definition. Devices that use IP packet addresses to determine forwarding steps, while using different MAC frame addresses on the different subnets they link, are called routers. In the case where the routers attach to the internet, these virtual devices are internet gateways (IGWs).

In cases where the source or destination IP address rules include IP network address translation (NAT), the packets are handled by a NAT gateway of some type (there are several types of NAT gateways). If some form of NAT is used, these are NAT gateways (NATGW or NGW).

In many cases of virtualized cloud networking, routing is very simple and can be handled by a small, static routing table that has a handful of destinations. More complex virtual environments require more complex logical routers, containing dynamic information that is updated periodically.

Logical routers inspect the packet's IP addresses. The IP addresses, source and destination, are looked up in a route table, called the *route rule table* in Oracle Private Cloud Appliance, and forwards the packet

to the next hop (another router) or destination (for local delivery) if the IP address rule allows this. If the route rules do not apply to the IP addresses, the packet is silently dropped and does not generate an error message (this is a security feature to prevent blocked probes from gathering information). However, this lack of error messages means that the route rules must be configured very carefully.

One IP address is special when it comes to route rules. This is the IPv4 address 0.0.0.0/0, which essentially matches any IP address at all. In some documents, the 0.0.0.0/0 is called an IP address CIDR block, but the same universal matching is true no matter what it is called.

For example, the following route rule allows a packet sent from inside the VCN to any IP address at all to reach a gateway to the internet:

Destination		Target
0.0.0.0/0	Internet Gateway	vcn-20210714-0910

Route rules do more than determine the destination of a packet. Route rules also form the basis for network firewalls.

4.7.2 Using Firewalls

Internet access is convenient, but brings concerns over vulnerability and security. Firewalls exist to limit the free passage of traffic between network elements and secure the network. A firewall, logical or physical, examines the traffic flow from a particular source to a particular destination and permits or blocks the packets based on the configured security rules in the route rule table.

Firewalls should be configured not only to allow or block traffic from or to external sources, but should also be configured to validate the traffic passing from subnet to subnet within the same VCN. Threats could be coming from external sources, but also from compromised VMs within the network.

4.7.3 Use of Network Segmentation

It is tempting to configure a virtual network as one big entity, with everything easily reachable by everything else. But this makes it relatively easy for attackers to compromise the network: once they are in, they are in everywhere. It is much better to use segmentation for the network and group resources and data into the various segments.

In Oracle Private Cloud Appliance, segments are essentially Tenant Groups.

Typically, you group data and resources into different Tenant Groups based on similarity or data sensitivity. For example, you can establish a group that examines all traffic received from the data center. Based on your security rules, this traffic can then be passed to a group of application servers, and then onto the database servers.

With this approach, firewalls between the groups secure the application and database servers from any compromised components of the data center.

4.7.4 Use of Tunneling

One complication of virtualized cloud networking is that there is no central authority to assign IP addresses to VCN subnets. Nothing stops one hypervisor from assigning, for example, IP address 192.168.1.6 to VM-1 in Subnet-1 of VCN-1, while another hypervisor assigns the same IP address 192.168.1.6 to VM-7 in Subnet-1 of VCN-1. Yet, if various tables are configured correctly, they can still communicate.

In order to effectively hide these network address complications, Oracle Private Cloud Appliance moves traffic between network components such as logical routers through IP tunnels. However, this use of tunnels, a common IP network practice, does not change the network scenarios. VMs still have IP

addresses, and these addresses are still assigned to a particular subnet and the subnets are assigned to virtual network interface cards (VNICs). All traffic from the various running VMs under the same a hypervisor is bonded together into an IP tunnel for transport to the next device between source and destination.

4.7.5 Use of Virtual Cloud Networks

It is easy to say that resources are gathered into one or more VCNs, which are private cloud networks running in a tenancy. But there is a lot more to VCNs than declaring a group of resources and creating a boundary for a VCN.

Planning your VCNs is a critical part of any deployment. VCNs serve as a foundation to structure your application servers, databases, and any other services provided. VCNs should take into account any needs such as redundancy, high availability, scalability, security, and more.

This section details the critical parts of a VCN.

4.7.5.1 IP Address Ranges

When planning VCNs, the first decision to make is which IP address CIDR block to use.

Note

To help with the calculation of CIDR blocks, a good resources is [IP Address Guide for CIDR](#)

The VCN network address range should be any VLSM between /16 and /30. This covers virtual networks for between 4 available IP addresses (/30) to 65,536 available IP addresses (/16), although the highest and lowest IP address are not useful for endpoint devices.

The size of the CIDR block chosen for the VCN is of critical importance. If the size is too large, then IP addresses are wasted that could be used in other places in the network. If the size is too small, the solution does not scale because there are not enough IP addresses for the VCN. You can always create another VCN and peer them together, but this is a complication that can be avoided through careful planning.

There are some important points about VCN CIDR blocks:

- VCNs should use one of the RFC1918 private address ranges: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16.
- VCNs use a contiguous IPv4 CIDR block. IPv6 is not supported.
- You cannot change the VCN and subnet sizes (such as 10.100.0.0/21) after their creation.
- The IP address range must not overlap with any VCN you want to peer with.
- Subnets within a VCN must not overlap with each other.
- Remember that the highest and lowest IP addresses in a range are reserved for network functions.

4.7.5.2 IP Subnets

A subnet is a subdivision of a VCN that uses a continuous IP address range as established by CIDR. Subnets group resources by IP address.

From the VCN perspective, subnets can be public or private. In the context of Oracle Private Cloud Appliance, private VCN addresses must be changed using NAT before they can interact with other VCNs

that might be using the same private address space. Public VCN addresses allow for connectivity to the data center network and are accessible from outside the rack.

It is, however, possible to have an RFC1918 address space subnetted into both public IP address spaces (which allows external data center connectivity) and private IP address spaces (which connect within the Oracle Private Cloud Appliance rack).

For example, it is possible for a VCN to do this with an IP CIDR block:

Subnet Name	Subnet Access	IP CIDR Block
Public_Subnet_01	Public	10.100.0.0/24
Private_Subnet_01	Private	10.100.1.0/24

Public_Subnet_01 has 256 IP addresses, from 10.100.0.0 to 10.100.0.255 (these two addresses are not often used for devices and have special uses in most IP networks). The netmask is 255.255.255.0

Private_Subnet_01 also has 256 IP addresses, from 10.100.1.0 to 10.100.1.255 (again, the low and high addresses are not often useful for devices). The netmask is 255.255.255.0.

Note that the two address ranges do not overlap. (If the public range was 10.100.0.0/23, then 256 devices would overlap with 10.100.1.0/24.)

4.7.5.3 Route Tables

A VCN uses a route table to hold the rules governing the sending and receiving of traffic by a VM. A VCN could be allowed or prevented from reaching destinations like:

- The global public internet
- An on-premise network, such as the data center network
- A peer VCN

Whenever a VCN is created, a virtual router with a default route table is also created.

For better security, it is preferable to create a dedicated route table for every subnet rather than use the default rules, which might not be adequate for the level of security needed. Dedicated route tables can more effectively manage the route rules that each subnet needs. For example, if the subnet is allowed to send traffic out onto the global public internet, then the route table for that subnet needs a rule for routing traffic to an internet gateway with the "universal destination" 0.0.0.0/0.

Note

The CIDR block 0.0.0.0/0 matches all addresses in the IPv4 address space. It means "any IPv4 address with any subnet mask."

A subnet not needing global public network access should not include that rule in the route table. In addition:

- Traffic with a source and destination within a VCN is not governed by any route table rules.
- If rules overlap (usually regarding different CIDR block or subnets), then the more specific rule applies (often seen as "the longest match").
- If there is no rule that applies to the traffic flow, then the traffic is silently dropped (no error message sent).

Every rule in the route table has a target. The targets are the functional network nodes for the common components of any network:

- Dynamic Routing Gateway (DRG: the route table rules are not statically configured, but use a routing protocol such as BGP to change them).
- Internet Gateway (IG) to connect to the global public internet.
- NAT gateway (NATG) for IP address translation.
- Service Gateway (SG) to reach a variety of services in other subnets.
- Local Peering Gateway (LPG) to connect to peer VCNs.
- Private IP address, which routes traffic to a specific instance inside the VCN.

As an example, a VCN could contain a route table with the following rules:

Subnet Name	Route Table	Target
Public_Subnet_01	Public_Subnet_01_Route_Table	IG
Private_Subnet_01	Private_Subnet_01_Route_Table	NAT Gateway

Route tables are the foundation of VCN security.

4.7.5.4 Security Lists and Network Security Groups

It might seem odd that an on-premises network needs firewall-like security. But security is important in every context, and not all threats come from outside an organization.

Oracle Private Cloud Appliance offers two security networking mechanisms that act like firewalls to control traffic at the packet level:

- Security lists, which are used like physical firewalls for subnets.
- Network security groups (NSGs), which act as firewalls for groups of instances across subnets.

You use a security list to define the rules that apply to all inbound (ingress) and outbound (egress) traffic of a subnet. You can associate up to five security lists per subnet. In the same way as route tables, there are default security lists and dedicated security lists. For better control and management, you should always use dedicated security lists for each subnet.

In contrast to security lists, NSGs let you build rules for groups of instances, even if the instances are in different subnets. For example, the same NSG can apply to all the database servers, or all the application servers running a certain application. Instead of applying security to a particular subnet, you create an NSG and then add the appropriate instances to the NSG.

There is no requirement to use either security lists or NSGs. You can use security lists without establishing NSGs, or create NSGs without creating any security lists. However, if you use both security lists and NSGs, the rules that apply to a VNIC are the union (both sets) of the rules that are in the security list for the VNIC and the rules specific to that VNIC from the NSG.

When creating a VCN, a default security list with three ingress rules and one egress rule is created. The default rules are stateful, which means that they know what connections are and that a request is followed by a response and that the rules should take this into consideration. In contrast, stateless rules are applied to every packet regardless of situation.

The default ingress and egress security rules look like this if you display them: First, the ingress rules (for a 10.0.0.0/16 CIDR block):

Stateless	Source	Source Type	IP Protocol	Source Port Range	Destination Port Range	Type and Code
No	10.0.0.0/16	CIDR Block	ICMP			3
No	0.0.0.0/0	CIDR Block	ICMP			3, 4
No	0.0.0.0/0	CIDR Block	TCP	22 - 22	22 - 22	

These rules, line by line, can be read as:

- There is a stateful rule that applies to traffic originating from the VCN 10.0.0.0/16 CIDR block using the ICMP protocol Type = 3 message format (Destination Unreachable) and all Codes. In other words, this rule allows devices in the 10.0.0.0/16 CIDR block to pass Destination Unreachable messages with any code (such as Net or Host Unreachable) back to the sender (another instance in the VCN subnet).
- There is a stateful rule that applies to traffic originating from any IP address (0.0.0.0/0 CIDR block) using the ICMP protocol Type = 3 message format (Destination Unreachable) and a Code = 4 (message must be fragmented, but the Do Not Fragment (DF) bit is set in the packet: these are Path MTU Discovery messages). In other words, this rule allows devices to notify sources that the DF bit is set on packets that need to be fragmented because the content exceeds the MTU size established for this VCN subnet.
- There is a stateful rule that applies to traffic originating from any IP address (0.0.0.0/0 CIDR block) using the connection-oriented TCP protocol and with the Source or Destination Port = 22 (SSH). In other words, this rule allows SSH for this VCN subnet.

This last rule makes allows you to create a new VCN and subnet, launch a Linux instance, and then use SSH to connect to that instance without writing any new security list rules.

Important

The default ingress security list does not include a rule to allow Remote Desktop Protocol (RDP) access. If you're using [Windows images](#), make sure to add a stateful ingress rule for TCP traffic on destination port 3389 from authorized source IP addresses and any source port. See [To enable RDP access](#) for more information.

The single egress rule is very simple:

Stateless	Source	Source Type	IP Protocol	Source Port Range	Destination Port Range	Type and Code
No	0.0.0.0/0	CIDR Block	All			

This can be read as: "There is a stateful rule that allows packets with any source address and for any IP protocol at all to leave the VCN." This basically says that anything can leave the VCN subnet without a problem.

The default security list comes with no stateless rules. However, you can always add or remove rules from the default security list.

To use these default rules in an NSG, let's give the default ingress and egress rules distinctive names, such as `Ingress_Security_List_Subnet01` and `Egress_Security_List_Subnet01`. Before you can add these rules to an NSG, you must first create the NSG. To create an NSG, you must give it a name and assign it to an existing compartment. You can also add tags, but these can be added later.

Chapter 5 Compute Images

Table of Contents

5.1 Accessing the Management Node Images	175
5.1.1 Initial User Account for Management Node Images	175
5.1.2 Importing Images from the Management Node (Direct Method)	176
5.1.3 Importing Images from the Management Node (Indirect Method)	176
5.1.4 Downloading an Image From the Management Node	177
5.1.5 Best Practices for Sharing an Image Across Tenancies	178
5.1.6 Importing an Image	178
5.1.7 Exporting an Image to Object Storage	180
5.2 Managing Images	182
5.2.1 Overview	182
5.2.2 Listing Images and Details	182
5.2.3 Creating an Image From an Instance	184
5.2.4 Editing the Image Name or Compatible Shapes	185
5.2.5 Moving an Image to a Different Compartment	187
5.2.6 Deleting an Image	187
5.3 Bring Your Own Image (BYOI)	188
5.3.1 Importing Custom Linux Images	188
5.3.2 Importing Custom Microsoft Windows Images	189

Before you can launch an instance, you need to have an image in a compartment you have access to.

Oracle Private Cloud Appliance provides these options for obtaining images:

- **Import images provided with Oracle Private Cloud Appliance:** These include Oracle Linux and Oracle Solaris images. See [Section 5.1, “Accessing the Management Node Images”](#).
- **Create custom images within the tenancy:** You can create a custom image of a compute instance's boot disk and use it to launch other compute instances. Instances you launch from your image include the customizations, configuration, and software installed when you created the image. See [Section 5.2.3, “Creating an Image From an Instance”](#)
- **Bring your own Images:** enables you to bring your own versions of operating systems to the cloud as long as the underlying hardware supports it. See [Section 5.3, “Bring Your Own Image \(BYOI\)”](#).

For conceptual information and important limitations, refer to *Compute Images* in the [Compute Instances](#) chapter of the [Oracle Private Cloud Appliance Concepts Guide](#).

5.1 Accessing the Management Node Images

Oracle Private Cloud Appliance includes Oracle Linux and Oracle Solaris images that you can import to your tenancies.

The images are located on management nodes. To make use of an image, you must import it from the management node into a compartment in a tenancy where you plan to launch instances. Then you select the image from the list of custom images during instance launch.

5.1.1 Initial User Account for Management Node Images

After you launch an instance from any of the images provided on the management node, you initially connect to the instance using `ssh` with this initial user account:

opc

The `ssh` connection is authenticated using your ssh key pair that is used during instance launch. For more information, see [Section 6.4, “Connecting to a Compute Instance”](#).

5.1.2 Importing Images from the Management Node (Direct Method)

This method imports an image directly into a tenancy from the https-hosted images on the management node. The advantage of this method is that it requires only a few steps and imports the image directly into a tenancy. Alternatively, you can import images to an Object Storage bucket using [Section 5.1.3, “Importing Images from the Management Node \(Indirect Method\)”](#)

1. Determine the URL of the image you plan to import.

```
Syntax:
https://<mgmt_vip_hostname>.<system_name>.<domain_name>:8079/images/<image_file_name>

Example:
https://mnvip.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.01.12_0.oci
```

where:

- `<mgmt_vip_hostname>.<system_name>.<domain_name>` is the fully qualified domain name for the management node VIP.

This information is available in the Service Enclave. If needed, consult with your Service Enclave administrator.

- `<image_file_name>` is the name of image file.

The image file names change as new versions of the OS are released. To find the specific image names for your appliance, refer to *Known Issues and Workarounds, Platform Issues*, in the [Oracle Private Cloud Appliance Release Notes](#).

2. Import the image.

See [Section 5.1.6, “Importing an Image”](#).

When you import, whether using the Compute Web UI or OCI CLI, select these import options:

- Import from an Object Storage URL.
- Include the image URL obtained in the previous step.

OCI CLI Example:

```
oci compute image import from-object-uri \
--uri https://mnvip.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.01.12_0.oci \
--display-name "Oracle Linux 8"
```

5.1.3 Importing Images from the Management Node (Indirect Method)

This method downloads the image to a local system, then uploads the image to an Object Storage bucket. The advantage of this method is that you can implement object versioning or pre-authenticated requests (for details see [Section 9.4, “Managing Object Versioning”](#) and [Section 9.5, “Using Pre-Authenticated](#)

Requests”). Alternatively, you can import images directly into a tenancy. See [Section 5.1.2, “Importing Images from the Management Node \(Direct Method\)”](#).

No.	Task	Links
1.	Download any of the images that are on management nodes to your local system.	Section 5.1.4, “Downloading an Image From the Management Node”
2.	Create an Object Storage bucket that will contain the image.	Section 9.2.5, “Creating a Bucket”
3.	Upload the image to the bucket.	Section 9.3.4, “Uploading an Object”
4.	Import the image so that it is available to use when launching instances.	Section 5.1.6, “Importing an Image”

5.1.4 Downloading an Image From the Management Node

You can download any of the Oracle Private Cloud Appliance images that are on management node to your local system.

Downloading an image provides these benefits:

- Provides a backup copy of the image.
- Enables you to upload the image to an Object Storage bucket where the image is available to tenancies for launching instances. See [Section 5.1.3, “Importing Images from the Management Node \(Indirect Method\)”](#).

1. Obtain this information:

- Determine the URL of the image you plan to download.

```
Syntax:
https://<mgmt_node_VIP_hostname>.<system_name>.<domain_name>:8079/images/<name_of_image_file>
Example:
https://mnvip.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.01.12_0.oci
```

where:

- `<mgmt_vip_hostname>.<system_name>.<domain_name>` is the fully qualified domain name for the management node VIP.

This information is available in the Service Enclave. If needed, consult with your Service Enclave administrator.

- `<image_file_name>` is the name of image file.

The image file names change as new versions of the OS are released. To find the specific image names for your appliance, refer to *Known Issues and Workarounds, Platform Issues*, in the [Oracle Private Cloud Appliance Release Notes](#).

Caution

The image file sizes are large. They range from 1.3 to 1.8 GB. Ensure that the system where you plan to download the images has enough space.

2. On the system you plan to download the image to, create a directory.

Example:

```
# sudo mkdir /root/images
```

3. Download an image from the management node to your local system.

Syntax:

```
wget --no-check-certificate -P <directory_to_download_to> https://<mgmt_node_VIP_hostname>.<system_name>.<co
```

Example:

```
wget --no-check-certificate -P . https://mnvip.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-
--2022-01-26 09:03:10-- https://mnvip.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.0
Resolving mnvip.myprivatecloud.example.com (mnvip.myprivatecloud.example.com)... 192.0.2.0
Connecting to mnvip.myprivatecloud.example.com (mnvip.myprivatecloud.example.com)|10.134.199.8|:8079... con
WARNING: cannot verify mnvip.myprivatecloud.example.com's certificate, issued by '/C=US/O=Oracle/CN=PCA Int
Self-signed certificate encountered.
WARNING: no certificate subject alternative name matches
requested host name 'mnvip.myprivatecloud.example.com'.
HTTP request sent, awaiting response... 200 OK
Length: 1727866880 (1.6G) [application/x-tar]
Saving to: './uln-pca-Oracle-Linux-8-2022.01.12_0.oci'

100%[=====
2022-01-26 09:03:14 (422 MB/s) - './uln-pca-Oracle-Linux-8-2022.01.12_0.oci' saved [1727866880/1727866880]
```

4. Verify that the image downloaded.

Example:

```
sudo ls -l /root/images
total 1687380
-rw-r--r-- 1 root root 1727866880 Jan 14 03:44 uln-pca-Oracle-Linux-8-2022.01.12_0.oci
```

5.1.5 Best Practices for Sharing an Image Across Tenancies

You can use image import/export to share images across tenancies, so you don't need to recreate the image manually in each tenancy. You must go through the steps required to create the image in one of the tenancies, but after this is done, you can export the image, making it available for import in additional tenancies.

These are the high-level tasks:

1. Export the image to an Object Storage bucket. For steps, see [Section 5.1.7, “Exporting an Image to Object Storage”](#).
2. Create a pre-authenticated request with read-only access for the image in the bucket. For steps, see [Section 9.5, “Using Pre-Authenticated Requests”](#).
3. In the destination tenancy, import the image. Use the pre-authenticated request URL as the Object Storage URL. For steps, see [Section 5.1.6, “Importing an Image”](#).

5.1.6 Importing an Image

You can import images into a tenancy or compartment in a tenancy from an Object Storage bucket or from a URL.

Images imported from the object store are always imported as QCOW2 images.

Before You Begin

Ensure that the bucket or URL contains the image. See [Section 9.2, “Managing Buckets”](#).

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Custom Images.
2. Click Import Image.
3. Enter the required information:
 - **Name:** Specify a name for the imported image.
 - **Create in Compartment:** Select the compartment where the image will be placed.
 - **Source Type:** Select one of the following:
 - **Import from an Object Storage bucket:** Select a bucket, then select the image from the object name menu.
 - **Import from an Object Storage URL:** Enter a URL.

Note – The URL does not need to be an Object Storage URL. It can be any URL that provides access to the image.
 - **Image Type:** Select one of the following options based on the type of image you are importing.
 - **VMDK:** Virtual machine disk file format (`.vmdk`), used for virtual machine disk images.
 - **QCOW2:** For disk image files (`.qcow2`) used by QEMU copy on write, and for OCI images (`.oci`).
 - **Launch Mode:** Paravirtualized is the default and cannot be changed.
 - **Tagging:** Optionally, add one or more tags to this image as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.
4. Click Import Image.

After you click Import Image, you'll see the imported image in the Custom Images list for the compartment, with a state of Importing.

To track the progress of the operation and troubleshoot errors that occur during instance creation, use the associated work request.

When the import completes successfully, the state changes to Available, and you can launch instances with it. See [Section 6.2.1, “Creating an Instance”](#)

If the state does not change, or no entry appears in the Custom Images list, the import failed. Ensure you have read access to the Object Storage object, and that the object contains a supported image.

Using the OCI CLI (Importing from a Bucket)

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
 - Object Storage bucket name (`oci os bucket list --compartment-id <compartment_OCID>`)
 - Name of the object identifying the image (`oci os object list --bucket-name <bucket_name>`)

- Object Storage namespace. See [Section 9.2.2, “Obtaining the Object Storage Namespace”](#).
2. Run this command.

Syntax (entered on a single line):

```
oci compute image import from-object
--compartment-id <compartment_OCID>
--bucket-name <bucket_name>
--name <image_object_name>
--namespace <namespace>
```

Using the OCI CLI (Importing From a URL)

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)
- The URL for the image.

The URL does not need to be an Object Storage URL. It can be any URL that provides access to the image.

2. Run this command.

Syntax (entered on a single line):

```
oci compute image import from-object-uri
--compartment-id <compartment_OCID>
--uri <URL_for_image>
```

5.1.7 Exporting an Image to Object Storage

You can export images to an Object Storage bucket or URL. You need write access to the to the export location.

Exported images are a copy of the boot volume and metadata when the image was created.

Perform one of the following procedures:

Export to an Object (OCI CLI)

1. Ensure that a bucket is available.

See [Section 9.2.3, “Listing Buckets”](#) and [Section 9.2.5, “Creating a Bucket”](#).

2. Gather the information that you need to run the command:

- Object Storage bucket name (`oci os bucket list --compartment-id <compartment_OCID>`)
- Image OCID (`oci compute image list --compartment-id <Compartment_OCID>`)
- Object Storage namespace. See [Section 9.2.2, “Obtaining the Object Storage Namespace”](#).
- The name you want to apply to the exported image.

The name of the exported image is in the following form. Specify the `.qcow2` extension.

```
namespace , bucketname , file/objectname . qcow2
```

3. Run this command.

Syntax (entered on a single line):

```
oci compute image export to-object
--bucket-name <bucket_name>
--image-id <image_OCID>
--namespace <namespace>
--name <exported_image_name>
```

Export to a URL (OCI CLI)

1. Ensure that a bucket with a pre-authenticated request is available, and that you have the request URL.

See:

- [Section 9.2.5, “Creating a Bucket”](#)
- [Section 9.5.3, “Creating a Pre-Authenticated Request for All Objects in a Bucket”](#)
- [Section 9.5.5, “Constructing the Pre-Authenticated Request URL”](#)

2. Gather the information that you need to run the command:

- Image OCID (`oci compute image list --compartment-id <Compartment_OCID>`)

3. Run this command.

Syntax (entered on a single line):

```
oci compute image export to-object-uri
--image-id <image_OCID>
--uri <URL_to_export_to>
```

Example:

Note – In some cases, you need to omit the slash at the end of the `<access-uri>` string. Refer to the [Oracle Private Cloud Appliance Release Notes](#).

```
oci compute image export to-object-uri
--image-id ocidl.image.....uniqueID
--uri https://objectstorage.mypca01.us.example.com/oci/p/MrxLFkKlFkIlNDhvhcZnrjBUAlsoeah/n/mynamespace
{
  "data": {
    "agent-features": null,
    "base-image-id": null,
    "compartment-id": "ocidl.tenancy.....uniqueID",
    "create-image-allowed": true,
    "defined-tags": null,
    "display-name": "PCA OL8 Image",
    "freeform-tags": null,
    "id": "ocidl.image.....uniqueID",
    "launch-mode": "PARAVIRTUALIZED",
    "launch-options": {
      "boot-volume-type": "PARAVIRTUALIZED",
      "firmware": "UEFI_64",
      "is-consistent-volume-naming-enabled": false,
      "is-pv-encryption-in-transit-enabled": false,
      "network-type": "PARAVIRTUALIZED",
```

```

    "remote-data-volume-type": "PARAVIRTUALIZED"
  },
  "lifecycle-state": "EXPORTING",
  "listing-type": null,
  "operating-system": "OracleLinux",
  "operating-system-version": "8",
  "size-in-mbs": 47694,
  "time-created": "2022-01-18T16:29:13.114742+00:00"
},
"etag": "5d24f645-b446-42f2-a777-112457f0cafe",
"opc-work-request-id": "ocidl.workrequest.AK00661530.scasg01.stor0ftb433j7tit"
}

```

5.2 Managing Images

5.2.1 Overview

You can create a custom image of a compute instance's boot disk and use it to launch other compute instances. Instances you launch from your image include the customizations, configuration, and software installed when you created the image.

Custom images do not include the data from any attached block volumes.

There are limitations and considerations. For example, custom images inherit the compatible shapes that are set by default from the base image. For additional details, refer to *Custom Images Created From Instances* in the [Compute Instances](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

5.2.2 Listing Images and Details

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Custom Images.
2. If needed, select the appropriate compartment.

The list of custom images is displayed.

3. To see the details, click the custom image that you are interested in.

The custom image details are displayed.

Using the OCI CLI

Listing All Images

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci compute image list
--compartment-id <Compartment_OCID>
```

Example:

```
oci compute image list \
--compartment-id ocidl.tenancy.....uniqueID
```



```
{
  "data": [
    {
      "agent-features": null,
      "base-image-id": "ocidl.bootvolume.....uniqueID",
      "billable-size-in-gbs": null,
      "compartment-id": "ocidl.tenancy.....uniqueID",
      "create-image-allowed": true,
      "defined-tags": {},
      "display-name": "Linux 8",
      "freeform-tags": {},
      "id": "ocidl.image.....uniqueID",
      "launch-mode": "PARAVIRTUALIZED",
      "launch-options": {
        "boot-volume-type": "PARAVIRTUALIZED",
        "firmware": "BIOS",
        "is-consistent-volume-naming-enabled": false,
        "is-pv-encryption-in-transit-enabled": false,
        "network-type": "PARAVIRTUALIZED",
        "remote-data-volume-type": "PARAVIRTUALIZED"
      },
      "lifecycle-state": "AVAILABLE",
      "listing-type": null,
      "operating-system": "CUSTOM",
      "operating-system-version": "CUSTOM",
      "size-in-mbs": 51200,
      "time-created": "2021-09-17T18:26:03.221604+00:00"
    },
  ],
}
```

Listing Image Details

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci compute image list
--compartment-id <Compartment_OCID>
```

Example:

```
oci compute image list \
--compartment-id ocidl.tenancy.....uniqueID
{
  "data": [
    {
      "agent-features": null,
      "base-image-id": "ocidl.bootvolume.....uniqueID",
      "billable-size-in-gbs": null,
      "compartment-id": "ocidl.tenancy.....uniqueID",
      "create-image-allowed": true,
      "defined-tags": {},
      "display-name": "Linux 8",
      "freeform-tags": {},
      "id": "ocidl.image.....uniqueID",
      "launch-mode": "PARAVIRTUALIZED",
      "launch-options": {
        "boot-volume-type": "PARAVIRTUALIZED",
        "firmware": "BIOS",
        "is-consistent-volume-naming-enabled": false,
        "is-pv-encryption-in-transit-enabled": false,

```

```
    "network-type": "PARAVIRTUALIZED",
    "remote-data-volume-type": "PARAVIRTUALIZED"
  },
  "lifecycle-state": "AVAILABLE",
  "listing-type": null,
  "operating-system": "CUSTOM",
  "operating-system-version": "CUSTOM",
  "size-in-mbs": 51200,
  "time-created": "2021-09-17T18:26:03.221604+00:00"
},
"etag": "31da36ff-b0d6-4327-a94e-fd3f2c3b41ec"
}
}
```

5.2.3 Creating an Image From an Instance

You can create a custom image of an instance's boot disk and use it to launch other instances. Instances you launch from your image include the customizations, configuration, and software installed on the boot disk when you created the image.

The instance you use as the basis, must be in a stopped state, as described in the following procedure.

Once the custom image reaches the Available state, you can use it to launch new instances the same way you launch any other instance. See [Section 6.2.1, "Creating an Instance"](#).

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the compartment where the source instance is located.
3. Click the name of instance that you want to use as the basis for the custom image.
4. Click Controls then click Stop.

Wait for the status to change to Stopped. The status is displayed above the icon of the object.

5. Click Controls then click Create Custom Image.
6. Enter this information in the dialog box:
 - **Name:** Replace the name with the name you want for the image.
 - **Create in Compartment:** Optionally, change the compartment where the image will be stored.
7. Click Create Custom Image.

The status of the instance changes to Creating Image. Creating a custom image takes a while to complete. The duration depends on the size of the instance's boot volume.

8. To see the progress, click Compute, then click Custom Images. Click on the image name. Under Resources click Work Requests.

If you don't see the custom image, ensure that the correct compartment is selected.

Once the instance status changes from Creating Image to Stopped, you can restart the instance.

Using the OCI CLI

1. Stop the instance that will provide the boot volume image as the basis for the custom image.

See [Section 6.2.4, “Stopping, Starting and Resetting an Instance”](#)

2. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
 - The OCID of the instance you want to use as the basis for the image. (`oci compute instance list`)
 - Display name of your choice
3. Run this command.

Syntax (entered on a single line):

```
oci compute image create
--compartment-id <Compartment_OCID>
--instance-id <base_instance_OCID>
--display-name <display_name>
```

Example:

```
oci compute image create \
--compartment-id ocid1.tenancy.....uniqueID \
--instance-id ocid1.instance.....uniqueID \
--display-name Linux v8 image
{
  "data": {
    "agent-features": null,
    "base-image-id": "ocid1.bootvolume.....uniqueID",
    "billable-size-in-gbs": null,
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "create-image-allowed": true,
    "defined-tags": {},
    "display-name": "Linux 7 image",
    "freeform-tags": {},
    "id": "ocid1.image.....uniqueID",
    "launch-mode": "PARAVIRTUALIZED",
    "launch-options": null,
    "lifecycle-state": "PROVISIONING",
    "listing-type": null,
    "operating-system": "Custom",
    "operating-system-version": "Custom",
    "size-in-mbs": 0,
    "time-created": "2021-09-17T18:26:03.221604+00:00"
  },
  "etag": "3c0e56a0-b58c-486b-b659-9f5b13f377ee",
  "opc-work-request-id": "ocid1.workrequest.....uniqueID"
}
```

Consider your next action:

- Create a virtual cloud network (VCN) for your instances. See [Section 4.1.1, “Creating a VCN”](#).
- Create an instance. See [Section 6.2.1, “Creating an Instance”](#).

5.2.4 Editing the Image Name or Compatible Shapes

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Custom Images.

2. If needed, select the appropriate compartment.
3. Click the custom image that you are interested in.
The custom image details are displayed.
4. Click Controls, then click Edit Details.
5. In the dialog box, change the name or the shapes supported by the image.
6. Click Save Changes.
7. If you added support for a shape, test the image on the shape.

Some images (especially Windows) might not be cross-compatible with other shapes because of driver or hardware differences.

Using the OCI CLI

Use this command to change the display name of the custom image.

1. Gather the information that you need to run the command:
 - Image OCID (`oci compute image list`)
2. Run this command.

Syntax (entered on a single line):

```
oci compute image update
--image-id <image_OCID>
--display-name <new-name>
```

Example:

```
oci compute image create \
--image-id ocidl.image.....uniqueID \
--display-name "Image version 2.1"
{
  "data": {
    "agent-features": null,
    "base-image-id": "ocidl.bootvolume.....uniqueID",
    "billable-size-in-gbs": null,
    "compartment-id": "ocidl.tenancy.....uniqueID",
    "create-image-allowed": true,
    "defined-tags": {},
    "display-name": "Linux v7",
    "freeform-tags": {},
    "id": "ocidl.image.....uniqueID",
    "launch-mode": "PARAVIRTUALIZED",
    "launch-options": {
      "boot-volume-type": "PARAVIRTUALIZED",
      "firmware": "BIOS",
      "is-consistent-volume-naming-enabled": false,
      "is-pv-encryption-in-transit-enabled": false,
      "network-type": "PARAVIRTUALIZED",
      "remote-data-volume-type": "PARAVIRTUALIZED"
    },
    "lifecycle-state": "PROVISIONING",
    "listing-type": null,
    "operating-system": "CUSTOM",
    "operating-system-version": "CUSTOM",
    "size-in-mbs": 51200,
```

```

    "time-created": "2021-09-17T18:26:03.221604+00:00"
  },
  "etag": "b91737ca-95fe-4c4d-9454-bd390f99535f"
}

```

5.2.5 Moving an Image to a Different Compartment

To move an image, you must use the OCI CLI.

Using the OCI CLI

1. Get the following information:

- The OCID of the current compartment, and the OCID of the destination compartment:

```
# oci iam compartment list [options]
```

- The OCID of the image that you want to move:

```
# oci compute image list --compartment-id current_compartment_OCID
```

2. Run the image change compartment command.

Syntax:

```
oci compute image change-compartment \
--compartment-id destination_compartment_OCID \
--image-id image_OCID
```

5.2.6 Deleting an Image

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Custom Images.
2. If needed, select the appropriate compartment.
3. For the image you want to delete, click the Actions menu (three dots), then click Delete image.

The custom image is deleted.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Image OCID (`oci compute image list`)

2. Run this command.

Syntax (entered on a single line):

```
oci compute image delete
--image-id <image_OCID>
```

Example:

```
oci compute image delete --image-id ocid1.image.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
{
  "etag": "bbb9a3df-8f9d-47df-a419-f9d2de912b57",
  "opc-work-request-id": "ocid1.workrequest.1742XC3024.broom14.storage-oyk39vma9bo496z9po8wbh9ori00us5v"
}
```

5.3 Bring Your Own Image (BYOI)

The Bring Your Own Image (BYOI) feature enables you to bring your own versions of operating systems to the appliance as long as the underlying hardware supports it. The Oracle Private Cloud Appliance services do not depend on the OS you run.

Important

You must comply with all licensing requirements when you upload and start instances based on OS images that you supply.

For more conceptual information, refer to *Bring Your Own Image (BYOI)* in the [Compute Instances](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

5.3.1 Importing Custom Linux Images

5.3.1.1 Preparing Linux VMs for Import

Before you import a custom Linux image, you must prepare the image to ensure that instances launched from the image can boot correctly and that network connections will work.

Perform these steps.

1. Review the requirements.

Refer to *Linux Source Image Requirements* in the [Compute Instances](#) chapter of the [Oracle Private Cloud Appliance Concepts Guide](#).

2. Create a backup of the root volume.
3. If the VM has remotely attached storage, such as NFS or block volumes, configure any services that rely on this storage to start manually. Remotely attached storage is not available the first time that an imported instance boots on the appliance.
4. Ensure that all network interfaces use DHCP, and that the MAC address and IP addresses are not hardcoded. See your system documentation for steps to perform network configuration for your system.
5. Stop the VM.
6. Clone the stopped VM as a VMDK or QCOW2 file, and then export the image from your virtualization environment.

Refer to the tools documentation for your virtualization environment.

5.3.1.2 Importing a Linux-Based Image

After you prepare a Linux image for import, follow these steps to import the image:

1. Upload the image file to an Object Storage bucket.

Ensure that you select a bucket where you have read and write access.

See [Section 5.1.7, "Exporting an Image to Object Storage"](#).

2. Import the image from the bucket to your tenancy.

See [Section 5.1.6, “Importing an Image”](#)

3. Complete the post-import tasks.

See [Section 5.3.1.3, “Post-Import Tasks for Linux Images”](#).

5.3.1.3 Post-Import Tasks for Linux Images

After you import a custom Linux-based image, perform these steps.

1. Create an instance based on the custom image.

See [Section 6.2.1, “Creating an Instance”](#).

2. If the instance requires any remotely attached storage, such as block volumes create and attach it.

See [Section 7.2, “Creating and Attaching Block Volumes”](#)

3. Create and attach any required secondary VNICs.

See [Section 4.4, “Configuring VNICs and IP Addressing”](#)

4. Test that all applications are working as expected.

5. Reconfigure any services that were set to start manually.

5.3.2 Importing Custom Microsoft Windows Images

The Compute service enables you to import Microsoft Windows images and use them to launch instances. You can import images that you create from Microsoft Windows systems that are running on your on-premises physical or virtual machines (VMs).

Perform the procedures in this section to prepare, create, export, import, and perform post-import tasks.

5.3.2.1 Preparing Microsoft Windows Systems for Import

The configuration described in this section is required so that Compute instances that are launched from the Microsoft Windows system image can boot correctly and network connections will work.

Important

The system drive configuration where the Microsoft Windows source system is installed will be imported to the image. All partitions on the drive will follow through the imported image. Any other drives will not be imported, and you must re-create them on the instance after they are launched from the image. You will then need to manually move the data on the non-system drives to storage on the instance.

You can perform this configuration on the running source system or after you have launched the Compute instance.

- [Preparing the Source System Prior to Creating the Image](#). This is the recommended method.
- [Preparing the Compute Instance After Instance Launch](#). If you have concerns about modifying the live source system, you can use this method. If you use this method, your Compute instance is not initially viable. After you launch your Compute instance, connect to the VNC console and use the VNC window to make the changes described in [Preparing the Source System Prior to Creating the Image](#).

Preparing the Source System Prior to Creating the Image

1. Review the requirements.

Refer to *Windows Source Image Requirements* in the [Compute Instances](#) chapter of the [Oracle Private Cloud Appliance Concepts Guide](#).

2. Follow your organization's security guidelines to ensure that the Microsoft Windows system is secured. This can include, but is not limited to the following tasks:
 - Install the latest security updates for the operating system and installed applications.
 - Enable the firewall, and configure it so that you only enable the rules that are needed.
 - Disable unnecessary privileged accounts.
 - Use strong passwords for all accounts.

3. Configure Remote Desktop Protocol (RDP) access to the image.

- a. Enable Remote Desktop connections to the image.

See [Section 6.4.4.1, "Enabling Remote Desktop Protocol \(RDP\) Access"](#).

- b. Modify the Microsoft Windows Firewall inbound port rule to allow RDP access for both Private and Public network location types. When you import the image, the Microsoft Windows Network Location Awareness service will identify the network connection as a Public network type.

4. Determine whether the current Microsoft Windows license type is a volume license by running the following command in PowerShell:

```
Get-CimInstance -ClassName SoftwareLicensingProduct | where {$_.PartialProductKey} | select ProductKeyChanr
```

If the license is not a volume license, after you import the image, you will update the license type.

5. If you plan to use this custom image to launch more than one instance, create a generalized image of the boot disk. A generalized image is cleaned of computer-specific information, such as unique identifiers. When you create instances from a generalized image, the unique identifiers are regenerated. This prevents two instances that are created from the same image from colliding on the same identifiers.
6. Create a backup of the root volume.
7. If the system has remotely attached storage, such as NFS or block volumes, configure any services that rely on this storage to start manually. Remotely attached storage is not available the first time an instance that was created from a custom image boots on Oracle Private Cloud Appliance.
8. Ensure that all network interfaces use DHCP, and that the MAC address and IP addresses are not hardcoded. See your system documentation for steps to perform network configuration for your system.
9. Install the Oracle VirtIO Drivers for Microsoft Windows.
 - a. [Downloading the Oracle VirtIO Drivers for Microsoft Windows](#)
 - b. [Installing the Oracle VirtIO Drivers for Microsoft Windows](#)
10. Perform the [Creating and Exporting an Image](#) procedure unless you already followed the [Preparing the Compute Instance After Instance Launch](#) procedure.

Creating and Exporting an Image

1. Stop the system.
2. Clone the stopped system as a VMDK or QCOW2 file.
Refer to the tools documentation for your system.
3. Export the image from your physical system or virtualization environment.
4. Perform the [Section 5.3.2.2, "Importing a Microsoft Windows Image"](#) procedure to import the image into Oracle Private Cloud Appliance.

Preparing the Compute Instance After Instance Launch

1. Perform as many of the [Preparing the Source System Prior to Creating the Image](#) steps as you are comfortable performing.
2. Perform the [Creating and Exporting an Image](#) procedure.
After importing the image, do *not* perform the [Section 5.3.2.3, "Post-Import Tasks for Microsoft Windows Images"](#) procedure.
3. Use the imported image to launch an instance.
For the image source, select Custom Images, and then select the image that you imported. See [Section 6.2.1, "Creating an Instance"](#).
4. Connect to the console as described in [Section 6.4.5, "Connecting to an Instance Using a Console Connection"](#).
5. Perform the [Preparing the Source System Prior to Creating the Image](#) procedure.
6. Perform the [Section 5.3.2.3, "Post-Import Tasks for Microsoft Windows Images"](#) procedure.

Downloading the Oracle VirtIO Drivers for Microsoft Windows

The Oracle VirtIO Drivers for Microsoft Windows are paravirtualized drivers for Microsoft Windows instances. These drivers improve performance for network and block (disk) devices on Microsoft Windows instances and resolve common issues.

Download the Oracle VirtIO Drivers for Microsoft Windows from [Oracle Software Delivery Cloud](#) or [My Oracle Support \(MOS\)](#).

Download the Oracle VirtIO Drivers for Microsoft Windows from Oracle Software Delivery Cloud

1. Sign in to the [Oracle Software Delivery Cloud](#) site.
2. In the All Categories list, select Release.
3. Type Oracle Linux 7.9 in the search box and click Search.
4. Click "REL: Oracle Linux 7.9.0.0.0" to add it to your cart.
5. At the top right of the page, to the right of your cart, click Continue.
6. In the Platforms/Languages list, select x86 64 bit. Click Continue.
7. Review and accept the license agreement (click "I reviewed and accept the Oracle License Agreement."). Click Continue.

8. Click the [v1009702-01.zip](#) filename to the left of "Oracle VirtIO Drivers Version for Microsoft Windows 1.1.7, 67.9 MB".
9. Follow the prompts to save the [v1009702-01.zip](#) file.

Download the Oracle VirtIO Drivers for Microsoft Windows from MOS

1. Sign in to [MOS](#).
2. Click the Patches & Updates tab.
3. In the Patch Search pane, in the Patch Name or Number field, enter [27637937](#). Click the Search button.
4. From the search results table, click the Patch Name to the left of "Oracle VirtIO driver version 1.1.7" for Release 7.9.0.0.0.

A more detailed description of the patch is shown.
5. In the box, click the Download button.
6. In the File Download window, follow the prompts to save the [p27637937_79000_MSWIN-x86-64.zip](#) file.

Installing the Oracle VirtIO Drivers for Microsoft Windows

To install the Oracle VirtIO Drivers for Microsoft Windows, configure Microsoft Windows policies and then run the installation program.

Configuring Policies for Device Installation

Configure Microsoft Windows policies to allow the installation of the Oracle VirtIO Drivers for Microsoft Windows, if these policies are not already configured.

1. Go to the Microsoft Windows system on which you want to install the Oracle VirtIO Drivers for Microsoft Windows.
2. From the Start menu, select Run.
3. Enter [gpedit.msc](#) and then click OK.

The Local Group Policy Editor is displayed.
4. From the Console Tree, display the list of Device Installation Restrictions as follows:
 - a. Expand Computer Configuration, and then expand Administrative Templates.
 - b. Expand System, and then expand Device Installation.
 - c. Select Device Installation Restrictions.
5. Edit the policy settings so that no device installation restrictions are configured.
6. Close the Local Group Policy Editor.
7. Restart the Microsoft Windows system.

After performing one of the procedures described in [Downloading the Oracle VirtIO Drivers for Microsoft Windows](#), the Microsoft Windows system should have a copy of the Oracle VirtIO Drivers for Microsoft Windows installation program, [Setup.exe](#).

You can use a graphical user interface (GUI) to install the drivers, or use the CLI to install the drivers by using a response file that you previously created.

The Oracle VirtIO Drivers for Microsoft Windows are installed in the following directories:

- On 32-bit systems: `C:\Program Files\Oracle Corporation\Oracle Windows VirtIO Drivers`
- On 64-bit systems: `C:\Program Files (x86)\Oracle Corporation\Oracle Windows VirtIO Drivers`

Installing the Oracle VirtIO Drivers for Microsoft Windows by Using the GUI

This procedure installs the drivers on a single Microsoft Windows system. You can optionally record your responses for use on other systems.

1. Run the `Setup.exe` driver installation program.
 - To install the drivers on only this system, double-click the `Setup.exe` file.
 - To record a response file for use on other systems, start the `Setup.exe` installer from the command line.
 - a. Open a command-line window.
 - b. Navigate to the directory where the `Setup.exe` file is located.
 - c. Run `Setup.exe -r` to start the installer and create a response file.

2. If prompted, select Yes in the User Account Control dialog to allow the installer to proceed.

The Welcome window is displayed.

3. Click Next.

The "Start to install Oracle VirtIO Drivers for Microsoft Windows Release 2.0" window is displayed with information about your selection.

4. Click Install to start the installation.

The installer copies the Oracle VirtIO Drivers for Microsoft Windows files and installs the drivers on the system.

5. Once the installation completes, click Finish.

The system is restarted.

Installing the Oracle VirtIO Drivers for Microsoft Windows by Using an Existing Response File

This procedure uses a response file that was created in the [Installing the Oracle VirtIO Drivers for Microsoft Windows by Using the GUI](#) procedure.

1. Locate the response file, `setup.iss`, in the `C:\Windows` directory.
2. Copy the response file to the same directory where the Oracle VirtIO Drivers for Microsoft Windows installation program, `Setup.exe`, is located.

Alternatively, you can specify the location of the response file at the command line.

3. Open a command-line window.
4. Run `Setup.exe -s` to install the drivers by using the response file.

The following additional options to the `Setup.exe -s` command are available:

- `-f1c:path_to\setup.iss` to specify the location of the `setup.iss` response file.
- `-f2c:path_to\setup.log` to specify the location of the `setup.log` log file.

By default, log files are written to the `C:\Windows` directory.

5.3.2.2 Importing a Microsoft Windows Image

After you prepare a Microsoft Windows image for import, follow these steps to import the image:

1. Upload the image file to an Object Storage bucket.

Ensure that you select a bucket where you have read and write access.

See [Section 5.1.7, “Exporting an Image to Object Storage”](#).

2. Import the image from the bucket to your tenancy.

See [Section 5.1.6, “Importing an Image”](#)

3. Complete the post-import tasks.

See [Section 5.3.2.3, “Post-Import Tasks for Microsoft Windows Images”](#).

5.3.2.3 Post-Import Tasks for Microsoft Windows Images

After you import a custom Microsoft Windows image, do the following:

1. Use the imported image to launch an instance.

For the image source, select Custom Images, and then select the image that you imported. See [Section 6.2.1, “Creating an Instance”](#).

2. Enable Remote Desktop Protocol (RDP) access to the Compute instance.

See [Section 6.4.4.1, “Enabling Remote Desktop Protocol \(RDP\) Access”](#).

3. Connect to the instance using RDP.

See [Section 6.4.4.2, “Connecting with an RDP Client”](#).

4. If the instance requires any remotely attached storage, such as block volumes create and attach it.

See [Section 7.2, “Creating and Attaching Block Volumes”](#)

5. Create and attach any required secondary VNICs.

See [Section 4.4, “Configuring VNICs and IP Addressing”](#)

6. Test that all applications are working as expected.

7. Reconfigure any services that were set to start manually.

8. Configure your instance to use the Network Time Protocol (NTP).

To avoid performing this post-launch configuration every time you launch an instance using this custom image, consider creating a new image from the fully configured instance. See [Section 5.2.3, "Creating an Image From an Instance"](#).

Chapter 6 Compute Instance Deployment

Table of Contents

6.1 Tutorial – Launching Your First Instance	197
6.1.1 Task Flow to Launch an Instance	198
6.1.2 Prerequisites	198
6.1.3 Log into Oracle Private Cloud Appliance	198
6.1.4 Create a Compartment	199
6.1.5 Import an Image	199
6.1.6 Create a Virtual Cloud Network (VCN)	201
6.1.7 Create a Subnet	202
6.1.8 Create an Internet Gateway and Configure Route Rules	203
6.1.9 Launch an Instance	204
6.1.10 Get the Instance IP Address	206
6.1.11 Connect to Your Instance	206
6.1.12 Add a Block Volume	208
6.1.13 Attach the Block Volume to an Instance	208
6.1.14 (Optional) Clean Up Resources	209
6.2 Working with Instances	211
6.2.1 Creating an Instance	212
6.2.2 Retrieving Instance Metadata from Within the Instance	218
6.2.3 Updating an Instance	219
6.2.4 Stopping, Starting and Resetting an Instance	220
6.2.5 Terminating an Instance	221
6.3 Working with Instance Configurations and Instance Pools	222
6.3.1 Creating an Instance Configuration	222
6.3.2 Updating an Instance Configuration	224
6.3.3 Moving an Instance Configuration to a Different Compartment	224
6.3.4 Deleting an Instance Configuration	225
6.3.5 Using an Instance Configuration to Launch an Instance	225
6.3.6 Creating an Instance Pool	226
6.3.7 Updating an Instance Pool	228
6.3.8 Stopping and Starting Instances in an Instance Pool	229
6.3.9 Deleting an Instance Pool	230
6.4 Connecting to a Compute Instance	231
6.4.1 Prerequisites	231
6.4.2 Managing Key Pairs	232
6.4.3 Connecting to a Linux or Oracle Solaris Instance	234
6.4.4 Connecting to a Windows Instance	236
6.4.5 Connecting to an Instance Using a Console Connection	238

6.1 Tutorial – Launching Your First Instance

In this tutorial you'll learn the basic features of Oracle Private Cloud Appliance by performing some guided steps to launch and connect to an instance. After your instance is up and running, this tutorial steps you through creating and attaching a block volume to your instance.

This tutorial also includes optional instructions for deleting all the resources you create.

6.1.1 Task Flow to Launch an Instance

No.	Task	Links
1.	Review the prerequisites.	Section 6.1.2, “Prerequisites”
2.	Log into the appliance.	Section 6.1.3, “Log into Oracle Private Cloud Appliance”
3.	Create a compartment for your resources.	Section 6.1.4, “Create a Compartment”
4.	Import an image into your compartment.	Section 6.1.5, “Import an Image”
5.	Create a Virtual Cloud Network (VCN).	Section 6.1.6, “Create a Virtual Cloud Network (VCN)”
6.	Create a subnet in the VCN.	Section 6.1.7, “Create a Subnet”
7.	Configure additional network parameters to enable instance connectivity.	Section 6.1.8, “Create an Internet Gateway and Configure Route Rules”
8.	Launch an instance	Section 6.1.9, “Launch an Instance”
9.	Get the instance IP address.	Section 6.1.10, “Get the Instance IP Address”
10.	Connect to your instance.	Section 6.1.11, “Connect to Your Instance”
11.	Add Storage to your instance.	Section 6.1.12, “Add a Block Volume” Section 6.1.13, “Attach the Block Volume to an Instance”
12.	(Optional) Clean up your resources.	Section 6.1.14, “(Optional) Clean Up Resources”

6.1.2 Prerequisites

To perform this tutorial, ensure that you have these items.

- The URL for your Oracle Private Cloud Appliance

For example, https://console.<pca_name>.example.com where *<pca_name>* is the name of your Oracle Private Cloud Appliance and *example.com* is your domain.

- An Oracle Private Cloud Appliance user account and password
- The name of your tenancy
- An SSH-2 RSA key pair. If you want to create a key pair for this tutorial, see [Section 6.4.2, “Managing Key Pairs”](#).
- The virtual IP address (VIP) or hostname of the Oracle Private Cloud Appliance management nodes

If you don't have these items, you might be able to get them from your Service Enclave administrator.

6.1.3 Log into Oracle Private Cloud Appliance

1. In a browser, enter the URL for your Oracle Private Cloud Appliance.

2. Enter your tenancy name and click Continue.

The Sign In page is displayed.

3. Enter your user name and password, and then click Sign In.

If this is the first time you've logged in, you are prompted to change your password.

For more information about using the Compute Web UI, see [Section 1.1, "Using the Compute Web UI"](#).

What's Next

Continue on with [Section 6.1.4, "Create a Compartment"](#)

6.1.4 Create a Compartment

Compartments help you organize and control access to your resources. A compartment is a collection of resources (such as cloud networks, compute instances, and block volumes) that can be accessed only by those groups that have been given permission by an administrator in your organization.

In a production environment, the compartment for the instance you plan to create might already exist, and you could use it instead of creating a new compartment. However, in this tutorial, you create a new compartment to learn how to do it, and to provide an empty compartment from which you can create your cloud network.

In this tutorial, you use one compartment for all your resources. However, when you are ready to create a production environment you can separate resources in different compartments. For example, you might place all instances in one compartment and all networking resources in another compartment.

For more information about compartments, refer to these resources:

- Conceptual information, see *Organizing Resources in Compartments* in the [Identity and Access Management Overview](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).
- Step-by-step instructions to manage compartments, see [Section 2.3, "Creating and Managing Compartments"](#).

Using the Compute Web UI

1. In the Navigation Menu, click Identity and then click Compartments.
2. Click Create Compartment.
3. Enter the following details:
 - **Name:** Enter Sandbox.
 - **Description:** Enter a description for the compartment.
 - **Create in Compartment:** Select the compartment in which to create this new compartment.
4. Click Create Compartment.

The new compartment is displayed on the Compartments page.

What's Next

Continue on with [Section 6.1.5, "Import an Image"](#).

6.1.5 Import an Image

Oracle Private Cloud Appliance includes Oracle Linux and Oracle Solaris images that can be used to launch instances.

To make use of these images, you must import them into a compartment in a tenancy where you plan to launch instances as described in this task.

This tutorial imports an Oracle Linux 8 image.

Before You Begin

You need to know the URL for the images you want to import. Images are imported from a web hosting on the management node. The image is reached using a URL that is formed as follows:

- URL of the image file on the management node uses this syntax:

```
Syntax:  
https://<mgmt_vip_hostname>.<system_name>.<domain_name>:8079/images/<image_file_name>  
  
Example:  
https://manager.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.01.12_0.oci
```

where:

- `<mgmt_vip_hostname>.<system_name>.<domain_name>` is the fully qualified domain name for the management node VIP.

This information is available in the Service Enclave. If needed, consult with your Service Enclave administrator.

- `<image_file_name>` is the name of the image file. In this tutorial use: `uln-pca-Oracle-Linux-8-2022.01.12_0.oci`

Note – The image file names change as new versions of the OS are released. To find the specific image names for your appliance, refer to *Known Issues and Workarounds*, *Platform Issues*, in the [Oracle Private Cloud Appliance Release Notes](#).

If you don't know these values, ask your Service Enclave administrator.

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Custom Images.
2. Click Import Image.
3. Enter the required information:
 - **Name:** Enter a descriptive name for the image.
 - **Create in Compartment:** Select the Sandbox compartment.
 - **Source Type:** Select Import from an Object Storage URL.
 - **Object Storage URL:** Enter the URL for the image. For this tutorial, enter the following URL with your management node VIP or hostname.

```
Syntax:  
https://<mgmt_vip_hostname>.<system_name>.<domain_name>:8079/images/<image_file_name>  
  
Example:  
https://manager.myprivatecloud.example.com:8079/images/uln-pca-Oracle-Linux-8-2022.01.12_0.oci
```

- **Image Type:** select OCI.

- **Launch Mode:** Select Paravirtualized Mode.
 - **Tagging:** Leave tagging blank for this tutorial.
4. Click Import Image.

The imported image appears in the Custom Images list for the compartment, with a state of Importing. When the import completes successfully, the state changes to Available, and you can launch instances with it.

While the image imports, continue to the next task.

What's Next

Continue on with [Section 6.1.6, "Create a Virtual Cloud Network \(VCN\)"](#).

6.1.6 Create a Virtual Cloud Network (VCN)

Before you can launch an instance, you need a virtual cloud network (VCN) and a subnet.

A virtual cloud network, or VCN, is a software-defined equivalent of a traditional network, with firewall rules and various types of communication gateways.

In a production environment, a VCN that you can use for the instance might already exist, and you could use it instead of creating a new VCN. However, in this tutorial, you create a new VCN to learn how to do it.

Important

This tutorial creates a simple cloud network to make it easy to launch an instance for learning purposes. When you create your production instances, ensure that you create appropriate security lists and route table rules to restrict network traffic to your instances.

For more information about VCNs, refer to these resources:

- For conceptual information, see *Virtual Cloud Network* in the [Virtual Networking Overview](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).
- For step-by-step instructions to manage VCNs, see [Section 4.1, "Managing VCNs and Subnets"](#).

Using the Compute Web UI

1. In the navigation menu, under Networking, click Virtual Cloud Networks.
2. At the top of the page, select the Sandbox compartment.
3. Click Create Virtual Cloud Network.
4. Enter the required information:
 - **Name:** Enter a descriptive name for your cloud network. Don't enter confidential information.
 - Create in **Compartment:** This field defaults to your current compartment. Ensure that the Sandbox compartment is selected.
 - **CIDR Block:** Enter a valid CIDR block for the VCN. For example 10.0.0.0/16.

- **Use DNS hostnames in this VCN:** Indicate whether you want to use DNS host names in the VCN.
- **DNS Label:** If you selected to use DNS, enter a DNS label or leave the field blank to let the system generate a DNS name for you.
- **Tagging:** Optionally, add one or more tags to this resource.

Leave blank. This tutorial does not use tags.

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create Virtual Cloud Network.

What's Next

Continue on with [Section 6.1.7, "Create a Subnet"](#).

6.1.7 Create a Subnet


A subnet is a subdivision of your VCN. The subnet directs traffic according to a route table.

For this tutorial, you'll access the instance over the internet using its public IP address, so your route table will direct traffic to an internet gateway. The subnet also uses a security list to control traffic in and out of the instance.

Using the Compute Web UI

1. Return to the Virtual Cloud Networks page.

A quick way to do this is to click the name of page in the breadcrumb that is in the top banner. For example:



Alternatively, you can open the navigation panel, click Networking then click Virtual Cloud Networks.

2. Click the name of the VCN you just created.

The VCN details page is displayed.

3. In the lower panel, click Create Subnet.

4. Enter the required information:

- **Name:** Enter a descriptive name for your subnet. Don't enter confidential information.
- **CIDR Block:** Enter a valid CIDR block for the subnet. The value must be within the VCN's CIDR block. For example, 10.0.0.0/24.
- **Create in Route Table:** For this tutorial, select the default route table.
- **Subnet Access:** For this tutorial, select Public Subnet to allow public IP addresses for instances in the subnet.

- **Use DNS hostnames in this VCN:** For this tutorial, leave this unselected.
- **DHCP Options:** Leave this unselected.
- Create in **Security Lists:** For this tutorial, click Add Security List and select the default security list.
- **Tagging:** Leave blank. This tutorial does not use tags.

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create Subnet.

What's Next

Continue on with [Section 6.1.8, "Create an Internet Gateway and Configure Route Rules"](#).

6.1.8 Create an Internet Gateway and Configure Route Rules

An internet gateway is an optional virtual router you can add to your VCN to enable access to your data center network.

The gateway supports connections initiated from within the VCN (egress) and connections initiated from the internet (ingress).

Security list rules control the types of traffic allowed in and out of resources in that subnet. Make sure to allow only the desired types of internet traffic.

Each public subnet that needs to use the internet gateway must have a route table rule that specifies the gateway as the target.

Using the Compute Web UI

1. If you aren't on your VCN's details page, navigate to it now.
2. Under the Resources panel, select Internet Gateways.
3. Click Create Internet Gateway.
4. Enter the required information:
 - **Name:** Enter a descriptive name for your subnet. Don't enter confidential information.
 - **Compartment:** This field defaults to your current compartment. Ensure that the Sandbox compartment is selected.
 - **Tagging:** Leave blank.
5. Click Submit.
6. Under Resources, click Route Tables.
7. Click the name of the default route table.
8. Click Add Route Rules.
9. Enter the required information:

- **Target Type:** From the drop-down menu, select Internet Gateway.
- **CIDR Block:** Enter: `0.0.0.0/0` (which means that all non-intra-VCN traffic that is not already covered by other rules in the route table will go to the target specified in this rule)
- **Internet Gateway:** From the drop-down menu, select the name of the Internet Gateway that you created.
- **Description:** An optional description of the rule.

10. Click Create Route Table Rule.

What's Next

Continue on with [Section 6.1.9, “Launch an Instance”](#).

6.1.9 Launch an Instance

In this task, launch an instance with an image and a basic shape.

A compute instance is a virtual machine (VM), which is an independent computing environment that runs on top of physical hardware. The virtualization makes it possible to run multiple compute instances that are isolated from each other.

A shape describes the instance resources such as the number of CPUs, amount of memory, and network resources. In a production environment, you would select a shape that best suits workload and application requirements for the instance.

For more information about instances, refer to these resources:

- Conceptual information, see the [Compute Instances](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).
- Step-by-step instructions to manage VCNs, see [Section 6.2, “Working with Instances”](#).

Before You Begin

Ensure that you have performed these tasks:


- [Section 6.1.4, “Create a Compartment”](#)
- [Section 6.1.5, “Import an Image”](#)
- [Section 6.1.6, “Create a Virtual Cloud Network \(VCN\)”](#)
- [Section 6.1.7, “Create a Subnet”](#)
- [Section 6.1.8, “Create an Internet Gateway and Configure Route Rules”](#)

Using the Compute Web UI

1. In the navigation menu, under Compute, click Instances.
2. Click Create Instance.
3. Enter the required information:

- **Name:** Enter a descriptive name for your compute instance. Don't enter confidential information.
 - **General Information:**
 - **Create in Compartment:** This field defaults to your current compartment. Ensure that the Sandbox compartment is selected.
 - **Fault Domain:** Leave the default so that the fault domain is assigned automatically.
 - **Source Image:**
 - **Source Type:** Select Custom Image.
 - **Compartment** Select the Sandbox compartment where the image is located.
 - **Operating System:** Leave set to Any.
 - **List of images:** Select the Oracle Linux 8 image that you imported.
 - **Shape:** Select one of the smaller shapes such as VM.PCAStandard1.1.
 - **Boot Volume:** Leave the check box empty so that the default boot volume size is created.
 - **Subnet:**
 - **VCN:** Select the VCN you created.
 - **Subnet:** Select the subnet you created.
 - **Public IP Address:** Ensure the check box is checked so that a public IP address is assigned to the instance.
 - **Private IP Address:** Leave the field blank.
 - **Hostname:** You can leave this field blank or enter a hostname.
 - **SSH Keys:** Provide your public SSH key. Either browse to (click the drag and drop box), or drag and drop, or select Paste the public SSH key and paste it (.pub) into the field.
 - **Initialization Script:** Leave this area as is.
 - **Network Security Group:** Leave the check box unchecked.
 - **Tagging:** Leave this area blank.
4. Click Launch Instance.
 5. Watch the state.

The state is displayed above the icon of the object. Example:



Your instance begins in the Provisioning state and in a short amount of time changes to a Running state. Once the instance is in the running state, you can connect to it.

What's Next

Continue on with [Section 6.1.10, "Get the Instance IP Address"](#).


6.1.10 Get the Instance IP Address

You can connect to the instance using SSH with the instance IP address.

Using the Compute Web UI

1. If you aren't on your instance's details page, navigate to it now.
From the Compute Instances page, click the name of your instance.
2. Select the Networking tab.

The tabs are displayed at the top of the details panel:



3. Make note of the Public IP Address.

What's Next

Continue on with [Section 6.1.11, "Connect to Your Instance"](#)

6.1.11 Connect to Your Instance

In most cases, you connect to a running instance using a Secure Shell (SSH) connection, but some instances support authenticating your connection with a password. This tutorial assumes that you used an image that creates an instance that authenticates your SSH connection with an SSH key pair.

For the system that you will be connecting from, most Linux and UNIX-like operating systems include an SSH client by default.

Windows 10 and Windows Server 2019 systems should include the [OpenSSH client](#), which you'll need if you created your instance using the SSH keys generated by Oracle Cloud Infrastructure.

For other Windows versions, you can download a free SSH client called PuTTY from <http://www.putty.org>.

Before You Begin

- Know the public IP address of your instance. See [Section 6.1.10, "Get the Instance IP Address"](#)
- Know the path to your private key file.
- Know the valid user name.

The name is determined by what is configured in the image used to launch the instance. If you launched an instance using one of the images provided on the appliance, the default user is `opc`. See [Section 5.1, "Accessing the Management Node Images"](#).

Perform one of the following tasks based on the type of system you are connecting from.

6.1.11.1 Connect from a UNIX-type System

1. Open a terminal window.
2. Use the SSH command to connect to your instance.

Syntax:

```
ssh -i <private_key_pathname> <username>@<public-ip-address>
```

- *<private_key_pathname>* is the full path name of the file that contains the private key associated with the instance you want to access.
- *<username>* is the default user name for the instance. For this tutorial, `opc` is the user name.
- *<public-ip-address>* is your instance IP address.

Example:

```
ssh -i ~/.ssh/my_private_ssh_key opc@192.0.2.0
```

3. If asked if you want to continue connecting, type `yes`.

You are now logged in to your instance.

What's Next

Continue on with [Section 6.1.12, "Add a Block Volume"](#).

6.1.11.2 Connect Using PuTTY

This connection method is commonly performed from Windows systems.

If the instance uses a key pair that you created using PuTTY Key Generator, use the following procedure.

1. Open PuTTY.
2. In the Category pane (on the left), select Session and enter the following:
 - **Host Name (or IP address):** *<username>@<public-ip-address>*
 - *<username>* is the default username for the instance. For this tutorial, the user name is `opc`.
 - *<public-ip-address>* is your instance IP address.
 - **Port:** 22
 - **Connection type:** SSH
3. In the Category pane, expand Window, and then select Translation.
4. In the Remote character set drop-down list, select UTF-8. The default locale setting on Linux-based instances is UTF-8, and this configures PuTTY to use the same locale.
5. In the Category pane, expand Connection, expand SSH, and then click Auth.

6. Click Browse, and then select your .ppk private key file.
7. Click Open to start the session.

If this is your first time connecting to the instance, you might see a message that the server's host key is not cached in the registry. Click Yes or Accept to continue the connection.

Tip – If the connection fails, you may need to update your PuTTY proxy configuration.

What's Next

Continue on with [Section 6.1.12, “Add a Block Volume”](#).

6.1.12 Add a Block Volume

You can use block volumes to expand the storage capacity of your compute instances.

After a block volume is created, you attach the volume to one or more instances. You can use the volume like a regular hard drive.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Click Create Block Volume.
3. Enter the required information:
 - **Name:** Enter a user-friendly name. Avoid entering confidential information.
 - **Create in Compartment:** This field defaults to your current compartment. Ensure that the Sandbox compartment is selected.
 - **Size:** Leave the default size (1024 GB).
 - **High Performance Volume:** Leave unselected.
 - **Backup Policy:** Do not select a backup policy.
 - **Tags:** Leave the tagging fields blank.
4. Click Create Block Volume.
5. View the state.

The state is displayed above the icon for the object (block volume in this case).

Initially, the block volume is in the provisioning state. When the volume changes to the Available state, you can attach it to your instance.

What's Next

Continue on with [Section 6.1.13, “Attach the Block Volume to an Instance”](#)

6.1.13 Attach the Block Volume to an Instance

Using the Compute Web UI

1. In the navigation menu, under Compute click Instances.
2. Click your instance name to view its details.

If you don't see your instance listed, ensure that the Sandbox compartment is selected at the top of the page.

3. In the Resources panel, click Attached Block Volumes.
4. Click Attach Block Volume.
5. Enter the required information:
 - **Select from Compartment:** Select the Sandbox compartment.
 - **Block Volume:** Select the block volume you created.
 - **Device Path:** Do not select a device path.
 - **Access:** Select Read/Write.
6. Click Attach to Instance.

The attachment process takes about a minute. You will know the volume is ready when the Attachment State for the volume is Attached.

Tip – Reload the web page if your block volume isn't displayed.

When a block volume is initially attached to an instance, the instance sees the volume as a new disk. To make the volume available to the instance OS, you need to perform administrative tasks such as give the volume a file system and mount it to the OS.

To learn about the block volume and how to make it available to the instance OS, refer to these sections outside of this tutorial:

- [Section 7.2.5, “Find Your Volume in the Instance”](#)
- [Section 7.2.6, “Configuring Volumes to Automatically Mount \(Linux Instances\)”](#)

What's Next

Continue on with [Section 6.1.14, “\(Optional\) Clean Up Resources”](#).

6.1.14 (Optional) Clean Up Resources

After you've finished with the resources you created for this tutorial, you can delete and release the resources you don't intend to continue working with.

6.1.14.1 Detach and Delete the Block Volume.

Caution

You cannot undo a termination. Any data on a volume will be permanently deleted once the volume is terminated.

Using the Compute Web UI

1. In the navigation menu and click Compute, then click Instances.
2. Select the Sandbox compartment.
3. Click the name of your instance.
4. In the Resources panel, click Attached Block Volumes.
5. Find your volume, click the Actions icon (three dots), and then click Detach. Confirm the detachment in the dialog box.

Note – You might need to refresh the web page to see that the block volume is no longer attached.

6. In the navigation menu, click Block Storage, then click Block Volumes.

If your block volume is not displayed, ensure that the Sandbox compartment is selected.

7. Delete the volume: Find your volume, click the Actions icon (three dots), and then click Terminate, and confirm the termination in the dialog box.

What's Next

Continue on with [Section 6.1.14.2, "Terminate the Instance"](#).

6.1.14.2 Terminate the Instance

You can permanently terminate (delete) instances that you no longer need. Any attached VNICs and volumes are automatically detached when the instance terminates. Eventually, the instance's public and private IP addresses are released and become available for other instances.

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the Sandbox compartment.
3. Find the instance you created and click the Actions icon (three dots), and select Terminate.
4. In the Confirm Instances termination dialog box, move the Permanently delete the attached boot volume selector to the right, and click Confirm.

Moving the selector to the right results in the boot volume being permanently deleted, which is appropriate for this tutorial.

In production, you can leave the selector in the left position to preserve the boot volume which can be used with another instance. This is convenient when you want to reuse a configured OS or data on the boot volume.

What's Next

Continue on with [Section 6.1.14.3, "Delete the Subnet, Internet Gateway, and VCN"](#)

6.1.14.3 Delete the Subnet, Internet Gateway, and VCN

Using the Compute Web UI

1. In the navigation menu, click Networking, then click Virtual Cloud Networks.
2. Select the Sandbox compartment.
3. Click the name of your VCN.
4. Under Resources, click Route Tables.
5. Click the name of the route table.
6. For the route rule you created, click the Actions icon (three dots), click Delete, and confirm the deletion.
The route rule is deleted.
7. In the breadcrumb path at the top of the page, click the name of your VCN.
The VCN details page is displayed.
8. Under Resources, click Internet Gateways.
9. For the Internet gateway that you created, click the Actions icon (three dots), click Delete, and confirm the deletion.
The Internet gateway is deleted.
10. Under Resources, click Subnets.
11. For the subnet you created, click the Actions icon (three dots), click Delete, and confirm the deletion.
The subnet is deleted.
12. On the VCN details page, click Terminate and confirm the termination.
The VCN is deleted.

What's Next

Continue on with

6.1.14.4 Delete the Compartment

You must remove all resources from a compartment before you can delete it, otherwise, the deletion action fails and returns to an Active state.

Using the Compute Web UI

1. In the navigation menu, click Identity, then click Compartments.
2. For the Sandbox compartment, click the Actions icon (three dots), and then click Delete.
3. Confirm the deletion in the dialog box.

6.2 Working with Instances

You can create compute instances as needed to meet your compute and application requirements. After you create an instance, you can access it securely from your computer, restart it, attach and detach volumes, and terminate it.

For general information about instances, see the [Compute Instances](#) chapter of the *Oracle Private Cloud Appliance Concepts Guide*.

6.2.1 Creating an Instance

See [Chapter 5, Compute Images](#) and [Section 6.1, “Tutorial – Launching Your First Instance”](#) for information about input you need to create an instance.

The following is the minimum information that you must provide to create an instance using the Compute Web UI:

- A name for the instance
- The compartment where you want to create the instance
- An image or boot volume
- A shape
- A subnet
- A public SSH key

To log in to the instance, users need either an SSH key or a password, depending on how the image was built. If the instance will require SSH keys for authentication, you must provide the public key when you create the instance. You cannot provide the public SSH key after the instance is created.

To create an instance using the OCI CLI, you need the same information as listed above for the Compute Web UI with the following differences:

- You need the name of the availability domain where you want to create the instance.
- You do not need an instance name. If you do not provide a name for the instance, the default name will be `instanceYYYYMMDDhhmmss`, where `YYYYMMDDhhmmss` is the creation date and time.

To modify launch options, see "Using the OCI CLI."

Using the Compute Web UI

1. Create or get the following resources and information:
 - An image or boot volume and the compartment where the image or boot volume is located
 - A virtual cloud network (VCN) and subnet and the compartment where the VCN and subnet are located
 - A public Secure Shell (SSH) key if users will connect to the instance using SSH
2. Open the Create an Instance dialog.
 - Using the navigation menu.
 - a. In the navigation menu, click Compute, and then click Instances.
 - b. Click the Create Instance button.
 - Using the Dashboard.
 - a. Click Dashboard.

- b. Use one of the following buttons:
 - Click the Compute/View Instances button. Click the Create Instance button.
 - Click the Compute/Create Virtual Machine Instance button. The Create an Instance dialog opens.
3. In the Create an Instance dialog, enter the following information:
 - **Name:** Enter a name for the instance. Instance names have the following characteristics:
 - Can be changed after the instance is created.
 - Do not need to be unique.
 - Can contain only alphanumeric characters and the hyphen (-) character.
 - Can be a maximum of 63 characters.
 - **Create in Compartment:** Select the compartment where you want to create the instance.
 - **Fault Domain:** (Optional) Select a fault domain. By default, the system automatically selects the best fault domain for the instance when the instance is created. You can change the fault domain after the instance is created.
 - **Source Image:** Select an image or boot volume.
 - a. Select the Source Type.
 - b. Select the compartment where the image or boot volume that you want to use is located..
 - c. Select an image or boot volume from the list. You can select an operating system to filter the list, or use the arrow buttons to view another page of the list.
 - **Shape:** Select a shape. For a description of each compute instance shape, see *Standard Shapes* in the [Compute Instances](#) chapter of the *Oracle Private Cloud Appliance Concepts Guide*.
 - **Boot Volume:** (Optional) Check the "Specify a custom boot volume size" box if you want to specify a custom boot volume size. The default value is 50 GB.
 - **Subnet:** Select a subnet.
 - a. Select a VCN from the list. You might need to change the compartment to the compartment where the VCN is located.
 - b. Select a subnet.
 - **Public IP Address:** To connect to the instance using SSH, check the Assign Public IP box to have a public IP address assigned to the instance. This box is checked by default if you specified a public subnet. If you do not check this box, or if you uncheck this box, and then want to assign a public IP address later, see [Section 4.4.2.3, "Assigning a Public IP Address to an Instance"](#) for instructions.
 - **Private IP Address:** (Optional) Specify an available private IP address from the subnet's CIDR. By default, a private IP address is automatically assigned.
 - **Hostname:** (Optional) Enter a hostname if you are using DNS within the cloud network. The hostname must be unique across all VNICs in the subnet.

By default, the instance name is used for the hostname. The hostname can also be configured in the OS after the instance is created.

If this is a UNIX-type instance, see [Section 8.1.3, “Creating a Mount Target”](#) and [Section 8.3.1, “Mounting Overview”](#) for more information about setting the host name correctly for mounting file systems.

- *SSH Keys*: To connect to the instance using SSH, provide a public SSH key.

Note

You cannot provide this SSH key after the instance is created.

- *Initialization Script*: (Optional) Provide an initialization script. This is a file of data to be used for custom instance initialization.
- *Network Security Group*: (Optional) Check the Enable Network Security Group box to add the primary VNIC for this instance to one or more network security groups (NSGs).

Select one or more NSGs.

See [Section 4.2.4, “Controlling Traffic with Network Security Groups”](#) for information about NSGs.

- *Tagging*: (Optional) Add defined or free-form tags for this instance as described in [Section 3.4.1, “Adding Tags at Resource Creation”](#). Tags can also be applied later.

4. Click the Launch Instance button.

Click Work Request(s) in the Resources box to check the status of the instance launch.

Using the OCI CLI

1. Create or get the following resources and information:

- The name of the availability domain that you want to use.

```
$ oci iam availability-domain list
```

- The OCID of the compartment where you want to create the instance.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- The name of the shape for this instance. Use the following command to list the available shapes and their characteristics. Use the OCID of the compartment where you want to create the instance. To list only shapes that are compatible with the image that you plan to use, specify the image OCID.

```
$ oci compute shape list --compartment-id compartment_OCID --image-id image_OCID
```

- The OCID of the subnet where the VNIC that is attached to this instance will be created.

```
$ oci compute vnic-attachment list --compartment-id compartment_OCID
```

- If you provide a value for the `--hostname-label` option, see the description of Hostname in the preceding Compute Web UI procedure.
- One of the following to specify either an image or a boot volume.
 - The OCID of the image used to boot the instance.


```
$ oci compute image list --compartment-id compartment_OCID
```

- The OCID of the boot volume used to boot the instance.

```
$ oci compute boot-volume-attachment list \
--availability-domain ad_name --compartment-id compartment_OCID
```

- A public Secure Shell (SSH) key to connect to the instance using SSH.

Note

You cannot provide this SSH key after the instance is created.

For a complete list of required and optional parameters, use the following command:

```
$ oci compute instance launch --help
```

See the Compute Web UI procedure for characteristics of the `--display-name` and `--hostname-label` values. See [Section 3.4.1, "Adding Tags at Resource Creation"](#) to add defined and free-form tags.

2. Construct an argument for the `--source-details` option.

The `--source-details` argument can be a JSON file or a command-line string. Use the following command to show the correct format of the JSON properties and values:

```
$ oci compute instance launch --generate-param-json-input source-details
[
  "This parameter should actually be a JSON object rather than an array - pick one of the
  following object variants to use",
  {
    "bootVolumeId": "string",
    "sourceType": "bootVolume"
  },
  {
    "bootVolumeSizeInGBs": 0,
    "imageId": "string",
    "kmsKeyId": "string",
    "sourceType": "image"
  }
]
```

3. (Optional) Construct an argument for the `--launch-options` option.

Only the `firmware` property can be changed. The default value is BIOS. You can alternatively specify `UEFI_64`. If you do not provide a correct value for `firmware`, the instance might not launch properly. You cannot update the value of the `firmware` property with the `instance update` command.

The following shows the default values:

```
{
  "bootVolumeType": "PARAVIRTUALIZED",
  "firmware": "BIOS",
  "isConsistentVolumeNamingEnabled": false,
  "is-pv-encryption-in-transit-enabled": false,
  "networkType": "PARAVIRTUALIZED",
  "remoteDataVolumeType": "PARAVIRTUALIZED"
```

```
}
```

To change the value of the `firmware` property, specify the following option:

```
--launch-options file://launch_options.json
```

Where the following is the content of the `launch_options.json` file:

```
{
  "bootVolumeType": "PARAVIRTUALIZED",
  "firmware": "UEFI_64",
  "isConsistentVolumeNamingEnabled": false,
  "is-pv-encryption-in-transit-enabled": false,
  "networkType": "PARAVIRTUALIZED",
  "remoteDataVolumeType": "PARAVIRTUALIZED"
}
```

4. (Optional) Construct an argument for the `--metadata` or `--extended-metadata` option.

Custom user data can be attached to the instance by using the `--metadata` and `--extended-metadata` options. Metadata key/value pairs are string/string maps in JSON format. Extended metadata can be nested JSON objects. Metadata and extended metadata have the following restrictions:

- Keys are limited to 255 characters.
- Most key values are limited to 255 characters.
 - The value of the `ssh_authorized_keys` key can be more than 255 characters. This value must be a valid public key in OpenSSH format.
 - The value of `user_data` can be a maximum of 16KB. This value is data that Cloud-Init can use to run custom scripts or provide custom Cloud-Init configuration.
- Metadata can have a maximum of 128 keys.
- The combined size of the metadata and extended metadata can be a maximum of 32,000 bytes.

SSH keys can alternatively be provided in the file argument of the `--ssh-authorized-keys-file` option, and user data can alternatively be provided in the file argument of the `--user-data-file` option. Use the `--help` option for more information.

In the example in the next step, the authorized keys file contains one or more public SSH keys in the format required by the SSH `authorized_keys` file. Use a newline character to separate multiple keys. SSH public keys can be provided as the value of the `ssh_authorized_keys` key in the `--metadata` option, or in the file argument of the `--ssh-authorized-keys-file` option. Use `--help` for more information.

5. Run the instance launch command.

Syntax:

```
oci compute instance launch --availability-domain availability_domain \
--compartment-id compartment_OCID --shape shape --subnet-id subnet_OCID \
```

```
--source-details file://image_info.json
```

Example:

If you are using a public subnet, a public IP address is assigned by default, or you can set the `--assign-public-ip` option value to `true`. If you need to assign a public IP address later, see [Section 4.4.2.3, "Assigning a Public IP Address to an Instance"](#) for instructions.

```
$ oci compute instance launch --availability-domain ad1 \
--compartment-id ocidl.compartment.unique_ID --display-name ops1 \
--shape VM.PCAStandard1.16 --subnet-id ocidl.subnet.unique_ID --source-details \
'{"bootVolumeSizeInGBs":100,"imageId":"ocidl.image.unique_ID","sourceType":"image"}' \
--assign-public-ip true --ssh-authorized-keys-file ./ssh/id_rsa.pub
{
  "data": {
    "agent-config": null,
    "availability-config": null,
    "availability-domain": "ad1",
    "capacity-reservation-id": null,
    "compartment-id": "ocidl.compartment.unique_ID",
    "dedicated-vm-host-id": null,
    "defined-tags": {},
    "display-name": "ops1",
    "extended-metadata": null,
    "fault-domain": "FAULT-DOMAIN-1",
    "freeform-tags": {},
    "id": "ocidl.instance.unique_ID",
    "image-id": "ocidl.image.unique_ID",
    "instance-options": null,
    "ipxe-script": null,
    "launch-mode": "PARAVIRTUALIZED",
    "launch-options": {
      "boot-volume-type": "PARAVIRTUALIZED",
      "firmware": "BIOS",
      "is-consistent-volume-naming-enabled": false,
      "is-pv-encryption-in-transit-enabled": false,
      "network-type": "PARAVIRTUALIZED",
      "remote-data-volume-type": "PARAVIRTUALIZED"
    },
    "lifecycle-state": "PROVISIONING",
    "metadata": {
      "ssh_authorized_keys": "ssh-rsa public_RSA_key"
    },
    "platform-config": null,
    "preemptible-instance-config": null,
    "region": "scasg01",
    "shape": "VM.PCAStandard1.1",
    "shape-config": null,
    "source-details": {
      "boot-volume-size-in-gbs": 100,
      "image-id": "ocidl.image.unique_ID",
      "kms-key-id": null,
      "source-type": "image"
    },
    "system-tags": null,
    "time-created": "2021-09-22T20:20:04.715304+00:00",
    "time-maintenance-reboot-due": null
  },
  "etag": "92180faa-3660-446c-9559-c12a6e6111f9",
  "opc-work-request-id": "ocidl.workrequest.unique_ID"
}
```

}

Use the `work-requests work-request get` command to monitor the status of the instance launch:

```
$ oci work-requests work-request get --work-request-id ocidl.workrequest.unique_ID
```

An alternative way to create an instance is to create an instance configuration and use that configuration to launch an instance, as described in [Section 6.3, “Working with Instance Configurations and Instance Pools”](#).

6.2.2 Retrieving Instance Metadata from Within the Instance

The instance metadata service (IMDS) serves information about a running instance to users who are logged in to that instance. IMDS also provides information to Cloud-Init that you can use for various system initialization tasks.

Note

To access IMDS metadata, use an instance image that is provided by Oracle.

The IMDS metadata includes instance information such as the following:

- The SSH public key that enables users to log in to the instance
- Instance attached VNICs, VNIC IDs
- Instance CIDR blocks

In general, the IMDS instance metadata includes the following:

- The same information that you see on the details page of an instance in the Compute Web UI and in the output of the instance `get` command in the OCI CLI.
- Custom information that you add to an instance by using the `--metadata`, `--extended-metadata`, `--ssh-authorized-keys-file`, and `--user-data-file` options of the instance `launch` command. This metadata cannot be updated after instance launch. For a user logged into the instance, the instance metadata is read-only.

To retrieve the IMDS instance metadata, follow these steps:

1. Log in to the instance.
2. Use a cURL command to retrieve the metadata information from the HTTP endpoint.

Information is provided through an HTTP endpoint that listens on 169.254.169.254. If an instance has multiple VNICs, you must send the request using the primary VNIC.

Use the `instance` command to retrieve the instance metadata. Use the `vnics` command to retrieve the VNIC data.

Example: Instance Metadata

```
$ curl -H "Authorization: Bearer Oracle" -L http://169.254.169.254/opc/v2/instance/
{
  "availabilityDomain": "PCA",
  "faultDomain": "FAULT-DOMAIN-1",
  "compartmentId": "ocidl.compartment.unique_ID",
  "displayName": "dev1",
  "hostname": "hostname",
```

```

    "id": "ocidl.instance.unique_ID",
    "image": "ocidl.image.unique_ID",
    "metadata": {
      "ssh_authorized_keys": "public_RSA_key"
    },
    "region": "PCA",
    "canonicalRegionName": "PCA",
    "ociAdName": "PCA",
    "regionInfo": null,
    "shape": "VM.PCAStandard1.1",
    "state": "RUNNING",
    "timeCreated": 1634943279000,
    "agentConfig": null
  }
}

```

To retrieve a single value, specify the key name as shown in the following example.

Example: VNIC Metadata

```

$ curl -H "Authorization: Bearer Oracle" -L http://169.254.169.254/opc/v2/vnics/
[
  {
    "vnicId": "ocidl.vnic.unique_ID",
    "privateIp": "privateIp",
    "vlanTag": 0,
    "macAddr": "00:13:97:9f:16:32",
    "virtualRouterIp": "virtualRouterIp",
    "subnetCidrBlock": "subnetCidrBlock"
  }
]

```

You can view all of the data for one of multiple VNICs by specifying the array index for that VNIC data, or you can retrieve a single value for that specified VNIC:

```

$ curl -H "Authorization: Bearer Oracle" -L http://169.254.169.254/opc/v2/vnics/0/privateIp
privateIp

```

6.2.3 Updating an Instance

In addition to updating the properties of an instance, you might want to attach additional block volumes or secondary VNICs. See [Section 6.1.13, “Attach the Block Volume to an Instance”](#), [Section 7.2, “Creating and Attaching Block Volumes”](#), [Section 4.4, “Configuring VNICs and IP Addressing”](#) and [Section 4.4.3.3, “Create and Attach a Secondary VNIC”](#). You can also specify `secondaryVnics` and `secondaryVnicSubnets` when you use an instance configuration to create an instance.

If you did not add a public IP address when you created the instance, and you want to assign a public IP address now, see [Section 4.4.2.3, “Assigning a Public IP Address to an Instance”](#) for instructions.

You can add, change, or remove tags as described in [Section 3.4.2, “Applying Tags to an Existing Resource”](#).

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instances.
2. If the instance that you want to update is not listed, use the Compartment drop-down menu above the instances list to select the correct compartment.
3. For the instance that you want to update, click the Actions menu, and click the Edit option.
4. In the Edit `instance_name` dialog, make the changes.

When you update an instance by using the Compute Web UI, you can change the following:

- The name of the instance
- The fault domain

The fault domain can be changed only if the instance is in the Running state. Changing the fault domain of a running instance migrates the instance to the new compute node.

- Tags

5. Click the Save Changes button.

Using the OCI CLI

1. Get the OCID of the instance that you want to update.

```
$ oci compute instance list --compartment-id compartment_OCID
```

2. Run the instance update command.

Syntax:

```
oci compute instance update --instance-id instance_OCID options_with_values_to_update
```

For descriptions of instance properties that you can change, enter the following command and scroll to Optional Parameters:

```
$ oci compute instance update --help
```

You can change the fault domain only if the instance `lifecycle-state` is RUNNING. Changing the fault domain of a running instance migrates the instance to the new compute node.

Example:

```
$ oci compute instance update \  
--instance-id ocid1.instance.unique_ID \  
--display-name new_name
```

6.2.4 Stopping, Starting and Resetting an Instance

By default, a reset operation restarts the instance by sending a shutdown command to the instance operating system. After waiting 15 minutes for the operating system shutdown to complete, the instance is powered off and then powered back on.

To avoid any issues with abruptly stopping applications that are running on the instance, use operating system commands on the instance to shut down the instance before you use the methods described in this section to stop the instance.

An instance that is Stopped cannot be migrated to a different compute node. See the following resources for more information:

- [Section 6.2.3, “Updating an Instance”](#)
- *Migrating Instances from a Compute Node* in the [Hardware Administration](#) chapter in the [Oracle Private Cloud Appliance Administrator Guide](#)

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instances.

2. If the instance that you want to manage is not listed, use the Compartment drop-down menu above the instances list to select the correct compartment.
3. For the instance that you want to manage, click the Actions menu, and click the Start, Stop, or Reset option.

Click Work Request(s) in the Resources box to check the status of the instance state change.

Using the OCI CLI

1. Get the OCID of the instance that you want to stop, start, or reset.

```
$ oci compute instance list --compartment-id compartment_OCID
```

2. Run the stop, start, or reset command.

Syntax:

```
oci compute instance action --instance-id instance_OCID \  
--action {START | STOP | RESET | SOFTSTOP | SOFTRESET}
```

For descriptions of these actions, enter:

```
$ oci compute instance action --help
```

Example:

```
$ oci compute instance action --instance-id ocid1.instance.unique_ID --action RESET
```

Use the `work-requests work-request get` command to monitor the status of the instance state change.

6.2.5 Terminating an Instance

By default, the boot volume of the instance is preserved when you terminate the instance. You can attach the boot volume to a different instance as a data volume, or use it to launch a new instance. If you no longer need the boot volume, you can permanently delete it as described in [Section 7.4.6, "Deleting a Boot Volume"](#).

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instances.
2. If the instance that you want to terminate is not listed, use the Compartment drop-down menu above the instances list to select the correct compartment.
3. Click the instance that you want to terminate.
4. On the instance details page, click the Controls menu and click the Terminate option.

Click Work Request(s) in the Resources box to check the status of the instance terminate.

Using the OCI CLI

1. Get the OCID of the instance that you want to terminate.

```
$ oci compute instance list --compartment-id compartment_OCID
```

2. Run the instance terminate command.

Example:

```
$ oci compute instance terminate --instance-id ocidl.instance.unique_ID
```

Use the `work-requests work-request get` command to check the status of the instance terminate.

6.3 Working with Instance Configurations and Instance Pools

Instance configurations and instance pools simplify the management of compute instances. An instance configuration contains settings that are used to create a compute instance. An instance pool defines a set of compute instances that is managed as a group.

6.3.1 Creating an Instance Configuration

An instance configuration contains settings that are used to create a compute instance.

To create an instance configuration, you must use the OCI CLI.

Using the OCI CLI

1. Get the following information:

- The OCID of the compartment where you want to create this instance configuration.
- The OCID of the compartment where you want instances that use this instance configuration to be created.
- The OCID of the subnet for instances that use this instance configuration.
- The OCID of the image or boot volume for instances that use this instance configuration.
- The name of the availability domain for instances that use this instance configuration.
- The name of the shape for instances that use this instance configuration.
- If you provide a value for the `hostnameLabel` property, see the description of Hostname in the Compute Web UI procedure in [Section 6.2.1, “Creating an Instance”](#).

2. Create the configuration file that is input to the configuration create command.

The configuration file is a JSON file of property/value pairs.

- The following command shows the correct syntax of the configuration file and names of properties:

```
$ oci compute-management instance-configuration create \  
--generate-param-json-input instance-details > instance_details.json
```

You do not need all of the data that is output by this command. Copy just the information you need, being careful to keep each property in its correct context.

You can specify secondary VNICs and subnets.

If you omit the fault domains specification, the system will automatically select the best fault domain. If you specify only a single fault domain, all instances will be placed in only that fault domain.

If a fault domain that you specify does not have enough resources, the instance could fail to launch:

- When you use the `launch-compute-instance` command as described in [Section 6.3.5, “Using an Instance Configuration to Launch an Instance”](#), and you specify a fault domain in the instance

configuration, only that specified fault domain will be used to launch the instance. Resource constraints could cause the instance launch to fail.

- When you create instances in a pool, fault domains specified in the placement configuration override fault domains specified in the instance configuration. See [Section 6.3.6, “Creating an Instance Pool”](#) for more information.

If you omit the `assignPublicIp` property, a public IP address is assigned by default if you specify a public subnet.

If users will use `ssh` to connect to the instance, specify the SSH public key as the value of the `ssh_authorized_keys` property in the `metadata` block. You cannot add the SSH public key after the instance is created.

The `displayName` property is used for the instance name when you use the `launch-compute-instance` command as described in [Section 6.3.5, “Using an Instance Configuration to Launch an Instance”](#). If you do not provide a value for the `displayName` property, the default name of instances will be `instanceYYYYMMDDhhmmss`, where `YYYYMMDDhhmmss` is the creation date and time.

The `displayName` property is ignored when you create instances in a pool as described in [Section 6.3.6, “Creating an Instance Pool”](#).

- The following command shows which properties are required to create an instance:

```
$ oci compute instance launch --help
```

Scroll to the Required Parameters section. Optional parameters are described below the required parameters.

The names of the properties in the configuration file are similar to, but different from, the names of the `instance launch` options. Also, some of the properties are organized into groups of properties, such as `createVnicDetails` and `sourceDetails`, as shown in the following example configuration file:

```
{
  "instanceType": "compute",
  "launchDetails": {
    "availabilityDomain": "availability_domain",
    "compartmentId": "compartment_OCID",
    "createVnicDetails": {
      "assignPublicIp": true,
      "freeformTags": {
        "ConfigType": "Configuration for an XYZ instance."
      },
      "subnetId": "subnet_OCID"
    },
  },
  "displayName": "instance_name",
  "metadata": {
    "ssh_authorized_keys": "ssh-rsa public_RSA_key"
  },
  "shape": "shape_name",
  "sourceDetails": {
    "bootVolumeSizeInGBs": 100,
    "imageId": "image_OCID",
    "sourceType": "image"
  }
}
```

```
}

```

To change the value of the `firmware` property, provide a value for the `launchOptions` property. The default value is BIOS. You can alternatively specify UEFI_64. Other properties in `launchOptions` cannot be changed.

```
"launchOptions": {
  "bootVolumeType": "PARAVIRTUALIZED",
  "firmware": "UEFI_64",
  "isConsistentVolumeNamingEnabled": false,
  "isPvEncryptionInTransitEnabled": false,
  "networkType": "PARAVIRTUALIZED",
  "remoteDataVolumeType": "PARAVIRTUALIZED"
}
```

3. Run the instance configuration create command.

Syntax:

```
oci compute-management instance-configuration create -c compartment_OCID \
--display-name IC_name --instance-details file://custom_config_file.json
```

The specified compartment is where this instance configuration will be created. This compartment could be different from the compartment specified in the instance details JSON file, which is where the instances will be created.

The specified display name is the name of the instance configuration. If you do not provide a value for the `--display-name` option, the default name of the instance configuration is `instanceconfigurationYYYYMMDDhhmmss`, where `YYYYMMDDhhmmss` is the creation date and time. (See Step 2 for a description of the display name specified in the instance details JSON file.)

The output of this command is the same as the output of the `instance-configuration get` command.

6.3.2 Updating an Instance Configuration

You can change the name of the instance configuration and change the tags. To change configuration such as the compartment, subnet, or image, create a new instance configuration.

Using the OCI CLI

1. Get the OCID of the instance configuration that you want to update.

```
$ oci compute-management instance-configuration list -c ocidl.compartment.unique_ID
```

2. Run the instance configuration update command.

Example:

```
$ oci compute-management instance-configuration update \
--instance-configuration-id ocidl.instanceConfiguration.unique_ID \
--defined-tags file://instcfgdeftags.json
```

6.3.3 Moving an Instance Configuration to a Different Compartment

You can move an instance configuration to a different compartment within the same tenancy. When you move an instance configuration to a different compartment, instances and instance pools created by using this instance configuration are not moved.

New instances and instance pools that are created using this instance configuration are created in the compartment specified in the instance configuration, not in the compartment to which the instance configuration has been moved.

To move an instance configuration, you must use the OCI CLI.

Using the OCI CLI

1. Get the following information:

- The OCID of the current compartment, and the OCID of the destination compartment:

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- The OCID of the instance configuration:

```
$ oci compute-management instance-configuration list -c current_compartment_OCID
```

2. Run the instance configuration change compartment command.

Syntax:

```
oci compute-management instance-configuration change-compartment \
--compartment-id destination_compartment_OCID \
--instance-configuration-id instance_configuration_OCID
```

6.3.4 Deleting an Instance Configuration

An instance configuration that is being used by any pool cannot be deleted.

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instance Configurations.
2. If the instance configuration that you want to delete is not listed, use the Compartment drop-down menu above the instance configurations list to select the correct compartment.
3. Click the name of the instance configuration that you want to delete.
4. On the instance configuration details page, click the Delete button.
5. Click the Confirm button.

Using the OCI CLI

1. Get the OCID of the instance configuration that you want to delete.

```
$ oci compute-management instance-configuration list -c ocid1.compartment.unique_ID
```

2. Run the instance configuration delete command.

Example:

```
$ oci compute-management instance-configuration delete \
--instance-configuration-id ocid1.instanceConfiguration.unique_ID
Are you sure you want to delete this resource? [y/N]: y
```

6.3.5 Using an Instance Configuration to Launch an Instance

This section shows how to use the instance configuration that you created in [Section 6.3.1, “Creating an Instance Configuration”](#) to launch a compute instance.

This method of launching a compute instance is an alternative to the method described in [Section 6.2.1, “Creating an Instance”](#).

The name of the instance will be one of the following:

- If the instance configuration specifies a value for the `displayName` property, the name of the instance will be `displayName`. If you use the same instance configuration with multiple `launch-compute-instance` commands, all instances will have the same name. Instance names are not required to be unique.
- If the instance configuration does not specify a value for the `displayName` property, the default name of the instance will be `instanceYYYYMMDDhhmmss`, where `YYYYMMDDhhmmss` is the creation date and time.

Using the OCI CLI

1. Get the OCID of the instance configuration that you want to use to launch the instance.

```
$ oci compute-management instance-configuration list \
--compartment-id ocidl.compartment.unique_ID
```

2. Run the instance configuration launch instance command.

Example:

```
$ oci compute-management instance-configuration launch-compute-instance \
--instance-configuration-id ocidl.instanceConfiguration.unique_ID
```

The output of this command is the same as the output of the `compute instance get` command with the addition of a work request OCID. Use the `work-requests work-request get` command to check the status of the instance launch.

6.3.6 Creating an Instance Pool

An instance pool is a group of compute instances within the same region.

Performing operations such as `reset` or `terminate` on the pool object performs that operation on all instances that are members of the pool. Performing these operations on an individual instance that is a member of the pool does not affect any other member instances.

Creating an instance pool requires an instance configuration and at least one placement configuration. Instances that are added to the pool in a pool update can be created with different instance and placement configurations.

For instances in a pool, the value of the `displayName` property in the instance configuration is ignored. Instances in a pool are named `inst-aaaaa-pool_name`, where `aaaaa` is five random alphanumeric characters.

Placement Configuration

In addition to an instance configuration, pool creation requires placement configurations. Values specified in a placement configuration override values specified in the instance configuration.

Specify one placement configuration for each availability domain. Each placement configuration can specify fault domains, primary subnet, and secondary VNIC subnets. To use the instance pool with a regional subnet, provide a placement configuration for each availability domain, and include the regional subnet in each placement configuration.

Fault Domains

If you do not specify a fault domain in either the instance configuration or the placement configuration, the system will automatically select the best fault domains for the pool instances. If you specify only a single fault domain, all instances will be placed in only that fault domain. If you specify more than one fault domain, pool instances will be placed in those fault domains evenly, providing better High Availability for the pool.

If one or more fault domains have been specified in the instance configuration or placement configuration, and resource constraints in that fault domain cause an instance launch to fail, the pool will remain in the Scaling state. Once `size` instances are launched, the pool can transition to the Running state.

The following are examples of actions you can take if a pool instance fails to launch because of resource constraints:

- Update the pool and reduce the "Number of instances" or `size` value.
- Update the pool and change the Fault Domains specification in the Compute Web UI or in a new instance or placement configuration.
- Update the pool to specify a new instance configuration that creates instances that require fewer resources.
- Terminate an instance that is not a member of a pool in the same fault domain where the pool instance is failing to launch because of resource constraints.

While the pool is in the Scaling state, pool instances that are in the Running state are available to use.

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instance Configurations.
2. If the instance configuration that you want to use to create this pool is not listed, use the Compartment drop-down menu above the instance configurations list to select the correct compartment.
3. Click the instance configuration that you want to use for the instances in this pool.
4. In the Resources box, click Attached Instance Pools, and then click the Create Instance Pool button.

In the Attach Instance Pool to `instance_configuration_name` dialog, enter the following information:

- *Name*: Enter a name for the instance pool. The name does not need to be unique. This name is used in the names of the created instances.
- *Create in Compartment*: Select a compartment for this instance pool definition. Note that the instances in the pool will be created in the compartment that is specified in the instance configuration.
- *Number of instances*: Specify the number of instances to create in this instance pool.
- *Pool Placement*: Select the Fault Domains, VCN, and Subnet for instances in this instance pool. You can select a different compartment from which to choose the VCN and Subnet. See the descriptions of Placement Configuration and Fault Domains at the beginning of this section.
- *Tagging*: (Optional) Add defined or free-form tags for this instance pool as described in [Section 3.4.1, "Adding Tags at Resource Creation"](#). Tags can also be applied later.

These tags are applied to the pool definition, not to the member instances.

5. Click the Create Instance Pool button.

Click Work Request(s) in the Resources box to check the status of the instance pool create.

Using the OCI CLI

1. Get the following information:

- The OCID of the compartment where you want to create the instance pool definition. Note that the instances in the pool will be created in the compartment that is specified in the instance configuration.

```
$ oci iam compartment list --compartment-id-in-subtree true
```

- The OCID of the instance configuration that you want to use.

```
$ oci compute-management instance-configuration list -c compartment_OCID
```

- The size of the instance pool. This is the number of compute instances in the instance pool.

2. Construct an argument for the `--placement-configurations` option.

See the descriptions of Placement Configuration and Fault Domains at the beginning of this section.

Use the following command to show the content of the placement configurations argument:

```
$ oci compute-management instance-pool create \
--generate-param-json-input placement-configurations > placement_configurations.json
```

3. Run the instance pool create command.

Syntax:

```
oci compute-management instance-pool create -c compartment_OCID \
--instance-configuration-id instance_configuration_OCID \
--placement-configurations file://placement_configuration.json \
--size number_of_instances
```

Example:

```
$ oci compute-management instance-pool create \
--compartment-id ocid1.compartment.unique_ID \
--display-name pool_name \
--instance-configuration-id ocid1.instanceConfiguration.unique_ID \
--placement-configurations file://placement_configurations.json \
--size 10
```

The value of the `--display-name` option is the name of the pool. The pool name is not required to be unique. The pool name is used in the names of the instances. Instances in a pool are named `inst-aaaaa-pool_name`, where `aaaaa` is five random alphanumeric characters.

If you do not provide a value for the `--display-name` option, the default name of the instance pool is `instancepoolYYYYMMDDhhmmss`, where `YYYYMMDDhhmmss` is the creation date and time.

The output of this command is the same as the output of the `instance-pool get` command.

6.3.7 Updating an Instance Pool

When you update an instance pool, you can change anything that you set when you created the instance pool, including the name of the pool.

If you specify a different instance configuration, new instances will be provisioned using the new instance configuration. The configuration of existing instances is not changed.

If you increase the size of the pool, new instances are provisioned. If you decrease the size of the pool, older instances are terminated. You cannot select which instances to terminate when you decrease the size of a pool. You can terminate individual instances that are members of the pool, as described in [Section 6.2.5, “Terminating an Instance”](#). When you directly terminate an instance in a pool, a new instance will be automatically provisioned, to keep the pool at the specified pool size.

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instance Pools.
2. If the instance pool that you want to update is not listed, use the Compartment drop-down menu above the instance pools list to select the correct compartment.
3. Click the name of the instance pool that you want to update.
4. On the instance pool details page, click the Controls menu and click the Edit option.
5. When you are finished editing, click the Update Instance Pool button.

Using the OCI CLI

1. Get the OCID of the instance pool that you want to update.

```
$ oci compute-management instance-pool list -c compartment_OCID
```

2. Run the instance pool update command.

Syntax:

```
oci compute-management instance-pool update \
--instance-pool-id instance_pool_OCID options_with_values_to_update
```

Example:

```
$ oci compute-management instance-pool update \
--instance-pool-id ocid1.instancePool.unique_ID \
--instance-configuration-id new_instance_configuration_OCID --size 20
```

The output of this command is the same as the output of the `instance-pool get` command.

6.3.8 Stopping and Starting Instances in an Instance Pool

Performing operations such as reset or stop on the pool object performs that operation on all instances that are members of the pool. Performing these operations on an individual instance that is a member of the pool does not affect any other member instances.

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instance Pools.
2. If the instance pool that you want to manage is not listed, use the Compartment drop-down menu above the instance pools list to select the correct compartment.
3. Click the name of the pool that you want to manage.
4. On the instance pool details page, click the Controls menu and click the Start, Stop, or Reboot option.

All of the instances in the pool are stopped, started, or rebooted. Click Attached Instances in the Resources box on the pool details page to view the status of the instances.

Click Work Request(s) in the Resources box to check the status of the instance pool stop, start, or reboot.

Using the OCI CLI

1. Get the OCID of the instance pool that you want to manage.

```
$ oci compute-management instance-pool list -c compartment_OCID
```

2. Run the instance pool stop, start, or reset command.

Syntax:

```
oci compute-management instance-pool {start | stop | reset | softreset} \  
--instance-pool-id instance_pool_OCID
```

For descriptions of these commands, enter:

```
$ oci compute-management instance-pool --help
```

Example:

```
$ oci compute-management instance-pool reset --instance-pool-id ocid1.instancePool.unique_ID
```

Use the `work-requests work-request get` command to check the status of the instance pool management change.

6.3.9 Deleting an Instance Pool

When you delete an instance pool, the resources that were created by the pool are permanently deleted, including associated instances, attached boot volumes, and block volumes.

Using the Compute Web UI

1. In the navigation menu, click Compute, and then click Instance Pools.
2. If the instance pool that you want to delete is not listed, use the Compartment drop-down menu above the instance pools list to select the correct compartment.
3. Click the name of the pool that you want to delete.
4. On the instance pool details page, click the Controls menu and click the Delete option.
5. On the confirmation dialog, click the Confirm button.

All of the instances in the pool are terminated. Terminated instances are not attached and therefore are not listed in Attached Instances in the Resources box on the pool details page.

Click Work Request(s) in the Resources box to check the status of the instance pool delete.

Using the OCI CLI

1. Get the OCID of the instance pool that you want to terminate.

```
$ oci compute-management instance-pool list -c compartment_OCID
```

2. Run the instance pool terminate command.

Example:


```
$ oci compute-management instance-pool terminate \  
--instance-pool-id ocidl.instancePool.unique_ID  
Are you sure you want to delete this resource? [y/N]: y  
{  
  "etag": "34153f54-0cc9-4e6b-bc02-328166efbb4a",  
  "opc-work-request-id": "ocidl.workrequest.unique_ID"  
}
```

Use the `work-requests work-request get` command to check the status of the instance pool terminate.

6.4 Connecting to a Compute Instance

6.4.1 Prerequisites

You need this information to connect to an instance:

- **The public IP address of the instance.**

You can get the address from the Instance Details page in the Web UI. Open the navigation menu and click Compute. Under Compute, click Instances. Then, select your instance, and click the Networking tab.

- **For UNIX-type instances: The full path to the private key portion of the SSH key pair that you used when you launched the instance.**

For more information about key pairs, see [Section 6.4.2, “Managing Key Pairs”](#).

- **The initial user name for the instance.**

The initial user name for an instance is determined by the image that was used to create the instance. Images fall into these categories:

- **Images provided with Oracle Private Cloud Appliance:**

If you used an image that is provided with the appliance such as Oracle Linux or Oracle Solaris to launch the instance, the user name is `opc`.

- **Custom images:**

The initial password depends on how the image was configured before it was imported as a custom image.

- **(In some circumstances) The initial user's password.**

The initial user password for an instance is determined by the image that was used to create the instance. Images fall into these categories:

- **Images provided with Oracle Private Cloud Appliance:**

Instances launched using an Oracle Linux or Oracle Solaris image that was provided by Oracle use SSH to authenticate a user, and there is no initial password required.

- **Custom images:**

The initial password depends on how the image was configured before it was imported as a custom image.

6.4.2 Managing Key Pairs

The method you use to log into an instance depends on how the image that was used to launch the instance was configured.

- **Images provided with Oracle Private Cloud Appliance** launch instances that use an SSH key pair instead of a password to authenticate a remote user. These images also include the `cloud-init` toolkit (required for SSH authentication) in launched instances.
- **Custom images** might be configured with the `cloud-init` toolkit and use SSH for authentication. Or, the image might be configured to use their own set of credentials to authenticate a user, for example they might require a password. If the latter is the case, you don't need to create an SSH key pair.

Note – Only instances that were created with the `cloud-init` toolkit can use SSH key pairs.

A key pair consists of a private key and public key. You keep the private key on your computer and provide the public key when you create an instance. When you connect to the instance using SSH, you provide the path to the private key in the SSH command.

You can have as many key pairs as you want, or you can keep it simple and use one key pair for all or several of your instances.

To create your own key pairs, you can use a third-party tool such as OpenSSH on UNIX-style systems (including Linux, Oracle Solaris, BSD, and OS X) or PuTTY Key Generator on Windows.

Required SSH Public Key Format

If you provide your own key pair, it must use the OpenSSH format.

A public key has the following format:

```
<key_type> <public_key> <optional_comment>
```

For example, an RSA public key looks like this:

```
ssh-rsa AAAAB3BzaC1yc2EAAAADAQABAAQD9BRwrUiLDki6P0+jZhwsjS2muM...  
...yXDus/5DQ== rsa-key-20201202
```

For images provided with the appliance, these SSH key types are supported: RSA, DSA, DSS, ECDSA, and Ed25519.

If you bring your own image, you're responsible for managing the SSH key types that are supported.

For RSA, DSS, and DSA keys, a minimum of 2048 bits is recommended. For ECDSA keys, a minimum of 256 bits is recommended.

Prerequisites

- If you're using a UNIX-style system, you probably already have the `ssh-keygen` utility installed. To determine whether it's installed, type `ssh-keygen` on the command line. If it's not installed, you can download OpenSSH for UNIX from <http://www.openssh.com/portable.html> and install it.
- If you're using a Windows operating system, you will need PuTTY and the PuTTY Key Generator. Download PuTTY and PuTTYgen from <http://www.putty.org> and install them.

6.4.2.1 Creating an SSH Key Pair on the Command Line

1. Open a shell or terminal for entering the commands.

2. At the prompt, enter `ssh-keygen` and provide a name for the key when prompted. Optionally, include a passphrase.

The keys will be created with the default values: RSA keys of 2048 bits.

3. Do one of the following:

- **On UNIX-type systems:**

Use this command to set the file permissions so that only you can read the private key file:

```
chmod 400 <private_key_file>
```

`<private_key_file>` is the full path and name of the file that contains the private key associated with the instance you want to access.

- **On a Windows system using OpenSSH:**

- a. In Windows Explorer, navigate to the private key file, right-click the file, and then click Properties.
- b. On the Security tab, click Advanced.
- c. Ensure that the Owner is your user account.
- d. Click Disable Inheritance, and then select Convert inherited permissions into explicit permissions on this object.
- e. Select each permission entry that is not your user account and click Remove.
- f. Ensure that the access permission for your user account is Full control.
- g. Save your changes.

6.4.2.2 Creating an SSH Key Pair Using PuTTY Key Generator

Using PuTTY Key Generator:

1. Open `puttygen.exe` on your computer.
For example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
2. Specify a key type of SSH-2 RSA and a key size of 2048 bits:
 - In the Key menu, confirm that the default value of SSH-2 RSA key is selected.
 - For the Type of key to generate, accept the default key type of RSA.
 - Set the Number of bits in a generated key to 2048 if it is not already set.
3. Click Generate.
4. Move your mouse around the blank area in the PuTTY window to generate random data in the key.
When the key is generated, it appears under Public key for pasting into OpenSSH `authorized_keys` file.
5. Leave the Key passphrase field blank.

6. Click Conversions, then click Export OpenSSH key.

When prompted to save this key without a passphrase, click Yes.

7. When prompted to save the private key, select a location and name of your choice.
8. Select all of the generated key that appears under Public key for pasting into OpenSSH authorized_keys file, copy it using Ctrl + C, paste it into a text file, and then save the file in the same location as the private key. (**Do not use Save public key because it does not save the key in the OpenSSH format.**)

You can name the key anything you want, but for consistency, use the same name as the private key and a file extension of `.pub`. For example, `mykey.pub`.

9. Do one of the following:

- **On UNIX-type systems:**

Use this command to set the file permissions so that only you can read the private key file:

```
chmod 400 <private_key_file>
```

`<private_key_file>` is the full path and name of the file that contains the private key associated with the instance you want to access.

- **On a Windows system:**

- a. In Windows Explorer, navigate to the private key file, right-click the file, and then click Properties.
- b. On the Security tab, click Advanced.
- c. Ensure that the Owner is your user account.
- d. Click Disable Inheritance, and then select Convert inherited permissions into explicit permissions on this object.
- e. Select each permission entry that is not your user account and click Remove.
- f. Ensure that the access permission for your user account is Full control.
- g. Save your changes.

10. Write down the names and location of your public and private key files. You will need the public key when launching an instance. You will need the private key to access the instance via SSH.

6.4.3 Connecting to a Linux or Oracle Solaris Instance

You can connect to a running instance by using a Secure Shell (SSH) or Remote Desktop connection. Most UNIX-style systems include an SSH client by default.

Note – If you created an instance without an SSH key, you can stop the instance, attach the boot volume to a new instance, and configure SSH on the new instance.

6.4.3.1 Connecting from a UNIX-Type System

1. Open a terminal window or shell.
2. Use this command to connect to the instance:

```
ssh -i <private_key_file> <username>@<public-ip-address>
```

- `<private_key_file>` is the full path and name of the file that contains the private key associated with the instance you want to access.
- `<username>` is the default username for the instance. See [Section 6.4.1, “Prerequisites”](#).
- `<public-ip-address>` is your instance IP address that you can get from the Web UI. See [Section 6.1.10, “Get the Instance IP Address”](#).

6.4.3.2 Connecting from Windows using OpenSSH

1. Open Windows PowerShell.
2. Use this command to connect to the instance:

```
ssh -i <private_key_file> <username>@<public-ip-address>
```

- `<private_key_file>` is the full path and name of the file that contains the private key associated with the instance you want to access.
- `<username>` is the default username for the instance. See [Section 6.4.1, “Prerequisites”](#).
- `<public-ip-address>` is your instance IP address that you can get from the Web UI. See [Section 6.1.10, “Get the Instance IP Address”](#).

6.4.3.3 Connecting from Windows using PuTTY

If the instance uses a key pair that you created using PuTTY Key Generator (see [Section 6.4.2.2, “Creating an SSH Key Pair Using PuTTY Key Generator”](#)), use the following procedure.

1. Open PuTTY.
2. In the Category pane (on the left), select Session and enter the following:
 - **Host Name (or IP address):** `<username>@<public-ip-address>`
 - `<username>` is the default username for the instance. For instances launched from images provided with Oracle Private Cloud Appliance, the default username is `opc`.
 - `<public-ip-address>` is your instance IP address.
 - **Port:** 22
 - **Connection type:** SSH
3. In the Category pane, expand Window, and then select Translation.
4. In the Remote character set drop-down list, select UTF-8. The default locale setting on Linux-based instances is UTF-8, and this configures PuTTY to use the same locale.
5. In the Category pane, expand Connection, expand SSH, and then click Auth.
6. Click Browse, and then select your .ppk private key file.
7. Click Open to start the session.

If this is your first time connecting to the instance, you might see a message that the server's host key is not cached in the registry. Click Yes to continue the connection.

Tip – If the connection fails, you may need to update your PuTTY proxy configuration.

Consider your next actions.

- Add storage. See [Chapter 7, Block Volume Storage](#), and [Chapter 9, Object Storage Chapter 8, File System Storage](#).
- Install software on the instance.
- Configure and enable additional users to connect to your instance.

The utilities you use to perform the administrative tasks vary depending on the type of OS in the instance. For additional administrative information, refer to the documentation for the version of the OS. These documentation libraries provide helpful information:

- Oracle OS Documentation: <https://docs.oracle.com/en/operating-systems/index.html>
- Oracle Virtualization Documentation: <https://docs.oracle.com/en/virtualization/index.html>

6.4.4 Connecting to a Windows Instance

You can connect to a Windows instance using a Remote Desktop connection. Most Windows systems include a Remote Desktop client by default.

6.4.4.1 Enabling Remote Desktop Protocol (RDP) Access

To enable Remote Desktop Protocol (RDP) access to the Windows instance, you need to add a stateful ingress security rule for TCP traffic on destination port 3389 from source 0.0.0.0/0 and any source port. You can implement this security rule in either a network security group that the Windows instance belongs to, or a security list that is used by the instance's subnet.

1. In the navigation menu, click Networking, then click Virtual Cloud Networks.
2. Choose a compartment you have permission to work in.
3. Click the VCN that you're interested in.
4. Perform one of the following actions:
 - **Add rules to an NSG:**
 - a. Under Resources, click Network Security Groups. Then click the network security group that you're interested in.
 - b. Click Add Rules.
 - c. Enter the following values for the rule:
 - **Stateless:** Leave the check box cleared.
 - **Source Type:** CIDR
 - **Source CIDR:** 0.0.0.0/0

- **IP Protocol:** TCP
 - **Source Port Range:** All (leave empty)
 - **Destination Port Range:** 3389
 - **Description:** An optional description of the rule.
- d. Click Add.
- **Add rules to a security list:**
 - a. Under Resources, click Security Lists. Then click the security list you're interested in.
 - b. Click Create Ingress Security Rule.
 - c. Enter the following values for the rule:
 - **Stateless:** Leave the check box empty.
 - **Source Type:** CIDR
 - **Ingress CIDR:** 0.0.0.0/0
 - **IP Protocol:** TCP
 - **Source Port Range:** All (leave empty)
 - **Destination Port Range:** 3389
 - **Description:** An optional description of the rule.
 - d. Click Add Ingress Rules.

6.4.4.2 Connecting with an RDP Client

Using the Compute Web UI

1. Open the Remote Desktop client.
2. In the Computer field, enter the public IP address of the instance. You can retrieve the public IP address from the Web UI. See [Section 6.1.10, "Get the Instance IP Address"](#).
3. The User name depends on how the image was configured. If you don't know the user name, consult with your administrator.

Note –Depending on the Remote Desktop client you are using, you might have to connect to the instance before you can enter this credential.

4. Click Connect to start the session.
5. Accept the certificate if you are prompted to do so.
6. If you are connecting to the instance for the first time, enter the initial password that was provided to you by your administrator when you launched the instance. You will be prompted to change the password as soon as you log in. Your new password must be at least 12 characters long and must comply with Microsoft's password policy.

Otherwise, enter the password that you created. If you are using a custom image, you might need to know the password for the instance that the image was created from.

7. Press Enter.

Consider your next actions.

- Add storage. See [Chapter 7, Block Volume Storage](#), and [Chapter 9, Object Storage Chapter 8, File System Storage](#).
- Install software on the instance.
- Configure and enable additional users to connect to your instance.

The utilities you use to perform the administrative tasks vary depending on the type of OS in the instance. For additional administrative information, refer to the documentation for the version of the OS.

6.4.5 Connecting to an Instance Using a Console Connection

Important

Instance console connections are for troubleshooting purposes only. To connect to a running instance for administration and general use, instead use a Secure Shell (SSH) or Remote Desktop connection. See [Section 6.4.3, “Connecting to a Linux or Oracle Solaris Instance”](#) and [Section 6.4.4, “Connecting to a Windows Instance”](#).

You can remotely troubleshoot malfunctioning instances using console connections. For example:

- An imported or customized image that does not complete a successful boot
- A previously working instance that stops responding

6.4.5.1 Prerequisites

Ensure that you have these items on the system you plan to use to connect to the instance console.

- RSA SSH key pair (See [Section 6.4.2, “Managing Key Pairs”](#))
- SSH client and command-line shell
- VNC viewer
- (Windows systems) `plink.exe` – the command link connection tool included with PuTTY. You can install PuTTY or install `plink.exe` separately. Refer to <https://www.putty.org/>.

6.4.5.2 Create an Instance Console Connection

Before you can make a connection to an instance serial console you need to create an instance console connection.

Note – Instance console connections are limited to one client at a time. If the client fails, the connection remains active for approximately five minutes. During this time, no other client can connect. After five minutes, the connection is closed, and a new client can connect. During the five-minute timeout, any attempt to connect a new client fails.

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the appropriate compartment where your instance resides.
3. Click the name of the instance you want to connect to.
4. Under Resources, click Console Connection.
5. Click Create Console Connection.
6. Provide the public key portion for the SSH key.

In the dialog box, enter your public SSH key in one of these ways.

- Drag and drop the public key into the Drag and Drop area.
 - Click the Drag and Drop area to browse to your public key.
 - Click Paste the key selection and paste the public key into the contents box.
7. Click Create Console Connection.
 8. Go to the next task.

See [Section 6.4.5.3, “Make a VNC Connection to the Serial Console”](#).

6.4.5.3 Make a VNC Connection to the Serial Console

After you create the console connection for the instance, you need to set up a secure tunnel to the VNC server on the instance, and then connect with a VNC client.

The VNC console connection uses SSH port forwarding to create a secure connection from your local system to the VNC server attached to your instance's console.

Caution

Although this method is a secure way to use VNC over the internet, owners of multiuser systems should know that opening a port on the local system makes it available to all users on that system until a VNC client connects. For this reason, we don't recommend using this product on a multiuser system unless you take proper actions to secure the port or you isolate the VNC client by running it in a virtual environment.

Use one of the following procedures based on the type of system you are using to connect to the instance's console.

Connecting From a Linux or Mac OS X System

This procedure sets up a secure tunnel to the VNC server on the instance using OPENSSH on Linux or Mac OS X. Mac OS X and most Linux and UNIX-like operating systems include the SSH client OpenSSH by default.

Using the Compute Web UI

Note – Remote management for Remote Desktop on OS X uses port 5900. Because VNC console connections in Oracle Private Cloud Appliance also use port 5900, VNC console connections are not compatible with remote management. To use VNC console connections, disable remote management

1. (If not already there) On the instance details page, under Resources, click Console Connection.
2. For the active console connection, click the Actions menu, and then click Copy VNC Connection for Linux/Mac.

This places an SSH command line string in your copy/paste buffer.

3. Paste the connection string into a terminal window, and then press Enter to set up the secure connection.
4. After the connection is established, open your VNC client and specify `localhost` as the host to connect to, and set the port to what was listed in the string.

Connecting From a PowerShell on Windows

Using the Compute Web UI

This procedure sets up a secure tunnel to a VNC server on the instance using PowerShell on Windows.

1. (If not already there) On the instance details page, under Resources, Under Resources, click Console Connection.
2. Click the Actions menu, and then click Copy VNC Connection for Windows.

This places a `plink.exe -ssh` command line string in your copy/paste buffer.

3. Check the copied connection string.

The copied connection string contains the parameter `-i` which specifies the location of the private key file. The default value for this parameter in the connection string references an environment variable which might not be configured on your Windows client, or it might not represent the location where the private key file is saved.

Tip – Paste the string into a text editor where you can view and edit the value for the `-i` parameter before proceeding to the next step.

4. Paste the connection string copied into a Windows Powershell, and then press Enter to set up the secure connection.
5. After the connection is established, open your VNC client and specify `localhost` as the host to connect to, and set the port to what was listed in the string.

Note – When you connect, you might see a warning from the VNC client that the connection is not encrypted. Because you are connecting through SSH, the connection is secure, so this warning is not an issue.

Chapter 7 Block Volume Storage

Table of Contents

7.1 Block Volumes	242
7.2 Creating and Attaching Block Volumes	242
7.2.1 Task Flow to Attach Block Volumes to Instances	242
7.2.2 Creating a Block Volume	242
7.2.3 Attaching a Volume	244
7.2.4 Attaching a Volume to Multiple Instances	246
7.2.5 Find Your Volume in the Instance	247
7.2.6 Configuring Volumes to Automatically Mount (Linux Instances)	249
7.3 Managing Block Volumes	250
7.3.1 Listing Block Volumes and Block Volume Details	250
7.3.2 Listing Block Volume Attachments	252
7.3.3 Editing a Volume's Configuration	254
7.3.4 Moving a Volume to a Different Compartment	255
7.3.5 Cloning a Volume	256
7.3.6 Detaching a Block Volume	258
7.3.7 Deleting a Block Volume	259
7.4 Managing Boot Volumes	259
7.4.1 Overview	260
7.4.2 Listing Boot Volumes	260
7.4.3 Listing Boot Volume Attachments	261
7.4.4 Detaching a Boot Volume	262
7.4.5 Reattaching a Boot Volume	262
7.4.6 Deleting a Boot Volume	263
7.5 Resizing Volumes	264
7.5.1 Resizing Overview	264
7.5.2 Online Volume Resizing	265
7.5.3 Offline Volume Resizing	267
7.6 Managing Volume Groups	271
7.6.1 Overview	271
7.6.2 Viewing the Volumes in a Volume Group	271
7.6.3 Creating a Volume Group from Existing Volumes	273
7.6.4 Adding Volumes to a Group	274
7.6.5 Removing Volumes from a Group	275
7.6.6 Creating a Clone of a Volume Group	276
7.6.7 Deleting a Volume Group	277
7.7 Backing Up Block Volumes	277
7.7.1 Block Volume Backups Overview	277
7.7.2 Viewing Volume Backups	278
7.7.3 Creating a Manual Boot or Block Volume Backup	279
7.7.4 Creating a Manual Backup of a Volume Group	281
7.7.5 Restoring a Backup to a New Volume	283
7.7.6 Restoring a Volume Group from a Volume Group Backup	284
7.8 Managing Backup Policies	286
7.8.1 Overview	286
7.8.2 Creating Backup Policies and Schedules	286
7.8.3 Accessing the Backups	289
7.8.4 Viewing Backup Policies	290
7.8.5 Editing a Backup Policy Schedule	291

7.8.6 Deleting a Backup Policy Schedule	292
7.8.7 Deleting a Backup Policy	293

7.1 Block Volumes

Block Volumes provide high-performance network storage capacity that supports a broad range of I/O intensive workloads.

You can use block volumes to expand the storage capacity of your compute instances, to provide durable and persistent data storage that can be migrated across compute instances, and to host large databases.

The Block Volume service lets you dynamically provision and manage block storage volumes. You can create, attach, connect, and move volumes to meet your storage, performance, and application requirements.

After a block volume is created, you attach the volume to one or more instances. You can use the volume like a regular hard drive. You can also disconnect a volume and attach it to another instance without the loss of data.

There are two types of volumes:

- **Block volume:** A detachable block storage device that allows you to dynamically expand the storage capacity of an instance.
- **Boot volume:** A detachable boot volume device that contains the image used to boot a Compute instance.

For more conceptual information, refer to the [Block Volume Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#)

This section provides instructions for managing Block Volumes.

7.2 Creating and Attaching Block Volumes

7.2.1 Task Flow to Attach Block Volumes to Instances

No.	Task	Links
1.	Create a block volume.	Section 7.2.2, "Creating a Block Volume"
2.	Attach the block volume to one or more instances.	Section 7.2.3, "Attaching a Volume" or Section 7.2.4, "Attaching a Volume to Multiple Instances"
3.	Identify the added block volume and perform administrative tasks.	Section 7.2.5, "Find Your Volume in the Instance"
4.	Configure the volume to automatically mount when the instance is rebooted.	Section 7.2.6, "Configuring Volumes to Automatically Mount (Linux Instances)"

7.2.2 Creating a Block Volume

Block volumes are created using the Block Volume service.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Click Create Block Volume.
3. Fill in the required volume information:
 - **Name:** Provide a name or description for the volume. Avoid entering confidential information.
 - **Compartment:** Select the compartment in which to create the block volume.
 - **Size (in GBs):** Must be between 50 GB and 32 TB and specified in 1 GB increments.
 - **High Performance Volume:** Optionally, choose to create a volume using the high performance feature. If not enabled, the volume uses balanced performance.
 - **High performance:** The Higher Performance elastic performance option is recommended for workloads with the highest I/O requirements, requiring the best possible performance, such as large databases.
 - **Balanced performance:** Suitable for most applications including boot volumes.

For more information, refer to the *Block Volume Performance Options* section in [Block Volume Storage](#) chapter in the Concepts Guide.

- **Backup Policy:** Optionally, you can enable the use of a backup policy for this volume. Select a backup policy from the drop-down list.

For more information about backup policies, see [Section 7.8, “Managing Backup Policies”](#).

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

4. Click Create Block Volume.

The volume is ready to attach to an instance once its icon lists the volume in the Available state.

See [Section 7.2.3, “Attaching a Volume”](#).

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
2. Run the `oci bv volume create` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci bv volume create
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
```

Example:

```
oci bv volume create \
--availability-domain MyAD \
--compartment-id ocid1.compartment.....uniqueID
{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "MyAD",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "volume20210106171509",
    "freeform-tags": {},
    "id": "ocid1.volume.....uniqueID",
    "is-auto-tune-enabled": null,
    "is-hydrated": null,
    "kms-key-id": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 50,
    "size-in-mbs": 51200,
    "source-details": null,
    "system-tags": null,
    "time-created": "2021-06-01T17:15:09+00:00",
    "volume-group-id": null,
    "vpus-per-gb": 0
  },
  "etag": "d507e0b7-cffa-4eae-bfdb-ed81bc132d5"
}
```

The volume is ready to attach to an instance once it is no longer in the PROVISIONING state. See [Section 7.2.3, “Attaching a Volume”](#).

7.2.3 Attaching a Volume

You can attach a volume to an instance to expand the available storage on the instance. You can attach volumes to more than one instance at a time, see [Section 7.2.4, “Attaching a Volume to Multiple Instances”](#).

You can also attach a boot volume that has been detached from its instance. This is convenient for troubleshooting a boot volume and for performing administrative activities while the boot volume is detached from its instance.

Important

Only attach Linux volumes to Linux instances and Windows volumes to Windows instances.

Important

If you are reattaching a volume that was detached, it might be associated with a different device name and the instance operating system might not recognize the volume.

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the compartment where the instance resides.
3. In the Instances list, click the instance that you want to attach a volume to.
4. In the lower left panel, under Resources, select Attached Block Volumes.
5. In the Attached Block Volumes panel, click Attach Block Volume.
6. Select the compartment where the block volume resides.
7. Select a Block Volume.
8. Select one of the access methods:
 - **Read/Write:** (Default) Configures the volume attachment with read/write capabilities. It cannot be shared with other instances. This enables attachment to a single instance only.
 - **Read/Write - Shareable:** Configures the volume attachment as read/write, shareable with other instances. This enables read/write attachment to multiple instances.
 - **Read Only - Shareable:** Configures the volume attachment as read-only, enabling attachment to multiple instances
9. Click Attach to Instance.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Instance OCID (`oci compute instance list --compartment-id <compartment_OCID>`)
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
2. Run the `oci compute volume-attachment attach` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci compute volume-attachment attach
--instance-id <instance_OCID>
--volume-id <volume_OCID>
--type paravirtualized
```

Example:

```
oci compute volume-attachment attach \
--instance-id ocid1.instance.....uniqueID \
--volume-id ocid1.volume.....uniqueID \
--type paravirtualized
{
  "data": {
    "attachment-type": "paravirtualized",
```

```

"availability-domain": "MyAD",
"compartment-id": "ocidl.compartment.....uniqueID",
"device": null,
"display-name": "volumeattachment.....uniqueID",
"id": "ocidl.volumeattachment.....uniqueID",
"instance-id": "ocidl.instance.....uniqueID",
"is-pv-encryption-in-transit-enabled": null,
"is-read-only": false,
"is-shareable": false,
"lifecycle-state": "ATTACHED",
"time-created": "2021-06-01T17:24:13+00:00",
"volume-id": "ocidl.volume.....uniqueID"
}
}

```

7.2.4 Attaching a Volume to Multiple Instances

The Block Volume service provides the capability to attach a block volume to multiple compute instances. With this feature, you can share block volumes across instances in read/write or read-only mode. Attaching block volumes as read/write and shareable enables you to deploy and manage cluster-aware solutions.

There are important limitations and considerations for attaching volumes to multiple instances. For more information, refer to the [Block Volume Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#)

Important

If you are attaching a volume that was detached, it might be associated with a different device name and the instance operating system might not recognize the volume.

Configuring Multiple Instance Volume Attachments with Read/Write Access

The Block Volume service does not provide coordination for concurrent write operations to volumes attached to multiple instances. To prevent data corruption from uncontrolled read/write operations you must install and configure a cluster aware system or solution such as Oracle Cluster File System version 2 (OCFS2) on top of the shared storage before you can use the volume.

This is a summary of the required steps:

1. Attach the block volume to an instance as Read/Write-Shareable using the Web UI, CLI, or API.
See [Section 7.2.3, "Attaching a Volume"](#)
2. Set up your OCFS2/O2CB cluster nodes.
3. Create your OCFS2 file system and mount point.

Configuring Multiple Instance Volume Attachments with Read-Only Mode

Once you attach a block volume to an instance as read-only, it can only be attached to other instances as read-only. If you want to attach the block volume to an instance as read/write, you need to detach the block volume from all instances and then you can reattach the block volume to instances as read/write.

1. Attach the block volume to an instance as read-only using the Web UI, CLI, or API.
See [Section 7.2.3, "Attaching a Volume"](#)
2. Attach the block volume to additional instances as read-only using the Web UI, CLI, or API.
See [Section 7.2.3, "Attaching a Volume"](#)

7.2.5 Find Your Volume in the Instance

When a block volume is initially attached to an instance, the instance sees the volume as a new disk, for example, as device `/dev/sdb`. This procedure describes how to list the disk devices in an instance so that you can find the volume.

For UNIX-based images, if you want to mount these volumes when an instance boots, you need to add the volume to the `/etc/fstab` file. See [Section 7.2.6, “Configuring Volumes to Automatically Mount \(Linux Instances\)”](#).

Optionally, you can perform a variety of administrative tasks to configure the storage to suit your storage requirements.

The utilities you use to perform the administrative tasks vary depending on the type of OS in the instance. For additional administrative information, refer to the documentation for the version of the OS that is on the instance. These documentation libraries provide access to helpful information:

- Oracle OS Documentation: <https://docs.oracle.com/en/operating-systems/index.html>
- Oracle Virtualization Documentation: <https://docs.oracle.com/en/virtualization/index.html>

These steps show you how to identify the boot volume and the attached block volume devices in the instance using Linux commands.

1. Log on to your instance as described in [Section 6.4, “Connecting to a Compute Instance”](#).
2. List the disk devices.

Important

On UNIX type operating systems, the order in which volumes are attached is non-deterministic, so it can change with each reboot. If you refer to a volume using the device name, such as `/dev/sdb`, and you have more than one non-root volume, there is no guarantee that the volume you intend to mount for a specific device name will be the volume mounted. When configuring the OS to recognize the block volume (for example, adding the volume to the `/etc/fstab` file), use the volume's SCSI ID as described in this procedure.

```
ls /dev/sd*
/dev/sda /dev/sda1 /dev/sda2 /dev/sdb
```

In this example, two devices are listed, `/dev/sda` and `/dev/sdb`.

3. Use the `fdisk -l` command to view configuration information about the devices.

In this example, `/dev/sda` is the boot volume. And `/dev/sdb` is the attached block volume.

```
sudo fdisk -l

Disk /dev/sda: 53.7 GB, 53687091200 bytes, 104857600 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 8192 bytes / 8192 bytes
Disk label type: dos
Disk identifier: 0x000af694

   Device Boot Start End Blocks Id System
/dev/sda1 * 2048 2099199 1048576 83 Linux
/dev/sda2 2099200 61442047 29671424 8e Linux LVM
```

```

Disk /dev/mapper/ol-root: 27.2 GB, 27229421568 bytes, 53182464 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 8192 bytes / 8192 bytes

Disk /dev/mapper/ol-swap: 3145 MB, 3145728000 bytes, 6144000 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 8192 bytes / 8192 bytes

Disk /dev/sdb: 1099.5 GB, 1099511627776 bytes, 2147483648 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 8192 bytes / 8192 bytes
    
```

This example output provides this information about `/dev/sda` and `/dev/sdb`:

- The size of `/dev/sda` is 53.7 GB (boot volume).
- `/dev/sda` has two partitions: `/dev/sda1` and `/dev/sda2`.
- The size of `/dev/sdb` is 1099.5 GB (the attached block volume), and does not have any partitions.

4. Identify the devices that have file systems and are mounted in the OS.

```

df -T
Filesystem          Type      1K-blocks    Used Available Use% Mounted on
devtmpfs            devtmpfs  16318164      0  16318164   0% /dev
tmpfs               tmpfs     16332596      0  16332596   0% /dev/shm
tmpfs               tmpfs     16332596    8744  16323852   1% /run
tmpfs               tmpfs     16332596      0  16332596   0% /sys/fs/cgroup
/dev/mapper/ol-root xfs       26578248 2907292  23670956  11% /
/dev/sda1           xfs       1038336 292512   745824   29% /boot
tmpfs               tmpfs     3266520      0   3266520   0% /run/user/0
    
```

In this example:

- `/dev/sda1` has an xfs file system and it is mounted on `/boot` (the boot volume).
- `/dev/sdb` is not listed because this block volume was just attached and hasn't had a file system created and is not mountable yet.

5. Find the SCSI ID for the newly attached volume.

```

ls -l /dev/disk/by-id
total 0
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 dm-name-ol-root -> ../../dm-0
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 dm-name-ol-swap -> ../../dm-1
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 dm-uuid-LVM-83pr2aUrW2ZdCbWgsN4ZRFqvsXGGNZ8J06il7j1YTWpywZeewYCiA6y
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 dm-uuid-LVM-83pr2aUrW2ZdCbWgsN4ZRFqvsXGGNZ8JsaUihE3RWozk5u4p5nOwG9s
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 lvm-pv-uuid-Dh9ydC-Rj90-chhj-tkwq-ZI0Z-mfop-Wtg5bh -> ../../sda2
lrwxrwxrwx. 1 root root  9 Dec  6 18:26 scsi-3600144f096933b92000061ae9bfc0025 -> ../../sda
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 scsi-3600144f096933b92000061ae9bfc0025-part1 -> ../../sda1
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 scsi-3600144f096933b92000061ae9bfc0025-part2 -> ../../sda2
lrwxrwxrwx. 1 root root  9 Dec  8 15:17 scsi-3600144f096933b92000061b1129e0037 -> ../../sdb
lrwxrwxrwx. 1 root root  9 Dec  6 18:26 wwn-0x600144f096933b92000061ae9bfc0025 -> ../../sda
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 wwn-0x600144f096933b92000061ae9bfc0025-part1 -> ../../sda1
lrwxrwxrwx. 1 root root 10 Dec  6 18:26 wwn-0x600144f096933b92000061ae9bfc0025-part2 -> ../../sda2
    
```

```
lrwxrwxrwx. 1 root root 9 Dec 8 15:17 wwn-0x600144f096933b92000061b1129e0037 -> ../../sdb
```

In this example, the following line shows the SCSI ID assigned to sdb:

```
lrwxrwxrwx. 1 root root 9 Dec 8 15:17 scsi-3600144f096933b92000061b1129e0037  
-> ../../sdb
```

where `scsi-3600144f096933b92000061b1129e0037` is the SCSI ID.

The SCSI ID is a persistent device name for `/dev/sdb` and is used when performing administrative operations on the device, such as partitioning, creating a file system, and mounting.

For more information about mounting a block volume file system to an instance, see [Section 7.2.6, “Configuring Volumes to Automatically Mount \(Linux Instances\)”](#).

6. Perform administrative tasks to configure the block volume to suit your storage requirements.

The specific tasks you perform depend on the type of OS that runs the instance and how you want the storage configured. Refer to your OS documentation for details.

7.2.6 Configuring Volumes to Automatically Mount (Linux Instances)

On Linux instances, if you want to automatically mount volumes during an instance boot, you need add the volumes to the `/etc/fstab` file.

Before You Begin

Get the SCSI ID for the block volume you plan to mount.

See [Section 7.2.5, “Find Your Volume in the Instance”](#)

On Linux operating systems, specify the volume SCSI ID in the `/etc/fstab` file instead of the device name (for example, `/dev/sdb`). This is an example of a Volume SCSI ID:

```
/dev/disk/by-id/scsi-3600144f096933b92000061b1129e0037
```

1. Prepare the newly attached block volume for mounting.

Use the disk administration utilities included with instance OS to perform tasks such as the following:

- Partition the volume
- Create file systems on the volume or partitions

Consult the documentation for your instance OS for details.

This is an example of creating an ext4 file system for a block volume attached to a Linux instance:

```
mkfs.ext4 /dev/disk/by-id/scsi-3600144f096933b92000061b1129e0037  
mke2fs 1.42.9 (28-Dec-2013)  
/dev/disk/by-id/scsi-3600144f096933b92000061b1129e0037 is entire device, not just one partition!  
Proceed anyway? (y,n) y  
Filesystem label=  
OS type: Linux  
Block size=4096 (log=2)  
Fragment size=4096 (log=2)  
Stride=2 blocks, Stripe width=2 blocks  
67108864 inodes, 268435456 blocks
```

```
13421772 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2415919104
8192 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

2. Create a mount point for each file system you plan to mount.

```
mkdir /mnt/volume1
```

3. Add the volume to the `/etc/fstab` file.

For this example, the following new line is added to the `/etc/fstab` file . . .

```
/dev/disk/by-id/scsi-3600144f096933b92000061b1129e0037 /mnt/volume1 ext4 _netdev,nofail 0 0
```

. . . using these field values:

- **Device:** Specified using the SCSI ID: `/dev/disk/by-id/scsi-3600144f096933b92000061b1129e003`.
- **Mount point:** The mount point created in the previous step: `/mnt/volume1`
- **Type:** The type of file system, `ext4` in this example.
- **Options:**
 - `_netdev` – Configures the mount process to initiate before the volumes are mounted.
 - `nofail` – No errors are reported for this device if it does not exist. This is a good option to use when an instance is used to create a custom image. Future instances created with that image will not include the block volume and might fail to boot without this option.
- **Dump:** `0` to not use the obsolete dump utility.
- **fsck:** `0` to not run fsck

4. Use this command to mount the volumes that are in the `/etc/fstab` file:

```
sudo mount -a
```

5. Verify that the file system is mounted:

```
mount | grep /mnt
/dev/sdb on /mnt/volume1 type ext4 (rw,relatime,seclabel,stripe=2,data=ordered,_netdev)
```

7.3 Managing Block Volumes

7.3.1 Listing Block Volumes and Block Volume Details

You can list all Block Volume volumes in a specific compartment, as well as detailed information about a single volume.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Select the appropriate compartment.
3. To view block volume details, click the name of the block volume.

The details are displayed.

Table 7.1

Detail Item	Description
Block volume icon (on the left)	Displays the status of the block volume.
Block volume name (at the top)	The name of the block volume.
Block Volume Information and Tags	Tabs that you can click to display: <ul style="list-style-type: none"> • General Information • Tags that have been applied to this object.
Created	The day and time that the volume was created.
Compartment	The compartment that the volume belongs to.
OCID	The Volume's Oracle Cloud ID.
Backup Policy	The backup policy assigned to the volume.
Size	The size of the volume.
Volume is Hydrated	Always no because this field does not apply to Oracle Private Cloud Appliance
High Performance Enabled	Indicates if the volume is configured as a high performance volume or not. If yes, the volume performance units (VPUs) per GB is displayed.
Effective VPUs when Idle	Displays the volume performance units (VPUs) per GB when the volume is idle.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume list --compartment <compartment_OCID>
```

Example:

```
oci bv volume list \
--compartment ocid1.compartment.....uniqueID
```

```

{
  "data": [
    {
      "auto-tuned-vpus-per-gb": null,
      "availability-domain": "MyAD",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "volume2",
      "freeform-tags": {},
      "id": "ocidl.volume.....uniqueID-2",
      "is-auto-tune-enabled": null,
      "is-hydrated": null,
      "kms-key-id": null,
      "lifecycle-state": "AVAILABLE",
      "size-in-gbs": 52,
      "size-in-mbs": 53248,
      "source-details": null,
      "system-tags": null,
      "time-created": "2021-06-01T17:33:24+00:00",
      "volume-group-id": null,
      "vpus-per-gb": 0
    },
    {
      "auto-tuned-vpus-per-gb": null,
      "availability-domain": "MyAD",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "volume20210106171509",
      "freeform-tags": {},
      "id": "ocidl.volume.....uniqueID-1",
      "is-auto-tune-enabled": null,
      "is-hydrated": null,
      "kms-key-id": null,
      "lifecycle-state": "AVAILABLE",
      "size-in-gbs": 50,
      "size-in-mbs": 51200,
      "source-details": null,
      "system-tags": null,
      "time-created": "2021-06-01T17:15:09+00:00",
      "volume-group-id":
    },
  ]
}

```

7.3.2 Listing Block Volume Attachments

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the compartment where the instance resides.
3. Click the instance name to display the details.
4. In the lower left corner, under Resources, select Attached Block Volumes.

All of the block volumes that are attached to this instance are displayed.

5. To see details for a block volume, click the block volume name.

Using the OCI CLI

- **Listing All Block Volumes in a Compartment**

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci compute volume-attachment list
--compartment-id <compartment_OCID>
```

Example:

```
oci compute volume-attachment list \
--compartment-id oocidl.compartment.....uniqueID
{
  "data": [
    {
      "attachment-type": "paravirtualized",
      "availability-domain": "MyAD",
      "compartment-id": "oocidl.compartment.....uniqueID",
      "device": null,
      "display-name": "volumeattachment20210106172413",
      "id": "oocidl.volumeattachment.....uniqueID-2",
      "instance-id": "oocidl.instance.....uniqueID",
      "is-pv-encryption-in-transit-enabled": null,
      "is-read-only": false,
      "is-shareable": false,
      "lifecycle-state": "ATTACHED",
      "time-created": "2021-06-01T17:24:13+00:00",
      "volume-id": "oocidl.volume.....uniqueID-2"
    },
    {
      "attachment-type": "paravirtualized",
      "availability-domain": "MyAD",
      "compartment-id": "oocidl.compartment.....uniqueID",
      "device": null,
      "display-name": "volumeattachment20210106175003",
      "id": "oocidl.volumeattachment.....uniqueID-1",
      "instance-id": "oocidl.instance.....uniqueID",
      "is-pv-encryption-in-transit-enabled": null,
      "is-read-only": false,
      "is-shareable": false,
      "lifecycle-state": "ATTACHED",
      "time-created": "2021-06-01T17:50:03+00:00",
      "volume-id": "oocidl.volume.....uniqueID-1"
    }
  ]
}
```

• Listing Block Volume Attachments for a Specific Instance

1. Gather the information that you need to run the command:

- Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci compute volume-attachment list
```

```
--volume-id <volume_OCID>
```

Example:

```
oci compute volume-attachment list \
--instance-id ocid1.instance.....uniqueID
{
  {
    "attachment-type": "paravirtualized",
    "availability-domain": "MyAD",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "device": null,
    "display-name": "volumeattachment20210106175003",
    "id": "ocid1.volumeattachment.....uniqueID",
    "instance-id": "ocid1.instance.....uniqueID",
    "is-pv-encryption-in-transit-enabled": null,
    "is-read-only": false,
    "is-shareable": false,
    "lifecycle-state": "ATTACHED",
    "time-created": "2021-06-01T17:50:03+00:00",
    "volume-id": "ocid1.volume.....uniqueID"
  }
}
```

7.3.3 Editing a Volume's Configuration

You can change settings for your block volumes and boot volumes while they are online, without any downtime.

You can change these volume settings:

- **Name:** Is the volume display name. The name does not have to be unique.
- **Size (in GB):** Is the capacity of the volume. You can only increase the capacity. Changing the size has repercussions. Before changing the size, see [Section 7.5, “Resizing Volumes”](#)
- **High Performance:** Enable or disable the higher performance option.
- **Backup Policy:** Select a backup policy from the drop-down list.
- **Tagging:** Add, remove, or change tags. For details about tagging, see [Resource Tag Management](#).

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Select the compartment where the block volume resides.
3. In the Actions menu (three dots), select Edit.
4. In the dialog, change the setting.
5. Click Save Changes.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)

2. Run the `oci bv volume update` command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

To change other volume parameters such as the volume size, you specify the parameters in a JSON format. For more information, refer to .

Changing a volume's display name

Syntax (entered on a single line):

```
oci bv volume update
--volume-id <volume_OCID>
--display-name <new-display-name>
```

Example:

```
oci bv volume update \
--volume-id ocidl.volume.....uniqueID \
--display-name volumeA
```

7.3.4 Moving a Volume to a Different Compartment

You can move Block Volume resources such as block volumes, boot volumes, clones, volume backups, volume groups, and volume group backups from one compartment to another.

When you move a volume resource to a new compartment, associated resources such as policies, are not moved. Inherent policies apply immediately and affect access to the resource.

Important

When moving Block Volume resources between compartments you need to ensure that the resource users have sufficient access permissions on the compartment the resource is being moved to.

- You can't move a block volume or boot volume from a security zone to a standard compartment.
- You can't move a block volume or boot volume from a standard compartment to a compartment that is in a security zone if the volume violates any security zone policies.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
 - Target Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume change-compartment
--volume-id <volume_OCID>
```

```
--compartment-id <destination_compartment_OCID>
```

Example:

```
oci bv volume change-compartment \
--volume-id ocidl.volume.....uniqueID \
--compartment-id ocidl.compartment.....uniqueID \
{
  "etag": "7e084c71-4729-4ddd-b131-d87bfc621e8c"
}
```

7.3.5 Cloning a Volume

You can create a clone from a volume using the Block Volume service. Cloning enables you to make a copy of an existing block volume without needing to go through the backup and restore process.

A cloned volume is a point-in-time direct disk-to-disk deep copy of the source volume, so all the data that is in the source volume when the clone is created is copied to the clone volume. Any subsequent changes to the data on the source volume are not copied to the clone.

The clone will be the same size as the source volume unless you specify a larger volume size when you create the clone.

The volume data is being copied in the background, and can take up to thirty minutes depending on the size of the volume. You can attach and use the cloned volume as a regular volume as soon as the state changes to available.

For more information about cloning volumes, refer to the *Block Storage Backups and Clones section* in the [Block Volume Storage](#) chapter.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Select the appropriate compartment.
3. For the volume you plan to clone, click the Actions icon (three dots), and select Create Clone.
4. In the dialog, enter this information:
 - **Name:** A name or description for the volume. Avoid entering confidential information.
 - **Compartment:** Select the compartment in which to clone the block volume.
 - **Size (in GBs):** You can keep the size the same, or increment the size up to 32 TB. You cannot decrease the size. The value is specified in 1 GB increments.
 - **High Performance Volume:** Optionally, choose to create a volume using the high performance feature. If not enabled, the volume uses balanced performance.
 - **High performance:** The Higher Performance elastic performance option is recommended for workloads with the highest I/O requirements, requiring the best possible performance, such as large databases.
 - **Balanced performance:** Suitable for most applications including boot volumes.

For more information, refer to the *Block Volume Performance Options section* in [Block Volume Storage](#) chapter in the Concepts Guide.

- **Backup Policy:** Optionally, you can enable the use of a backup policy for this volume by specifying the following items:

- **Compartment:** Choose the compartment where the backup policy resides.
- **Backup Policy:** Choose a backup policy from the drop-down list.

For more information about backup policies, see [Section 7.8, “Managing Backup Policies”](#).

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create Clone.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Availability Domain Name (`oci iam availability-domain list`)
- Compartment OCID that contains the source volume (`oci iam compartment list`)
- Volume OCID of the volume to clone (`oci bv volume list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume create
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
--source-volume-id <source_volume_OCID>
--display-name <display_name>
```

Example:

```
oci bv volume create \
--availability-domain ad1 \
--compartment-id ocid1.compartment.....uniqueID \
--source-volume-id ocid1.volume.....uniqueID \
--display-name "MyVolumeClone"
{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "ad1",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyVolumeClone",
    "freeform-tags": {},
    "id": "ocid1.volume.....uniqueID",
    "is-auto-tune-enabled": null,
    "is-hydrated": null,
    "kms-key-id": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 1024,
    "size-in-mbs": 51200,
    "source-details": {
```

```

    "id": "ocidl.volume.....uniqueID",
    "type": "volume"
  },
  "system-tags": null,
  "time-created": "2021-07-02T21:12:35+00:00",
  "volume-group-id": null,
  "vpus-per-gb": 0
},
"etag": "13864f86-cd1c-49f7-b414-4c4800103b0c",
"opc-work-request-id": "ocidl.workrequest.....uniqueID"
}

```

7.3.6 Detaching a Block Volume

When an instance no longer needs access to a volume, you can detach the volume from the instance without affecting the volume's data.

Caution

If you later reattach the detached volume, it might be associated with a different device name and the instance operating system might not recognize the volume.

Using the Compute Web UI

1. Perform administrative tasks to remove any dependencies that any instances have for the block volume.

For example, ensure no applications are accessing the volume. Unmount the volume and remove it from the `/etc/fstab` file, and so on.

2. In the navigation menu, click Compute, then click Instances.
3. Select the compartment where the instance resides.
4. In the Instances list, click the instance that has the volume attached.
5. In the lower left corner, under Resources, select Attached Block Volumes.
6. Click the Actions icon (three dots) next to the volume you want to detach, and then click Detach.
7. Confirm when prompted.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume attachment OCID (`oci compute volume-attachment list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci compute volume-attachment detach
--volume-attachment-id <volume-attachment_OCID>
```

Example:

```
oci compute volume-attachment detach \
```

```
--volume-attachment-id ocid1.volumeattachment.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

7.3.7 Deleting a Block Volume

You can delete a volume that is no longer needed.

Caution

You cannot undo this operation. Any data on a volume will be permanently deleted once the volume is deleted.

Caution

All policy-based backups will eventually expire, so if you want to keep a volume backup indefinitely, you need to create a manual backup before deleting the volume. See [Section 7.7, “Backing Up Block Volumes”](#).

Using the Compute Web UI

1. Perform administrative tasks to remove any dependencies that any instances have for the block volume.

For example, ensure no applications are accessing the volume. Unmount the volume and remove it from the `/etc/fstab` file, and so on.

2. In the navigation menu, under Block Storage, click Block Volumes.
3. Select the compartment that contains the block volume that you plan to delete.
4. Click the Action menu (three dots) for the volume you plan to delete, and select Terminate.
5. Confirm the termination when prompted.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume delete
--volume-id <volume_OCID>
```

Example using the minimum required parameters:

```
oci bv volume delete --volume-id ocid1.volume.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
{
  "etag": "bd576de7-3193-4171-9792-uniqueID"
}
```

7.4 Managing Boot Volumes

7.4.1 Overview

When you launch an instance, a new boot volume for the instance is created in the same compartment. That boot volume is associated with that instance until you terminate the instance. When you terminate the instance, you can preserve and reuse the boot volume and its data.

Boot volumes are encrypted by default.

For more conceptual information, refer to the [Block Volume Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

This section provides instructions for managing Boot Volumes.

7.4.2 Listing Boot Volumes

Using the Compute Web UI

1. In the navigation menu, under Compute, click Boot Volumes.
2. Select the appropriate compartment.
A list of boot volumes is displayed.
3. To view details about a boot volume, click the boot volume name.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv boot-volume list
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
```

Example:

```
oci bv boot-volume list
--availability-domain MyAD
--compartment-id ocid1.compartment.....uniqueID
{
  "data": [
    {
      "auto-tuned-vpus-per-gb": null,
      "availability-domain": "MyAD",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "instance20211705214217(Boot Volume)",
      "freeform-tags": {},
      "id": "ocid1.bootvolume.....uniqueID",
      "image-id": "ocid1.image.....uniqueID",
      "is-auto-tune-enabled": null,
      "is-hydrated": null,
```

```

    "kms-key-id": null,
    "lifecycle-state": "AVAILABLE",
    "size-in-gbs": 50,
    "size-in-mbs": 51200,
    "source-details": null,
    "system-tags": null,
    "time-created": "2021-05-17T21:42:17+00:00",
    "volume-group-id": null,
    "vpus-per-gb": 0
  }

```

7.4.3 Listing Boot Volume Attachments

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the appropriate compartment.
3. Click the name of the instance for which you want to view the boot volume attachment.
4. Under Resources, click Attached Boot Volumes.

The boot volume attachments are displayed.

5. To view the details about an attachment, click the boot volume attachment name.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```

oci compute boot-volume-attachment list
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>

```

Example:

```

oci compute boot-volume-attachment list \
--availability-domain MyAD \
--compartment-id ocidl.compartment.....uniqueID
{
  "data": [
    {
      "availability-domain": "MyAD",
      "boot-volume-id": "ocidl.bootvolume.....uniqueID",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "display-name": "bootvolumeattachment20211705214514",
      "id": "ocidl.bootvolumeattachment.....uniqueID",
      "instance-id": "ocidl.instance.....uniqueID",
      "is-pv-encryption-in-transit-enabled": null,
      "lifecycle-state": "ATTACHED",
      "time-created": "2021-05-17T21:45:14+00:00"
    }
  ]
}

```

7.4.4 Detaching a Boot Volume

If you think a boot volume issue is causing a compute instance problem, you can stop the instance and detach the boot volume using the steps described in this topic. Then you can attach it to another instance as a data volume to troubleshoot it.

Note – The instance must be in a stopped state to detach a boot volume (described in this procedure). You must reattach a boot volume before you can start the instance.

Using the Compute Web UI

1. Stop the instance:
 - a. In the navigation menu, under Compute, click Instances.
 - b. Select the appropriate compartment.
 - c. Click the name of the instance.
 - d. Click Controls, then select Stop.

Wait for the status to change from Stopping to Stopped.

2. Under Resources, click Boot Volumes.
3. Click the Actions menu (three dots) for the boot volume, then select Detach.
4. Confirm when prompted.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Boot volume attachment OCID (`oci compute boot-volume-attachment list`)
2. Stop the instance.

See [Section 6.2.4, “Stopping, Starting and Resetting an Instance”](#).
3. Detach the boot volume by running this command.

Syntax (entered on a single line):

```
oci compute boot-volume-attachment detach
--boot-volume-attachment-id <boot_volume_attachment_OCID>
```

Example:

```
oci compute boot-volume-attachment detach \
--boot-volume-attachment-id ocid1.bootvolumeattachment.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

7.4.5 Reattaching a Boot Volume

A boot volume is automatically attached to an instance when the instance is launched. However, there are circumstances when you might need to detach and reattach a boot volume as a data volume for troubleshooting purposes. This procedure describes how to reattach a boot volume as a boot volume.

If a boot volume has been detached from the instance for troubleshooting purposes, you can attach the boot volume to another instance as a data volume. To do this, see [Section 7.2.3, “Attaching a Volume”](#).

Using the Compute Web UI

1. In the navigation menu, click Compute, then click Instances.
2. Select the compartment where the instance resides.
3. Click the instance name.
4. Under Resources, click Attached Boot Volumes.
5. Click Attach Boot Volume.
6. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Boot volume OCID (see [Section 7.4.2, “Listing Boot Volumes”](#))
 - Instance OCID (`oci compute instance list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci compute boot-volume-attachment attach
--boot-volume-id <boot_volume_OCID>
--instance-id <instance_OCID>
```

Example:

```
oci compute boot-volume-attachment attach \
--boot-volume-id ocid1.bootvolume.....uniqueID \
--instance-id ocid1.instance.....uniqueID
{
  "data": {
    "availability-domain": "MyAD",
    "boot-volume-id": "ocid1.bootvolume.....uniqueID",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "display-name": "bootvolumeattachment20212405205526",
    "id": "ocid1.bootvolumeattachment.....uniqueID",
    "instance-id": "ocid1.instance.....uniqueID",
    "is-pv-encryption-in-transit-enabled": null,
    "lifecycle-state": "ATTACHED",
    "time-created": "2021-05-24T20:55:26+00:00"
  }
}
```

7.4.6 Deleting a Boot Volume

When you terminate an instance, you choose to delete or preserve the associated boot volume. You can also delete a boot volume if it has been detached from the associated instance. See [Section 7.4.4, “Detaching a Boot Volume”](#).

Caution

You cannot undo this operation. Any data on a volume will be permanently deleted once the volume is deleted. You will also not be able to restart the associated instance.

Using the Compute Web UI

1. In the navigation menu, under Compute, click Boot Volumes.
2. Select the appropriate compartment.
3. In the Boot Volumes list, find the volume you want to delete.
4. Click the Actions icon (three dots) for the boot volume and select Delete.
5. Confirm when prompted.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Boot volume OCID (see [Section 7.4.2, “Listing Boot Volumes”](#))
2. Run this command.

Syntax (entered on a single line):

```
oci bv boot-volume delete
--boot-volume-id <boot_volume_OCID>
```

You can obtain the `<boot_volume_OCID>` by listing boot volumes. See [Section 7.4.2, “Listing Boot Volumes”](#).

Example:

```
oci bv boot-volume delete \
--boot-volume-id ocid1.bootvolume.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

7.5 Resizing Volumes

7.5.1 Resizing Overview

The Block Volume service lets you expand the size of block volumes and boot volumes. You have several options to increase the size of your volumes:

- Expand an existing volume in place with online resizing.
- Restore from a volume backup to a larger volume.
- Clone an existing volume to a new, larger volume.
- Expand an existing volume in place with offline resizing.

You cannot decrease the size of a volume.

For more conceptual information, refer to the [Block Volume Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#)

7.5.2 Online Volume Resizing

With online resizing, you can expand the volume size without detaching the volume from an instance. Online resizing requires you to rescan the disk and extend the partition.

7.5.2.1 Online Block Volume Resizing

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes.
2. Select the appropriate compartment.
3. For the volume you plan to resize, click the Actions icon (three dots), and select Edit.
4. Change the size:
 - **Size (in GBs):** You can keep the size the same, or increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.
5. Click Save Changes.
6. Rescan the disk.

For details, consult the OS documentation for the OS type and version running in the instance.

7. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume update
-volume-id <volume_OCID>
--size-in-gbs <size_in_GBS>
```

`<size_in_GBS>` is the size of the block volume. You can increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.

Example:

```
oci bv volume update \
-volume-id ocidl.volume.....uniqueID \
--size-in-gbs 72{
  "data": {
    "auto-tuned-vpus-per-gb": null,
```

```

"availability-domain": "ad1",
"compartment-id": "ocidl.compartment.....uniqueID",
"defined-tags": {},
"display-name": "clone-w-policy",
"freeform-tags": {},
"id": "ocidl.volume.....uniqueID",
"is-auto-tune-enabled": null,
"is-hydrated": null,
"kms-key-id": null,
"lifecycle-state": "PROVISIONING",
"size-in-gbs": 72,
"size-in-mbs": 71424,
"source-details": {
  "id": "ocidl.volume.....uniqueID",
  "type": "volume"
},
"system-tags": null,
"time-created": "2021-07-02T20:48:20+00:00",
"volume-group-id": null,
"vpus-per-gb": 0
},
"etag": "58851b71-236d-4d99-8175-b27835d6b34f"
}

```

3. Rescan the disk.

For details, consult the OS documentation for the OS type and version running in the instance.

4. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

7.5.2.2 Online Boot Volume Resizing

Using the Compute Web UI

1. **Boot Volume:** In the navigation menu, under Compute, click Boot Volumes.
2. Select the appropriate compartment.
3. For the volume you plan to resize, click the Actions icon (three dots), and select Edit.
4. Change the size:
 - **Size (in GBs):** You can keep the size the same, or increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.
5. Click Save Changes.
6. Rescan the disk.

For details, consult the OS documentation for the OS type and version running in the instance.

7. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Boot volume OCID (`oci bv boot-volume list --availability-domain <availability_domain_name> --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci bv boot-volume update
--boot-volume-id <volume_OCID>
--size-in-gbs <size_in_GBS>
```

`<size_in_GBS>` is the size of the boot volume. You can increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.

Example:

```
oci bv boot-volume update \
--boot-volume-id ocidl.bootvolume.....uniqueID \
--size-in-gbs 1024

{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "adl",
    "compartment-id": "ocidl.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "MyInstance(Boot Volume)",
    "freeform-tags": {},
    "id": "ocidl.bootvolume.....uniqueID",
    "image-id": "ocidl.image.....uniqueID",
    "is-auto-tune-enabled": null,
    "is-hydrated": null,
    "kms-key-id": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 1024,
    "source-details": null,
    "system-tags": null,
    "time-created": "2021-08-10T20:14:03.053300+00:00",
    "volume-group-id": null,
    "vpus-per-gb": 0
  },
  "etag": "bd0677e3-c542-45f3-bf04-c473b184c795"
}
```

3. Rescan the disk.

For details, consult the OS documentation for the OS type and version running in the instance.

4. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

7.5.3 Offline Volume Resizing

With offline resizing, you detach the volume from an instance before you expand the volume size. Once the volume is resized and reattached, you need to extend the partition, but you do not need to rescan the disk.

7.5.3.1 Considerations When Resizing an Offline Volume

Whenever you detach and reattach volumes, there are complexities and risks for both UNIX-based and Windows-based instances. Keep the following points in mind when resizing volumes:

- Before you resize a volume, you should create a full backup of the volume.
- When you reattach a volume to an instance after resizing, if you are not using consistent device paths, or if the instance does not support consistent device paths, device order and path may change. If you are using a tool such as Logical Volume Manager (LVM), you may need to fix the device mappings.

7.5.3.2 Offline Block Volume Resizing

Using the Compute Web UI

1. Detach the block volume.

See [Section 7.3.6, “Detaching a Block Volume”](#).
2. In the navigation menu, click Block Storage, then click Block Volumes.
3. Select the appropriate compartment.
4. For the volume you plan to resize, click the Actions icon (three dots), and select Edit.
5. Change the size:
 - **Size (in GBs):** You can keep the size the same, or increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.
6. Click Save Changes.
7. Reattach the volume.

See [Section 7.2.3, “Attaching a Volume”](#)
8. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

Using the OCI CLI

1. Detach the block volume.

See [Section 7.3.6, “Detaching a Block Volume”](#).
2. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
3. Run this command.

Syntax (entered on a single line):

```
oci bv volume update
-volume-id <volume_OCID>
--size-in-gbs <size_in_GBS>
```

`<size_in_GBS>` is the size of the block volume. You can increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.

Example:

```
oci bv volume update \
-volume-id ocidl.volume.....uniqueID \
--size-in-gbs 72{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "ad1",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "clone-w-policy",
    "freeform-tags": {},
    "id": "ocidl.volume.....uniqueID",
    "is-auto-tune-enabled": null,
    "is-hydrated": null,
    "kms-key-id": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 72,
    "size-in-mbs": 71424,
    "source-details": {
      "id": "ocidl.volume.....uniqueID",
      "type": "volume"
    },
    "system-tags": null,
    "time-created": "2021-07-02T20:48:20+00:00",
    "volume-group-id": null,
    "vpus-per-gb": 0
  },
  "etag": "58851b71-236d-4d99-8175-b27835d6b34f"
}
```

4. Reattach the volume.

See [Section 7.2.3, “Attaching a Volume”](#).

5. Extend the partition.

For details, consult the OS documentation for the OS type and version running in the instance.

7.5.3.3 Offline Boot Volume Resizing

Using the Compute Web UI

1. Stop the instance.

See [Section 6.2.4, “Stopping, Starting and Resetting an Instance”](#).

2. Detach the boot volume.

See [Section 7.4.4, “Detaching a Boot Volume”](#).

3. In the navigation menu, under Compute, click Boot Volumes.
4. Select the appropriate compartment.
5. For the volume you plan to resize, click the Actions icon (three dots), and select Edit.
6. Change the size:
 - **Size (in GBs):** You can keep the size the same, or increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.
7. Click Save Changes.

8. Attach the boot volume to a second instance as a data volume.
See [Section 7.2.3, "Attaching a Volume"](#).
9. Extend the partition and grow the file system.
For details, consult the OS documentation for the OS type and version running in the instance.
10. Detach the data volume.
See [Section 7.3.6, "Detaching a Block Volume"](#).
11. Reattach the boot volume.
See [Section 7.4.5, "Reattaching a Boot Volume"](#).
12. Restart the instance.
See [Section 6.2.4, "Stopping, Starting and Resetting an Instance"](#).

Using the OCI CLI

1. Stop the instance.
See [Section 6.2.4, "Stopping, Starting and Resetting an Instance"](#).
2. Detach the boot volume.
See [Section 7.4.4, "Detaching a Boot Volume"](#).
3. Gather the information that you need to run the command:
 - Boot volume OCID (`oci bv boot-volume list --availability-domain <availability_domain_name> --compartment-id <compartment_OCID>`)
4. Run this command.

Syntax (entered on a single line):

```
oci bv boot-volume update
--boot-volume-id <volume_OCID>
--size-in-gbs <size_in_GBS>
```

`<size_in_GBS>` is the size of the boot volume. You can increment the size. You cannot decrease the size. The value must be between 50 GB and 32 TB and specified in 1 GB increments.

Example:

```
oci bv boot-volume update \
--boot-volume-id ocidl.bootvolume.....uniqueID \
--size-in-gbs 1024

{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "ad1",
    "compartment-id": "ocidl.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "MyInstance(Boot Volume)",
    "freeform-tags": {},
```



```
"id": "ocidl.bootvolume.....uniqueID",
"image-id": "ocidl.image.....uniqueID",
"is-auto-tune-enabled": null,
"is-hydrated": null,
"kms-key-id": null,
"lifecycle-state": "PROVISIONING",
"size-in-gbs": 1024,
"source-details": null,
"system-tags": null,
"time-created": "2021-08-10T20:14:03.053300+00:00",
"volume-group-id": null,
"vpus-per-gb": 0
},
"etag": "bd0677e3-c542-45f3-bf04-c473b184c795"
}
```

5. Attach the boot volume to a second instance as a data volume.

See [Section 7.2.3, “Attaching a Volume”](#).

6. Extend the partition and grow the file system.

For details, consult the OS documentation for the OS type and version running in the instance.

7. Detach the data volume.

See [Section 7.3.6, “Detaching a Block Volume”](#).

8. Reattach the boot volume.

See [Section 7.4.5, “Reattaching a Boot Volume”](#).

9. Restart the instance.

See [Section 6.2.4, “Stopping, Starting and Resetting an Instance”](#).

7.6 Managing Volume Groups

7.6.1 Overview

The Block Volume service provides you with the capability to group together multiple block volumes in a volume group.

You can use volume groups to create volume group backups and clones that are point-in-time and crash-consistent.

For additional details about volume group considerations and limitations, refer to the *Volume Group* section in [Block Volume Storage](#).

7.6.2 Viewing the Volumes in a Volume Group

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the appropriate compartment.
3. In the Volume Groups list, click the volume group you want to view.

4. To view the block volumes for the volume group, in Resources, click Volumes.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group list
--compartment-id <compartment_OCID>
```

Example:

```
oci bv volume-group list \
--compartment-id ocidl.compartment.....uniqueID
{
  "data": [
    {
      "availability-domain": "MyAD",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "myVolumeGroup",
      "freeform-tags": {},
      "id": "ocidl.volumeGroup.....uniqueID",
      "is-hydrated": null,
      "lifecycle-state": "AVAILABLE",
      "size-in-gbs": 150,
      "size-in-mbs": 153600,
      "source-details": {
        "type": "volumeIds",
        "volume-ids": [
          "ocidl.volumeuniqueID-1",
          "ocidl.volume.....uniqueID-2",
          "ocidl.volume.....uniqueID-3"
        ]
      }
    },
    "time-created": "2021-05-26T20:47:06+00:00",
    "volume-ids": [
      "ocidl.volume.....uniqueID-1",
      "ocidl.volume.....uniqueID-2",
      "ocidl.volume.....uniqueID-3"
    ]
  },
  {
    "availability-domain": "MyAD",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "anotherVolumeGroup",
    "freeform-tags": {},
    "id": "ocidl.volumeGroup.....uniqueID",
    "is-hydrated": null,
    "lifecycle-state": "AVAILABLE",
    "size-in-gbs": 100,
    "size-in-mbs": 102400,
    "source-details": {
      "type": "volumeIds",
      "volume-ids": [
        "ocidl.volume.....uniqueID-4",
        "ocidl.volume.....uniqueID-5"
      ]
    }
  }
},
}
```

```

    "time-created": "2021-05-25T19:08:55+00:00",
    "volume-ids": [
      "ocidl.volume.....uniqueID-4",
      "ocidl.volume.....uniqueID-5"
    ]
  }
]
}

```

7.6.3 Creating a Volume Group from Existing Volumes

This procedure assumes that the volumes you plan to group are created. See [Section 7.2.2, “Creating a Block Volume”](#).

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the compartment where you want the volume group created.
3. Click Create Volume Group.
4. Fill in the required volume information:
 - **Name:** A user-friendly name or description for the group.
 - **Compartment:** The compartment for the volume group.
 - **Volumes:** For each volume you want to add, select the compartment containing the volume, select the volume to add. Click **+ Add Volume** to add additional volumes.

5. **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

6. Click Create Volume Group.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
 - Source details – a JSON string or file that lists the all of the volumes you want in the group.

For details about the JSON format, run this command:

```

oci bv volume-group create \
--generate-param-json-input source-details

```

2. Run this command.

Syntax (entered on a single line):

```

oci bv volume-group create

```

```
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
--source-details <json_string> or file://<path_to_JSON_file>
```

Example :

```
oci bv volume-group create \
--availability-domain MyAD \
--compartment-id ocid1.compartment.....unique_ID \
--source-details '{"type": "volumeIds", "volumeIds": ["ocid1.volume.....uniqueID-1", \
"ocid1.volume.....uniqueID-2"]}'
{
  "data": {
    "availability-domain": "MyAD",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "volumegroup20212605212205",
    "freeform-tags": {},
    "id": "ocid1.volumeGroup.....uniqueID",
    "is-hydrated": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 100,
    "size-in-mbs": 102400,
    "source-details": {
      "type": "volumeIds",
      "volume-ids": [
        "ocid1.volume.....uniqueID-1",
        "ocid1.volume.....uniqueID-2"
      ]
    },
    "time-created": "2021-05-26T21:22:05+00:00",
    "volume-ids": []
  },
  "etag": "c7053513-6819-49ad-8785-dd3e2a45272a"
}
```

7.6.4 Adding Volumes to a Group

Note

You cannot add a volume with an existing backup policy assignment to a volume group with a backup policy assignment. You must first remove the backup policy assignment from the volume before you can add it to the volume group.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the compartment that contains the volume group.
3. In the Volume Groups list, click the volume group you want to add the volume to.
4. Click Add Block Volumes.
5. For each block volume you want to add, select the compartment containing the volume and then select the volume to add. Click + Add Volume to add additional volumes.
6. Once you have selected all the block volumes to add to the volume group, click Update Volume Group.

Using the OCI CLI

Caution

Updates to volume-ids replace any existing values. This means that you need to specify the volume IDs for all of the volumes in the volume group each time you update the volume group.

1. Gather the information that you need to run the command:
 - Volume group OCID (`oci bv volume-group list --compartment-id <compartment_OCID>`)
 - Volume OCIDs – a JSON string or file that lists the all of the volumes (existing and new) that you want in the group.

For details about the JSON format, run this command:

```
oci bv volume-group update --generate-param-json-input volume-ids
```

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group update
--volume-group-id <volume_group_OCID>
--volume-ids <volume_OCIDs_JSON>
```

Example:

In this example, the first two volume IDs are already in the volume group, and the third volume ID is added to the group.

```
oci bv volume-group update \
--volume-group-id ocid1.volumeGroup.....uniqueID \
--volume-ids ["ocid1.volume.....uniqueID-1", "ocid1.volume.....uniqueID-2", "ocid1.volume.....uniqueID-3"]
```

7.6.5 Removing Volumes from a Group

When you remove the last volume in a volume group the volume group is deleted.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the appropriate compartment.
3. In the Volume Groups list, click the volume group that contains the volume you plan to remove.
4. In Actions menu (three dots) for the block volume you want to remove, click Remove.
5. Confirm the removal.

Using the OCI CLI

Caution

Updates to volume-ids replace any existing values. This means that you need to specify the volume IDs for all of the volumes in the volume group each time you update the volume group.

You remove a volume from a volume group using the `oci bv volume-group update` command with valid JSON formatting. For details about the syntax and how to specify the volume OCIDs in a valid JSON file, refer to [volume group update in the OCI CLI Command Reference](#).

1. Gather the information that you need to run the command:

- Volume group OCID (`oci bv volume-group list --compartment-id <compartment_OCID>`)
- Volume OCIDs – a JSON string or file that lists the volumes that you want to keep in the group.

For details about the JSON format, run this command:

```
oci bv volume-group update --generate-param-json-input volume-ids
```

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group update
--volume-group-id <volume_group_OCID>
--volume-ids <volume_OCIDs_JSON>
```

Example:

In this example, the first two volume IDs remain in the volume group. Omitting a volume ID removes it from the group.

```
oci bv volume-group update \
--volume-group-id ocid1.volumeGroup.....uniqueID \
--volume-ids '["ocid1.volume.....uniqueID-1", "ocid1.volume.....uniqueID-2"]'
```

7.6.6 Creating a Clone of a Volume Group

You can create a clone from a volume group. Cloning enables you to make a copy of an existing volume group without needing to go through the backup and restore process.

A cloned volume group is a point-in-time direct disk-to-disk deep copy of the source volume group, so all the data that is in the source volume group is copied to the clone volume group.

Any subsequent changes to the data on the source volume group are not copied to the clone.

For additional details about clones and how they differ from backups, refer to the *Block Storage Backups and Clones* section in the [Block Volume Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the appropriate compartment.
3. Click the name of the Volume Group you plan to clone.
4. Under Resources, click Volume Group Clones.
5. Click Create Volume Group Clone.
6. Enter the required information:

- **Volume Group Clone Name:** Enter a descriptive name for the clone.
- **Create in Compartment:** Select the compartment where the clone will be created.

7. Click Create Volume Group Clone.

7.6.7 Deleting a Volume Group

When you delete a volume group the individual volumes in the group are not deleted, only the volume group is deleted.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Groups.
2. Select the appropriate compartment.
3. In the Volume Groups list, click the volume group you want to delete.
4. On the Volume Group Details page, click Terminate.
5. Confirm the termination.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume group OCID (`oci bv volume-group list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group delete
--volume-group-id <volume_group_OCID>
```

Example:

```
oci bv volume-group delete \
--volume-group-id ocid1.volumegroup.....unique_ID
```

7.7 Backing Up Block Volumes

This section describes how to manually back up block volumes, boot volumes, and volume groups.

7.7.1 Block Volume Backups Overview

The backups feature for the Block Volume service lets you make a point-in-time snapshot of the data on a block volume. You can make a backup of a volume when it is attached to an instance or while it is detached. These backups can then be restored to new volumes either immediately after a backup or at a later time that you choose.

Backups can be restored as new volumes.

There are two ways to initiate a backup:

- **Manual Backups:** Are on-demand one-off backups that you can launch immediately. The backup is a full backup. See [Section 7.7, “Backing Up Block Volumes”](#).
- **Policy-Based Backups:** Are automated scheduled backups as defined by the backup policy assigned to the volume. See [Section 7.8, “Managing Backup Policies”](#).

For more information about Block Volume backups, such as the differences between backups and clones, refer to the *Block Storage Backups and Clones* section in the [Block Volume Storage](#) chapter in the Concepts Guide.

7.7.2 Viewing Volume Backups

Use these steps to view volume backups and volume group backups.

Using the Compute Web UI

1. Perform one of these actions based on the type of backup you want to view:
 - **Block Volume Backups:** In the navigation menu, under Block Storage, click Block Volume Backups.
 - **Volume Group Backups:** In the navigation menu, under Block Storage, click Volume Group Backups.

The backups are listed.

2. If you don't see your backup listed, ensure you are viewing the correct compartment which is displayed at the top of the page.
3. To view the details of a backup, click the backup name.

Using the OCI CLI

• Viewing a Block Volume Backup

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv backup list
--compartment-id <Compartment_OCID>
```

Example:

```
oci bv backup list --compartment-id ocid1.....uniqueID
{
  "data": [
    {
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "Vol2-manual-backup",
      "expiration-time": null,
      "freeform-tags": {},
      "id": "ocid1.volumebackup.....uniqueID",
      "kms-key-id": null,
```



```

    "lifecycle-state": "AVAILABLE",
    "size-in-gbs": 0,
    "size-in-mbs": null,
    "source-type": "MANUAL",
    "source-volume-backup-id": null,
    "system-tags": null,
    "time-created": "2021-06-29T19:02:16.000001+00:00",
    "time-request-received": "2021-06-29T19:02:04.000001+00:00",
    "type": "FULL",
    "unique-size-in-gbs": 0,
    "unique-size-in-mbs": null,
    "volume-id": "ocidl.volume.....uniqueID"
  }
]
}

```

- **Viewing a Volume Group Backup**

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```

oci bv volume-group-backup list
--compartment-id <Compartment_OCID>

```

Example:

```

oci bv volume-group-backup list \
--compartment-id ocidl.compartment.....uniqueID
{
  "data": [
    {
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": null,
      "display-name": "MyVolGrpBackup",
      "freeform-tags": null,
      "id": "ocidl.volumegroupbackup.....uniqueID",
      "lifecycle-state": "AVAILABLE",
      "size-in-gbs": 0,
      "size-in-mbs": 0,
      "time-created": "2021-06-29T21:37:29+00:00",
      "time-request-received": "2021-06-29T21:37:18+00:00",
      "type": "FULL",
      "unique-size-in-gbs": null,
      "unique-size-in-mbs": 0,
      "volume-backup-ids": [
        "ocidl.volumebackup.....uniqueID-1",
        "ocidl.volumebackup.....uniqueID-2"
      ],
      "volume-group-id": "ocidl.volumegroup.....uniqueID"
    }
  ]
}

```

7.7.3 Creating a Manual Boot or Block Volume Backup

This procedure describes how to backup both types of volumes:

- Block Volumes

- Boot Volumes

Using the Compute Web UI

1. Locate the volume you plan to back up:
 - **Block Volume:** In the navigation menu, under Block Storage, click Block Volumes.
 - **Boot Volume:** In the navigation menu, under Compute, click Boot Volumes.
2. If you don't see your volume listed, ensure you are viewing the correct compartment which is displayed at the top of the page.
3. For the volume you plan to backup, click the Action Menu (three dots), then click Create Manual Backup.
4. In the dialog, enter this information:
 - **Name:** Enter a descriptive name for the backup.
 - **Tagging:** Optionally, add one or more tags to this resource.
If you are not sure whether to apply tags, skip this option (you can apply tags later).
For more information about tagging resources, see [Tagging Overview](#).
5. Click Create Manual Backup or Submit.

Using the OCI CLI

- **Creating a Block Volume Backup**

1. Gather the information that you need to run the command:
 - Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv backup create
--volume-id <block_volume_OCID>
--display-name <Backup_Name>
--type FULL
```

Example:

```
oci bv backup create \
--volume-id ocid1.volume.....uniqueID.xyz \
--display-name "ABC-Full-Backup" \
--type FULL
{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "ABC-Full-Backup",
    "expiration-time": null,
    "freeform-tags": {},
    "id": "ocid1.volumeBackup.oc1.....uniqueID.abc",
    "kms-key-id": null,
```

```

"lifecycle-state": "CREATING",
"size-in-gbs": null,
"size-in-mbs": null,
"source-type": "MANUAL",
"source-volume-backup-id": null,
"system-tags": null,
"time-created": "2021-05-19T21:50:11+00:00",
"time-request-received": "2021-05-19T21:50:11+00:00",
"type": "FULL",
"unique-size-in-gbs": null,
"unique-size-in-mbs": null,
"volume-id": "ocid1.volume.....uniqueID.xyz"
},
"etag": "616112e8-728c-43d6-b0d1-c6cfcc1a46e6"
}

```

• Creating a Boot Volume Backup

1. Gather the information that you need to run the command:

- Volume OCID (`oci bv volume list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```

oci bv boot-volume-backup create
--volume-id <block_volume_OCID>
--display-name <Backup_Name>
--type FULL

```

Example:

```

oci bv backup create \
--volume-id ocid1.bootvolume.....uniqueID.123 \
--display-name "mybootvolume-20210424" \
--type FULL

{
  "data": {
    "boot-volume-id": "ocid1.bootvolume.oc1.....uniqueID.123",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "mybootvolume-20210424",
    "expiration-time": null,
    "freeform-tags": {},
    "id": "ocid1.bootvolumeBackup.oc1.....uniqueID.456",
    "image-id": "ocid1.image.oc1.....uniqueID",
    "kms-key-id": null,
    "lifecycle-state": "CREATING",
    "size-in-gbs": null,
    "source-boot-volume-backup-id": null,
    "source-type": "MANUAL",
    "system-tags": null,
    "time-created": "2021-04-24T21:40:13+00:00",
    "time-request-received": "2021-04-24T21:40:13+00:00",
    "type": "FULL",
    "unique-size-in-gbs": null
  },
  "etag": "123a12b3-daa8-4557-8c83-uniqueID"
}

```

7.7.4 Creating a Manual Backup of a Volume Group

Using the Compute Web UI

1. In the navigation menu, click Block Storage, then click Volume Groups.
2. Select the appropriate compartment.
3. Click the name of the Volume Group you plan to back up.
4. Under Resources, click Volume Group Backups.
5. Click Create Volume Group Backup.
6. In the dialog, enter this information:
 - **Name:** Enter a descriptive name for the backup.
7. Click Create Volume Group Backup.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Volume group OCID (`oci bv volume-group list --compartment-id <compartment_OCID>`)
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group-backup create
--volume-group-id <volume_group_OCID>
--compartment-id <compartment_OCID>
--display-name <display_name>
--type FULL
```

Example:

```
oci bv volume-group-backup create \
--compartment-id ocid1.compartment.....uniqueID \
--volume-group-id ocid1.volumeGroup.....uniqueID \
--display-name vol-grp-backup-2205 \
--type FULL
{
  "data": {
    "compartment-id": "cid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "vol-grp-backup-2205",
    "freeform-tags": {},
    "id": "ocid1.volumeGroupBackup.....uniqueID",
    "lifecycle-state": "CREATING",
    "size-in-gbs": null,
    "size-in-mbs": null,
    "time-created": "2021-05-26T21:44:36+00:00",
    "time-request-received": "2021-05-26T21:44:36+00:00",
    "type": "FULL",
    "unique-size-in-gbs": null,
    "unique-size-in-mbs": null,
    "volume-backup-ids": [],
    "volume-group-id": "ocid1.volumeGroup.....uniqueID"
  },
```

```
"etag": "04761386-6ec5-4cfa-b88e-a085ad833eac"
}
```

7.7.5 Restoring a Backup to a New Volume

Using the Compute Web UI

1. In the navigation menu, click Block Storage, then click Block Volume Backups.
2. If you don't see your backup listed, ensure you are viewing the correct compartment which is displayed at the top of the page.
3. Click the Actions icon (three dots) for the block volume backup you want to restore, and click Create Block Volume.
4. Fill in the required volume information:

- **Name:** A name or description for the volume. Avoid entering confidential information.
- **Compartment:** Select the compartment in which to create the block volume.
- **Size (in GBs):** Must be between 50 GB and 32 TB. You can specify a value in 1 GB increments.
- **High Performance Volume:** Optionally, choose to create a volume using the high performance feature. If not enabled, the volume uses balanced performance.
 - **High performance:** The Higher Performance elastic performance option is recommended for workloads with the highest I/O requirements, requiring the best possible performance, such as large databases.
 - **Balanced performance:** Suitable for most applications including boot volumes.

For more information, refer to the *Block Volume Performance Options* section in [Block Volume Storage](#) chapter in the Concepts Guide.

- **Backup Policy:** Optionally, you can enable the use of a backup policy for this volume by specifying the following items:
 - **Compartment:** Choose the compartment where the backup resides.
 - **Backup Policy:** Choose a backup policy from the drop-down list.

For more information about backup policies, see [Section 7.8, "Managing Backup Policies"](#).

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create Block Volume.

The volume is ready to attach once its icon no longer lists it as PROVISIONING in the volume list.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Availability Domain Name (`oci iam availability-domain list`)
- Source volume backup OCID – the OCID of the volume backup from which the data should be restored on the newly created volume.

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume create
--availability-domain <availability_domain_name>
--volume-backup-id <source_volume_backup_OCID>
```

Example:

```
oci bv volume create \
--availability-domain ad1 \
--volume-backup-id ocid1.....uniqueID{
  "data": {
    "auto-tuned-vpus-per-gb": null,
    "availability-domain": "ad1",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "volume20212906194712",
    "freeform-tags": {},
    "id": "ocid1.volume.....uniqueID",
    "is-auto-tune-enabled": null,
    "is-hydrated": null,
    "kms-key-id": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 50,
    "size-in-mbs": 51200,
    "source-details": {
      "id": "ocid1.volumebackup.....uniqueID",
      "type": "volumeBackup"
    },
    "system-tags": null,
    "time-created": "2021-06-29T19:47:12+00:00",
    "volume-group-id": null,
    "vpus-per-gb": 0
  },
  "etag": "13864f86-cd1c-49f7-b414-4c4800103b0c",
  "opc-work-request-id": "ocid1.workrequest.....f4s6nne9dbi4e"
}
```

7.7.6 Restoring a Volume Group from a Volume Group Backup

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Volume Group Backups.
2. Select the appropriate compartment.
3. Click the Actions icon (three dots) for the block volume backup you want to restore, and click Create Volume Group.
4. Fill in the required volume information:
 - **Name:** Enter a descriptive name for the group.
 - **Compartment:** Select the compartment for the volume group.

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create Volume Group.

Using the OCI CLI

You restore a volume group backup using the `oci bv volume-group create` command with valid JSON formatting. For details about the syntax and how to specify the volume OCIDs in a valid JSON file, refer to [volume group create in the OCI CLI Command Reference](#).

1. Gather the information that you need to run the command:

- Availability Domain Name (`oci iam availability-domain list`)
- Compartment OCID (`oci iam compartment list`)
- Source details – a JSON string or file that lists the volume group backup that is used to create the restored volume group.

For details about the JSON format, run this command:

```
oci bv volume-group create --generate-param-json-input source-details
```

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-group create
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
--source-details <json_string> or file://<path_to_JSON_file>
```

Example:

```
oci bv volume-group create
--availability-domain MyAD
--compartment-id ocid1.compartment.....uniqueID
--display-name VolumeGroupBackup-Monday
--source-details '{"type": "volumeGroupBackupId-1", "volumeGroupBackupId-2": ["ocid1.volumeGroupBackupId-1", "ocid1.volumeGroupBackupId-2"]}'

{
  "data": {
    "availability-domain": "MyAD",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "VolumeGroupBackup-Monday",
    "freeform-tags": {},
    "id": "ocid1.volumeGroup.....new-uniqueID",
    "is-hydrated": null,
    "lifecycle-state": "PROVISIONING",
    "size-in-gbs": 0,
    "size-in-mbs": 0,
    "source-details": {
      "type": "volumeGroupBackupId",
      "volume-group-backup-id": [
        "ocid1.volumeGroupBackup.....uniqueID"
      ]
    }
  }
}
```

```

    },
    "time-created": "2021-05-26T22:44:07+00:00",
    "volume-ids": []
  },
  "etag": "c7053513-6819-49ad-8785-dd3e2a45272a"
}

```

7.8 Managing Backup Policies

7.8.1 Overview

Block Volume service provides you with the capability to perform volume backups and volume group backups automatically on a schedule and retain them based on the selected backup policy.

For more conceptual information, refer to *Block Storage Backups and Clones* in the [Block Volume Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

7.8.2 Creating Backup Policies and Schedules

7.8.2.1 Creating a Backup Policy

Using the Compute Web UI

Creating a backup policy using the Web UI is a two task process. First create the backup policy, then create a backup policy schedule, as described in this procedure.

1. In the navigation menu, under Block Storage, click Backup Policies.
2. Click Create Backup Policy.
3. Enter the required information:
 - **Name:** Enter a descriptive name for the policy.
 - **Create in Compartment:** Select the compartment where the policy will reside.
 - **Tagging:** Optionally, add one or more tags to this resource.
If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

4. Click Create Backup Policy

The policy is empty until you add a schedule to the policy.

5. Click the name of the policy to display the details page.
6. Click Add Schedule.
7. In the dialog box, enter the schedule parameters:
 - **Schedule Type options:** Specify the backup frequency by selecting Daily, Weekly, Monthly, or Yearly.
 - Configure the additional schedule options.

Depending on the schedule type, the additional schedule options include one or more of the following choices:

- Hour of the day
- Day of the week
- Day of the month
- Month of the year
- **Retention Time:** Depending on the schedule type, specify the days, weeks, months, or years.
- **Time Zone:** Select the time zone to base the schedule settings on, either UTC or Regional Data Center Time.

8. Click Add Schedule.

The policy can now be applied to one or more volumes and volume groups. See [Section 7.8.2.2, “Assigning a Backup Policy to a Volume or Volume Group”](#)

Using the OCI CLI

Creating a backup policy using the CLI creates the policy and the schedule using one command, as described in this procedure.

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)
- Schedule items – A valid JSON string on the command line, or a valid JSON file with all of the required schedule items plus any optional schedule items of your choice.

For details about the schedule items and the proper JSON format, run this command:

```
oci bv volume-backup-policy create --generate-param-json-input schedules
```

For additional details about required and optional schedule items, refer to the [VolumeBackupSchedule Reference](#) in the Oracle Cloud Infrastructure Documentation.

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-backup-policy create \
--compartment-id <compartment_OCID> \
--display-name <discriptive_name> \
--schedule <json_string> or file://<path_to_JSON_file>
```

Example using a JSON string:

```
oci bv volume-backup-policy create \
--compartment-id ocid1.compartment.....uniqueID \
--display-name daily-full-backup \
--schedules '{"backupType": "FULL", "period": "ONE_DAY", "hourOfDay": 23, "retentionSeconds": 172800, "timeZone": "UTC"}' \
--data '{"compartment-id": "ocid1.compartment.....uniqueID", "defined-tags": {}}'
```

```

"destination-region": null,
"display-name": "daily-full-backup",
"freeform-tags": {},
"id": "ocidl.volumebackuppolicy.....uniqueID",
"schedules": [
  {
    "backup-type": "FULL",
    "day-of-month": null,
    "day-of-week": null,
    "hour-of-day": 23,
    "month": null,
    "offset-seconds": null,
    "offset-type": null,
    "period": "ONE_DAY",
    "retention-seconds": 172800,
    "time-zone": "UTC"
  }
],
"time-created": "2021-10-04T23:22:08.688884+00:00"
},
"etag": "0e1d293e-42fe-4fa9-88f7-a0ab30cc8256",
"opc-work-request-id": "ocidl.workrequest.....uniqueID"
}

```

3. Assign the policy to a volume.

See [Section 7.8.2.2, “Assigning a Backup Policy to a Volume or Volume Group”](#).

7.8.2.2 Assigning a Backup Policy to a Volume or Volume Group

You can assign a backup policy to a volume at these stages:

- During volume creation. See [Section 7.2.2, “Creating a Block Volume”](#).
- After the volume is created, you can add a backup policy to a volume or volume group by following the steps in this section.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Block Volumes or Volume Groups.
2. Select the appropriate compartment.
3. Click the volume or volume group name.
4. Click Edit.
5. Select a backup policy from the drop-down menu.
6. Click Save Changes.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Asset ID – the target volume OCID (`oci bv volume list --compartment <compartment_OCID>`)
 - Policy ID – the policy OCID (`oci bv volume-backup-policy list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-backup-policy-assignment create
--asset-id <volume_OCID>
--policy-id <backup-policyABC>
```

Example:

```
oci bv volume-backup-policy-assignment create \
--asset-id ocidl.volume.....uniqueID123 \
--policy-id ocidl.volumeBackupPolicy.....uniqueID456
{
  "data": {
    "asset-id": "ocidl.volume.....uniqueID123",
    "id": "ocidl.backupPolicyAssignment.....new-uniqueID",
    "policy-id": "ocidl.volumeBackupPolicy.....uniqueID456",
    "time-created": "2021-05-26T23:31:53+00:00"
  },
  "etag": "57ddd89d-14bc-4ecf-b164-8a6cc9dc9014"
}
```

7.8.3 Accessing the Backups

After a manual or automatic backup takes place, follow these steps to locate the backups.

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click one of the following items:

- Block Volume Backups
- Volume Group Backups

The manual and automatic backups are displayed.

2. To see the details of a backup, click the backup name.

Using the OCI CLI

1. Gather the information that you need to run the command:

-
-

2. Run this command.

Syntax (entered on a single line):

```
oci bv backup list
--compartment-id <compartment_OCID>
```

Example:

```
oci bv backup list --compartment-id ocidl.compartment.....uniqueID
{
  "data": [
    {
```

```

    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "volumebackup20212909223152",
    "expiration-time": null,
    "freeform-tags": {},
    "id": "ocidl.volumebackup.....uniqueID",
    "kms-key-id": null,
    "lifecycle-state": "AVAILABLE",
    "size-in-gbs": 0,
    "size-in-mbs": null,
    "source-type": "MANUAL",
    "source-volume-backup-id": null,
    "system-tags": null,
    "time-created": "2021-09-29T22:31:55.000001+00:00",
    "time-request-received": "2021-09-29T22:31:52.000001+00:00",
    "type": "FULL",
    "unique-size-in-gbs": 0,
    "unique-size-in-mbs": null,
    "volume-id": "ocidl.volume.....uniqueID"
  }
]
}

```

7.8.4 Viewing Backup Policies

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Backup Policies.
2. Select the appropriate compartment.
3. To see details about a policy, click the policy name.

Using the OCI CLI

- **Listing All the Backup Policies in a Compartment**

1. Gather the information that you need to run the command:
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-backup-policy list --compartment-id <compartment_OCID>
```

Example:

```

oci bv volume-backup-policy list \
--compartment-id ocidl.compartment.oc1.....unique_ID
{
  "data": [
    {
      "compartment-id": "ocidl.compartment.oc1.....unique_ID",
      "defined-tags": {},
      "destination-region": null,
      "display-name": "backup-policyABC",
      "freeform-tags": {},
      "id": "ocidl.volumeBackupPolicy.....uniqueID-1",
      "schedules": null,
      "time-created": "2021-05-17T22:40:17+00:00"
    }
  ]
}

```

```

    }
  {
    "compartment-id": "ocidl.compartment.oc1.....unique_ID",
    "defined-tags": {},
    "destination-region": null,
    "display-name": "backup-policy1",
    "freeform-tags": {},
    "id": "ocidl.volumeBackupPolicy.....uniqueID-2",
    "schedules": null,
    "time-created": "2021-05-17T22:10:45+00:00"
  }
}

```

• **Listing a Specific Backup Policy**

1. Gather the information that you need to run the command:

- Backup Policy OCID (`oci bv volume-backup-policy list --compartment-id <compartment_OCID>`)

2. Run this command.

Syntax (entered on a single line):

```
oci bv volume-backup-policy get --policy-id <backup_policy_OCID>
```

Example:

```

oci bv volume-backup-policy get \
--policy-id ocidl.volumeBackupPolicy.....uniqueID-1
{
  "data": [
    {
      "compartment-id": "ocidl.compartment.oc1.....unique_ID",
      "defined-tags": {},
      "destination-region": null,
      "display-name": "backup-policyABC",
      "freeform-tags": {},
      "id": "ocidl.volumeBackupPolicy.....uniqueID-1",
      "schedules": null,
      "time-created": "2021-05-17T22:40:17+00:00"
    }
  ]
}

```

7.8.5 Editing a Backup Policy Schedule

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Backup Policies.
2. Select the appropriate compartment.
3. Click the backup policy name that has the schedule you want to edit.
4. In Schedules, find the schedule you want to edit, click the Actions icon (three dots), and then click Edit.
5. After making your changes to the schedule, click Update Schedule.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Backup policy OCID (oci bv volume-backup-policy list --compartment-id *<compartment_OCID>*)
- Schedule items – A valid JSON string on the command line, or a valid JSON file with all of the required schedule items plus any optional schedule items of your choice.

For details about the JSON items and format, run this command:

```
oci bv volume-backup-policy create --generate-param-json-input schedules
```

For additional details about required and optional schedule items, refer to the [VolumeBackupSchedule Reference](#) in the Oracle Cloud Infrastructure Documentation.

2. Run this command.

Caution

Running this command replaces all schedule items with the items you specify in the --schedule JSON.

Syntax (entered on a single line):

```
oci bv volume-backup-policy update
--policy-id <backup_policy_OCID>
--schedule <json_string> or file://<path_to_JSON_file>
```

Example:

```
oci bv volume-backup-policy update \
--policy-id ocid1.volumebackuppolicy.....uniqueID \
--schedules '{"backupType": "FULL", "period": "ONE_DAY", "hourOfDay": 3, "retentionSeconds": 172800, "timeZone": "UTC"}'

WARNING: Updates to schedules and defined-tags and freeform-tags will replace any existing values.
Are you sure you want to continue? [y/N]: y
{
  "data": {
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "destination-region": null,
    "display-name": "policy-daily-full",
    "freeform-tags": {},
    "id": "ocid1.volumebackuppolicy.....uniqueID",
    "schedules": [
      {
        "backup-type": "FULL",
        "day-of-month": null,
        "day-of-week": null,
        "hour-of-day": 3,
        "month": null,
        "offset-seconds": null,
        "offset-type": null,
        "period": "ONE_DAY",
        "retention-seconds": 172800,
        "time-zone": "UTC"
      }
    ],
    "time-created": "2021-10-05T21:01:45.153443+00:00"
  },
  "etag": "fb61d20c-92b2-465e-a545-68f6a5463b8e"
}
```

7.8.6 Deleting a Backup Policy Schedule

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Backup Policies.
2. Select the appropriate compartment.
3. Click the backup policy that has the schedule you want to delete.
4. In Schedules, for the schedule you want to delete, click the Actions icon (three dots), and then click Delete.
5. Confirm the deletion.

7.8.7 Deleting a Backup Policy

Using the Compute Web UI

1. In the navigation menu, under Block Storage, click Backup Policies.
2. Select the appropriate compartment.
3. Click the name of the policy you want to delete.
4. Click Delete.
5. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Backup Policy OCID (`oci bv volume-backup-policy list --compartment-id <compartment_OCID>`)
2. Run the command.

Syntax (entered on a single line):

```
oci bv volume-backup-policy delete --policy-id <policy_OCID>
```

Example:

```
oci bv volume-backup-policy delete
--policy-id ocid1.volumeBackupPolicy.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
{
  "etag": "63a6b74f-e86e-423c-9948-123456789012"
}
```

Chapter 8 File System Storage

Table of Contents

8.1 Creating a File System, Mount Target, and Export	295
8.1.1 Overview	295
8.1.2 Task Flow to Create and Export A File System	296
8.1.3 Creating a Mount Target	296
8.1.4 Creating a File System	299
8.1.5 Creating an Export for a File System	300
8.2 Controlling Access to File Storage	303
8.2.1 Overview	303
8.2.2 Configuring VCN Security Rules for File Storage	303
8.2.3 Adding File Storage to a Network Security Group	304
8.2.4 Setting NFS Export Options	305
8.3 Mounting File Systems on UNIX-Type Instances	307
8.3.1 Mounting Overview	307
8.3.2 Obtaining the Mount Target IP Address	308
8.3.3 Mounting a File System on Linux, RedHat, or CentOS	309
8.3.4 Mounting a File System on Ubuntu or Debian	311
8.3.5 Configuring a File System to Automatically Mount (Linux Instances)	312
8.4 Mounting File Systems On Windows Instances	313
8.4.1 Mounting a File System On a Windows Instance Using NFS	313
8.4.2 Mounting a File System on a Window Instance Using SMB	316
8.5 Managing Mount Targets and Exports	318
8.5.1 Overview	318
8.5.2 Listing Mount Targets and Viewing Details	319
8.5.3 Changing the Mount Target Name	320
8.5.4 Listing Exports	321
8.5.5 Listing Export Sets	322
8.5.6 Deleting an Export	323
8.5.7 Moving a Mount Target to a Different Compartment	323
8.5.8 Deleting a Mount Target	324
8.6 Managing File Systems	324
8.6.1 Listing and Viewing the Details of a File System	325
8.6.2 Changing the File System Name	326
8.6.3 Moving a File System to a Different Compartment	327
8.6.4 Deleting a File System	328
8.7 Managing Snapshots	328
8.7.1 Snapshots Overview	328
8.7.2 Listing and Getting Snapshot Details	328
8.7.3 Creating a Snapshot	330
8.7.4 Accessing a Snapshot on the Mounted File System	331
8.7.5 Restoring a Snapshot (UNIX-Type Instances)	332
8.7.6 Deleting a Snapshot	332

8.1 Creating a File System, Mount Target, and Export

8.1.1 Overview

The File Storage service provides scalable and secure shared network file systems.

The File Storage service encrypts all file system and snapshot data at rest.

You can mount a File Storage service file system on any compute instance in your Virtual Cloud Network (VCN).

For more conceptual information, refer to the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

8.1.2 Task Flow to Create and Export A File System

To create a file system and make it mountable from an NFS client, perform the tasks listed in the table (Specific steps are described in subsequent procedures).

Note – Alternatively, you can mount the file system from a Windows SMB client.

No.	Description	Links to Procedures
1.	<p>Ensure there is a mount target available that is assigned to the VCN and subnet of your choice. Create a mount target if one doesn't exist.</p> <p>Only one mount target can be created per VCN. A mount target can be used for many file systems.</p> <p>Note – the file system and mount target must be in the same compartment when you create an export.</p>	Section 8.1.3, “Creating a Mount Target”
2.	Create the file system.	Section 8.1.4, “Creating a File System”
3.	Create a file system export in the mount target.	Section 8.1.5, “Creating an Export for a File System”
4.	Enable Security Rules for File Storage.	Section 8.2, “Controlling Access to File Storage”
5.	Change NFS export options to control access to the file system.	Section 8.2.4, “Setting NFS Export Options”

Once the file system is exported, on the NFS client, perform these tasks to mount the file system:

1. (If needed) Install NFS client software.
2. Create a mount point.
3. On the client, mount the file system to the mount point.
4. On the client, add whatever files, directories, and data that you want in the file system.

For more information about mounting file systems, see [Section 8.3, “Mounting File Systems on UNIX-Type Instances”](#).

8.1.3 Creating a Mount Target

A mount target is an NFS endpoint assigned to a subnet of your choice. The mount target provides the IP address or DNS name that is used in the mount command when connecting NFS clients to a file system.

For an instance to mount a file system, the instance's VCN must have a Mount Target.

You can only create one mount target per VCN. If a mount target is already created in the VCN you want to use, do not create a new mount target. Instead, use the mount target that is already available.

You can reuse the same mount target to make many file systems available on the network. To reuse the same mount target for multiple file systems, create an export in the mount target for each file system.

Caution

Do not use /30 or smaller subnets for mount target creation because they might not have sufficient available IP addresses.

Important

When more than one file system is exported to the same mount target, you must export to the mount target with the smallest network (largest CIDR number) first. For detailed information and instructions, refer to the My Oracle Support [Doc ID 2823994.1](#)

Before you can create a mount target, ensure that these items are configured:

- At least one Virtual Cloud Network (VCN) in the compartment where the file system will be created. See [Section 4.3, “Configuring VCN Gateways”](#)
- An internet gateway with a route rule in the VCN. See [Section 4.2, “Configuring VCN Rules and Options”](#).
- (Optional) Security rules for the file system mount target. Security rules can be created in the security list for the mount target subnet, or in a Network Security Group (NSG) that you add the mount target to. See [Section 8.2, “Controlling Access to File Storage”](#).

Note – You don't need security rules to create a mount target, but you need the rules to eventually mount files systems that are associated with this mount target.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Target.

If a mount target is listed, you can use the existing mount target if it is on the subnet you were planning to assign the mount target. Click on the mount target name to see the details. If the mount target meets your needs, skip this procedure and go to [Section 8.1.4, “Creating a File System”](#).

2. Click Create Mount Target.
3. Enter the mount target information:
 - **Name:** It doesn't have to be unique. An Oracle Cloud Identifier (OCID) uniquely identifies the mount target. Avoid entering confidential information.

Note

The mount target name is different than the DNS hostname.

- **Create in Compartment:** Specify the compartment.
- **VCN:** Select the VCN where you want to create the new mount target.
- **Subnet:** Select a subnet to attach the mount target to.
- **Enable Network Security Groups:** Select this option to add this mount target to an NSG you've created.

Important

Rules for the NSG you select must be configured to allow traffic to the mount target's VNIC using specific protocols and ports. For more information, see [Section 8.2, “Controlling Access to File Storage”](#) Configuring VCN Security Rules for File Storage.

- **IP Address:** Optionally, you can specify an unused IP address in the subnet you selected for the mount target. If left blank, an IP address is automatically assigned.
- **Hostname:** Optionally, you can specify a hostname you want to assign to the mount target.

Note

The File Storage service constructs a fully qualified domain name (FQDN) by combining the hostname with the FQDN of the subnet the mount target is located in.

For example, `myhostname.subnet123.dnslabel.examplevcn.com`.

- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

4. Click Create Mount Target.

Next, create a file system. See [Section 8.1.4, “Creating a File System”](#).

Using the OCI CLI

1. Gather the information that you need to run the command:

- Availability Domain Name (`oci iam availability-domain list`)
- Compartment OCID (`oci iam compartment list`)
- Subnet OCID (`oci network subnet list`)
- (Optional) Display Name you wanted assigned to this mount target.

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci fs mount-target create
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
--subnet-id <subnet_OCID>
--display-name <name_to_assign_to_mount-target>
```

Example:

```
oci fs mount-target create \
--availability-domain MyAD \
--compartment-id ocid1.compartment.....uniqueID \
--subnet-id ocid1.subnet.....uniqueID \
--display-name MyMountTarget2
{
  "data": {
    "availability-domain": "pca",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyMountTarget2",
    "export-set-id": "ocid1.exportset.....uniqueID",
    "freeform-tags": {},
    "id": "ocid1.mounttarget.....uniqueID",
    "lifecycle-details": null,
    "lifecycle-state": "CREATING",
    "nsg-ids": null,
    "private-ip-ids": null,
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": null
  },
  "etag": "2d278b37-a74a-4fec-b74a-fd9e9a1c72de"
```

3. Next, create a file system. See [Section 8.1.4, “Creating a File System”](#).

8.1.4 Creating a File System

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. In the left panel, select File Systems.
3. Click Create File System.
4. Enter this information:

- **File System Information:**

- **Name:** It doesn't have to be unique. An Oracle Cloud Identifier (OCID) uniquely identifies the file system. Avoid entering confidential information.
- **Create in Compartment:** Select the compartment where the file system will be created.
- **Tagging:** Optionally, add one or more tags to this resource.

If you are not sure whether to apply tags, skip this option (you can apply tags later).

For more information about tagging resources, see [Tagging Overview](#).

5. Click Create File System.

The file system is created.

Next, create an export for the file system. See [Section 8.1.5, “Creating an Export for a File System”](#).

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
 - File System Name: The display name you want assigned to this file system
2. Run this command.

Syntax (entered on a single line):

```
oci fs file-system create
--availability-domain <availability_domain_name>
--compartment-id <compartment_id>
--display-name <fs_display_name>
```

Example:

```
oci fs file-system create \
--availability-domain MyAD \
--compartment-id ocidl.compartment.....uniqueID \
--display-name MyFileSystem

{
  "data": {
    "availability-domain": "pca",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyFileSystem",
    "freeform-tags": {},
    "id": "ocidl.filesystem.....uniqueID",
    "kms-key-id": null,
    "lifecycle-state": "CREATING",
    "metered-bytes": 0,
    "time-created": null
  },
  "etag": "58dec47e-4732-4730-9e18-6b5db1ac30d6"
}
```

3. Next, create an export for the file system. See [Section 8.1.5, “Creating an Export for a File System”](#).

8.1.5 Creating an Export for a File System

Exports control how NFS clients access file systems when they connect to a mount target.

A file system must have at least one export in one mount target for instances to mount the file system.

Important

When more than one file system is exported to the same mount target, you must export to the mount target with the smallest network (largest CIDR number) first. For detailed information and instructions, refer to the My Oracle Support [Doc ID 2823994.1](#)

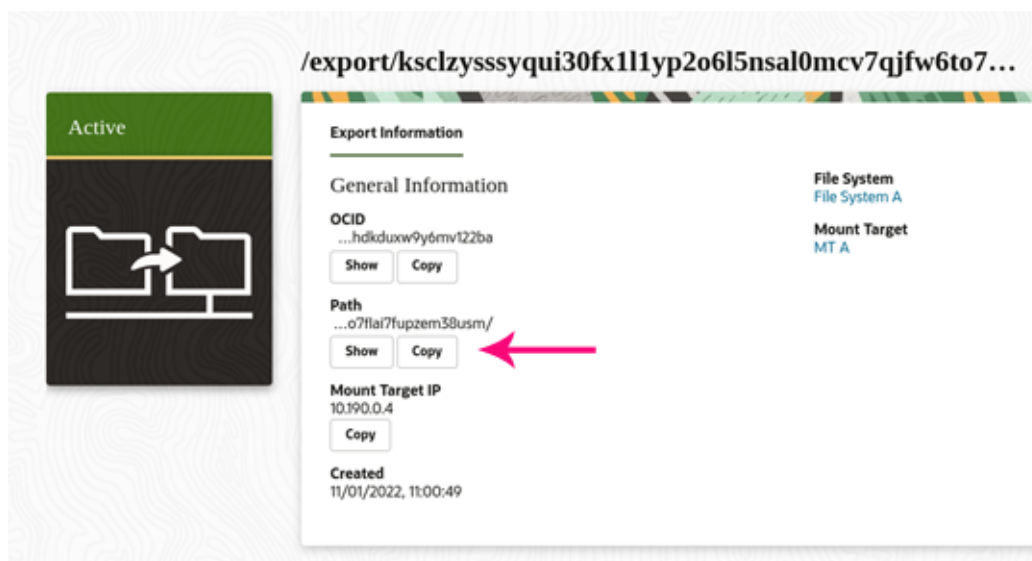
Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. In the left panel, select File Systems.

3. Click the name of the file system that you plan to create an export for.
4. In the lower panel, click Create Export.
5. Enter the required information:
 - Mount Target: Select a mount target from the drop-down list.
6. Click Create Export.

The file system export is created and the export details page is displayed.

7. In the export details page, make note of the export path. The export path is used to mount the file system on an instance. Example:



8. In the lower panel, review the NFS Export Options.

The NFS export options for that file system are set to the default values, which allow full access for all NFS client source connections. These defaults must be changed if you want to restrict access.

9. Consider your next action:
 - Mount the file system from an NFS client. See [Section 8.3, “Mounting File Systems on UNIX-Type Instances”](#).
 - Configure NFS options to secure the exported file system. See [Section 8.2.4, “Setting NFS Export Options”](#).

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Export set OCID (`oci fs export-set list --availability-domain <name> --compartment-id <compartment_OCID>`)
 - File system OCID (`oci fs file-system list --availability-domain <name> --compartment-id <compartment_OCID>`)

- (Required) Export path of your choice. Note that the system will assign an auto-generated path to the export. The auto-generated path is eventually used to mount the file system. The path you enter here is recorded, but not used.

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci fs export create
--export-set-id <export_set_OCID>
--file-system-id <file_system_OCID>
--path "</pathname>"
```

Example:

```
oci fs export create \
--export-set-id ocid1.exportset.....uniqueID \
--file-system-id ocid1.filesystem.....uniqueID \
--path "/export/departmentA"
{
  "data": {
    "export-options": [
      {
        "access": "READ_WRITE",
        "anonymous-gid": 65534,
        "anonymous-uid": 65534,
        "identity-squash": "NONE",
        "require-privileged-source-port": false,
        "source": "0.0.0.0/0"
      }
    ],
    "export-set-id": "ocid1.exportset.....uniqueID",
    "file-system-id": "ocid1.filesystem.....uniqueID",
    "id": "ocid1.export.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "path": "/export/181t6v4drhddiz2mn7vwmqt7mjiz3kfbw4reqaew33y50pdrj35p4ef5p04x",
    "time-created": "2021-09-02T22:41:36.284348+00:00"
  },
  "etag": "a0842b0b-b27b-4c98-a1ff-da85ae4bf150"
}
```

3. In the output, make note of the value for "path". The path value is used to mount the file system.

Example:

```
...
  "path": "/export/181t6v4drhddiz2mn7vwmqt7mjiz3kfbw4reqaew33y50pdrj35p4ef5p04x",
  "time-created": "2021-09-01T19:23:15.774764+00:00"
...
```

4. In the output, review the export options.

In this example, the NFS export options for the file system are set to the default values, which allow full access for all NFS client source connections. These defaults must be changed if you want to restrict access

5. Next, control access to the file system.

See [Section 8.2, “Controlling Access to File Storage”](#).

8.2 Controlling Access to File Storage

8.2.1 Overview

Before you can mount a file system, you must configure security rules to allow traffic to the mount target's VNIC using specific protocols and ports. Security rules enable traffic for the following protocols:

- Open Network Computing Remote Procedure Call (ONC RPC) rpcbind utility protocol
- Network File System (NFS) protocol
- Network File System (MOUNT) protocol

For more conceptual information, refer to the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

8.2.2 Configuring VCN Security Rules for File Storage

You can add the required rules to a pre-existing security list associated with a subnet, such as the default security list that is created along with the VCN.

For specific information about the what security rules are required for the File Storage service, refer to [File Storage Network Ports](#) in the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

For more information about managing VCNs and subnets, see [Section 4.1, “Managing VCNs and Subnets”](#).

Using the Compute Web UI

1. In the navigation menu, under Networking, click Virtual Cloud Networks.
2. Select the compartment where the VCN is located.
3. Click the name of the VCN.
4. Under Resources, click Security Lists.
5. Click the name of the security.
6. Under Resources, click Ingress Rules.
7. Click Create Ingress Security Rule, and enter the required information:
 - **Stateless check box:** Specify that it's a stateful rule by leaving the check box unchecked.
 - **Ingress CIDR:** Enter the CIDR block for the subnet. For example, 10.0.0.0/24.
 - **IP Protocol:** Choose the protocol. For example, TCP.
 - **Description:** Enter a meaningful description for the rule.
8. Click Create Security List Rule.
9. Under Resources, click Egress Rules.

10. Click Create Egress Security Rule and enter the required information:

- **Stateless check box:** Specify that it's a stateful rule by leaving the check box unchecked.
- **Egress Type:** To allow traffic from the subnet, select CIDR.
- **Egress CIDR:** Enter the CIDR block for the subnet. For example, 10.0.0.0/24.
- **IP Protocol:** Choose the protocol. For example, TCP.
- **Description:** Enter a meaningful description for the rule.

11. Click Create Security List Rule.

8.2.3 Adding File Storage to a Network Security Group

8.2.3.1 Process Overview

This is the general process for setting up NSGs that work with File Storage:

No.	Description	Links to Procedures
1.	Create an NSG with the required security rules. (Alternatively, you can add them to a previously existing NSG.)	Section 4.2.4, "Controlling Traffic with Network Security Groups"
2.	Add the mount target (or more specifically, the mount target's VNIC) to the NSG. You can do this when you create the mount target, or you can update the mount target and add it to one or more NSGs that contain the required security rules.	Section 8.2.3.2, "Adding a Mount Target to a Network Security Group"
3.	If you're setting up a mount target and instance in different subnets, you'll also have to add the instance to an NSG that contains the required security rules. Adding an existing instance to an NSG means adding its primary VNIC to the NSG.	Section 4.4.3.5, "Add or Remove a VNIC from a Network Security Group"

For specific information about the what security rules are required for the File Storage service, refer to *File Storage Port Configurations* in the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

8.2.3.2 Adding a Mount Target to a Network Security Group

You can add the mount target to one or more Network Security Groups (NSGs). File storage requires specific rules to be configured for NSGs that are associated with mount targets.

Using the Compute Web UI

1. Ensure that an NSG with ingress and egress rules has been configured.

See [Section 4.2, "Configuring VCN Rules and Options"](#).

2. Ensure that a mount target is created.

See [Section 4.1, "Managing VCNs and Subnets"](#).

3. In the navigation menu, under File Storage, click Mount Targets.
4. Click the mount target name to see the details page.
5. Click Edit.
6. Enable Network Security Groups.
7. Select the NSG from the drop-down list.
8. Click Save Changes.

Using the OCI CLI

1. Ensure that an NSG with ingress and egress rules has been configured.

See [Section 4.2, “Configuring VCN Rules and Options”](#).

2. Ensure that a mount target is created.

See [Section 4.1, “Managing VCNs and Subnets”](#).

3. Gather the information that you need to run the command:

- Mount target OCID (`oci fs mount-target list`)
- NSG OCIDs (`oci network nsg list`)

4. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target update
--mount-target-id <mount_target_OCID>
--nsg-ids '["<nsg1_OCID>","<nsg2_OCID>"]'
```

Example:

```
oci fs export update \
--mount-target-id ocid1.mounttarget.....uniqueID \
--nsg-ids '["ocid1.networksecuritygroup.....uniqueID-01","ocid1.networksecuritygroup.....uniqueID-02"]'
```

8.2.4 Setting NFS Export Options

When you create a file system and export, the NFS export options for that file system are set to the defaults listed in this table. The default values allow full access for all NFS client source connections. These defaults must be changed if you want to restrict access:

Caution

If a file system is mounted by any clients, creating, deleting, or editing the Source value can disrupt file system I/O operations.

Export Option in the Web UI	Export Option in the CLI	Default Value
Source:	<code>source</code>	0.0.0.0/0
Ports:	<code>require-privileged-source-port</code>	Any

Export Option in the Web UI	Export Option in the CLI	Default Value	De
			• V • C
Access:	<code>access</code>	Read/Write	Sp • I • I
Squash:	<code>identity-squash</code>	None	De po • I • I
Squash UID/GID:	<code>anonymous-uid</code> and <code>anonymous-gid</code>	65534	Th to

Note – If you change the RW/RO permissions of an export option for an SMB share, the changes are only enforced for newly network-mapped drives of that share. Any previously mapped drives of the same share will retain the original permissions. To have the changed permissions enforced on previously mapped drives on SMB clients, disconnect the shares and map them again.

For more information about configuring the options to suit various access scenarios, refer to the section titled *NFS Access Control and Export Options* in the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the file system name.
4. Under Resources, select Exports.
5. Click on the export's export path.

The NFS Export Options are displayed.

6. Click Edit Options.
7. In the NFS Export Options dialog, configure the NFS options.
8. Click Update Options.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Export ID (`oci fs export list --all --compartment-id <compartment_OCID>`)
 - Export options, listed in json format, in a json file or as a string on the command line.

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci fs export update
--export-id <export_id>
--export-options <file://json_file or json_string>
```

Note – The `require-privileged-source-port` option can only be set to `false`.

This example sets the export options for file system A to allow Read_Write access only to Client A, who is assigned to CIDR block 10.0.0.0/24. Client B and Client C are not included in this CIDR block, and cannot access the file system:

```
oci fs export update \
--export-id File_system_A_export_ID \
--export-options \
'[{ "source": "10.0.0.0/24", "require-privileged-source-port": "false", "access": "READ_WRITE", "identity-squa
WARNING: Updates to export-options will replace any existing values. Are you sure you want to continue?
{
  "data": {
    "export-options": [
      {
        "access": "READ_WRITE",
        "anonymous-gid": 65534,
        "anonymous-uid": 65534,
        "identity-squash": "NONE",
        "require-privileged-source-port": false,
        "source": "10.0.0.0/24"
      }
    ],
    "export-set-id": "ocidl.exportset.....uniqueID",
    "file-system-id": "ocidl.filesystem.....uniqueID",
    "id": "ocidl.export.ocl.pca.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "path": "/export/85aiiadclw81s8id63knxdq22nt95pe63sgs9c45yp3qovhut14cq9r6eqhn",
    "time-created": "2021-09-27T20:20:34.231009+00:00"
  },
  "etag": "bc660e11-644a-4043-9ad7-622d9581da9b"
}
```

8.3 Mounting File Systems on UNIX-Type Instances

8.3.1 Mounting Overview

Instance users of UNIX-type operating systems, such as Linux and Oracle Solaris, can use OS commands to mount and access file systems.

Mount targets serve as network access points for file systems. After your mount target is assigned an IP address, you can use it together with the export path to mount the file system.

On the instance from which you want to mount the file system, you need to install an NFS client package and create a mount point. When you mount the file system, the mount point effectively represents the root directory of the File Storage file system, allowing you to write files to the file system from the instance.

Prerequisites

- The file system must be created and have at least one export in a mount target. See [Section 8.1, “Creating a File System, Mount Target, and Export”](#).
- The mount target must have correctly configured security rules or be assigned to an NSG. See [Section 8.2.2, “Configuring VCN Security Rules for File Storage”](#).

Note

Only for NFSv4 Mounts in Oracle Linux instances – If you find that the file system owner is assigned as `nobody` instead of the actual user who mounts the file system, and if you have not set identity squash, you might need to edit the `/etc/idmapd.conf` file. In the file, set the DOMAIN entry to either `localdomain` or to the Active Directory domain name, if applicable. After the change, run `service rpcidmapd restart` to restart the `rpcidmapd` service.

Defining settings in the `/etc/idmapd.conf` file is specific to Oracle Linux, and there are other ways to configure the domain depending on the OS in use. Consult your operating system documentation.

For more conceptual information, refer to the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

8.3.2 Obtaining the Mount Target IP Address

To mount a file system, you need to know the private IP address of the mount target that has the export for the file system.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Target.
2. Click the Mount Target name to see the details page.

The IP address is displayed.

Using the OCI CLI

1. Gather the information that you need to run the commands:
 - Mount Target ID (`oci fs mount-target list --availability-domain <availability_domain_name> --compartment-id <compartment_OCID>`)
2. Run this command to get the mount target IP ID.

Syntax (entered on a single line):

```
oci fs mount-target get
--mount-target-id <mount_target_OCID>
```

Example:

```
oci fs mount-target get \
```

```

--mount-target-id ocid1.mounttarget.....uniqueID
{
  "data": {
    "availability-domain": "ad1",
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "defined-tags": {
      "Finance": {
        "CostCenter": "admin"
      }
    },
    "display-name": "mount-target01",
    "export-set-id": "ocid1.exportset.....uniqueID",
    "freeform-tags": {},
    "id": "ocid1.mounttarget.....uniqueID",
    "lifecycle-details": null,
    "lifecycle-state": "ACTIVE",
    "nsg-ids": [],
    "private-ip-ids": [
      "ocid1.privateip.....uniqueID"
    ],
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-09-01T18:45:25.251048+00:00"
  },
  "etag": "c2f84c0b-d0b5-422c-9761-9e43d7fc4214"
}

```

3. Run this command to get the mount target IP address.

Syntax (entered on a single line):

```

oci network private-ip get
--private-ip-id <mount_target_IP_OCID>

```

Example:

```

oci network private-ip get \
--private-ip-id ocid1.....uniqueID{
  "data": {
    "availability-domain": "ad1",
    "compartment-id": "ocid1.tenancy.....uniqueID",
    "defined-tags": {},
    "display-name": "privateip20210901184525",
    "freeform-tags": {},
    "hostname-label": null,
    "id": "ocid1.privateip.....uniqueID",
    "ip-address": "10.200.0.3",
    "is-primary": false,
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-09-01T18:45:25.406808+00:00",
    "vlan-id": null,
    "vnic-id": "ocid1.vnic.....uniqueID"
  },
  "etag": "c98377e4-ae89-46cf-9c61-52aea68a3476"
}

```

8.3.3 Mounting a File System on Linux, RedHat, or CentOS

1. Log into the instance where you want to mount the file system.

See [Section 6.4, “Connecting to a Compute Instance”](#).

Example:

```

ssh user@192.0.2.0

```

2. Install the NFS client using this command:

```
sudo yum install nfs-utils
```

3. Create a directory that will be used as the mount point.

Replace `<yourmountpoint>` with a directory name of your choice. Example: `/mnt/mountpoint-A`

```
sudo mkdir -p <yourmountpoint>
```

4. Mount the file system.

Caution

Omitting the `-o nosuid` option can allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.

Example:

```
sudo mount -t nfs -o nfsvers=<version>,nosuid <10.x.x.x>:<fs-export-path> <yourmountpoint>
```

- Replace `<version>` with one of the following, based on the NFS protocol version you want to use:
 - 3,noacl
 - 4.0
 - 4.1
- Replace `<10.x.x.x>` with the mount target's private IP address. See [Section 8.3.2, "Obtaining the Mount Target IP Address"](#).
- Replace `<fs-export-path>` with the export path that was generated when the export was created. See [Section 8.1.5, "Creating an Export for a File System"](#).
- Replace `<yourmountpoint>` with the full path to the local mount point.

5. View the mounted file system.

```
df -h
```

6. Write a file to the file system.

Replace `<yourmountpoint>` with the path to the local mount point and `<filename>` with your file name.

```
sudo touch /mnt/<yourmountpoint>/<filename>
```

7. Verify that you can access the file system and view the file.

Replace `yourmountpoint` with the path to the local mount point.

```
cd <yourmountpoint>  
ls
```

8. Add the file system mount information to the appropriate mount file for your OS.

So far, the file system is manually mounted to the client. If the client is rebooted, the file system will not remount unless you add it to the mount file (for example the `/etc/fstab` or `/etc/vfstab` file).

8.3.4 Mounting a File System on Ubuntu or Debian

Operating Systems and versions of operating systems differ in the way software is added. Consult the documentation for our specific operating system for details.

1. On the NFS client, open a command window, and install the NFS client using this command:

```
sudo apt-get install nfs-common
```

2. Create a directory that will be used as the mount point.

Replace `<yourmountpoint>` with a directory name of your choice. Example: `/mnt/mountpoint-A`

```
sudo mkdir -p <yourmountpoint>
```

3. Mount the file system.

Caution

Omitting the `-o nosuid` option may allow unprivileged users to escalate their permissions to 'root'. The `nosuid` option disables set-user-identifier or set-group-identifier bits within the mounted system, which are rarely used.

Example:

```
sudo mount -t nfs -o nfsvers=<version>,nosuid <10.x.x.x>:<fs-export-path> <yourmountpoint>
```

- Replace `<version>` with one of the following, based on the NFS protocol version you want to use:
 - 3,noacl
 - 4.0
 - 4.1
 - Replace `<10.x.x.x>` with the mount target's private IP address. See [Section 8.3.2, "Obtaining the Mount Target IP Address"](#).
 - Replace `<fs-export-path>` with the export path that was generated when the export was created. See [Section 8.1.5, "Creating an Export for a File System"](#).
 - Replace `<yourmountpoint>` with the full path to the local mount point.
4. View the file system.

```
df -h
```

5. Write a file to the file system.

Replace `<yourmountpoint>` with the path to the local mount point and `<filename>` with your file name.

```
sudo touch /mnt/<yourmountpoint>/<filename>
```

6. Verify that you can access the file system and view the file.

Replace yourmountpoint with the path to the local mount point.

```
cd <yourmountpoint>
ls
```

7. Add the file system mount information to the appropriate mount file for your OS.

So far, the file system is manually mounted to the client. If the client is rebooted, the file system will not remount unless you add it to the mount file (for example the `/etc/fstab` or `/etc/vfstab` file).

8.3.5 Configuring a File System to Automatically Mount (Linux Instances)

On Linux instances, if you want to automatically mount exported file systems during an instance boot, you need to add the mount information in the `/etc/fstab` file.

1. Log into the instance where the file system will be mounted.

See [Section 6.4, “Connecting to a Compute Instance”](#).

2. Create a mount point, if one has not been created.

Example:

```
mkdir /mnt/fs01
```

3. Open the `/etc/fstab` file in an editor and add a line for the nfs file systems you want automatically mounted.

This is an example of an `/etc/fstab` file entry.

```
192.0.2.0:/export/3ywflz8hhqfde81miewqwjfd049zju69502t9ouo6shzidr4dndaz1hd6qfi /mnt/fs01 nfs nfsvers=4.1,no
```

The `/etc/fstab` file space-separated fields are specified with these entries:

- **Field 1:** Device to mount. For network file systems, specify: `<mount target IP>:<export_path>`

See [Section 8.3.2, “Obtaining the Mount Target IP Address”](#) and [Section 8.1.5, “Creating an Export for a File System”](#).

- **Field 2:** Full path of the mount point on the instance.
- **Field 3:** File system type. In this case, specify `nfs`.
- **Field 4:** NFS mount options separated with commas, such as:

```
nfsvers=<version>,nosuid,nofail
```

- `nfsvers=` where `<version>` is one of the following:
 - `3,noacl`
 - `4.0`
 - `4.1`
- `nosuid` – prevents unprivileged users from escalating their permissions to root.
- `nofail` – Ensures that an unavailable file system does not cause the instance reboot process to fail.

In this case, use the same options as described in [Section 8.3.3, “Mounting a File System on Linux, RedHat, or CentOS”](#). Each option is separated by a comma (no spaces).

- **Field 5:** Obsolete option for dump backups. Specify 0 (zero) for no dump backup.
 - **Field 6:** File system check (fsck) order. Specify 0 (zero) for no check.
4. Use this command to mount the volumes that are in the `/etc/fstab` file:

```
sudo mount -a
```

If you get any error messages, fix the cause before proceeding.

5. Verify that the file systems are mounted:

```
mount | grep nfs
```

6. To verify that the file system will automatically mount, reboot the instance.

```
sudo reboot
```

7. After the reboot, log into the instance and check to see if the nfs file system is mounted.

```
mount | grep nfs
```

8.4 Mounting File Systems On Windows Instances

You can make file systems available to Windows instances by mapping a network drive to the mount target IP address and export path provided by the File Storage service. You can accomplish this task using NFS or SMB protocols.

Using the SMB protocol requires that the Windows instances and Oracle Private Cloud Appliance belong to the same Active Directory domain. For more information about configuring Active Directory in the Service Enclave, refer to [Hardware Administration](#) in the [Oracle Private Cloud Appliance Administrator Guide](#).

For more conceptual information, refer to the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

8.4.1 Mounting a File System On a Windows Instance Using NFS

Prerequisites

- The file system must be created and have at least one export in a mount target. See [Section 8.1, “Creating a File System, Mount Target, and Export”](#).
- The mount target must have correctly configured security rules or be assigned to an NSG. See [Section 8.2.2, “Configuring VCN Security Rules for File Storage”](#).
- You need to know the mount target's IP address. See [Section 8.3.2, “Obtaining the Mount Target IP Address”](#).
- You must be able to log into the Windows OS on the instance with superuser or administrator privileges.

Before You Begin

The following tasks are included in this procedure, and you might want to be aware of them before you begin.

- **Installation of the Windows NFS Client** – This service must be installed on the instance from which you want to mount the file system. Installing the client often requires a restart of the instance.
- **The `AnonymousGid` and `AnonymousUid` identity values must be configured to allow write access.** – Access to NFS file systems requires UNIX-style user and group identities, which are not the same as Windows user and group identities. By default, file systems write permissions are only granted to the root user. To enable user access to NFS shared resources, the Windows client for NFS accesses file systems anonymously, using `AnonymousGid` and `AnonymousUid`.

Caution

Updating the `AnonymousGid` and `AnonymousUid` values require registry changes to your instance.

Choose one the following methods:

- [Using the Windows Command Prompt](#)
- [Using Windows File Explorer](#)

Using the Windows Command Prompt

1. Log into your Windows instance.
See [Section 6.4, “Connecting to a Compute Instance”](#).
2. Open Windows PowerShell and run as Administrator:
 - a. Go to Start and open Windows PowerShell.
 - b. In Windows PowerShell, type the following to run as Administrator:

```
Start-Process powershell -Verb runAs
```

- c. In the User Account Control window, click Yes. A new Administrator: PowerShell window opens. You can close the standard PowerShell window to avoid confusing them.
3. In Administrator: PowerShell, get the NFS client and update the registry by typing the following:

```
Install-WindowsFeature -Name NFS-Client  
Set-ItemProperty HKLM:\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default -Name AnonymousUid -Value 0  
Set-ItemProperty HKLM:\SOFTWARE\Microsoft\ClientForNFS\CurrentVersion\Default -Name AnonymousGid -Value 0  
Stop-Service -Name NfsClnt  
Restart-Service -Name NfsRdr  
Start-Service -Name NfsClnt
```

4. Open a standard Command Prompt Window.

Important

NFS file systems mounted as Administrator are not available to standard users.

5. From the Command Prompt window, mount the file system.

See the cautions and notes below the example.

In the following example, replace:

- `10.x.x.x` with the mount point IP address (see [Section 8.3.2, “Obtaining the Mount Target IP Address”](#))

- `fs-export-path` with the file system export path (see [Section 8.1.5, “Creating an Export for a File System”](#))
- `X` with the drive letter of any available drive you want to map the file system to.

Example:

```
mount 10.x.x.x:/fs-export-path X:
```

6. Verify that you can access and write to the file system.

- a. Access the file system.

In the example, replace `X` with the drive letter you used to mount the file system.

```
X:
```

- b. Write a file.

```
echo > myfile.txt
```

- c. Verify that you can view the file.

```
dir
```

Using Windows File Explorer

1. Log into your Windows instance.

See [Section 6.4, “Connecting to a Compute Instance”](#).

2. Open Windows PowerShell and run as Administrator:

- a. Go to Start and open Windows PowerShell.

- b. In Windows PowerShell, type the following to run as Administrator:

```
Start-Process powershell -Verb runAs
```

- c. In the User Account Control window, click Yes. A new Administrator: PowerShell window opens. You can close the standard PowerShell window to avoid confusing them.

3. In Administrator: PowerShell, get the NFS client by typing the following:

```
Install-WindowsFeature -Name NFS-Client
```

4. If necessary, restart your system.

5. Open the registry editor (`regedit`) to map the `AnonymousGid` and `AnonymousUid` to the root user.

Caution

User identity mapping requires changes to your system registry.

- a. Click Windows Search.
- b. Enter `regedit` in the Search field and press Enter.
- c. Click Yes to allow changes to your device.

- d. Click `HKEY_LOCAL_MACHINE`. Then, browse to: `Software\Microsoft\ClientForNFS\CurrentVersion\Default`.
6. Add a new DWORD32 registry entry for `AnonymousGid`:
 - a. Click Edit, and select New DWORD (32 bit) Value.
 - b. In the Name field, enter `AnonymousGid`. Leave the value at 0.
7. Repeat the previous step to add a second DWORD32 registry entry named `AnonymousUid` with a value of 0.
8. Open Windows Command Line (CMD) and run as Administrator:
 - a. Go to Start and scroll down to Apps.
 - b. In the Windows System section, press Ctrl+Shift and click Command Prompt.
9. In the Windows Command Line (CMD) window, restart the NFS Client by typing the following:

```
nfsadmin client stop
```

```
nfsadmin client start
```
10. Open File Explorer and select This PC. In the Computer tab, select Map network drive.
11. Select the Drive letter that you want to assign to the file system.
12. In the Folder field, enter the following line, replacing:
 - `10.x.x.x` with the mount point IP address (see [Section 8.3.2, "Obtaining the Mount Target IP Address"](#))
 - `fs-export-path` with the file system export path (see [Section 8.1.5, "Creating an Export for a File System"](#))Line:

```
\\10.x.x.x\fs-export-path
```
13. Click the Finish button when complete.

8.4.2 Mounting a File System on a Window Instance Using SMB

General Prerequisites

- The file system must be created and have at least one export in a mount target. See [Section 8.1, "Creating a File System, Mount Target, and Export"](#).
- The mount target must have correctly configured security rules or be assigned to an NSG. See [Section 8.2.2, "Configuring VCN Security Rules for File Storage"](#).
- You need to know the mount target's IP address. See [Section 8.3.2, "Obtaining the Mount Target IP Address"](#).
- You must be able to log into the Windows OS on the instance with superuser or administrator privileges.

Specific Prerequisites for SMB Support

SMB support for the File Storage service requires that both Oracle Private Cloud Appliance and the client Windows instances belong to the same Active Directory (AD) domain.

This procedure assumes that the AD service is already configured in your data center infrastructure.

To add a Windows instance to your AD service, perform the necessary administrative tasks according to the documentation for your version of Windows OS.

To add the appliance to your AD service, an administrator with privileges to the Oracle Private Cloud Appliance Service Enclave must add the AD domain name to the appliance's Active Directory Domain configuration. For information on how to perform this task, refer to [Hardware Administration](#) in the [Oracle Private Cloud Appliance Administrator Guide](#).

Relaxing File System Permissions Before Network Mapping with SMB

By default, write permissions to a file system are limited to the UNIX-style superuser and group identity. To provide write permission to AD domain users, the permissions need to be relaxed.

1. Mount the network drive using NFS protocol.

See [Section 8.4.1, "Mounting a File System On a Windows Instance Using NFS"](#).

2. Relax the file system permissions:

- a. Open File Explorer, select the mapped drive and right click on it, then select Properties.
- b. Select the NFS Attributes tab.
- c. Change File permissions by checking all RWX check boxes to relax the permissions for Owner, Group, and Other.
- d. Click OK.

3. Disconnect the NFS mounted drive.

Now that the file system permissions are relaxed, you can mount the file system using the SMB protocol.

Mounting a File System Using SMB

1. Log into your Windows instance.

See [Section 6.4, "Connecting to a Compute Instance"](#).

2. Open File Explorer and select This PC.
3. In the Computer tab, select Map network drive.
4. In the Folder field, enter the following line and replace these items:

- `10.x.x.x` with the mount target IP address.
- `fs-export-path-ID` with the file system export path (see [Section 8.1.5, "Creating an Export for a File System"](#))

Note – Do not include `\export` in the `fs-export-path-ID` string when mounting using SMB.

```
\\10.x.x.x\fs-export-path-ID
```

Example:

```
\\192.0.2.0\39u2lbtystm8xlaxizezb9a3lfnpzjho98evi3ij450i96vj0a8jpf36au26
```

5. select the 'Drive' letter of any available drive you want to map the file system to.
6. If needed, select the Connect using different credentials checkbox.
7. Click Finish.
8. When prompted, provide the user name and password of the AD domain user used for mapping the network drive.
9. Click OK.
10. In a Command Prompt window (cmd), verify that the drive is properly mapped using this command:

```
C:\>net use
New connections will be remembered.
Status      Local      Remote                                          Network
-----
OK           Z:         \\10.0.0.2\uvjliw6ytyecqijcbdgpy7ec15mgsv044i7609giqx7ukfn6t2pwgfgot0ma
                                                Microsoft Windows Network
The command completed successfully.
C:\>
```

8.5 Managing Mount Targets and Exports

8.5.1 Overview

A mount target is an NFS endpoint assigned to a VCN subnet of your choice and provides network access for file systems. The mount target provides the IP address or DNS name that is used together with a unique export path to mount the file system.

For an instance to mount a file system, the instance's VCN must have a Mount Target. A VCN can only have one mount target.

You can reuse the same mount target to make as many file systems available on the network as you wish. To reuse the same mount target for multiple file systems, create an export in the mount target for each file system.

Important

When more than one file system is exported to the same mount target, you must export to the mount target with the smallest network (largest CIDR number) first. For detailed information and instructions, refer to the My Oracle Support [Doc ID 2823994.1](#)

For instructions to create a mount target, see [Section 8.1, "Creating a File System, Mount Target, and Export"](#).

For more conceptual information, refer to the [File Storage](#) chapter in the [Oracle Private Cloud Appliance Concepts Guide](#).

This section provides instructions for administering mount targets.

8.5.2 Listing Mount Targets and Viewing Details

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Targets.
2. Select the compartment where the mount target resides.

The mount targets are displayed.

3. To see the mount target details, click on the mount target name.

Using the OCI CLI

• Listing Mount Targets

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target list
--availability-domain <availability_domain_name>
--compartment-id <compartment_id>
```

Example:

```
oci fs mount-target list \
--availability-domain MyAD \
--compartment-id ocid1.compartment.....uniqueID

{
  "data": [
    {
      "availability-domain": "MyAD",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "MyMountTarget",
      "export-set-id": "ocid1.exportset.....uniqueID",
      "freeform-tags": {},
      "id": "ocid1.mounttarget.....uniqueID",
      "lifecycle-state": "ACTIVE",
      "nsg-ids": null,
      "private-ip-ids": [
        "ocid1.privateip.....uniqueID"
      ],
      "subnet-id": "ocid1.subnet.....uniqueID",
      "time-created": "2021-07-16T22:56:57+00:00"
    },
    {
      "availability-domain": "MyAD",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "AnotherMountTarget",
      "export-set-id": "ocid1.exportset.....uniqueID",
      "freeform-tags": {},
```

```

    "id": "ocidl.mounttarget.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "nsg-ids": [],
    "private-ip-ids": [
      "ocidl.privateip.....uniqueID"
    ],
    "ocidl.privateip.....uniqueID"
    "subnet-id": "ocidl.subnet.....uniqueID",
    "time-created": "2021-06-16T22:56:57+00:00"
  }
]
}

```

- **Getting Mount Target Details**

1. Gather the information that you need to run the command:
 - Mount target ID (`oci fs mount-target list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target get
--mount-target-id <mount_target_OCID>
```

Example:

```

oci fs mount-target get \
--mount-target-id ocidl.mounttarget.....uniqueID
{
  "data": {
    "availability-domain": "MyAD",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyMountTarget",
    "export-set-id": "ocidl.exportset.....uniqueID",
    "freeform-tags": {},
    "id": "ocidl.mounttarget.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "nsg-ids": null,
    "private-ip-ids": [
      "ocidl.privateip.....uniqueID"
    ],
    "subnet-id": "ocidl.subnet.....uniqueID",
    "time-created": "2021-07-16T22:56:57+00:00"
  }
}

```

8.5.3 Changing the Mount Target Name

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Targets.
2. Select the compartment where the mount target resides.
3. Click the Action menu (three dots) for the mount target, and select Edit.
4. Change the name.
5. Click Save.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Mount target ID (`oci network subnet list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target update
--mount-target-id <mount_target_OCID>
--display-name "<New_Mount_Target_Name>"
```

Example:

```
oci fs mount-target update \
--mount-target-id ocid1.mounttarget.....uniqueID \
--display-name "MyMountTarget"

{
  "data": {
    "availability-domain": "pca",
    "compartment-id": "ocid1.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyMountTarget",
    "export-set-id": "ocid1.exportset.....uniqueID",
    "freeform-tags": {},
    "id": "ocid1.mounttarget.....uniqueID",
    "lifecycle-details": null,
    "lifecycle-state": "ACTIVE",
    "nsg-ids": null,
    "private-ip-ids": [
      "ocid1.privateip.....uniqueID"
    ],
    "subnet-id": "ocid1.subnet.....uniqueID",
    "time-created": "2021-06-17T19:01:37+00:00"
  },
  "etag": "b7efb0d7-d5fb-45d8-8bdd-a4a2f3f0371d"
}
```

8.5.4 Listing Exports

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Targets.
2. Select the compartment where the mount target resides.
3. Click on the mount target name.

The exports are display at the bottom of the page.

4. To see the export details, click the export name.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs export list
--compartment-id <compartment_id>
```

Example:

```
oci fs export list \
--compartment-id ocid1.....uniqueID
{
  "data": [
    {
      "export-set-id": "ocid1.exportset.....uniqueID",
      "file-system-id": "ocid1.filesystem.....uniqueID",
      "id": "ocid1.export.....uniqueID-1",
      "lifecycle-state": "ACTIVE",
      "path": "/export/8g0afgj16nuwx77a4wublc3ekkdaekeflbct2zt8qcbukfscnxmkp9su0ys",
      "time-created": "2021-06-17T21:15:44+00:00"
    },
    {
      "export-set-id": "ocid1.exportset.....uniqueID",
      "file-system-id": ".....uniqueID",
      "id": "ocid1.export.....uniqueID-2",
      "lifecycle-state": "ACTIVE",
      "path": "/export/8g0afgj16nuwx77a4wublc3ekkdaekeflbct2zt8qcbukfscnxmkp9su0ys",
      "time-created": "2021-06-17T21:20:55+00:00"
    }
  ]
}
```

8.5.5 Listing Export Sets

Using the OCI CLI

1. Gather the information that you need to run the command:

- Availability Domain Name (`oci iam availability-domain list`)
- Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs export-set list
--availability-domain <availability_domain_name>
--compartment-id <compartment_id>
```

Example:

```
oci fs export-set list \
--availability-domain pca \
--compartment-id ocid1.compartment.....uniqueID
{
  "data": [
    {
      "availability-domain": "pca",
      "compartment-id": "ocid1.compartment.....uniqueID",
      "display-name": "MyMountTarget2 - export set",
    }
  ]
}
```

```

    "id": "ocid1.exportset.....uniqueID6",
    "lifecycle-state": "ACTIVE",
    "time-created": "2021-06-17T19:01:37+00:00",
    "vcn-id": "ocid1.vcn.....uniqueID"
  }
]
}

```

8.5.6 Deleting an Export

Deleting an export deletes the file system path that clients use to mount the file system. Deleting an export does not delete any file systems.

Caution

When you delete an export, you can no longer mount the file system using the file path specified in the deleted export. Any clients that use the export path to mount a file system will not be able to access the file system.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the name of a file system that uses the export you plan to delete.
4. Click the Action menu (three dots) for the export and select Delete.
5. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:

- export OCID (`oci fs file-system list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs export delete
--export-id <export_OCID>
```

Example:

```
oci fs export delete --export-id ocid1.export.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

8.5.7 Moving a Mount Target to a Different Compartment

Using the OCI CLI

1. Gather the information that you need to run the command:

- Mount target OCID (`oci fs mount-target list`)
- Destination Compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target change-compartment
--mount-target-id <mount_target_OCID>
--compartment-id <destination_compartment_OCID>
```

Example:

```
oci fs mount-target change-compartment \
--mount-target-id ocid1.....uniqueID \
--compartment-id ocid1.compartment.....uniqueID
{
  "etag": "864d51bd-ed69-44bc-8c54-2a65d55fe07b"
}
```

8.5.8 Deleting a Mount Target

Caution

Deleting a mount target deletes all the exports that are associated with the mount target.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click Mount Targets.
2. Select the compartment where the mount target resides.
3. Click the Action menu (three dots) for the mount target you plan to delete.
4. Select Delete.
5. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - Mount target OCID (`oci fs mount-target list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs mount-target delete
--mount-target-id <mount_target_OCID>
```

Example:

```
oci fs mount-target delete \
--mount-target-id ocid1.mounttarget.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

8.6 Managing File Systems

A file system in the File Storage service represents a network file system that is mounted by one or more clients. File systems are associated with a single compartment. File systems must have at least one export

in one mount target for any client to mount and use the file system. Data is added to a file system from the client.

This section describes how to manage file systems after they are created. For instructions to create a file system, see [Section 8.1, “Creating a File System, Mount Target, and Export”](#).

8.6.1 Listing and Viewing the Details of a File System

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.

The file systems for the compartment are listed.

3. To see file system details, click the name of the file system.

Using the OCI CLI

• Listing File Systems

1. Gather the information that you need to run the command:
 - Availability Domain Name (`oci iam availability-domain list`)
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs file-system list
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
```

Example:

```
oci fs file-system list \
--availability-domain MyAD \
--compartment-id ocidl.compartment.....uniqueID
{
  "data": [
    {
      "availability-domain": "pca",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "MyFileSystem",
      "freeform-tags": {},
      "id": "ocidl.filesystem.....uniqueID-1",
      "kms-key-id": null,
      "lifecycle-state": "ACTIVE",
      "metered-bytes": 180224,
      "time-created": "2021-06-16T19:48:18+00:00"
    },
    {
      "availability-domain": "pca",
      "compartment-id": "ocidl.compartment.....uniqueID",
      "defined-tags": {},
      "display-name": "pluto",
      "freeform-tags": {},
```

```

    "id": "ocidl.filesystem.....uniqueID-2",
    "kms-key-id": null,
    "lifecycle-state": "ACTIVE",
    "metered-bytes": 147456,
    "time-created": "2021-06-17T23:16:43+00:00"
  }
]
}

```

• **Getting the File System Details**

1. Gather the information that you need to run the command:

- File System OCID (`oci fs file-system list`)

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```

oci fs file-system get
--file-system-id <file_system_OCID>

```

Example:

```

oci fs file-system get \
--file-system-id ocidl.filesystem.....uniqueID-1
{
  "data": {
    "availability-domain": "pca",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "MyFileSystem",
    "freeform-tags": {},
    "id": "ocidl.filesystem.....uniqueID-1",
    "kms-key-id": null,
    "lifecycle-state": "ACTIVE",
    "metered-bytes": 180224,
    "time-created": "2021-06-16T19:48:18+00:00"
  },
  "etag": "58dec47e-4732-4730-9e18-6b5db1ac30d6"
}

```

8.6.2 Changing the File System Name

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the Action menu (three dots) for the file system, and select Edit.
4. Enter a new name in the name field.
5. Click Save Changes.

Using the OCI CLI

1. Gather the information that you need to run the command:

- File System OCID (`oci fs file-system list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs file-system update
--file-system-id <file_system_OCID>
--display-name <new_file-system_name>
```

Example:

```
oci fs file-system update \
--file-system-id ocidl.filesystem.....uniqueID-2 \
--display-name neptune

{
  "data": {
    "availability-domain": "pca",
    "compartment-id": "ocidl.compartment.....uniqueID",
    "defined-tags": {},
    "display-name": "neptune",
    "freeform-tags": {},
    "id": "ocidl.filesystem.....uniqueID-2",
    "kms-key-id": null,
    "lifecycle-state": "ACTIVE",
    "metered-bytes": 147456,
    "time-created": "2021-06-17T23:16:43+00:00"
  },
  "etag": "6536c835-51bc-4288-a907-ae37d1af080b"
}
```

8.6.3 Moving a File System to a Different Compartment

Using the OCI CLI

1. Gather the information that you need to run the command:

- File System OCID (`oci fs file-system list`)
- Destination compartment OCID (`oci iam compartment list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs file-system change-compartment
--file-system-id <file-system_OCID>
--compartment-id <destination_compartment_OCID>
```

Example:

```
oci fs file-system change-compartment \
--file-system-id ocidl.filesystem.....uniqueID \
--compartment-id ocidl.compartment.....destination-uniqueID

{
  "etag": "0acc73ca-839d-451e-b079-4013889c233a"
}
```

8.6.4 Deleting a File System

A file system that has an export cannot be deleted. To delete the export, see [Section 8.5.6, “Deleting an Export”](#).

You cannot delete file systems that have dependencies. For example, if you have created a snapshot of this file system and then created a new file system from the snapshot, you cannot delete the source file system.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the Action menu (three dots) for the file system and select Delete.
4. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - File System OCID (`oci fs file-system list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs file-system delete
--file-system-id <file-system_OCID>
```

Example:

```
oci fs file-system delete \
--file-system-id ocidl.filesystem.....uniqueID
Are you sure you want to delete this resource? [y/N]: y
```

8.7 Managing Snapshots

8.7.1 Snapshots Overview

The File Storage service supports snapshots for data protection of your file system.

Snapshots are a consistent, point-in-time view of your file systems. Snapshots are copy-on-write, and scoped to the entire file system. The File Storage service encrypts all file system and snapshot data at rest. You can take as many snapshots as you need.

For more conceptual information, refer to *Snapshots* in the [File Storage](#) chapter of the [Oracle Private Cloud Appliance Concepts Guide](#).

This section provides instructions for managing file system snapshots.

8.7.2 Listing and Getting Snapshot Details

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the file system name.
4. In the Resources panel, click Snapshots.

The file system snapshots are listed.

5. To get the details for a specific snapshot, click the snapshot name.

Using the OCI CLI

- **Listing Snapshots**

1. Gather the information that you need to run the command:
 - File system OCID (`oci fs file-system list`)
2. Run this command.

Syntax (entered on a single line):

```
oci fs snapshot list
--file-system-id <file-system_OCID>
```

Example:

```
oci fs snapshot list \
--file-system-id ocid1.filesystem.....uniqueID
{
  "data": [
    {
      "defined-tags": {},
      "file-system-id": "ocid1.filesystem.....uniqueID",
      "freeform-tags": {},
      "id": "ocid1.snapshot.....uniqueID-1",
      "lifecycle-state": "ACTIVE",
      "name": "MySnapshot",
      "time-created": "2021-06-21T17:12:37+00:00"
    },
    {
      "defined-tags": {},
      "file-system-id": "ocid1.filesystem.....uniqueID",
      "freeform-tags": {},
      "id": "ocid1.snapshot.....uniqueID-2",
      "lifecycle-state": "ACTIVE",
      "name": "MySnapshot2",
      "time-created": "2021-06-21T17:31:18+00:00"
    }
  ]
}
```

- **Getting a Specific Snapshot**

1. Gather the information that you need to run the command:
 - Snapshot OCID (`oci fs snapshot list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs snapshot get \
--snapshot-id <snapshot_OCID>
```

Example:

```
oci fs snapshot get --snapshot-id ocidl.snapshot.....uniqueID
{
  "data": {
    "defined-tags": {},
    "file-system-id": "ocidl.filesystem.....uniqueID",
    "freeform-tags": {},
    "id": "ocidl.snapshot.....uniqueID",
    "lifecycle-state": "ACTIVE",
    "name": "MySnapshot",
    "time-created": "2021-06-21T17:12:37+00:00"
  },
  "etag": "f38aa070-0f3e-407f-a0b4-9bc841ff3fa4"
}
```

8.7.3 Creating a Snapshot

You can create a snapshot of a file system. A snapshot is a point-in-time view of the file system. The snapshot is accessible at `.zfs/snapshot/name`.

Using the Compute Web UI

1. In the navigation menu, under File Storage, click File Systems.
2. Select the appropriate compartment.
3. Click the file system name.
4. In the Resources panel, click Snapshots.
5. Click Create Snapshot.
6. Enter a name for the snapshot.

The name is limited to 64 characters and it must be unique among all other snapshots for this file system. The name can't be changed. Avoid entering confidential information.

7. Click Create Snapshot.

The snapshot is accessible under the root directory of the file system at `.zfs/snapshot/name`.

Using the OCI CLI

1. Gather the information that you need to run the command:
 - File system OCID (`oci fs file-system list`)
 - Snapshot name of your choice. The name is limited to 64 characters and it must be unique among all other snapshots for this file system. The name can't be changed. Avoid entering confidential information.

2. Run this command.

Note

This procedure shows the minimum required parameters for this command. For information about optional parameters, run the command with the `--help` option.

Syntax (entered on a single line):

```
oci fs snapshot create
--file-system-id <file-system_OCID>
--name <snapshot_name>
```

Example:

```
oci fs snapshot create \
--file-system-id ocidl.filesystem.....uniqueID \
--name "MySnapshot"
{
  "data": {
    "defined-tags": {},
    "file-system-id": "ocidl.filesystem.....uniqueID",
    "freeform-tags": {},
    "id": "ocidl.snapshot.....uniqueID",
    "lifecycle-state": "CREATING",
    "name": "MySnapshot",
    "time-created": null
  },
  "etag": "f38aa070-0f3e-407f-a0b4-9bc841ff3fa4"
}
```

8.7.4 Accessing a Snapshot on the Mounted File System

When a file system snapshot is created, it is placed in the file system. If the file system is mounted in a client system, you can access the snapshot on the client system.

The snapshot is accessible in this directory path: `<mount-point>/zfs/snapshot/<snapshot-name>`.

Using a UNIX-type OS.

1. Log into the instance OS that has the mounted the file system from which the snapshot was made.
2. List the snapshots.

Syntax:

```
ls -la <mount-point>/zfs/snapshot/
```

Example:

```
ls -la /mnt/MyMountPoint/zfs/snapshot
total 17
dr-xr-xr-x. 4 root root 4 Sep  8 15:54 .
dr-xr-xr-x. 4 root root 4 Sep  1 17:27 ..
drwxr-xr-x. 4 root root 7 Sep  8 15:53 file-system-FS-snapshot-02
drwxr-xr-x. 4 root root 6 Sep  1 18:12 file-system-FS-snapshot-01
```

3. Change to the directory of a snapshot.

Example:

```
cd /mnt/MyMountPoint/.zfs/snapshot/file-system-FS-snapshot-02
```

- List the contents of the snapshot.

```
ls -la
total 3027
drwxr-xr-x. 4 root root      7 Sep  8 15:53 .
dr-xr-xr-x. 4 root root      4 Sep  8 15:54 ..
-rwxr-xr-x. 1 root root    429 Sep  8 15:53 example1
drwxr-x---. 2 root sys       3 Sep  1 17:28 .$EXTEND
drwxr-xr-x. 2 root root      2 Sep  1 18:10 ABC-directory
-rw-r--r--. 1 root root      0 Sep  1 18:10 xyz-file
-rw-r--r--. 1 root root 3073219 Sep  1 18:12 zap.zip
```

8.7.5 Restoring a Snapshot (UNIX-Type Instances)

You can restore individual snapshot files or an entire snapshot using the `cp` command.

Note

Optionally, you can use `rsync`, `tar`, or another tool that supports NFS to copy your data to another remote location.

Using the Instance OS.

- Log into the instance OS that has the mounted the file system from which the snapshot was made.
- List the snapshots.

Syntax:

```
ls -la <mount-point>/.zfs/snapshot/
```

Example:

```
ls -la /mnt/MyMountPoint/.zfs/snapshot
total 17
dr-xr-xr-x. 4 root root 4 Sep  8 15:54 .
dr-xr-xr-x. 4 root root 4 Sep  1 17:27 ..
drwxr-xr-x. 4 root root 7 Sep  8 15:53 file-system-FS-snapshot-02
drwxr-xr-x. 4 root root 6 Sep  1 18:12 file-system-FS-snapshot-01
```

- Use the `cp` command to copy individual snapshot files, or the entire snapshot to a location of your choice.

Use the `-r` option when restoring a snapshot that contains subdirectories.

Example:

```
cp -r /mnt/MyMountPoint/.zfs/snapshot/<snapshot_name>/ * <destination_directory>
```

8.7.6 Deleting a Snapshot

You cannot delete snapshots that have dependencies. For example, if you have created a file system from the snapshot, you cannot delete the source snapshot.

Using the Compute Web UI

- In the navigation menu, under File Storage, click File Systems.

2. Select the appropriate compartment.
3. Click the name of the file system where the snapshot resides.
4. In the Resources panel, click Snapshots.
5. Click the Actions icon (three dots), and then click Delete.
6. Confirm the deletion.

Using the OCI CLI

1. Gather the information that you need to run the command:

- Snapshot OCID (`oci fs snapshot list`)

2. Run this command.

Syntax (entered on a single line):

```
oci fs snapshot delete  
--snapshot-id <snapshot_OCID>
```

Example:

```
oci fs snapshot delete \  
--snapshot-id ocid1.snapshot.....uniqueID  
Are you sure you want to delete this resource? [y/N]: y
```

Chapter 9 Object Storage

Table of Contents

9.1 Object Storage	336
9.2 Managing Buckets	336
9.2.1 Object Storage Buckets	336
9.2.2 Obtaining the Object Storage Namespace	336
9.2.3 Listing Buckets	336
9.2.4 Viewing Bucket Details	337
9.2.5 Creating a Bucket	338
9.2.6 Moving a Bucket to a Different Compartment	340
9.2.7 Deleting a Bucket	341
9.3 Managing Storage Objects	342
9.3.1 Objects	342
9.3.2 Viewing Objects in a Bucket	342
9.3.3 Creating a Folder or SubFolder	344
9.3.4 Uploading an Object	345
9.3.5 Performing a Multi-Part Upload	346
9.3.6 Listing the Parts of an Unfinished or Failed Multi-part Upload	347
9.3.7 Canceling a Multi-Part Upload	347
9.3.8 Performing a Bulk Object Upload	348
9.3.9 Copying an Object to a Different Bucket	349
9.3.10 Downloading an Object	350
9.3.11 Performing a Multi-Part Download	351
9.3.12 Performing a Bulk Download	352
9.3.13 Deleting an Object	352
9.3.14 Performing a Bulk Delete of All Objects in a Bucket	353
9.4 Managing Object Versioning	354
9.4.1 Object Versioning	354
9.4.2 Enabling Versioning During Bucket Creation	355
9.4.3 Enabling, Disabling, or Suspending Versioning (After Bucket Creation)	356
9.4.4 Viewing Object Versions and Details	357
9.4.5 Deleting the Previous Version of an Object	358
9.4.6 Recovering a Deleted Object Version	358
9.5 Using Pre-Authenticated Requests	359
9.5.1 Pre-Authenticated Requests	359
9.5.2 Listing Pre-Authenticated Requests	359
9.5.3 Creating a Pre-Authenticated Request for All Objects in a Bucket	361
9.5.4 Creating a Pre-Authenticated Request for a Specific Object	362
9.5.5 Constructing the Pre-Authenticated Request URL	364
9.5.6 Deleting a Pre-Authenticated Request	364
9.5.7 Listing Objects for Pre-Authenticated Requests	365
9.5.8 Uploading an Object Using a Pre-Authenticated Request	365
9.5.9 Downloading an Object Using a Pre-Authenticated Request	365
9.6 Defining Retention Rules	366
9.6.1 Retention Rules	366
9.6.2 Viewing Retention Rules and Details	366
9.6.3 Creating a Retention Rule	368
9.6.4 Modifying a Retention Rule	370
9.6.5 Deleting a Retention Rule	371

9.1 Object Storage

The Object Storage service is a storage platform that offers reliable and cost-efficient data durability.

The Object Storage service stores unstructured data of any content type, including analytic data and rich content, like images and videos.

The data is stored as an object in a bucket. Buckets are associated with a compartment within a tenancy.

An Object Storage namespace serves as the top-level container for all buckets and objects. At account creation time, each tenant is assigned one unique system-generated and immutable Object Storage namespace name. The namespace spans all compartments.

With Object Storage, you can safely and securely store or retrieve data directly from the internet or from within the cloud appliance.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#)

This chapter provides instructions for managing Object Storage.

9.2 Managing Buckets

9.2.1 Object Storage Buckets

A bucket is a container for storing objects in a compartment within an Object Storage namespace.

A bucket is associated with a single compartment. The compartment has policies that indicate what actions you can perform on a bucket and all of the objects in the bucket.

A bucket cannot contain other buckets.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#).

9.2.2 Obtaining the Object Storage Namespace

An Object Storage namespace serves as the top-level container for all buckets and objects. Each tenant is assigned one unique system-generated and immutable Object Storage namespace name. The namespace spans all compartments.

Use this step to view your Object Storage namespace name.

Using the Compute Web UI

1. Click your user name (upper right corner), and select Tenancy.

The namespace string is listed under Object Storage Settings.

9.2.3 Listing Buckets

Using the Compute Web UI

1. In the navigation menu, under Object Storage, click Object Storage.

A list of the buckets in the compartment you're viewing is displayed.

2. If you don't see the bucket you're looking for, ensure that you're viewing the correct compartment (select from the list at the top of the page).

The page shows only the resources in that compartment.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, "Obtaining the Object Storage Namespace"](#))
 - Compartment OCID (`oci iam compartment list`)
2. Run this command.

Syntax (entered on a single line):

```
oci os bucket list
--namespace-name <object_storage_namespace>
--compartment-id <compartment_OCID>
```

Example:

```
oci os bucket list \
--namespace-name examplnamespace \
--compartment-id ocid.compartment.....uniqueID

{
  "data": [
    {
      "compartment-id": "ocid.compartment.....uniqueID",
      "created-by": "ocidl.user.....uniqueID",
      "defined-tags": null,
      "etag": "cdb5bc11561e476cb0d8aa5b8f8668f6",
      "freeform-tags": null,
      "name": "MyBucket",
      "namespace": "export/examplnamespace",
      "time-created": "2021-05-04T18:56:39+00:00"
    },
    {
      "compartment-id": "ocid.compartment.....uniqueID",
      "created-by": "ocidl.user.....uniqueID",
      "defined-tags": null,
      "etag": "aa7642fec45729ce7cb8b321d3ee1463",
      "freeform-tags": null,
      "name": "JoesBucket",
      "namespace": "export/examplnamespace",
      "time-created": "2021-05-04T20:26:33+00:00"
    }
  ]
}
```

9.2.4 Viewing Bucket Details

Use this task to view bucket details.

Using the Compute Web UI

1. In the navigation menu, under Object Storage, click Object Storage.

A list of the buckets in the compartment you are viewing is displayed.

2. From the list at the top of the page, select the compartment where the bucket resides.
3. Click the bucket name to display the details.
4. Click View or Copy.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run this command.

Syntax (entered on a single line):

```
oci os bucket get
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

The OCID is identified as `id` in the output.

Example:

```
oci os bucket get \
--namespace-name examplnamespace \
--bucket-name MyBucket

{
  "data": {
    "approximate-count": null,
    "approximate-size": null,
    "compartment-id": "ocid.compartment.....uniqueID",
    "created-by": "ocid1.user.....uniqueID",
    "defined-tags": null,
    "etag": "cdb5bc11561e476cb0d8aa5b8f8668f6",
    "freeform-tags": null,
    "id": ocid.bucket.....uniqueID,
    "is-read-only": null,
    "kms-key-id": null,
    "metadata": null,
    "name": "MyBucket",
    "namespace": "export/exemplnamespace",
    "object-events-enabled": null,
    "object-lifecycle-policy-etag": null,
    "public-access-type": "NoPublicAccess",
    "replication-enabled": null,
    "storage-tier": "Standard",
    "time-created": "2021-05-04T18:56:39+00:00",
    "versioning": null
  },
  "etag": "cdb5bc11561e476cb0d8aa5b8f8668f6"
}
```

9.2.5 Creating a Bucket

Use this procedure to create an Object Storage bucket.

When you create a bucket, it is created with no public access. To make the bucket publicly available, see [Section 9.5, “Using Pre-Authenticated Requests”](#).

Using the Compute Web UI

1. In the navigation menu, click Object Storage, then click Object Storage.
2. Click Create Bucket.
3. Enter the following details:
 - **Name:** Enter a name for the bucket.
Specify a name that is unique within your tenancy Object Storage namespace.
 - **Create in Compartment:** Select the compartment in which to create this bucket.
 - **Enable Object Versioning:** Optionally, you can enable object versioning.
For more information, refer to [Object Storage](#).
 - **Tagging:** Optionally, add one or more tags to this resource.
If you are not sure whether to apply tags, skip this option (you can apply tags later).
For more information about tagging resources, see [Tagging Overview](#).
4. Click Create Bucket.
The bucket is created immediately and you can start uploading objects. See [Section 9.3.4, “Uploading an Object”](#).

Using the OCI CLI

1. Gather the information you need to run the next command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Compartment OCID (`oci iam compartment list`)
 - Bucket name: The name you want for this bucket.
2. Run this command.

Syntax (entered on a single line):

```
oci os bucket create
--namespace-name <object_storage_namespace>
--compartment-id <compartment_OCID>
--name <bucket_name>
```

The bucket is created immediately and you can start uploading objects. See [Section 9.3.4, “Uploading an Object”](#).

Example:

```
oci os bucket create \
--namespace-name examplnamespace \
--compartment-id ocid.compartment.....uniqueID \
```

```
--name MyBucket

{
  "data": {
    "approximate-count": null,
    "approximate-size": null,
    "compartment-id": "ocidl.compartment.....uniqueID",
    "created-by": "ocidl.user.....uniqueID",
    "defined-tags": null,
    "etag": "b78d4193ab3eb2270b1373aa52b443a1",
    "freeform-tags": null,
    "id": null,
    "is-read-only": null,
    "kms-key-id": null,
    "metadata": null,
    "name": "MyBucket",
    "namespace": "export/exemplenamespace",
    "object-events-enabled": null,
    "object-lifecycle-policy-etag": null,
    "public-access-type": "NoPublicAccess",
    "replication-enabled": null,
    "storage-tier": "Standard",
    "time-created": "2021-06-11T20:11:02+00:00",
    "versioning": null
  },
  "etag": "b78d4193ab3eb2270b1373aa52b443a1"
}
```

9.2.6 Moving a Bucket to a Different Compartment

You can move a bucket from one compartment to another as long as both the source and target compartments are in the same tenancy. This includes moving a bucket from one compartment level down to a sub-level within the source compartment.

Using the CLI

1. Gather the information you need for the next command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Compartment OCID of the compartment you are moving the bucket to (`oci iam compartment list`)
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run this command to move the bucket.

Syntax (entered on a single line):

```
oci os bucket update
--namespace-name <object_storage_namespace>
--compartment-id <target_compartment_id>
--bucket-name <bucket_name>
```

Example:

```
oci os bucket update \
--namespace-name exemplenamespace \
--compartment-id ocidl.compartment.....target-compartmentID \
--bucket-name MyBucket
{
  "data": {
    "approximate-count": null,
```

```

"approximate-size": null,
"compartment-id": "ocidl.compartment.....target-compartmentID",
"created-by": "ocidl.user.....uniqueID",
"defined-tags": null,
"etag": "5d72fb7ac4385e24f42ac830bc6490ca",
"freeform-tags": null,
"id": null,
"is-read-only": null,
"kms-key-id": null,
"metadata": null,
"name": "MyBucket",
"namespace": "export/exemplenamespace",
"object-events-enabled": null,
"object-lifecycle-policy-etag": null,
"public-access-type": "NoPublicAccess",
"replication-enabled": null,
"storage-tier": "Standard",
"time-created": "2021-06-02T20:44:57+00:00",
"versioning": null
},
"etag": "5d72fb7ac4385e24f42ac830bc6490ca"
}

```

3. Run this command to verify that the bucket moved to the correct compartment:

Syntax (entered on a single line):

```

oci os bucket list
--namespace-name <object_storage_namespace>
--compartment-id <target_compartment_OICD>

```

Example:

```

oci os bucket list \
--namespace-name exemplenamespace \
--compartment-id ocidl.compartment.....target-compartmentID
{
  "data": [
    {
      "compartment-id": "ocidl.compartment.....target-compartmentID",
      "created-by": "ocidl.user.....uniqueID",
      "defined-tags": null,
      "etag": "5d72fb7ac4385e24f42ac830bc6490ca",
      "freeform-tags": null,
      "name": "MyBucket",
      "namespace": "export/exemplenamespace",
      "time-created": "2021-06-02T20:44:57+00:00"
    }
  ]
}

```

9.2.7 Deleting a Bucket

Caution

You cannot recover a deleted bucket.

You can permanently delete an empty bucket. You cannot delete a bucket that contains any of the following:

- Any objects
- Previous versions of an object
- A multipart upload in progress

- A pre-authenticated request

Tip

When you delete an object in a version-enabled bucket, a previous version of that object is created. Select Show Deleted Objects to display the object versions that might prevent you from deleting the bucket. For more information, see [Section 9.4, “Managing Object Versioning”](#).

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run this command.

Syntax (entered on a single line):

```
oci os bucket delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

Example:

```
oci os bucket delete \
--namespace-name examplnamespace \
--bucket-name MyBucket
Are you sure you want to delete this resource? [y/N]: y
```

9.3 Managing Storage Objects

9.3.1 Objects

In the Object Storage service, an object is a file or unstructured data you upload to a bucket within a compartment within an Object Storage namespace.

The object can be any type of data, for example, multimedia files, data backups, static web content, or logs. You can store objects that are up to 10 TiB. Objects are processed as a single entity. You can't edit or append data to an object, but you can replace the entire object.

Object Storage is not tied to any specific compute instance. You can access data from anywhere inside or outside the context of the Oracle Private Cloud Appliance, as long you have internet connectivity, access to the Object Storage endpoint, and authorization.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#).

9.3.2 Viewing Objects in a Bucket

Using the Compute Web UI

1. In the navigation menu, under Object Storage, click Object Storage.
2. Choose the compartment that contains the bucket that contains your object.

A list of buckets is displayed.

3. Click the bucket name that contains your object.
4. Click Objects under Resources.

Using the OCI CLI

• Listing Objects in a bucket

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Enter this command.

Syntax (entered on a single line):

```
oci os object list
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

Example:

```
oci os object list \
--namespace-name examplnamespace \
--bucket-name MyBucket

{
  "data": [
    {
      "etag": null,
      "md5": "Ucf+fZbCK/RN5gGsE17G5w=",
      "name": "eventslogreference.htm",
      "size": 1363,
      "time-created": "2021-06-01T17:57:16+00:00",
      "time-modified": null
    }
  ],
  "prefixes": []
}
```

• Listing object details

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object name (`oci os object list`), see previous example
2. Run this command.

Syntax (entered on a single line):

```
oci os object head
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

```
--name <object_name>
```

Example:

```
oci os object head \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--name eventslogreference.htm

{
  "access-control-allow-credentials": "true",
  "access-control-allow-methods": "POST,PUT,GET,HEAD,DELETE",
  "access-control-allow-origin": "*",
  "access-control-expose-headers": "Content-Type, Etag, last-modified, Content-MD5, Content-Length, opc-client",
  "connection": "Keep-Alive",
  "content-length": "1363",
  "content-md5": "Ucf+fZbCK/RN5gGsEl7G5w==",
  "content-type": "application/octet-stream",
  "date": "Tue, 01 Jun 2021 18:05:32 GMT",
  "etag": "33edlaff724eac56f00616552fc61f3e",
  "keep-alive": "timeout=5, max=100",
  "last-modified": "2021-06-01T17:57:16.000Z",
  "opc-client-request-id": "8965F8B5A9B84F00B51D4C965F029230",
  "opc-request-id": "txae7c2c9aa7094f16adee8-0060b676ec",
  "server": "Apache",
  "x-content-type-options": "nosniff"
}
```

9.3.3 Creating a Folder or SubFolder

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object file location
 - Object name
2. Run this command.

Syntax:

```
oci os object put
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--file <file_location>
--name <object_name>
```

Example:

```
oci os object put \
--namespace-name examplnamespace \
--bucket-name Bucket1_objv-enabl \
--file /home/log_files/install.log \
--name /home/log_files/install.log
```

```
oci os object put \
--namespace-name examplnamespace \
--bucket-name Bucket1_objv-enabl \
```

```

--file myfile \
--name /home/log_files/install.log

oci os object put \
--namespace-name examplnamespace \
--bucket-name Bucket1_objv-enabl \
--file /home/log_files/install.log \
--name /home/log_files/install.log

Uploading object [#####] 100%
{
  "etag": "bae04836d4ea5d521c23cbee70566cf2",
  "last-modified": "2021-05-13T15:37:18.000Z",
  "opc-content-md5": "GWZbZ8CXPCjLcPxBs6cPCQ=="
}

```

9.3.4 Uploading an Object

Using the OCI CLI

An object can be uploaded as a single part or as multiple parts. Use the `--no-multipart` option to upload as a single part. For detailed information on multi-part uploads, see [Section 9.3.5, “Performing a Multi-Part Upload”](#).

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object file location
2. Run this command.

Syntax (entered on a single line):

```

oci os object put
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--file <file_location>

```

`<file_location>` is the source directory path of the object being uploaded, such as `C:\workspace\Uploads\MyFile.txt` or `/home/user/Documents/Uploads/MyFile.txt`.

Example:

```

oci os object put \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--file /home/user/Documents/Uploads/MyFile.txt

Uploading object [#####] 100%
{
  "etag": "33ed1aff724eac56f00616552fc61f3e",
  "last-modified": "2021-06-01T17:57:16.000Z",
  "opc-content-md5": "Ucf+fZbCK/RN5gGsE17G5w=="
}

```

Example of uploading an object using the original file names:

```

oci os object put \
--namespace-name examplnamespace \

```

```
--bucket-name MyBucket \
--file C:\workspace\Uploads\MyFile.txt \
--no-multipart
{
  "etag": "cadb9f8a-3292-45e6-ale8-f075699fb619",
  "last-modified": "Fri, 11 Dec 2020 14:04:19 GMT",
  "opc-content-md5": "9P6l0SaYe4fXxaeK8siuDw=="
}
```

9.3.5 Performing a Multi-Part Upload

With multi-part uploads, individual parts of an object can be uploaded in parallel to reduce the amount of time you spend uploading.

Multi-part uploads accommodate objects that are too large for a single upload operation. Object parts must be no larger than 50 GiB.

You can pause between the uploads of individual parts, and resuming the upload when your schedule and resources allow.

Using the OCI CLI

To upload an object, run `oci os object put` with the `--part-size` flag. The `--part-size` value represents the size of each part in mebibytes (MiBs). Object Storage waives the minimum part size restriction for the last uploaded part. The `--part-size` value must be an integer.

Optionally, you can use the `--parallel-upload-count` flag to set the maximum number of parallel uploads allowed.

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object file location
2. Run the command.

Syntax (entered on a single line):

```
oci os object put
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--file <file_location>
--parallel-upload-count <maximum_number_parallel_uploads>
--part-size <upload_part_size_in_MB>
--force
```

Example:

```
oci os object put \
--namespace-name examplnamespace \
--file /boot/intramfs-0-rescue-e542c19f0fbf4e41a41428d933a7357f.img \
--parallel-upload-count 5 \
--part-size 15 \
--force

Upload ID: a21bba2c-8922-4b9c-a98a-9ef3569c0138
Split file into 6 parts for upload.
Uploading object [#####] 100%
{
```

```

"etag": "0964effc8dc4394fd317f03a025ae5d0",
"last-modified": "2021-05-11T21:35:19",
"opc-multipart-md5": "UIVRhiwSHY6o0E4pi/yfGg==6"
}

```

9.3.6 Listing the Parts of an Unfinished or Failed Multi-part Upload

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run this command.

Syntax (entered on a single line):

```

oci os multipart list
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>

```

Example:

```

oci os multipart list
--namespace-name examplnamespace \
--bucket-name MyBucket \
{
  "data": [
    {
      "bucket": "MyBucket",
      "namespace": "examplnamespace",
      "object": "MyObject",
      "time-created": "2019-07-25T21:55:21.973000+00:00",
      "upload-id": "0b7abd48-9ff2-9d5f-2034-63a02fdd7afa"
    },
    {
      "bucket": "MyBucket",
      "namespace": "examplnamespace",
      "object": "MyObject",
      "time-created": "2019-07-25T21:53:09.246000+00:00",
      "upload-id": "1293ac9d-83f8-e055-a5a7-d1e13277b5c0"
    },
    {
      "bucket": "MyBucket",
      "namespace": "examplnamespace",
      "object": "MyObject",
      "time-created": "2019-07-25T21:46:34.981000+00:00",
      "upload-id": "33e7a875-9e94-c3bc-6577-2ee5d8226b53"
    }
  ]
}
...

```

9.3.7 Canceling a Multi-Part Upload

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)

- Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
- Upload ID (`oci os multipart list`), see [Section 9.3.6, “Listing the Parts of an Unfinished or Failed Multi-part Upload”](#)

2. Run this command.

Syntax (entered on a single line):

```
oci os multipart abort
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--object-name <object_name>
--upload-id <upload_ID>
```

Example:

```
oci os multipart abort \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--object-name MyObject \
--upload-id 22d5f6d2-8e03-48ca-8593-0192d25770b8

"data": [
{
"etag": "dd434179cfbc22458a9739096ec43226",
"md5": "PBrT093rZrcSDwQsKh9azQ==",
"part-number": 13,
"size": 15728640
}
],
"opc-next-page": "00013"
}
WARNING: Are you sure you want to permanently remove this incomplete upload? [y/N]: y
```

9.3.8 Performing a Bulk Object Upload

Bulk operations at a specific level of the hierarchy do not affect objects in any level above.

Using the OCI CLI

1. Gather the information you need to run the command.

- Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
- Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
- Source directory location – is the upload directory path, such as `C:\workspace\Upload\` or `/home/user/Documents/Upload`. If your source directory has subdirectories, the subdirectory names are prepended to the names of the files stored in those subdirectories, delimited with a forward slash (/) character. For example, if a file named `maple.jpg` is stored in the subdirectory `trees`, when the file is uploaded, Object Storage assigns the name `trees/maple.jpg` to the object.

2. Run this command.

Syntax (entered on a single line):

```
oci os object bulk-upload
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

```
--src-dir <source_directory_location>
```

Example:

```
oci os object bulk-upload
--namespace-name examplnamespace \
--bucket-name MyBucket \
--src-dir /home/log-dir/

Uploaded Jan-logs [#####] 100%
Uploaded Feb-logs [#####] 100%
Uploaded Mar-logs [#####] 100%
Uploaded Apr-logs [#####] 100%

{
  "skipped-objects": [],
  "upload-failures": {},
  "uploaded-objects": {
    "Jan-logs": {
      "etag": "33ed1aff724eac56f00616552fc61f3e",
      "last-modified": "2021-06-01T20:42:50.000Z",
      "opc-content-md5": "Ucf+fZbCK/RN5gGsE17G5w=="
    },
    "Feb-logs": {
      "etag": "e1875449257cc6ac6ab93cc9c7921c87",
      "last-modified": "2021-06-01T20:42:50.000Z",
      "opc-content-md5": "1B2M2Y8AsgTpgAmY7PhCfg=="
    },
    "Mar-logs": {
      "etag": "c784ac5216d889f55138ecfb428eee3c",
      "last-modified": "2021-06-01T20:42:51.000Z",
      "opc-content-md5": "1B2M2Y8AsgTpgAmY7PhCfg=="
    },
    "Apr-logs": {
      "etag": "3b4571c73bdb9e44bec0512a5e48fba7",
      "last-modified": "2021-06-01T20:42:51.000Z",
      "opc-content-md5": "1B2M2Y8AsgTpgAmY7PhCfg=="
    }
  }
}
```

9.3.9 Copying an Object to a Different Bucket

You can copy an object to a different bucket as long as the target bucket is located in the same region.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Name of the Source object (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
 - Name of the destination bucket (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Name of the object in the new destination

2. Run this command.

Syntax (entered on a single line):

```
oci os object copy
--namespace-name <object_storage_namespace>
--bucket-name <source_bucket_name>
--source-object-name <source_object>
--destination-bucket <destination_bucket_name>
--destination-object-name <destination_object_name>
```

Example:

```
oci os object copy
--namespace-name examplnamespace
--bucket-name MyBucket
--source-object-name Compute_Logs.tar.gz
--destination-bucket Bucket-log-backups
--destination-object-name Compute_Logs.tar.gz.backup
```

Verify that the copied object is in the bucket.

```
oci os object list
--namespace-name examplnamespace
--bucket-name Bucket-log-backups

{
  "data": [
    {
      "etag": null,
      "md5": "XzYkstrjaprhbZyemalRbQ==",
      "name": "Compute_Logs.tar.gz.backup",
      "size": 132631,
      "time-created": "2021-04-01T21:00:55+00:00",
      "time-modified": null
    }
  ],
  "prefixes": []
}
```

9.3.10 Downloading an Object

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
 - Object file location
2. Run this command.

Syntax (entered on a single line):

```
oci os object get
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--name <object_name>
--file <file_location>
```

`<file_location>` is the destination path for the file being downloaded, such as `C:\workspace\Downloads\MyFile.txt` or `/home/user/Documents/Downloads/MyFile.txt`.

Example:

```
oci os object get \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--name photos \
--file /home/photos_backup

Downloading object [#-----] 100%

# ls -l
total 8
-rw-r--r-- 1 root root 1363 Jun 1 17:56 photo1
-rw-r--r-- 1 root root 1363 Jun 1 21:40 photo1_backup
-rw-r--r-- 1 root root 0 Jun 1 20:42 photo2
-rw-r--r-- 1 root root 0 Jun 1 20:42 photo3
-rw-r--r-- 1 root root 0 Jun 1 20:42 photo4
```

9.3.11 Performing a Multi-Part Download

Using the OCI CLI

- Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
 - Object file location
 - The byte-range for the download. Multi-part object downloading is available using the byte-range request standard defined in [RFC 7233, section 2.1](#)
- Run the command.

Syntax (entered on a single line):

```
oci os object get
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--name <object_name>
--file <file_location>
--range bytes=<byte_range>
```

Example:

```
oci os object get \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--name MyObject.mp4 \
--file c:\workspace\Downloads\MyObject.mp4 \
--range bytes=0-50

cusobjstorenamespace --range bytes=0-50
Downloading object [#-----] 3%

# ls -l
total 12
-rw-r--r-- 1 root root 1363 Jun 1 17:56 abc.mp41
-rw-r--r-- 1 root root 51 Jun 1 21:50 def.mp4
```

```
-rw-r--r-- 1 root root 1363 Jun 1 21:40 ghi.mp4
-rw-r--r-- 1 root root 0 Jun 1 20:42 jkl.mp4
-rw-r--r-- 1 root root 0 Jun 1 20:42 mno.mp4
-rw-r--r-- 1 root root 0 Jun 1 20:42 pqr.mp4
```

9.3.12 Performing a Bulk Download

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Download directory. `<download_directory_location>` is the destination path for the objects being downloaded, such as `C:\workspace\Downloads\` or `/home/user/Documents/Downloads/`. If the directory does not exist, Object Storage creates the directory when the command runs.
2. Run the command.

Syntax (entered on a single line):

```
oci os object bulk-download
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--download-dir <download_directory_location>
```

Example:

```
oci os object bulk-download \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--download-dir c:\workspace\Downloads

Downloaded MyFile.txt [#####] 100%
Downloaded logFile.log [#####] 100%

{
  "download-failures": {},
  "skipped-objects": []
}
```

9.3.13 Deleting an Object

You can permanently delete an object from a bucket or folder. You cannot, however, recover a deleted object unless you have object versioning enabled. See [Section 9.4, “Managing Object Versioning”](#) for details.

You cannot delete an object that has an active retention rule. See [Section 9.6, “Defining Retention Rules”](#) for details.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))

- Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
- Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)

2. Syntax (entered on a single line):

```
oci os object delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--object-name <object_name>
```

Example:

```
oci os object delete \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--object-name MyFile.txt

Are you sure you want to delete this resource? [y/N]: y
```

9.3.14 Performing a Bulk Delete of All Objects in a Bucket

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. To see a list of the files impacted by a bulk delete command without actually deleting the files, use the `--dry-run` option.

Syntax (entered on a single line):

```
oci os object bulk-delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--dry-run
```

Example:

```
oci os object bulk-delete \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--dry-run
{
  "delete-failures": {},
  "deleted-objects": [
    "MyFile.txt",
    "logfile.log"
  ]
}
```

3. To perform the bulk deletion:

Syntax (entered on a single line):

```
oci os object bulk-delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

Example:

```
oci os object bulk-delete \
--namespace-name examplnamespace \
--bucket-name MyBucket

WARNING: This command will delete 2 objects. Are you sure you wish to continue? [y/N]:y

Deleted MyRenamedFile.txt [#####] 100%
Deleted logFile.log [#####] 100%

{
  "delete-failures": {},
  "deleted-objects": [
    "MyFile.txt",
    "logFile.log"
  ]
}
```

9.4 Managing Object Versioning

9.4.1 Object Versioning

Object versioning provides data protection against accidental or malicious object updates, overwrites, or deletions.

Object versioning is enabled at the bucket level. Versioning directs Object Storage to automatically create an object version each time a new object is uploaded, an existing object is overwritten, or when an object is deleted. You can enable object versioning at bucket creation time or later. A bucket that is versioning-enabled can have many versions of an object. There is always one latest version of the object and zero or more previous versions.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#).

9.4.1.1 Determining the Object Versioning Status for a Bucket

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, "Obtaining the Object Storage Namespace"](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, "Listing Buckets"](#)
2. Run the command.

Syntax (entered on a single line):

```
oci os bucket get
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

The "versioning" line lists the state. Example:

```
oci os bucket get \
--namespace-name examplnamespace \
--bucket-name bucket-4-versioning
```

```

{
  "data": {
    "approximate-count": null,
    "approximate-size": null,
    "compartment-id": "ocidl.compartment.....uniqueID",
    "created-by": "ocidl.user.....uniqueID",
    "defined-tags": null,
    "etag": "00b4edbb27012ae78a912428ad1e630c",
    "freeform-tags": null,
    "id": null,
    "is-read-only": null,
    "kms-key-id": null,
    "metadata": null,
    "name": "bucket-4-versioning",
    "namespace": "export/examplnamespace",
    "object-events-enabled": null,
    "object-lifecycle-policy-etag": null,
    "public-access-type": "NoPublicAccess",
    "replication-enabled": null,
    "storage-tier": "Standard",
    "time-created": "2021-06-10T18:39:12+00:00",
    "versioning": "Enabled"
  },
  "etag": "00b4edbb27012ae78a912428ad1e630c"
}

```

9.4.2 Enabling Versioning During Bucket Creation

Object versioning provides data protection against accidental or malicious object updates and deletions.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, "Obtaining the Object Storage Namespace"](#))
 - Compartment OCID (`oci iam compartment list -all`)
 - Bucket name: The name you want for this bucket.
2. Syntax (entered on a single line):

```

oci os bucket create
--namespace-name <object_storage_namespace>
--compartment-id <target_compartment_id>
--name <bucket_name>
--versioning enabled

```

Example:

```

oci os bucket create \
--namespace-name examplnamespace \
--compartment-id ocid.compartment.....exampleuniqueID \
--name MyStandardBucket \
--versioning enabled

{
  "data": {
    "approximate-count": null,
    "approximate-size": null,
    "compartment-id": "ocidl.compartment.....uniqueID",
    "created-by": "ocidl.user.....uniqueID",
    "defined-tags": null,
    "etag": "00b4edbb27012ae78a912428ad1e630c",

```

```

"freeform-tags": null,
"id": null,
"is-read-only": null,
"kms-key-id": null,
"metadata": null,
"name": "bucket-4-versioning",
"namespace": "export/examplnamespace",
"object-events-enabled": null,
"object-lifecycle-policy-etag": null,
"public-access-type": "NoPublicAccess",
"replication-enabled": null,
"storage-tier": "Standard",
"time-created": "2021-06-10T18:39:12+00:00",
"versioning": "Enabled"
},
"etag": "00b4edbb27012ae78a912428ad1e630c"
}

```

9.4.3 Enabling, Disabling, or Suspending Versioning (After Bucket Creation)

Object versioning provides data protection against accidental or malicious object updates and deletions.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Compartment OCID (`oci iam compartment list -all`)
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run the command.

Syntax (entered on a single line):

```

oci os bucket update
--namespace-name <object_storage_namespace>
--compartment-id <target_compartment_id>
--bucket-name <bucket_name>
--versioning <enabled | disabled | suspended>

```

For `--versioning`, choose one of the options: `enabled`, `disabled`, or `suspended`.

Example of enabling object versioning:

```

oci os bucket update \
--namespace-name examplnamespace \
--compartment-id ocid1.compartment.....uniqueID \
--bucket-name MyBucket \
--versioning Enabled
{
  "data": {
    "approximate-count": null,
    "approximate-size": null,
    "compartment-id": "ocid1.compartment.....uniqueID",
    "created-by": "ocid1.user.....uniqueID",
    "defined-tags": null,
    "etag": "117f0608bdf83b9c7ea393db556a0ee4",
    "freeform-tags": null,
    "id": null,
    "is-read-only": null,
    "kms-key-id": null,

```

```

"metadata": null,
"name": "MyBucket",
"namespace": "export/exemplnamespace",
"object-events-enabled": null,
"object-lifecycle-policy-etag": null,
"public-access-type": "ObjectRead",
"replication-enabled": null,
"storage-tier": "Standard",
"time-created": "2021-06-02T17:06:18+00:00",
"versioning": "Enabled"
},
"etag": "117f0608bdf83b9c7ea393db556a0ee4"
}

```

9.4.4 Viewing Object Versions and Details

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, "Obtaining the Object Storage Namespace"](#))
 - Compartment OCID (`oci iam compartment list -all`)
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, "Listing Buckets"](#)
2. Run the command.

Syntax (entered on a single line):

```

oci os object list-object-versions
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>

```

Example:

```

oci os object list-object-versions \
--namespace-name exemplnamespace \
--bucket-name MyBucket
{
  "data": [
    {
      "etag": null,
      "is-delete-marker": false,
      "md5": "3DI5GbLmKiRxY/ozWxyXHQ==",
      "name": "bucket-data",
      "size": 103,
      "time-created": "2021-06-02T22:20:25+00:00",
      "time-modified": null,
      "version-id": null
    },
    {
      "etag": null,
      "is-delete-marker": false,
      "md5": "VIic5JncRWwDQj6CnsZlWw==",
      "name": "compute.log",
      "size": 4878456,
      "time-created": "2021-06-10T19:03:26+00:00",
      "time-modified": null,
      "version-id": "5f4ce7e8-656f-409a-b70a-ebfedddcfeda"
    }
  ],
  "prefixes": []
}

```

9.4.5 Deleting the Previous Version of an Object

When versioning is enabled, deleting an object without targeting a specific version creates a delete marker and previous version of the object that can be recovered. However, deleting a previous version of an object is a permanent deletion.

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Compartment OCID (`oci iam compartment list -all`)
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
2. Syntax:

Note

If an object has a `version-id` of `null`, this indicates that there is only one version of the object. To delete this object, omit the `--version-id` argument.

```
oci os object delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--version-id <bucket_version_id>
--object-name <object_name>
```

Example:

```
oci os object delete
--namespace-name examplnamespace \
--bucket-name MyBucket \
--version-id 7f1f537d-ec9c-4706-867a-b1dae355c263 \
--object-name compute.log
```

9.4.6 Recovering a Deleted Object Version

Recovering a deleted object version is as simple as deleting the delete marker that was created when you deleted the latest version of an object. The previous version of the object listed just below the delete marker is recovered and becomes the latest version of the object.

Using the OCI CLI

1. List the objects in the bucket. See [Section 9.4.4, “Viewing Object Versions and Details”](#). In the output, locate the object version that has `"is-delete-marker": true`.

Use the version-id of that object with the delete command to delete the delete marker.

Note

If an object has a `version-id` of `null`, this indicates that there is only one version of the object. To delete this object marker, omit the `--version-id` argument.

2. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Compartment OCID (`oci iam compartment list -all`)
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Object name (`oci os object list`), see [Section 9.3.2, “Viewing Objects in a Bucket”](#)
 - Version ID (see previous step)
3. Syntax:

```
oci os object delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--object-name <object_name>
--version-id <bucket_version_id>
```

Example:

```
oci os object delete
--namespace-name examplnamespace \
--bucket-name MyBucket
--object-name application.log
--version-id 6ce3eb93-8850-4732-8949-cb6e67b722b0
Are you sure you want to delete this resource? [y/N]: y
```

9.5 Using Pre-Authenticated Requests

9.5.1 Pre-Authenticated Requests

Pre-authenticated requests provide a way to let users access a bucket or an object without having their own credentials, as long as the request creator has permissions to access those objects.

For example, you can create a request that lets an operations support user upload backups to a bucket without owning API keys. Or, you can create a request that lets a business partner update shared data in a bucket without owning API keys.

When you create a pre-authenticated request, a unique URL is generated. Anyone you provide this URL to can access the Object Storage resources identified in the pre-authenticated request, using standard HTTP tools like curl and wget.

Important

Assess the business requirement for and the security ramifications of pre-authenticated access to a bucket or objects.

A pre-authenticated request URL gives anyone who has the URL access to the targets identified in the request. Carefully manage the distribution of the URL.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#).

9.5.2 Listing Pre-Authenticated Requests

Use this procedure to obtain information about pre-authenticated requests, such as obtaining the pre-authenticated requests id that you might need for other commands.

Note

Listing pre-authenticated requests does not display the unique URL provided by the system when you created a pre-authenticated request. The URL is displayed only at the time of creation and cannot be retrieved later.

Using the OCI CLI

- **Listing All the Pre-Authenticated Requests in a Bucket**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
2. Run the command.

Syntax (entered on a single line):

```
oci os preauth-request list
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

Example:

```
oci os preauth-request list \
--namespace-name examplnamespace \
--bucket-name MyBucket
{
  "data": [
    {
      "access-type": "ObjectRead",
      "id": "5299a6f9-55c7-4805-88ca-b270c9a9e94f",
      "name": "PAR_ObjRead",
      "object-name": "compute.log",
      "time-created": "2021-06-10T20:34:01+00:00",
      "time-expires": "2021-07-30T23:55:00+00:00"
    },
    {
      "access-type": "AnyObjectWrite",
      "id": "783cd56b-9df5-4518-aacf-f523deae5102",
      "name": "PAR-all-objectsRW",
      "object-name": null,
      "time-created": "2021-06-10T20:49:11+00:00",
      "time-expires": "2021-07-30T23:54:59+00:00"
    },
    {
      "access-type": "ObjectRead",
      "id": "2ea48624-16ed-4d81-95ca-b23ea750ed3d",
      "name": "PAR-OS-READ",
      "object-name": "backup.log",
      "time-created": "2021-06-10T21:16:47+00:00",
      "time-expires": "2021-07-30T23:55:00+00:00"
    }
  ]
}
```

- **Getting the Details for a Specific Pre-Authenticated Request**

1. Gather the information you need to run the command.

- Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Preauth ID (`oci os preauth-request list`), see [Section 9.5.2, “Listing Pre-Authenticated Requests”](#)
2. Run the command.

Syntax (entered on a single line):

```
oci os preauth-request list
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--par-id <preauth-id>
```

Example:

```
oci os preauth-request get \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--par-id 5299a6f9-55c7-4805-88ca-b270c9a9e94f
{
  "data": {
    "access-type": "ObjectRead",
    "id": "5299a6f9-55c7-4805-88ca-b270c9a9e94f",
    "name": "PAR_ObjRead",
    "object-name": "compute.log",
    "time-created": "2021-06-10T20:34:01+00:00",
    "time-expires": "2021-07-30T23:55:00+00:00"
  }
}
```

9.5.3 Creating a Pre-Authenticated Request for All Objects in a Bucket

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Name for this pre-authenticated request.
 - Access type is one of these items:
 - `AnyObjectRead` permits reads on all objects in the bucket.
 - `AnyObjectWrite` permits writes to all objects in the bucket.
 - `AnyObjectReadWrite` permits reads and writes to all objects in the bucket.

Note

Listing objects in a bucket is denied by default. If the `--access-type` is `AnyObjectRead` or `AnyObjectReadWrite`, you can specify the optional

`--bucket-listing-action ListObjects` parameter when creating the pre-authenticated request that lets users list the objects in the bucket.

- Timestamp is a required argument and must be an [RFC 3339](#) timestamp. For example: `2017-09-01T00:09:51.000+02:00`.

2. Run the command.

Syntax (entered on a single line):

```
oci os preauth-request create
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--name <preauthenticated_request_name>
--access-type <access_value>
--time-expires <timestamp>
```

This example creates a pre-authenticated request that allows reads and writes to all objects in the bucket:

```
oci os preauth-request create \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--name PAR-all-objectsRW \
--access-type AnyObjectWrite \
--time-expires '2021-07-30 23:55'
{
  "data": {
    "access-type": "AnyObjectWrite",
    "access-uri": "/oci/p/KOCRWzqBilJmIsaBbJNlKLWcOxwRLq/n/exemplnamespace/b/MyBucket/o/",
    "id": "783cd56b-9df5-4518-aacf-f523deae5102",
    "name": "PAR-all-objectsRW",
    "object-name": null,
    "time-created": "2021-06-10T20:49:11+00:00",
    "time-expires": "2021-07-30T23:54:59+00:00"
  }
}
```

3. **Important** – Copy the access-uri to durable storage.

The unique `access-uri` provided by the system is the only way to construct a URL that a user can use to access the bucket or object specified as the request target.

The `access-uri` is displayed only at the time of creation and cannot be retrieved later.

4. Construct a URL from the unique `access-uri`.

See [Section 9.5.5, “Constructing the Pre-Authenticated Request URL”](#).

9.5.4 Creating a Pre-Authenticated Request for a Specific Object

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Name for this pre-authenticated request.

- Access type is one of these items:
 - `AnyObjectRead` permits reads on all objects in the bucket.
 - `AnyObjectWrite` permits writes to all objects in the bucket.
 - `AnyObjectReadWrite` permits reads and writes to all objects in the bucket.

Note

Listing objects in a bucket is denied by default. If the `--access-type` is `AnyObjectRead` or `AnyObjectReadWrite`, you can specify the optional `--bucket-listing-action ListObjects` parameter when creating the pre-authenticated request that lets users list the objects in the bucket.

- Timestamp is a required argument and must be an [RFC 3339](#) timestamp. For example: `2017-09-01T00:09:51.000+02:00`.
- Object name, or `null`

2. Syntax (entered on a single line):

```
oci os preauth-request create
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--name <preauthenticated_request_name>
--access-type <access_value>
--time-expires <timestamp>
-on <object_name_or_null>
```

Example:

```
oci os preauth-request create
--namespace-name examplnamespace \
--bucket-name MyBucket
--name PAR-OS-READ
--access-type ObjectRead
--time-expires '2021-07-30 23:55'
-on compute.log

{
  "data": {
    "access-type": "ObjectRead",
    "access-uri": "/oci/p/eWvgyLcDthhvVUNkVaejymgDToILHli/n/exemplnamespace/b/MyBucket/o/compute.log",
    "id": "2ea48624-16ed-4d81-95ca-b23ea750ed3d",
    "name": "PAR-OS-READ",
    "object-name": "compute.log",
    "time-created": "2021-06-10T21:16:47+00:00",
    "time-expires": "2021-07-30T23:55:00+00:00"
  }
}
```

3. **Important** – Copy the access-uri to durable storage.

The unique `access-uri` provided by the system is the only way to construct a URL that a user can use to access the bucket or object specified as the request target.

The `access-uri` is displayed only at the time of creation and cannot be retrieved later.

4. Construct a URL from the unique `access-uri`.

See [Section 9.5.5, “Constructing the Pre-Authenticated Request URL”](#).

9.5.5 Constructing the Pre-Authenticated Request URL

Once you have a unique `access-uri`, you can construct the access URL that enables users to access pre-authenticated objects.

1. Construct the URL using this syntax.

Syntax:

```
https://objectstorage.<pca_fqdn>/oci/<access-uri>
```

where:

- `<pca_fqdn>` is the fully qualified domain name of your appliance.
- `<access-uri>` is the access URI that was obtained from one of these procedures:
 - [Section 9.5.3, “Creating a Pre-Authenticated Request for All Objects in a Bucket”](#)
 - [Section 9.5.4, “Creating a Pre-Authenticated Request for a Specific Object”](#)

Example:

```
https://objectstorage.mypca01.us.example.com/oci/p/MrxLFkKlFkI1NDhvhcZnrjBUAlsoeah/n/mynamespace/b/my-bucket
```

Note – In some cases, you need to omit the slash at the end of the `<access-uri>` string. Refer to the [Oracle Private Cloud Appliance Release Notes](#).

9.5.6 Deleting a Pre-Authenticated Request

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Preauth ID (`oci os preauth-request list`), see [Section 9.5.2, “Listing Pre-Authenticated Requests”](#)
2. Syntax (entered on a single line):

```
oci os preauth-request delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--par-id <preauthenticated_request_id>
```

Example:

```
oci os preauth-request delete \
--namespace-name examplnamespace \
```


9.6 Defining Retention Rules

9.6.1 Retention Rules

Retention rules provide immutable storage options for data written to Object Storage for data governance, regulatory compliance, and legal hold requirements. Retention rules can also protect your data from accidental or malicious writes or deletion. Retention rules can be locked to prevent rule modification and data deletion or modification even by administrators.

Retention rules are configured at the bucket level and are applied to all individual objects in the bucket.

For more conceptual information, refer to the [Object Storage](#) section in the [Oracle Private Cloud Appliance Concepts Guide](#).

9.6.2 Viewing Retention Rules and Details

Using the OCI CLI

- **Listing the Retention Rules for a Bucket**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, "Obtaining the Object Storage Namespace"](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, "Listing Buckets"](#)
2. Run the command.

Syntax:

```
oci os retention-rule list
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

Example:

```
oci os retention-rule list \
--namespace-name examplnamespace \
--bucket-name MyBucket
{
  "data": {
    "items": [
      {
        "display-name": "RegulatoryCompliance",
        "duration": {
          "time-amount": 5,
          "time-unit": "YEARS"
        },
        "etag": "72be3a47de931cd50ad9d93c077def64",
        "id": "72be3a47de931cd50ad9d93c077def64",
        "time-created": "2021-06-10T22:24:21+00:00",
        "time-modified": "2021-06-10T22:24:21+00:00",
        "time-rule-locked": "2021-06-30T17:00:00+00:00"
      },
      {
        "display-name": "TempHold",
        "duration": {
          "time-amount": 30,
          "time-unit": "DAYS"
        },
        "etag": "344a9c205187408699b51c7769dc1bb4",
```



```

    "id": "344a9c205187408699b51c7769dc1bb4",
    "time-created": "2021-06-10T22:17:50+00:00",
    "time-modified": "2021-06-10T22:17:50+00:00",
    "time-rule-locked": null
  },
  {
    "display-name": "LegalHold",
    "duration": null,
    "etag": "bd8d8efb964d1025f4305c86de630a4f",
    "id": "bd8d8efb964d1025f4305c86de630a4f",
    "time-created": "2021-06-10T22:13:37+00:00",
    "time-modified": "2021-06-10T22:13:37+00:00",
    "time-rule-locked": null
  }
]
}
}

```

• Getting Details for a Specific Retention Rule

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Retention rule ID (`oci os retention-rule list`), see [Section 9.6.2, “Viewing Retention Rules and Details”](#)
2. Run the command.

Syntax:

```

oci os retention-rule get
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--retention-rule-id <retention_rule_identifier>

```

Example:

```

oci os retention-rule get \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--retention-rule-id 72be3a47de931cd50ad9d93c077def64

{
  "data": {
    "display-name": "RegulatoryCompliance",
    "duration": {
      "time-amount": 5,
      "time-unit": "YEARS"
    },
    "etag": "72be3a47de931cd50ad9d93c077def64",
    "id": "72be3a47de931cd50ad9d93c077def64",
    "time-created": "2021-06-10T22:24:21+00:00",
    "time-modified": "2021-06-10T22:24:21+00:00",
    "time-rule-locked": "2021-06-30T17:00:00+00:00"
  }
}

```

9.6.3 Creating a Retention Rule

Using the OCI CLI

- **Creating an Indefinite Retention Rule**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Display name: The name you want to apply to this retention rule.
2. Run this command.

Syntax:

```
oci os retention-rule create
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--display-name <name_displayed_for_rule>
```

Example:

```
oci os retention-rule create \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--display-name LegalHold
{
  "data": {
    "display-name": "LegalHold",
    "duration": null,
    "etag": "bd8d8efb964d1025f4305c86de630a4f",
    "id": "bd8d8efb964d1025f4305c86de630a4f",
    "time-created": "2021-06-10T22:13:37+00:00",
    "time-modified": "2021-06-10T22:13:37+00:00",
    "time-rule-locked": null
  }
}
```

- **Creating a time-bound, Unlocked Retention Rule**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Display name: The name you want to apply to this retention rule.
 - Time and unit(days|years). For example, 30 days or 5 years.
2. Run this command.

Syntax:

```
oci os retention-rule create
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
```

```
--display-name <display_name>
--time-amount <time_integer>
--time-unit <days|years>
```

Example:

```
oci os retention-rule create \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--display-name TempHold \
--time-amount 30 \
--time-unit days
{
  "data": {
    "display-name": "TempHold",
    "duration": {
      "time-amount": 30,
      "time-unit": "DAYS"
    },
    "etag": "344a9c205187408699b51c7769dc1bb4",
    "id": "344a9c205187408699b51c7769dc1bb4",
    "time-created": "2021-06-10T22:17:50+00:00",
    "time-modified": "2021-06-10T22:17:50+00:00",
    "time-rule-locked": null
  }
}
```

• Creating a Time-Bound, Locked Retention Rule

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Display name: The name you want to apply to this retention rule.
 - Time and unit (days|years). For example, 30 days or 5 years.
 - Date and time to lock the rule.
2. Run this command.

Syntax:

```
oci os retention-rule create
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--display-name <display_name>
--time-amount <time_integer>
--time-unit <days|years>
--time-rule-locked <date and time>
```

Example:

```
oci os retention-rule create \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--display-name RegulatoryCompliance \
--time-amount 5 \
--time-unit years \
--time-rule-locked "2021-06-30 17:00"
{
```

```

"data": {
  "display-name": "RegulatoryCompliance",
  "duration": {
    "time-amount": 5,
    "time-unit": "YEARS"
  },
  "etag": "72be3a47de931cd50ad9d93c077def64",
  "id": "72be3a47de931cd50ad9d93c077def64",
  "time-created": "2021-06-10T22:24:21+00:00",
  "time-modified": "2021-06-10T22:24:21+00:00",
  "time-rule-locked": "2021-06-30T17:00:00+00:00"
}
    
```

9.6.4 Modifying a Retention Rule

Using the OCI CLI

- **Updating a Retention Rule**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Retention rule ID (`oci os retention-rule list`), see [Section 9.6.2, “Viewing Retention Rules and Details”](#)
2. Run this command.

Syntax:

```

oci os retention-rule update
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--retention-rule-id <retention_rule_id>
    
```

Followed by the retention rule item that you plan to change:

```

--time-amount <time_integer>
--time-unit <days/years>
    
```

Example:

```

oci os retention-rule update \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--retention-rule-id 344a9c205187408699b51c7769dc1bb4 \
--time-amount 60 \
--time-unit days
{
  "data": {
    "display-name": "TempHold",
    "duration": {
      "time-amount": 60,
      "time-unit": "DAYS"
    },
    "etag": "344a9c205187408699b51c7769dc1bb4",
    "id": "344a9c205187408699b51c7769dc1bb4",
    "time-created": "2021-06-10T22:17:50+00:00",
    "time-modified": "2021-06-10T22:45:16+00:00",
    "time-rule-locked": null
  }
}
    
```

```
}
}
```

- **Removing a Retention Rule Lock During the Delay Period**

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Retention rule ID (`oci os retention-rule list`), see [Section 9.6.2, “Viewing Retention Rules and Details”](#)
2. Run this command.

Syntax:

```
oci os retention-rule update
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--retention-rule-id <retention_rule_id>
--time-rule-locked ""
```

Example:

```
oci os retention-rule update
--namespace-name examplnamespace \
--bucket-name MyBucket \
--retention-rule-id bla6c84c-57c4-416c-b006-f864b0904c9e
--time-rule-locked ""
{
  "data": {
    "display-name": "RegulatoryCompliance",
    "duration": {
      "time-amount": 6,
      "time-unit": "YEARS"
    },
  },
  "etag": "5b4fa526-faec-47d4-9162-4acdf1813ee0",
  "id": "bla6c84c-57c4-416c-b006-f864b0904c9e",
  "time-created": "2020-03-25T15:11:44.423000+00:00",
  "time-modified": "2020-03-25T22:02:43.745000+00:00",
  "time-rule-locked": null
},
"etag": "5b4fa526-faec-47d4-9162-4acdf1813ee0"
}
```

9.6.5 Deleting a Retention Rule

Using the OCI CLI

1. Gather the information you need to run the command.
 - Namespace (see [Section 9.2.2, “Obtaining the Object Storage Namespace”](#))
 - Bucket name (`oci os bucket list`), see [Section 9.2.3, “Listing Buckets”](#)
 - Retention rule ID (`oci os retention-rule list`), see [Section 9.6.2, “Viewing Retention Rules and Details”](#)
2. Syntax:

Deleting a Retention Rule

```
oci os retention-rule delete
--namespace-name <object_storage_namespace>
--bucket-name <bucket_name>
--retention-rule-id <retention_rule_identifier>
```

Example:

```
oci os retention-rule delete \
--namespace-name examplnamespace \
--bucket-name MyBucket \
--retention-rule-id 344a9c205187408699b51c7769dc1bb4
Are you sure you want to delete this resource? [y/N]: y
```