

Oracle® Cloud

Exadata Database Service on Dedicated Infrastructure



G39062-29
January 2026



Oracle Cloud Exadata Database Service on Dedicated Infrastructure,

G39062-29

Copyright © 2022, 2026, Oracle and/or its affiliates.

Primary Authors: James Spiller, Neil Hebert, Nirmal Kumar

Contributors: Bryce Cracco, Sanjay Narvekar, Pravin Jha , Aneesh Khandelwal, Alan Williams, Bethany Lapaglia, Chandrashekhar Garud, Dileep Thiagarajan, Guruprasad Hegde, Jai Krishnani, Jeffery Wright, Vineet Kakani

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

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

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

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

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

1 Oracle Exadata Database Service on Dedicated Infrastructure Overview

Oracle Exadata Database Service on Dedicated Infrastructure Description	1
About Exadata Cloud Infrastructure	2
Roles and Responsibilities for Exadata Cloud Infrastructure Software Maintenance	3
Licensing Considerations for Oracle Exadata Database Service on Dedicated Infrastructure	4
Supported Database Edition and Versions for Exadata Cloud Infrastructure	4
Subscription Types	5
Metering Frequency and Per-Second Billing	5
Technical Architecture of Exadata Cloud Infrastructure Systems	5
Scaling Options	5
Scaling CPU cores within an Exadata Cloud Infrastructure instance	6
Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations	7
System and Shape Configuration Options	7
Exadata Shape Configuration	8
Support Requirements	12
Exadata Cloud Management Interfaces	13
Introduction to Exadata Cloud Management Interfaces	13
OCI Control Plane Interfaces for Exadata Cloud Infrastructure	13
Local VM Command-Line Interfaces	15
Exadata Database Service on Dedicated Infrastructure Technical Architecture	16

2 What's New in Oracle Exadata Database Service on Dedicated Infrastructure

Support for Private Service Access (PSA)	4
Enhanced Out-of-Place Restore Options for ExaDB-D	4
Support for Concurrent Data Guard Operations in Multiple Standby Environments	5
Enhanced Data Guard Health Insights	5
AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS	5
Convert a Physical Standby Database to a Snapshot Standby Database and vice versa	6
Exadata Exascale with Oracle Exadata Database Service on Dedicated Infrastructure	6
Exadata Database Service for Developers	8

Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud	9
Cross-Service Data Guard Between ExaDB-D and ExaDB-XS	9
Enhancements to Serial Console Functionality	10
Change in Backup Destination for New Tenancies in Select OCI Regions (Effective August 06, 2025)	10
Exadata System Software 25.1.0.0.0	11
Expanded Tagging Support Across Database Workflows	11
Amazon S3 Support for Oracle-Managed Backups on Exadata Database Service on Oracle Database@AWS	12
Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure	12
Oracle Data Guard Setup with Precheck Validation	12
Schedule VM Cloud Automation Updates	13
Classify and Control Access to Critical OCI REST APIs	13
Oracle Key Vault (OKV) Integration with ExaDB-D to Manage Transparent Data Encryption (TDE) Keys	14
Remove Provisioned Storage Servers	14
Enhancements to support concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations	14
Dual Stack (IPv4 and IPv6) Network Support	16
Long-Term Retention Backup (LTR)	17
Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution	18
Granular Permissions for VM Cluster Update Operations	18
Multiple Standby Databases	19
X11M System Support	20
Microsoft Entra ID (MS-EI) Integration with Oracle Exadata Database Service on Dedicated Infrastructure	21
Delegate Access Control for ExaDB-D	21
Different RUs for Primary and Standby DB Homes in Data Guard Associations, Switchover, and Failover Operations	22
Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing	22
Assign New Key Versions to KMS-Based Container Databases (CDBs) and Pluggable Databases (PDBs)	23
Enhancement to Backup and Restore from a Standby Database in a Data Guard Environment	24
VM Cluster on a Single VM	24
Multicloud Oracle Database Backup Support	25
Restore a Backup to Create a Database Across Regions	25
Cost and Usage Attribution for Pluggable Databases (PDBs)	26
Use the Same Custom Software Image Across OCI Regions	26
Ability to Increase the Size of Guest VM Local File Systems	27
Create and Use Custom Software Images	27
Manage Serial Console Access to Oracle Exadata Database Service on Dedicated Infrastructure Systems	27

Oracle AI Database 26ai on Exadata Database Service on Dedicated Infrastructure	28
Enable Unified Auditing While Creating a Database Home	29
Provision a VM Cluster with Either an OL7 or OL8-Based Image	29
Enhancement to the OCI Console to Remove Database and Storage Servers	30
Enable Data Guard Across Different VCNs or Compartments in the Same OCI Region	30
Enhancement to Pluggable Database (PDB) Management	30
Manage Administrator (SYS User) and TDE Wallet Passwords	31
Backup and Restore from a Standby Database with OCI Object Storage in a Data Guard Environment	31
Cancel a Running Full or Incremental Backup	33
Autonomous Recovery Service as the Default Backup Destination	33
Exadata Fleet Update	33
Update Guest VM (domU) Operating System to Oracle Linux 8	34
Use a Backup to Create a Database Across Availability Domains within the Same Region	34
Interim Software Updates	35
Enhanced Controls to Configure Automatic Full (L0) and Incremental (L1) Backups	35
Configure Oracle Database Autonomous Recovery Service as a Backup Destination	36
Application VIP Support	37
Monthly ExaDB-D Infrastructure Security Maintenance	38
Identity and Access Management (IAM) Integration with Oracle Exadata Database Service on Dedicated Infrastructure	38
Exadata Cloud Infrastructure: Private DNS	39
Enhanced Infrastructure Maintenance Controls	39
Database Management Support for Pluggable Databases in Oracle Exadata Database Service on Dedicated Infrastructure	40
Microsoft Azure Active Directory Integration with Oracle Cloud Infrastructure Databases	40
Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting	40
VM Cluster and Database Health and Performance Metrics in the OCI Console	41
Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Dedicated Infrastructure	42
Automatic Diagnostic Collection	42
Exadata Database on Dedicated Infrastructure: Key Management Service for Cross Region Data Guard	43
Concurrently Create or Terminate Oracle Databases in a VM Cluster	44
VM Guest Exadata OS Image Major Version Update	44
Database Service Events capability for Exadata Database	44
Exadata Database on Dedicated Infrastructure: 'Create database from backup' now available for databases using customer-managed encryption	45
Support for DB Home Minor Version Selection (N-3)	46
Oracle Cloud Infrastructure Operations Insights Support for Oracle Cloud Databases	46
Specify the Same SID for Primary and Standby Databases in Data Guard Association	47
Exadata Cloud Infrastructure: Pluggable database lifecycle support	47

Exadata Cloud Infrastructure: Set DB_UNIQUE_NAME and Oracle SID prefix during database creation	47
Elastic Expansion	47
Oracle Database: Encryption key options updated for Exadata Cloud Infrastructure databases	48
Performance Hub Exadata Tab	48
Exadata Cloud Infrastructure: custom SCAN listener port for VM cluster	49
Performance Hub & metrics available for databases running in Exadata Cloud Infrastructure	49
Maintenance advisory contacts for Exadata infrastructure	49
Data Guard protection mode enhancement for Exadata Cloud Infrastructure instances	49
Exadata Cloud Infrastructure: Non-rolling infrastructure patching option now available	49
Customer-managed encryption keys available with Oracle Data Guard-enabled databases in Exadata Cloud Infrastructure	50
ExaDB-D OS/DomU Patching Project	50
Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure	50
Exadata Cloud Infrastructure: Oracle Database 19c upgrade feature available	51
Create custom database software images for Exadata Cloud Infrastructure instances	51
Exadata Cloud Infrastructure: grid infrastructure upgrade for cloud VM clusters	51
Exadata Cloud Infrastructure: the flexible X8M shape now available	51
Exadata Cloud Infrastructure: use an existing Database Home when setting up a Data Guard standby database	52

3 Preparing for Exadata Cloud Infrastructure

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Dedicated Infrastructure	1
Required IAM Policy for Exadata Cloud Infrastructure	2
Network Setup for Exadata Cloud Infrastructure Instances	3
VCN and Subnets	3
Option 1: Public Client Subnet with Internet Gateway	4
Option 2: Private Subnets	6
Requirements for IP Address Space	7
Configuring a Static Route for Accessing the Object Store	8
Setting Up DNS for an Exadata Cloud Infrastructure Instance	8
DNS: Short Names for the VCN, Subnets, and Exadata Cloud Infrastructure instance	8
DNS: Between On-Premises Network and VCN	9
Configure Private DNS	9
Node Access to Object Storage: Static Route	10
Object Storage IP allocations	11
To configure a static route for Object Storage access	11
Service Gateway for the VCN	12
Option 1: Service Gateway Access to OCI Service for Backup Subnet	12
Option 2: Service Gateway Access to OCI Service for Both the Client and Backup Subnets	13

Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure	15
Rules Required for Both the Client Network and Backup Network	18
Rules Required Specifically for the Client Network	20
Rule Required Specifically for the Backup Network	22
Rules Required for Events Service	22
Rules Required for Monitoring Service	23
Ways to Implement the Security Rules	23
If you use network security groups	23
If you use security lists	24
Network Requirements for Oracle Database Autonomous Recovery Service	25
Create a Service Gateway to Object Storage	25
Creating Protection Policies	26
Storage Configuration Requirements for Oracle Exadata Database Service on Dedicated Infrastructure	26
Configuring Exadata ASM Storage	27
Impact of Configuration Settings on ASM Storage	27
Configuring Exadata Exascale Storage	28
About Exascale Storage Configuration for Oracle Exadata Database Service on Dedicated Infrastructure	28
Using the Console to Configure Exascale Storage on ExaDB-D Infrastructure	29
Using the Console to Create an Exascale Storage Vault	30
Using the Console to Scale an Exascale Storage Vault	31
Using the Console to Move an Exascale Storage Vault to Another Compartment	31
Using the Console to Delete an Exascale Storage Vault	32
Cross-Service Data Guard Between ExaDB-D and ExaDB-XS	32

4 Getting Started with Exadata Cloud Infrastructure Deployment

Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources	1
Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure	7
The Exadata Cloud Infrastructure Resource Model	7
The Cloud Exadata Infrastructure Resource	7
The Cloud VM Cluster Resource	8
Additional Exadata Cloud Infrastructure Instance Resources	8
The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes	8
The New Exadata Cloud Infrastructure Resource Model	10
Creating an Exadata Cloud Infrastructure Instance	10
Resources to Be Created	11
Prerequisites for Creating an Cloud Exadata Infrastructure Instance	11
Default Options for the Initial Database	12
Using the Console to Create Infrastructure Resources	12
To create a Cloud Exadata infrastructure resource	13

To create an ASM cloud VM cluster	17
VM Cloud Automation Software Update Management	26
Configuring Network Resources for Recovery Service	28
Autonomous Recovery Service Checklist	34
Connecting to an Exadata Cloud Infrastructure Instance	36
Prerequisites	36
SCAN Listener Port Setting	37
Connecting to a Virtual Machine with SSH	37
Connecting from a Unix-Style System	38
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	38
Accessing a Database After You Connect to the Virtual Machine	39
Using Oracle Net Services to Connect to a Database	40
Prerequisites for Connecting to a Database with Oracle Net Services	41
Connecting to a Database with SQL Developer	41
Connecting to a Database Using SCAN	43
Connecting to a Database Using a Node Listener	45
Using the Console to Create API Access Control	46
Best Practices for Exadata Cloud Infrastructure Instances	46
Moving to Oracle Cloud Using Zero Downtime Migration	46

5 How-to Guides

Manage Database Security with Oracle Data Safe	2
About Oracle Data Safe	2
Get Started	3
Using Oracle Data Safe	4
Connecting to an Exadata Cloud Infrastructure Instance	5
Prerequisites	6
About Connecting to a Compute Node with SSH	6
Connecting from a Unix-Style System	7
Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY	7
To access a database after you connect to the compute node	8
Connect to the Exadata Cloud Infrastructure Service	9
Connecting to a Database with SQL Developer	9
Connecting to a Database with Oracle Net Services	9
Manage Exadata Cloud Infrastructure	13
Using the Console to Provision Exadata Cloud Infrastructure	13
Lifecycle Management Operations	14
Network Management Operations	18
Management Tasks for the Oracle Cloud Infrastructure Platform	18
Oracle Database License Management Tasks	21
Scaling Resources within an Exadata Infrastructure Instance	22

Using the API to Create Infrastructure Components	28
Using the API to Manage Exadata Cloud Infrastructure Instance	29
Configure Oracle-Managed Infrastructure Maintenance	30
About Oracle-managed Exadata Cloud Infrastructure Maintenance	32
Overview of the Quarterly Infrastructure Maintenance Process	32
Overview of Monthly Security Maintenance	34
Understanding Monthly and Quarterly Maintenance in the Same Month	35
Maintenance Scheduling Policy	37
View the List of Maintenance Scheduling Policy	38
Create a Maintenance Scheduling Policy	38
View Details of a Maintenance Scheduling Policy	39
Edit a Maintenance Scheduling Policy	39
Edit a Maintenance Scheduling Policy that's in Needs Attention State	40
View a Maintenance Window of a Maintenance Scheduling Policy	40
Edit a Maintenance Window of a Maintenance Scheduling Policy	40
Add Additional Maintenance Windows to a Maintenance Scheduling Policy	41
View Resources Associated with a Policy	41
Delete a Maintenance Window of a Maintenance Scheduling Policy	41
Move a Maintenance Scheduling Policy to a Different Compartment	42
Add Tags to a Maintenance Scheduling Policy	42
Delete a Maintenance Scheduling Policy	42
Using the Console to Configure Oracle-Managed Infrastructure Updates	43
To view or edit quarterly maintenance preferences	44
Manage Quarterly Maintenance Plan using Scheduling Policy	47
To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure	52
View and Edit Maintenance While Maintenance is In Progress	55
View and Edit Maintenance While Maintenance is Waiting for Custom Action	56
To view or edit a scheduled security maintenance	56
To view the maintenance history of an Exadata Cloud Infrastructure resource	57
To set the node patching order for a scheduled infrastructure maintenance run	58
Manage Quarterly Maintenance Run Created From Scheduling Plan	58
View Maintenance Windows Associated with a Maintenance Run	58
Edit Maintenance Window Associated with a Maintenance Run	59
View Maintenance Actions Associated with a Maintenance Run	60
Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run	61
View and Edit Maintenance While Maintenance is In Progress	63
View and Edit Maintenance While Maintenance is Waiting for Custom Action	64
Cancel a Maintenance Run While In Progress	64
View Maintenance Activity in a Compartment	66
View the Details of a Maintenance Activity	66

Review and Respond to Unplanned Maintenance Activity	66
Monitor Infrastructure Maintenance Using Lifecycle State Information	67
Receive Notifications about Your Infrastructure Maintenance Updates	68
Managing Infrastructure Maintenance Contacts	68
To manage maintenance contacts in an Exadata Cloud Infrastructure	68
Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls	69
Manage VM Clusters	70
Introduction to Scale Up or Scale Down Operations	71
Scale VM Resources in Multi VM Enabled Infrastructure	71
Resizing Memory and Large Pages	73
Calculating the ASM Storage	73
Estimating How Much Local Storage You Can Provision on Your VMs	75
Scaling Local Storage	76
Overview of VM Cluster Node Subsetting	78
Add a VM to a VM Cluster	79
Terminate or Remove a VM from a VM Cluster	80
Overview of Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR)	80
Private Service Access for OCI Service Access from VCN	81
About Application VIP	82
Using the Console to Manage VM Clusters on Exadata Cloud Infrastructure	84
To create an ASM cloud VM cluster	85
To create an Exascale cloud VM cluster	94
To add security attributes to an Exadata VM Cluster	103
To edit a security attribute	103
To remove a security attribute	103
To add database server or storage server capacity to a cloud VM cluster	104
To Enable, Partially Enable, or Disable Diagnostics Collection	105
To Update the License Type on a VM Cluster	106
To add SSH keys to a VM cluster	106
To Add SSH Keys After Creating a VM Cluster	107
To Stop, Start, or Reboot a VM Cluster Virtual Machine	107
To Check the Status of a VM Cluster Virtual Machine	107
To Move a VM Cluster to Another Compartment	108
To change the VM cluster display name	109
To Terminate a VM Cluster	109
To view details about private DNS configuration	110
To Attach a Virtual IP Address	110
To Detach a Virtual IP Address	111
Migrate from a Single Stack (IPv4) Exadata VM Cluster to a Dual Stack (IPv4/IPv6) Exadata VM Cluster with Data Guard Synchronization	111
Overview of Automatic Diagnostic Collection	111
Incident Logs and Trace Files	112

Health Metrics	117
Using the API to Manage Exadata Cloud Infrastructure Instance	122
Troubleshooting Virtual Machines Using Console Connections	123
Required IAM Policies	124
Prerequisites	124
Create the Virtual Machine Serial Console Connection	127
Make an SSH Connection to the Serial Console	128
Using Cloud Shell to Connect to the Serial Console	131
Displaying the Console History for a Virtual Machine	132
Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems	135
Exiting the Virtual Machine Serial Console Connection	137
Manage Software Images	138
Using Software Images in Oracle Cloud Infrastructure	138
Creation and Storage of Software Images	138
Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home	139
Using a Software Image with an Exadata Cloud Infrastructure Instance	140
Using the Console for Software Images	140
To view the list of software images	140
To create a database software image	141
To create a Grid Infrastructure software image	141
To create a database software image from a Database Home	142
To view the image details of a software image	142
To move a software image to a different compartment	143
To update database software using custom database software image	143
To update Grid Infrastructure software using custom Grid Infrastructure software image	144
To delete a software image	144
Using the API to manage database software images	144
Create Oracle Database Homes on an Exadata Cloud Infrastructure System	145
About Creating Oracle Database Homes on an Exadata Cloud Infrastructure System	145
To create a new Database Home in an existing Exadata Cloud Infrastructure instance	146
To create a database software image from a Database Home	147
Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure	148
Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance	148
Manage Database Home Using the Console	149
To view information about a Database Home	149
To delete a database home	149
To manage tags for your Database Home	150
To move a database to another Database Home	150
Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure	151
Manage Databases on Exadata Cloud Infrastructure	151

Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure	152
Prerequisites for Oracle Database Autonomous Recovery Service Cross Region Restore (Same Tenancy)	152
Prerequisites for Oracle Database, Object Storage Cross Region Restore (Same Tenancy)	154
Oracle Database Releases Supported by Oracle Exadata Database Service on Dedicated Infrastructure	155
Provisioning and Managing Exadata Databases	156
Database Memory Initialization Parameters	157
Customer-Managed Keys in Exadata Cloud Infrastructure	157
Using the Console to Manage Databases on Oracle Exadata Database Service on Dedicated Infrastructure	162
Known Issues in Exadata Cloud Infrastructure	179
Using the API to manage Databases	180
Create and Manage Exadata Pluggable Databases	180
Limitations for Pluggable Database Management	182
Creating an Exadata Pluggable Database	183
Managing an Exadata Pluggable Database	187
Cloning an Exadata Pluggable Database	193
Concurrently Create or Delete Pluggable Databases (PDBs)	200
Restoring an Exadata Pluggable Database	201
To perform an in-place restore of a pluggable database (PDB)	201
To perform an out-of-place restore of a pluggable database (PDB)	201
Cost and Usage Attribution for Pluggable Databases (PDBs)	205
Generate Attributed Cost Analysis Report for Pluggable Databases	206
Changing the Database Passwords	207
To Change the SYS Password for an Exadata Cloud Infrastructure Database	207
To Change Database Passwords in a Data Guard Environment	207
To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database	207
Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure	208
Oracle Recommended Options to Perform Backup and Recovery Operations	209
Managing Exadata Database Backups	210
Managed Backup Types and Usage Information	211
Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console	212
Long-Term Retention Backup with Recovery Service	213
Default Backup Channel Allocation	214
Prerequisites for Backups on Exadata Cloud Infrastructure	214
Using the Console to Manage Backups	215
To configure automatic backups for a database	216
To create an on-demand backup of a database	219
To view backup status	220

To cancel a backup	220
To delete full backups from Object Storage	221
To delete standalone backups from Object Storage	221
To change the retention period of an LTR backup with Recovery Service	221
To designate Autonomous Recovery Service as a Backup Destination for an Existing Database	222
Recovering an Exadata Database from Backup Destination	223
Using the Console to restore a database	223
Managing Exadata Database Backups by Using dbaascli	225
Default Backup Configuration	225
To get the default backup configuration for a freshly provisioned database	225
To create a backup configuration file	226
To create an on-demand backup	232
To remove the backup configuration	233
To delete a backup in Object Storage	234
Using the API to Manage Backup and Recovery	234
Using the API to manage backups	234
Alternative Backup Methods	234
Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management	235
Recovering a Database Using Oracle Recovery Manager (RMAN)	236
Patch and Update an Exadata Cloud Infrastructure System	236
User-Managed Maintenance Updates	236
Patching and Updating an Exadata Cloud Infrastructure System	237
Patching and Updating VM Cluster's GI and Database Homes	238
Updating an Exadata Cloud VM Cluster Operating System	252
Upgrading Exadata Grid Infrastructure	254
Upgrading Exadata Databases	257
Database Health Checks During Infrastructure and Guest VM OS Update Operations	263
Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade	264
Patching and Updating an Exadata Cloud Infrastructure System Manually	265
Patching Oracle Database and Oracle Grid Infrastructure Software Manually	265
Updating the Exadata Cloud VM Cluster OS Manually	266
Updating Tooling on an Exadata Cloud Infrastructure Instance	274
Interim Software Updates	274
Create an Interim Software Update	275
Download an Interim Software Update	275
Delete an Interim Software Update	276
Move an Interim Software Update Resource to Another Compartment	276
Using the API to Manage Interim Software Updates	276
Use Oracle Data Guard with Exadata Cloud Infrastructure	277
About Using Oracle Data Guard with Exadata Cloud Infrastructure	278

Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure	278
Network Requirements for Data Guard	279
Password Requirements	281
Known Issues for Exadata Cloud Infrastructure and Data Guard	281
Adding a Node to a VM Cluster	281
Removing a Node from a VM Cluster	282
Working with Data Guard	282
Switchover	282
Failover	283
Reinstate	283
Support for Concurrent Data Guard Operations in Multiple Standby Environments	283
Enhanced Data Guard Health Status Reporting	284
Using the Console to Manage an Oracle Data Guard Group	285
To Enable Data Guard on an Exadata Cloud Infrastructure System	286
To view Data Guard Group details of databases in a Cloud VM Cluster	291
To enable automatic backups on a standby database	291
To perform a database switchover	293
To edit the Oracle Data Guard Group details	294
To perform a database failover	295
To reinstate a database	296
To terminate a Data Guard Group on an Exadata Cloud Infrastructure instance	297
Using the API to manage Data Guard associations	297
Using the API to manage Data Guard Group	298
Configure Oracle Database Features for Exadata Cloud Infrastructure	299
Using Oracle Multitenant on an Exadata Cloud Infrastructure Instance	299
To determine if you need to create and activate an encryption key for the PDB	300
To create and activate the master encryption key in a PDB	300
To export and import a master encryption key	302
Managing Tablespace Encryption	302
Managing Huge Pages	304
Managing Exadata Cloud Infrastructure I/O Resource Management (IORM)	305
About IORM	305
Using the Console to Manage IORM	306
To enable IORM on your Exadata cloud VM cluster	306
To modify the IORM configuration on your cloud VM cluster	307
Using the API to manage the I/O resources of an Exadata cloud VM cluster	308
Managing Encryption Keys on External Devices	308
Customer-Managed Keys in Oracle Exadata Database Service on Dedicated Infrastructure	309
About Oracle Key Vault	310
Overview of Key Store	310

Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Dedicated Infrastructure	310
Tagging Resources	311
Moving Resources to a Different Compartment	311
Setting Up Your Oracle Exadata Database Service on Dedicated Infrastructure to Work With Oracle Key Vault	311
Step 1: Create a Vault in OKV Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password	313
Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault	313
Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store	314
Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service	314
Step 5: Create Key Store	315
Managing Your Key Store	315
View Key Store Details	316
Edit Key Store Details	316
Move a Key Store to Another Compartment	316
Delete a Key Store	317
View Key Store Associated Container Database Details	317
Using the API to Manage Key Store	317
Administer Transparent Data Encryption (TDE) Keys	318
How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV)	320
Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server	320
Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server	330
How to Upgrade Oracle Key Vault (OKV) Home in Oracle Exadata Database Service on Dedicated Infrastructure	340
Migrate to Exadata Cloud Infrastructure	345
Moving to Oracle Cloud Using Zero Downtime Migration	345
Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Dedicated Infrastructure	345
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database	346
About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database	346
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication	347
Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication	348
Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database	350
Prerequisites for IAM Authentication on Oracle Database	350
Disable External Authentication Scheme	351

Configure TLS to Use IAM Tokens	351
Enabling the Database and Clients for IAM Integration	352
Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure	352
About Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure	353
Configuring the Oracle Database for Microsoft Entra ID (MS-EI) Integration	354
Prerequisites for Microsoft Entra ID (MS-EI) Authentication	354
Networking Prerequisites for Microsoft Entra ID (MS-EI) Authentication	355
Configure TLS to Use Microsoft Entra ID (MS-EI) Tokens	356
Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure	357
Prerequisites	358
Network Requirements for Creating an Identity Connector and KMS Resources	358
Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure	360
Create an Identity Connector from the OCI Console	361
View the details of an Identity Connector	362
Enable or Disable the Azure Key Management	363
Create Azure Key Vault (Managed HSM, Premium, and Standard) and Assign Required Permissions	363
Register Azure Key Vaults in the OCI console	364
Create a Database and Use Azure Key Vault as the Key Management Solution	365
Change the Key Management from Oracle Wallet to Azure Key Vault	367
Rotate the Keys Managed by Azure Key Vault for a Container Database	367
Rotate the Keys Managed by Azure Key Vault for a Pluggable Database	368
Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure	368
Updating the Multicloud PKCS#11 Driver	369
Release Notes	372
Release 0.2 (250910.0501)	372
Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud	372
Prerequisites	372
Using the Console to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud	373
To create an ASM cloud VM cluster	374
Verify the Default Identity Connector Attached to the VM Cluster	383
Create a Key Ring in Google Cloud Console	383
Create a Key in Google Cloud Console	384
Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI)	385
Register GCP Key Ring in Oracle Cloud Infrastructure (OCI)	386
Enable or Disable Google Cloud Key Management	387

Create a Database and Use GCP Customer-Managed Encryption Key (CMEK) as the Key Management Solution	387
Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key (CMEK)	388
Rotate the GCP Customer Managed Encryption Key of a Container Database (CDB)	388
Rotate the GCP Customer Managed Encryption Key of a Pluggable Database (PDB)	389
Using the API to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud	389
Updating the Multicloud PKCS#11 Driver	390
Release Notes	392
Release 2.185 (251201.0551)	393
AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS	393
Prerequisites	393
Configure OCI Identity Domain	393
Create or Use an Existing ODB Network	394
Create an Exadata VM Cluster	394
Configure Identity Provider	394
Associate an IAM Role to an Exadata VM Cluster	395
Create a Key	396
Register AWS KMS Key	398
OCI IAM Policy Requirements	398
Using the Console to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS	399
Enable or Disable AWS Key Management	399
Create a Database and Use AWS KMS as the Key Management Solution	399
Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS	400
Rotate the AWS KMS Key of a Container Database (CDB)	400
Rotate the AWS KMS Key of a Pluggable Database (PDB)	401
(Optional) Disable a CDB or PDB Key	401
Using the API to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS	402
OracleDbAwsIdentityConnector	402
OracleDbAwsKey	403
Updating the Multicloud PKCS#11 Driver	403
Release Notes	406
Release 2.185 (251201.0551)	406
Database Multicloud Integration for Oracle Database Cloud Services	406
Using the Console to Manage Database Multicloud Integration for Oracle Database Cloud Services	406
Integrate Oracle Exadata with Azure Cloud services	406
Create an Identity Connector from the OCI Console	407
View the list of Identity Connector Resources	408

View the details of an Identity Connector Resource	408
Create a Storage Container Resource	408
Discover Azure Storage Containers	409
View the list of Storage Container Resources	409
View the details of a Storage Container Resource	409
Create a Storage Mount Resource	409
View the list of Storage Mount Resources	410
View the details of a Storage Mount Resource	410
Using the API to Manage Multi-Cloud Data Integration for Oracle Database Cloud Services	410
oracle-db-azure-connectors	410
oracle-db-azure-blob-containers	411
oracle-db-azure-blob-mounts	411
oracle-db-mci-work-requests	412
multi-cloud-resource-discoveries	412
oracle-db-azure-vaults	413
oracle-db-azure-keys	413
oracle-db-azure-vault-associations	413

6 Reference Guides for Exadata Cloud Infrastructure

Using the dbaascli Utility on Exadata Cloud Infrastructure	1
About Using the dbaascli Utility on Exadata Cloud Infrastructure	2
Creating Databases Using dbaascli	3
Listing Available Software Images and Versions for Database and Grid Infrastructure	3
Creating Oracle Database Home	4
Creating Oracle Database In the Specified Oracle Database Home	5
Changing the Database Passwords	8
To Change the SYS Password for an Exadata Cloud Infrastructure Database	8
To Change Database Passwords in a Data Guard Environment	9
To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database	9
Managing Exadata Cloud Infrastructure Software Images Using the Dbaascli Utility	9
Listing Available Software Images and Versions for Database and Grid Infrastructure	10
To download a software image	11
Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli	12
Patching Databases using dbaascli	12
Patching Oracle Grid Infrastructure	14
Listing Available Software Images and Versions for Database and Grid Infrastructure	15
Performing a Precheck Before Patching Databases and Grid Infrastructure	17
Resuming or Rolling Back a Patching Operation	19
Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli	21
Collecting Tooling Log Data Examples	21

Performing a Health Check Examples	25
Updating Cloud Tooling Using dbaascli	26
Creating a Duplicate Database	27
Using dbaascli to Duplicate a Cloud Database	27
Considerations When Using OCI Vault for the Key Management	28
Duplicate an On-Premises Database	29
Release Notes	30
Release 25.4.1.0.0 (251107)	31
Release 25.3.1.0.0 (250826)	31
Release 25.2.1.0.0 (250522)	32
Release 25.1.2.0.0 (250325)	32
Release 25.1.1.0.0 (250107)	32
Release 24.4.1.0.0 (241104)	32
Release 24.3.2.0.0 (240828)	32
Release 24.3.1.0.0 (240711)	32
Release 24.2.1.0.0 (240530)	32
Release 24.1.2.0.0 (240306)	32
Release 24.1.1.0.0 (231219)	32
Release 23.4.1.0.0 (231102)	33
Release 23.3.2.0.0 (230921)	33
Release 23.3.1.0.0 (230712)	33
Release 23.2.1.0.0 (230503)	34
Release 23.1.2.0.0 (230305)	34
Release 23.1.1.0.1 (230113)	34
Release 22.4.1.0.1 (221122)	35
Release 22.3.1.1.0 (221003)	35
Release 22.3.1.0.1 (220721)	36
Release 22.2.1.1.0 (220623)	36
Release 22.2.1.1.0 (220609)	37
Release 22.2.1.0.1 (220423)	37
Release 22.1.1.2.0 (220405)	38
Release 22.1.1.1.0 (220317)	38
Release 22.1.1.0.1 (220223)	38
Release 21.4.1.1.0 (220209)	39
Release 21.4.1.1.0	39
Release 21.3.1.2.0	40
Release 21.3.1.1.0	40
Release 21.3.1.0.1	41
Release 21.2.1.x.x	41
dbaascli Command Reference	42
dbaascli admin updateAHF	48
dbaascli admin updateStack	48

dbaascli cswlib deleteLocal	49
dbaascli cswlib download	49
dbaascli cswlib listLocal	50
dbaascli cswlib showImages	51
dbaascli database addInstance	51
dbaascli database backup	52
dbaascli database bounce	54
dbaascli database changepassword	54
dbaascli database convertToPDB	55
dbaascli database create	56
dbaascli database delete	58
dbaascli database deleteInstance	58
dbaascli database duplicate	59
dbaascli database getDetails	61
dbaascli database getPDBs	61
dbaascli database modifyParameters	61
dbaascli database move	62
dbaascli database recover	64
dbaascli database runDatapatch	65
dbaascli database createTemplate	65
dbaascli database start	66
dbaascli database status	67
dbaascli database stop	67
dbaascli database upgrade	68
dbaascli dataguard prepareStandbyBlob	69
dbaascli dataguard updateDGConfigAttributes	69
dbaascli dataguard failover	69
dbaascli dataguard reinstate	70
dbaascli dataguard switchover	72
dbaascli dataguard prepareForStandby	73
dbaascli dataguard configureStandby	75
dbaascli dataguard registerStandby	77
dbaascli dataguard deregisterStandby	79
dbaascli dataguard configureAWR	80
dbaascli dataguard updateConfiguration	81
dbaascli dbhome create	83
dbaascli dbHome delete	83
dbaascli dbhome getDatabases	84
dbaascli dbHome getDetails	84
dbaascli dbHome patch	85
dbaascli dbimage purge	86
dbaascli diag collect	87

dbaascli diag healthCheck	88
dbaascli gridHome create	88
dbaascli grid configureTCPS	89
dbaascli grid patch	90
dbaascli grid removeTCPSCert	91
dbaascli grid rotateTCPSCert	92
dbaascli grid upgrade	93
dbaascli job getStatus	93
dbaascli patch db apply	94
dbaascli patch db prereq	95
dbaascli pdb backup	95
dbaascli pdb bounce	96
dbaascli pdb close	97
dbaascli pdb getConnectString	98
dbaascli pdb create	98
dbaascli pdb delete	100
dbaascli pdb getDetails	101
dbaascli pdb list	101
dbaascli pdb localClone	102
dbaascli pdb open	103
dbaascli pdb recover	104
dbaascli pdb refresh	105
dbaascli pdb relocate	106
dbaascli pdb remoteClone	108
dbaascli system getDBHomes	110
dbaascli system getGridHomes	110
dbaascli tde changePassword	110
dbaascli tde addSecondaryHsmKey	111
dbaascli tde enableWalletRoot	112
dbaascli tde encryptTablespacesInPDB	113
dbaascli tde fileToHsm	114
dbaascli tde getHsmKeys	115
dbaascli tde getMkidForKeyVersionOCID	115
dbaascli tde getPrimaryHsmKey	116
dbaascli tde hsmToFile	116
dbaascli tde listKeys	118
dbaascli tde removeSecondaryHsmKey	118
dbaascli tde rotateMasterKey	121
dbaascli tde setKeyVersion	121
dbaascli tde setPrimaryHsmKey	122
dbaascli tde status	123
Monitoring and Managing Exadata Storage Servers with ExaCLI	124

About the ExaCLI Command	124
Exadata Storage Server Username and Password	125
ExaCLI Command Syntax	125
Connecting to a Storage Server with ExaCLI	127
ExaCLI Command Reference	128
Monitor Metrics for VM Cluster Resources	133
Prerequisites for Using Metrics	133
View Metrics for VM Cluster	134
View Metrics for a Database	135
View Metrics for VM Clusters in a Compartment	137
View Metrics for Databases in a Compartment	138
Manage Oracle Trace File Analyzer	139
Manage Database Service Agent	140
Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service	140
Oracle Exadata Database Service on Dedicated Infrastructure Events	146
About Event Types on Exadata Cloud Infrastructure	147
Prerequisites for Event Service	147
Oracle Exadata Database Service on Dedicated Infrastructure Event Types	148
Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types	160
Exadata Cloud Infrastructure Critical and Information Event Types	167
Exascale DB Storage Vaults Event Types	170
Exadata Cloud Infrastructure VM Cluster Event Types	170
VM Node Subsetting Event Types	183
Data Guard Event Types	187
Oracle Database Home Event Types	190
Database Event Types	191
Pluggable Database Event Types	193
Database Service Events	200
Overview of Database Service Events	201
Database Service Event Types	204
Receive Notifications about Database Service Events	215
Temporarily Restrict Automatic Diagnostic Collections for Specific Events	215
Remediation	219
Application VIP Event Types	241
Interim Software Updates Event Types	253
Serial Console Connection Event Types	259
Viewing Audit Log Events	264
Serial Console History Event Types	265
Policy Details for Exadata Cloud Infrastructure	270
About Resource-Types	270
Resource-Types for Exadata Cloud Service Instances	270

Supported Variables	272
Details for Verb + Resource-Type Combinations	272
Database-Family Resource Types	273
Permissions and API operation details for Cloud Exadata Infrastructures	274
Permissions and API operation details for VM Clusters	274
Permissions and API operation details for Exascale DB Storage Vaults	276
Permissions and API operation details for DB Nodes	277
Permissions and API operation details for DB Node Console Connection	277
Permissions and API operation details for DB Node Console History	278
Permissions and API operation details for DB Homes	278
Permissions and API operation details for DB Servers	279
Permissions and API operation details for Database Software Image	280
Permissions and API operation details for Pluggable Databases (PDBs)	280
Permissions and API operation details for Databases (CDBs)	282
Permissions and API operation details for DB Backups	283
Permissions and API operation details for Data Guard Association	283
Permissions and API operation details for Data Guard Group	284
Permissions and API operation details for Key Stores	284
Permissions and API operation details for Application VIPs	285
Permissions and API operation details for Interim Software Updates	285
Permissions and API operation details for Scheduling Policies	286
Permissions and API operation details for Scheduling Windows	286
Permissions and API operation details for Scheduling Plan	287
Permissions and API operation details for Scheduled Action	287
Permissions and API operation details for Execution Windows	288
Permissions and API operation details for Execution Action	288
Permissions and API operation details for Oracle DB Azure Connectors	289
Permissions and API operation details for Oracle DB Azure Blob Mounts	290
Permissions and API operation details for Oracle DB Azure Blob Containers	290
Permissions and API operation details for Oracle DB MCI Work Requests	290
Permissions and API operation details for Multicloud Resource Discoveries	291
Permissions and API operation details for Oracle DB Azure Vaults	291
Permissions and API operation details for Oracle DB Azure Keys	292
Permissions and API operation details for Oracle DB Azure Vault Associations	292
Permissions and API operation details for Oracle DB GCP Connectors	293
Permissions and API operation details for Oracle DB GCP Keyrings	293
Permissions and API operation details for Oracle DB GCP Keys	294
Permissions and API operation details for Oracle DB AWS Identity Connectors	294
Permissions and API operation details for Oracle DB AWS Key	295
Permissions Required for Each API Operation	295
Managing Exadata Resources with Oracle Enterprise Manager Cloud Control	304
Overview of Oracle Enterprise Manager Cloud Control	305

Features of Enterprise Manager Cloud Control	305
Analyzing Exadata Database Service Database Performance	306
Observability and Management for Exadata Database Service on Dedicated Infrastructure	306
Metrics for Exadata Cloud Infrastructure in the Database Management Service	307
Oracle Cloud Infrastructure Operations Insights	319
Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases	319
About Database Management	320
Using the Console to Enable Database Management for a Container Database (CDB)	320
Using the Console to Enable Database Management for a Pluggable Database (PDB)	323
Using the Console to Edit Database Management for a Pluggable Database (PDB)	326
Using the Console to Disable Database Management for a Pluggable Database (PDB)	328
Using the Console to View Performance Hub for a Container Database (CDB)	329
Using the Console to View Performance Hub for a Pluggable Database (PDB)	330
Using the API to Enable, Disable, or Update Database Management Service	330
Oracle Cloud Database Metrics	331
Security Guide for Oracle Exadata Database Service on Dedicated Infrastructure	333
Part 1: Security Configurations and Default Enabled Features	333
Responsibilities	333
Infrastructure Security	334
Guiding Principles Followed for Security Configuration Defaults	337
Security Features	338
Guest VM Default Fixed Users	339
Default Security Settings: Customer VM	341
Default Processes on Customer VM	342
Default Database Security Configuration	346
Default Backup Security Configuration	347
Operator Access to Customer System and Customer Data	348
Compliance Requirements	348
Break Glass Procedure for Accessing Customer's Guest VM	349
Part 2: Additional Procedures for Updating Security Posture	350
Customer Responsibilities	351
Enabling Additional Security Capabilities	351
EU Data Act Compliance	352
Troubleshooting Exadata Cloud Infrastructure Systems	353
Known Issues for Exadata Cloud Infrastructure	353
CPU Offline Scaling Fails	353
Adding a VM to a VM Cluster Fails	354
Troubleshoot Network Connectivity	354
Backup Failures in Exadata Database Service on Dedicated Infrastructure	356
Determining the Problem	356

Database Service Agent Issues	357
Object Store Connectivity Issues	357
Host Issues	358
Database Issues	358
TDE Wallet and Backup Failures	362
Troubleshooting Oracle Data Guard	365
Troubleshooting Data Guard using logfiles	366
Troubleshooting the Data Guard Setup Process	368
Patching Failures on Exadata Cloud Infrastructure Systems	370
Determining the Problem	371
Troubleshooting and Diagnosis	371
Obtaining Further Assistance	372
Collecting Cloud Tooling Logs	373
Collecting Oracle Diagnostics	373
Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup	373
Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel	374

1

Oracle Exadata Database Service on Dedicated Infrastructure Overview

This topic is an overview of the Exadata Cloud Infrastructure formerly Exadata Cloud Service.

- [Oracle Exadata Database Service on Dedicated Infrastructure Description](#)
Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Dedicated Infrastructure
- [Exadata Cloud Management Interfaces](#)
Exadata Cloud Infrastructure provides a variety of management interfaces to fit your use case and automation needs.
- [Exadata Database Service on Dedicated Infrastructure Technical Architecture](#)

Oracle Exadata Database Service on Dedicated Infrastructure Description

Learn how you can leverage the combined capabilities of Oracle Exadata and Oracle Cloud Infrastructure with Oracle Exadata Database Service on Dedicated Infrastructure

- [About Exadata Cloud Infrastructure](#)
Exadata Cloud Infrastructure allows you to leverage the power of Exadata in the cloud.
- [Licensing Considerations for Oracle Exadata Database Service on Dedicated Infrastructure](#)
Subscription to Exadata Cloud Infrastructure can include all of the required Oracle Database software licenses, or you can choose to bring Oracle Database software licenses that you already own to Oracle Exadata Database Service on Dedicated Infrastructure.
- [Supported Database Edition and Versions for Exadata Cloud Infrastructure](#)
Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.
- [Subscription Types](#)
Available subscription types.
- [Metering Frequency and Per-Second Billing](#)
Per-Second billing, minimums and limitations on billing.
- [Technical Architecture of Exadata Cloud Infrastructure Systems](#)
Exadata Cloud Infrastructure systems integrate Oracle's Exadata Database Machine hardware with the networking resources needed to securely connect to your organization's on-premise network and to other services in the Oracle cloud.
- [Scaling Options](#)
Introduction to the Scaling option on Exadata Cloud Infrastructure.
- [System and Shape Configuration Options](#)
Review the list of Exadata System Shapes
- [Support Requirements](#)

About Exadata Cloud Infrastructure

Exadata Cloud Infrastructure allows you to leverage the power of Exadata in the cloud.

You can provision flexible X8M and X9M systems that allow you to add database compute servers and storage servers to your system as your needs grow. X8M and X9M systems offer RDMA over Converged Ethernet (RoCE) networking for high bandwidth and low latency, persistent memory (PMEM) modules, and intelligent Exadata software. X8M and X9M systems can be provisioned using a shape equivalent to a quarter rack X8 or X9M system, and then database and storage servers can be added at any time after provisioning. For more information on X8M and X9M systems, see [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#).

Note

The [RDMA](#) software allows computers in a network to exchange data in main memory without involving the processor, cache, or OS of either computer. RDMA can improve throughput and performance because it frees up resources, and it can also facilitate a faster data transfer rate. RDMA over Converged Ethernet (RoCE) is the network protocol that allows RDMA over an Ethernet network.

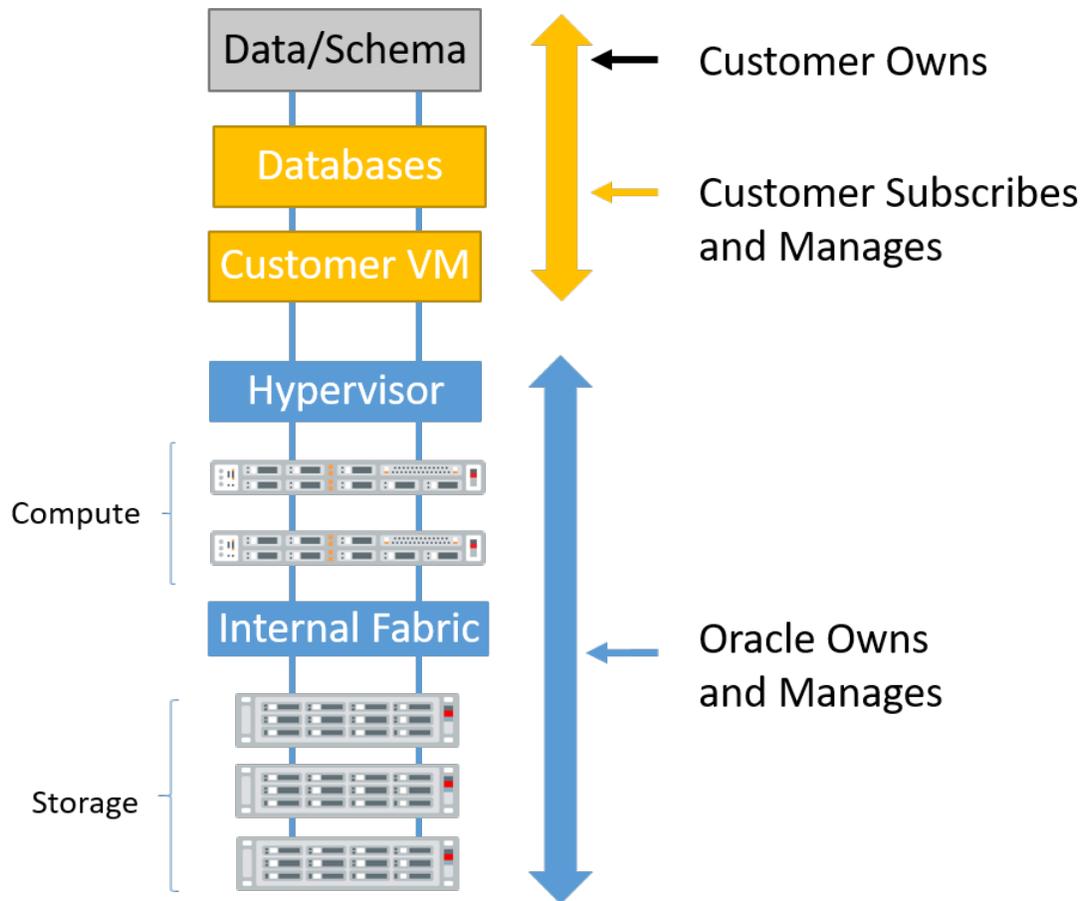
X8 and X7 systems are also available in fixed-shapes (quarter, half, and full rack systems). These systems use InfiniBand networking, and do not have the ability to scale database and storage servers. You can also provision an Exadata base system, which has a smaller capacity than a quarter rack system.

For all Exadata Cloud Infrastructure instances, you can configure automatic backups, optimize for different workloads, and scale the OCPU and storage allocations as needed.

- [Roles and Responsibilities for Exadata Cloud Infrastructure Software Maintenance](#)
Oracle is responsible for the base OS and hardware. The customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.

Roles and Responsibilities for Exadata Cloud Infrastructure Software Maintenance

Oracle is responsible for the base OS and hardware. The customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.



Customer Responsibilities

Customer is responsible for the Guest VM OS, Grid Infrastructure, and the database software maintenance.

The customer owns everything inside the database: data schema, and encryption keys.

Customer subscribes to database services

- Customer manages VMs and Databases using Cloud Automation (UI / APIs)
- Automation to create, delete, patch, backup, scale up/down, and so on
- Runs supported Oracle Database versions 19c and 26ai
- Customer controls access to customer VM
- Customer can install and manage additional software in customer VM
- Oracle staff are not authorized to access customer VM

Oracle Responsibilities

Oracle owns and manages infrastructure. Oracle is responsible for the base OS and hardware.

- Hypervisor, physical database and storage servers, storage network

- Patching, security scans, security updates
- Monitoring and maintenance
- Customer is not authorized to access Oracle infrastructure.

Licensing Considerations for Oracle Exadata Database Service on Dedicated Infrastructure

Subscription to Exadata Cloud Infrastructure can include all of the required Oracle Database software licenses, or you can choose to bring Oracle Database software licenses that you already own to Oracle Exadata Database Service on Dedicated Infrastructure.

If you choose to include Oracle Database software licenses in your Oracle Exadata Database Service on Dedicated Infrastructure subscription, then the included licenses contain all of the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC). Exadata Cloud Infrastructure also comes with cloud-specific software tools that assist with administration tasks, such as backup, recovery, and patching.

Supported Database Edition and Versions for Exadata Cloud Infrastructure

Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.

The Enterprise Edition - Extreme Performance provides all the features of Oracle Database Enterprise Edition, plus all the database enterprise management packs and all the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC).

Exadata Cloud Infrastructure supports the following database versions:

- Oracle AI Database 26ai
- Oracle Database 19c
- Oracle Database 12c Release 2 (12.2) (**Upgrade Support Required**)
- Oracle Database 12c Release 1 (12.1) (**Upgrade Support Required**)
- Oracle Database 11g Release 2 (11.2) (**Upgrade Support Required**)

Note

- Earlier database versions are supported on a 19c cloud VM cluster and can be created at anytime. Cloud VM clusters created with earlier Oracle Database versions will not automatically support Oracle Database 19c.
- For information on upgrading an existing database, see [Upgrading Exadata Databases](#).
- To use Autonomous Recovery Service as a backup destination, your target database must have a minimum compatibility level of 19.0 (the `COMPATIBLE` initialization parameter must be set to 19.0.0 or higher).

For Oracle Database release and software support timelines, see Release Schedule of Current Database Releases (Doc ID 742060.1) in the My Oracle Support portal.

Related Topics

- [Release Schedule of Current Database Releases \(Doc ID 742060.1\)](#)

Subscription Types

Available subscription types.

The available purchase models are:

- **Pay As You Go**
Pay As You Go (PAYG) pricing lets customers quickly provision services with no commitment, and they're only charged for what they use. There's no upfront commitment and no minimum service period. Any cloud infrastructure (IaaS) and platform (PaaS) services consumed are metered and billed based on that consumption. If, during the services period of your order, Oracle makes new IaaS and PaaS services available within your cloud services account, Oracle will notify you of any fees that would apply to their activation and use. For more details, see our complete price list.
- **Annual Universal Credits**
Oracle Annual Universal Credits enables customers to have the flexibility to use any Oracle Cloud Infrastructure and platform services at any time, in any region, to deliver faster time to market. Customers can commit to an amount of Oracle Annual Universal Credits that can be applied towards the future usage of eligible Oracle IaaS and PaaS cloud services. This payment option offers a significant savings across cloud services, combining cost reduction and a predictable monthly spend with a ramp up period as you onboard your workloads.

See the [Universal Credit Pricing FAQ](#) for more information.

Metering Frequency and Per-Second Billing

Per-Second billing, minimums and limitations on billing.

For each Exadata Cloud Infrastructure instance you provision, you are billed for the infrastructure for a minimum of 48 hours, and then by the second after that. Each OCPU you add to the system is billed by the second, with a minimum usage period of 1 minute. If you terminate the cloud VM cluster and do not terminate the cloud Exadata infrastructure resource, billing will continue for the infrastructure resource.

Technical Architecture of Exadata Cloud Infrastructure Systems

Exadata Cloud Infrastructure systems integrate Oracle's Exadata Database Machine hardware with the networking resources needed to securely connect to your organization's on-premise network and to other services in the Oracle cloud.

For a complete architectural overview of the components that make up an Exadata Cloud Infrastructure system, see [Oracle Exadata Cloud Service \(ExaCS\) Technical Architecture](#). This interactive reference guides you through the key hardware and networking resources in the system, and provides technical specifications for the database (compute) and storage servers to help you plan for your deployment.

Scaling Options

Introduction to the Scaling option on Exadata Cloud Infrastructure.

Introduction to the Scaling option on Exadata Cloud Infrastructure.

Two kinds of scaling operations are supported for an Exadata Cloud Infrastructure:

- For X8M, X9M, and X11M systems, the flexible shape allows you to add additional database and storage servers to the cloud Exadata infrastructure resource as needed. See [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#).

For more information on each type of scaling, see *Scaling an Exadata Cloud Infrastructure Instance*.

- [Scaling CPU cores within an Exadata Cloud Infrastructure instance](#)
If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores (ECPUs for X11M) symmetrically across all the nodes in the system as follows:
- [Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations](#)
Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

Related Topics

- [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#)
Oracle Cloud Infrastructure scalable X8M, X9M, and X11M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- [Scaling Resources within an Exadata Infrastructure Instance](#)
If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

Scaling CPU cores within an Exadata Cloud Infrastructure instance

If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores (ECPUs for X11M) symmetrically across all the nodes in the system as follows:

The options for each of the shapes are:

You can scale the infrastructure by adding DB servers or Storage Servers, up to the infrastructure limits. For more information on adding compute and storage resources to an X8M or X9M system enabled for MVM, see [Scaling Exadata X8M or X9M Compute and Storage](#).

You can scale CPU cores (ECPUs for X11M) in multiples of the number of database servers currently provisioned for the cloud VM cluster. For example, if you have 6 database servers provisioned, you can add CPU cores in multiples of 6. In the case of X11M, if you have 6 database servers provisioned, you can add ECPUs in multiples of 24. At the time of provisioning, X8M, X9M, and X11M systems have as few as 2 database servers or up to 32 database servers. For more information on adding compute and storage resources to an X8M or X9M system, see [Scaling Exadata X8M or X9M Compute and Storage](#) (replace this text with "Scaling Exadata X8M, X9M, and X11M Compute and Storage")

All systems that are not X8M, X9M, or X11M are fixed shape systems. For a base system or an X7 or X8 quarter rack, you can scale in multiples of 2 across the 2 database compute nodes. For an X7 or X8 half rack, you can scale in multiples of 4 across the 4 database compute nodes. For an X7 or X8 full rack, you can scale in multiples of 8 across the 8 database compute nodes.

For non-metered service instances, you can temporarily modify the compute node processing power (bursting) or add compute node processing power on a more permanent basis. For a metered service instance, you can simply modify the number of enabled CPU cores.

You can provision an Exadata Cloud Infrastructure instance with zero CPU cores, or scale the service instance down to zero cores after you provision it. With zero cores, you are billed only for the infrastructure until you scale up the system. For detailed information about pricing, see [Exadata Cloud Service Pricing](#).

Note

OCPU (ECPUs for X11M) scaling activities are done online with no downtime.

For information on CPU cores per configuration, see Exadata Shape Configurations. To learn how to scale a system, see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster*.

Related Topics

- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.
- [Exadata Shape Configuration](#)
This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.
- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)

Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations

Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.

This is useful when a database deployment requires:

- Processing power that is beyond the capacity of the current system configuration.
- Storage capacity that is beyond the capacity of the current system configuration.
- A performance boost that can be delivered by increasing the number of available compute nodes.
- A performance boost that can be delivered by increasing the number of available Exadata Storage Servers.

You can move your workloads to a larger fixed shape (X7 and X8 hardware shapes) or to the flexible X8M shape, which allows for easy expansion of compute and storage resources as your workloads grow.

To assist with moving your database deployments between Exadata Cloud Infrastructure instances, you can restore a backup to a different service instance with more capacity, create a Data Guard association for your database in a service instance with more capacity, and then perform a switchover so that your new standby database assumes the primary role. To start the process, contact Oracle and request a service limit increase so that you can provision the larger service instance needed by your database.

System and Shape Configuration Options

Review the list of Exadata System Shapes

- [Exadata Shape Configuration](#)
This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.

Exadata Shape Configuration

This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.

Exadata X11M, X9M, and X8M shapes start with 2 database and 3 storage servers. Compute and/or storage servers can be added independently to these shapes up to a total of 32 DB servers and 64 storage servers. The initial minimum configuration for Exadata X6, X7, and X8 also starts with 2 database and 3 storage servers similar to the quarter rack shape. They are also offered in half and full rack shapes.

See the following sections for shape specifications:

- [Exadata X11M](#)
The values in the table that follows represent the specifications for an X11M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.
- [Exadata X9M](#)
The values in the table that follows represent the specifications for an X9M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.
- [Exadata X8M](#)
The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.
- [Exadata X8 Shapes](#)
The values in the table that follows represent the specifications for an X8 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- [Exadata X7 Shapes](#)
The values in the table that follows represent the specifications for an X7 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- [Exadata X6 Shapes](#)
The values in the table that follows represent the specifications for an X6 cloud instance with fixed Quarter, Half, and Full Rack shapes.
- [Exadata Base System](#)
An Exadata Base System is a fixed shape similar in size to a Quarter Rack with some differences in capacity.

Exadata X11M

The values in the table that follows represent the specifications for an X11M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.

Independently add compute and/or storage servers up to a total of 32 database servers and 64 storage servers.

- A single database (DB) server contains 760 ECPUs and 1390 GB of memory.
- A single storage server contains 80 TB of usable disk storage capacity.

Property	Minimum Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3

Property	Minimum Configuration
Minimum (Default) Number of Enabled ECPUs	0
Total Usable ECPUs in DB Servers per System	1520
Total Memory Available for VMs (GB)	2780
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable File System Size Per VM (GB)	900
VM Image size minimum and default (GB)	244
Max Number of VM Clusters per System	8
Max Number of VMs per DB server	8
Total Flash Capacity (TB)	81.6
Total Usable Disk Storage Capacity (TB)	240

VM Image size minimum and default includes 60 GB for /u02.

A maximum of 8 VM Clusters can be created on a single infrastructure. For more information, see [Estimating How Much Local Storage You Can Provision to Your VMs](#) and [Scaling Local Storage](#).

Exadata X9M

The values in the table that follows represent the specifications for an X9M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.

Independently add compute and/or storage servers up to a total of 32 DB servers and 64 storage servers.

- A single DB server contains 126 usable cores and 1390 GB memory.
- A single storage server contains 63.6 TB of usable disk storage capacity.

Property	Minimum Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum (Default) Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	252
Total Memory Available for VMs (GB)	2780
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable File System Size Per VM (GB)	900
VM Image size minimum and default (GB)	244
Max Number of VM Clusters per System	8
Max Number of VMs per DB server	8
Total Flash Capacity (TB)	76.8
Total Usable Disk Storage Capacity (TB)	190

VM Image size minimum and default includes 60 GB for /u02.

A maximum of 8 VM Clusters can be created on a single infrastructure. For more information, see [Estimating How Much Local Storage You Can Provision on Your VMs](#) and [Scaling Local Storage](#).

Exadata X8M

The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.

Independently add compute and/or storage servers up to a total of 32 DB servers and 64 storage servers.

- A single DB server contains 50 usable cores and 1390 GB memory.
- A single storage server contains 49.9 TB of usable disk storage capacity.

Property	Minimum Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum (Default) Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	100
Total Memory Available for VMs (GB)	2780
Max Usable Local Storage Per DB Server (GB)	2243
Max Usable File System Size Per VM (GB)	900
VM Image size minimum and default (GB)	244
Max Number of VM Clusters per System	8
Max Number of VMs per DB server	8
Total Flash Capacity (TB)	76.8
Total Usable Disk Storage Capacity (TB)	149

VM Image size minimum and default includes 60 GB for /u02.

A maximum of 8 VM Clusters can be created on a single infrastructure. For more information, see [Estimating How Much Local Storage You Can Provision on Your VMs](#) and [Scaling Local Storage](#).

Exadata X8 Shapes

The values in the table that follows represent the specifications for an X8 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter3.100	Exadata.Half3.200	Exadata.Full3.400
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	0	0	0
Total Usable Cores in DB Servers per System	100	200	400
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	700	700	700

Property	Quarter Rack	Half Rack	Full Rack
Total Flash Capacity (TB)	76.8	179.2	358.4
Total Usable Disk Storage Capacity (TB)	149	299	598

Exadata X7 Shapes

The values in the table that follows represent the specifications for an X7 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter2.92	Exadata.Half2.184	Exadata.Full2.368
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	0	0	0
Total Usable Cores in DB Servers per System	92	184	368
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	1000	1000	1000
Total Flash Capacity (TB)	76.8	153.6	307.2
Total Usable Disk Storage Capacity (TB)	106	212	424

Exadata X6 Shapes

The values in the table that follows represent the specifications for an X6 cloud instance with fixed Quarter, Half, and Full Rack shapes.

Property	Quarter Rack	Half Rack	Full Rack
Shape Name	Exadata.Quarter1.84	Exadata.Half1.168	Exadata.Full1.336
Number of DB servers per System	2	4	8
Number of Storage Servers per System	3	6	12
Minimum Number (Default) of Enabled CPU Cores	22	44	88
Total Usable Cores in DB Servers per System	84	168	336
Total Memory Available (GB)	1440	2880	5760
Max Usable Local Storage (GB)	200	200	200

Property	Quarter Rack	Half Rack	Full Rack
Total Flash Capacity (TB)	38.4	76.8	153.6
Total Usable Disk Storage Capacity (TB)	73	168	336

Note

Exadata X6 shapes must be provisioned using the License Included option. Bring-Your-Own-License (BYOL) is not supported with the X6 shapes.

Exadata Base System

An Exadata Base System is a fixed shape similar in size to a Quarter Rack with some differences in capacity.

Property	Configuration
Number of DB servers per System	2
Number of Storage Servers per System	3
Minimum Number of Enabled CPU Cores	0
Total Usable Cores in DB Servers per System	48
Total Memory Available (GB)	720
Max Usable Local Storage (GB)	900
Total Flash Capacity (TB)	38.4
Total Usable Disk Storage Capacity (TB)	73

For information on provisioning an Exadata Cloud Infrastructure instance, see [Creating an Exadata Cloud Infrastructure Instance](#).

Support Requirements

Oracle Database Support is required to run a database service on Exadata Database Service. Customers running the license included option have support through their database cloud subscription. Customers bringing their own licenses to Exadata Database Service must have a valid support contract covering the database licenses they are running in the database cloud service. Customers are responsible for applying updates to their databases in a timely manner to protect from product issues and security vulnerabilities.

Oracle Applications Unlimited (specifically Oracle E-Business Suite, PeopleSoft, JD Edwards, Siebel and Hyperion applications) require ongoing updates to ensure continued interoperability. As such, Oracle requires customers to have an active Oracle support contract for any such programs running on Oracle Exadata Database Service. Note that third party support does not satisfy this requirement. This requirement shall not be interpreted as superseding anything in the Oracle Technical Support Policies, including, but limited to, the Matching Service Level requirement.

Exadata Cloud Management Interfaces

Exadata Cloud Infrastructure provides a variety of management interfaces to fit your use case and automation needs.

- [Introduction to Exadata Cloud Management Interfaces](#)
The Exadata Cloud resources on Oracle Cloud Infrastructure (OCI) are created and managed through a variety of interfaces provided to fit your different management use cases.
- [OCI Control Plane Interfaces for Exadata Cloud Infrastructure](#)
The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tool and plugins provided to facilitate development and simplify the management of OCI resources.
- [Local VM Command-Line Interfaces](#)
In addition to the OCI REST-based APIs, CLI utilities located on the VM guests, provisioned as part of the VM clusters on the Exadata Cloud Infrastructure, are available to perform various lifecycle and administration operations.

Introduction to Exadata Cloud Management Interfaces

The Exadata Cloud resources on Oracle Cloud Infrastructure (OCI) are created and managed through a variety of interfaces provided to fit your different management use cases.

The various interfaces include:

- OCI Console interface and automation tools, see *Using the Console*
- Application Programming Interfaces (APIs)
- Command-Line Interfaces (CLIs)

The management interfaces are grouped into two primary categories:

- OCI Control Plane Interfaces
- Local Exadata Cloud VM CLIs

Note

For more information and best practices on how these interfaces align for various Exadata Cloud database management use cases, refer to My Oracle Support note: *Exadata Cloud API/CLI Alignment Matrix (Doc ID 2768569.1)*.

Related Topics

- [Oracle Database console overview](#)
- [Using the Console](#)
- <https://support.oracle.com/epmos/faces/DocContentDisplay?id=2768569.1>

OCI Control Plane Interfaces for Exadata Cloud Infrastructure

The OCI control plane accepts input from the OCI APIs, the OCI Console, and custom interfaces built with kits, tool and plugins provided to facilitate development and simplify the management of OCI resources.

The OCI APIs are typical REST APIs that use HTTPS requests and responses. The OCI Console, an intuitive, graphical interface for creating and managing your Exadata Cloud and other OCI resources, is one of the interfaces to the OCI APIs. When looking to develop automation utilizing the OCI APIs, a number of additional interfaces including: kits, tools and plug-ins, are provided to facilitate development and simplify the management of OCI resources. A subset of these APIs applies to Exadata Cloud resources and the containing infrastructure. Each of these various interfaces provide the same functionality, all calling the OCI APIs, and are provided to enable flexibility and choice depending on preference and use case.

- **Command Line Interface (CLI):** The OCI CLI is a small footprint tool that you can use on its own or with the Console to perform Exadata Cloud resource tasks and other OCI tasks. The CLI provides the same core functionality as the Console, plus additional commands. Some of these, such as the ability to run scripts, extend the Console's functionality.
- **Software Development Kits (SDK):** OCI provides SDKs to enable you to develop custom solutions for your Exadata Cloud and other OCI based services and applications.
- **DevOps Tools and Plug-ins:** These tools can simplify provisioning and managing infrastructure, enable automated processes and facilitate development. Tools include the OCI Terraform Provider used with Resource Manager and OCI Ansible Collection.
- **Cloud Shell:** Cloud Shell is a free-to-use, browser-based terminal, accessible from the OCI Console, that provides access to a Linux shell with pre-authenticated OCI CLI and other useful developer tools. You can use the shell to interact with Exadata Cloud and other OCI resources, follow labs and tutorials, and quickly run OCI CLI commands.
- **Documentation: Appendix and Reference:** This general reference shows how to configure the SDKs and other developer tools to integrate with Oracle Cloud Infrastructure services.
- **Documentation: REST APIs:** This complete reference provides details on the Oracle Cloud Infrastructure REST APIs, including descriptions, syntax, endpoints, errors, and signatures. Exadata Cloud Infrastructure specific OCI REST APIs can be found throughout the documentation in the *Using the API* sections specific to each service:
 - *Using the API to Create Infrastructure Components*
 - *Using the API to Manage Exadata Cloud Service Instance*
 - *Using the API to manage database software images*
 - *Using the API to Create Oracle Database Home on Exadata Cloud Service*
 - *Using the API to Manage Oracle Database Home*
 - *Using the API to manage Databases*
 - *Using the API to Update the Grid Infrastructure on a VM Cluster Resources*
 - *Using the API to manage the I/O resources of an Exadata cloud VM cluster*
 - *Using the API to Patch an Exadata Cloud Service Instance*
 - *Using the API to upgrade Databases*
 - *Using the API to Manage Data Guard Associations*
 - *Using the API to manage backups*

Related Topics

- [Command Line Interface \(CLI\)](#)
- [Software Development Kits](#)
- [DevOps Tools and Plug-ins](#)

- [Terraform Provider](#)
- [Resource Manager](#)
- [Ansible Collection](#)
- [Cloud Shell](#)
- [Appendix and Reference](#)
- [REST APIs](#)
- [Using the API to Create Infrastructure Components](#)
- [Using the API to Manage Exadata Cloud Infrastructure Instance](#)
- [Using the API to manage database software images](#)
Use these API operations to manage database software images:
- [Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure](#)
To create an Oracle Database home, review the list of API calls.
- [Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure](#)
Review the list of API calls to manage Oracle Database home.
- [Using the API to manage Databases](#)
- [Using the API to Upgrade the Grid Infrastructure in a VM Cluster](#)
- [Using the API to manage the I/O resources of an Exadata cloud VM cluster](#)
- [Using the API to Patch an Exadata Cloud Infrastructure Instance](#)
Use these API operations to manage patching the following Exadata resources: cloud VM clusters, databases, and Database Homes.
- [Using the API to upgrade Databases](#)
Use the following APIs to manage database upgrades:
- [Using the API to manage Data Guard associations](#)
Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:
- [Using the API to manage backups](#)

Local VM Command-Line Interfaces

In addition to the OCI REST-based APIs, CLI utilities located on the VM guests, provisioned as part of the VM clusters on the Exadata Cloud Infrastructure, are available to perform various lifecycle and administration operations.

The best practice is to use these utilities only when a corresponding Console command or OCI API is not available.

The utilities include:

- **dbaascli:** Use the `dbaascli` utility to perform various database lifecycle and administration operations on the Exadata Cloud Infrastructure such as
 - changing the password of a database user
 - starting a database
 - managing pluggable databases (PDBs)
- **bkup_api:** Use the `bkup_api` utility to perform various backup and recovery operations on the Exadata Cloud Infrastructure such as creating an on-demand backup of a complete

database or an individual pluggable database (PDB), or to *customize backup settings* used by the automatic backup configuration

Note

`bkup_api` is deprecated. Use `dbaascli database backup`, `dbaascli pdb backup`, or `dbaascli pdb recover`

instead.

- **ExaCLI:** Use the ExaCLI command-line utility to perform monitoring and management functions on Exadata storage servers in the Exadata Cloud.

These utilities are provided in addition to, and separate from, the OCI API-based interfaces listed above. To use the local VM command-line utilities, you must be connected to a virtual machine in an Exadata Cloud VM cluster and use the VM operating system user security, not the OCI user security, for execution. Most operations executed by these utilities sync their changes back to the OCI control plane using a process called `DB Sync`. However, there can be operations not synced with the control plane.

The cloud tooling software on the virtual machines, containing these CLI utilities, is automatically updated by Oracle on a regular basis. If needed, the tooling can be updated manually by following the instructions in *Updating Cloud Tooling Using dbaascli*.

Related Topics

- [About Using the dbaascli Utility on Exadata Cloud Infrastructure](#)
You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.
- [Managing Exadata Database Backups by Using dbaascli](#)
- [To create a backup configuration file](#)
- [Monitoring and Managing Exadata Storage Servers with ExaCLI](#)
The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.
- [Updating Cloud Tooling Using dbaascli](#)
To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

Exadata Database Service on Dedicated Infrastructure Technical Architecture

Explore the architecture of Oracle Exadata Database Service on Dedicated Infrastructure, which leverages Exadata racks within an OCI data center and includes one or more virtual machine clusters. Interactive diagrams showcase secure connectivity through Site-to-Site VPN and OCI FastConnect, while detailing client access to databases and VM clusters. The overview also highlights management capabilities via the OCI Console, CLI, and REST APIs, and emphasizes Oracle Cloud Operations' role in maintaining the infrastructure. For in-depth information, explore the [Exadata Database Service on Dedicated Infrastructure Technical Architecture](#).

2

What's New in Oracle Exadata Database Service on Dedicated Infrastructure

Oracle is constantly adding new capabilities to Exadata Cloud Infrastructure.

- [Support for Private Service Access \(PSA\)](#)
- [Enhanced Out-of-Place Restore Options for ExaDB-D](#)
- [Support for Concurrent Data Guard Operations in Multiple Standby Environments](#)
- [Enhanced Data Guard Health Insights](#)
- [AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS](#)
- [Convert a Physical Standby Database to a Snapshot Standby Database and vice versa](#)
- [Exadata Exascale with Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [Exadata Database Service for Developers](#)
- [Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud](#)
- [Cross-Service Data Guard Between ExaDB-D and ExaDB-XS](#)
- [Enhancements to Serial Console Functionality](#)
- [Change in Backup Destination for New Tenancies in Select OCI Regions \(Effective August 06, 2025\)](#)
- [Exadata System Software 25.1.0.0.0](#)
- [Expanded Tagging Support Across Database Workflows](#)
- [Amazon S3 Support for Oracle-Managed Backups on Exadata Database Service on Oracle Database@AWS](#)
- [Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure](#)
- [Oracle Data Guard Setup with Precheck Validation](#)
- [Schedule VM Cloud Automation Updates](#)
- [Classify and Control Access to Critical OCI REST APIs](#)
- [Oracle Key Vault \(OKV\) Integration with ExaDB-D to Manage Transparent Data Encryption \(TDE\) Keys](#)
- [Remove Provisioned Storage Servers](#)
- [Enhancements to support concurrent Data Guard, Container Database \(CDB\), and Pluggable Database \(PDB\) Operations](#)
- [Dual Stack \(IPv4 and IPv6\) Network Support](#)
- [Long-Term Retention Backup \(LTR\)](#)
- [Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution](#)
- [Granular Permissions for VM Cluster Update Operations](#)

- [Multiple Standby Databases](#)
- [X11M System Support](#)
- [Microsoft Entra ID \(MS-EI\) Integration with Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [Delegate Access Control for ExaDB-D](#)
- [Different RUs for Primary and Standby DB Homes in Data Guard Associations, Switchover, and Failover Operations](#)
- [Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Assign New Key Versions to KMS-Based Container Databases \(CDBs\) and Pluggable Databases \(PDBs\)](#)
- [Enhancement to Backup and Restore from a Standby Database in a Data Guard Environment](#)
- [VM Cluster on a Single VM](#)
- [Multicloud Oracle Database Backup Support](#)
- [Restore a Backup to Create a Database Across Regions](#)
- [Cost and Usage Attribution for Pluggable Databases \(PDBs\)](#)
- [Use the Same Custom Software Image Across OCI Regions](#)
- [Ability to Increase the Size of Guest VM Local File Systems](#)
- [Create and Use Custom Software Images](#)
- [Manage Serial Console Access to Oracle Exadata Database Service on Dedicated Infrastructure Systems](#)
- [Oracle AI Database 26ai on Exadata Database Service on Dedicated Infrastructure](#)
- [Enable Unified Auditing While Creating a Database Home](#)
- [Provision a VM Cluster with Either an OL7 or OL8-Based Image](#)
- [Enhancement to the OCI Console to Remove Database and Storage Servers](#)
- [Enable Data Guard Across Different VCNs or Compartments in the Same OCI Region](#)
- [Enhancement to Pluggable Database \(PDB\) Management](#)
- [Manage Administrator \(SYS User\) and TDE Wallet Passwords](#)
- [Backup and Restore from a Standby Database with OCI Object Storage in a Data Guard Environment](#)
- [Cancel a Running Full or Incremental Backup](#)
- [Autonomous Recovery Service as the Default Backup Destination](#)
- [Exadata Fleet Update](#)
- [Update Guest VM \(domU\) Operating System to Oracle Linux 8](#)
- [Use a Backup to Create a Database Across Availability Domains within the Same Region](#)
- [Interim Software Updates](#)
- [Enhanced Controls to Configure Automatic Full \(L0\) and Incremental \(L1\) Backups](#)
- [Configure Oracle Database Autonomous Recovery Service as a Backup Destination](#)
- [Application VIP Support](#)
- [Monthly ExaDB-D Infrastructure Security Maintenance](#)

- [Identity and Access Management \(IAM\) Integration with Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [Exadata Cloud Infrastructure: Private DNS](#)
- [Enhanced Infrastructure Maintenance Controls](#)
- [Database Management Support for Pluggable Databases in Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [Microsoft Azure Active Directory Integration with Oracle Cloud Infrastructure Databases](#)
- [Create and Manage Multiple Virtual Machines per Exadata System \(MultiVM\) and VM Cluster Node Subsetting](#)
- [VM Cluster and Database Health and Performance Metrics in the OCI Console](#)
- [Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [Automatic Diagnostic Collection](#)
- [Exadata Database on Dedicated Infrastructure: Key Management Service for Cross Region Data Guard](#)
- [Concurrently Create or Terminate Oracle Databases in a VM Cluster](#)
- [VM Guest Exadata OS Image Major Version Update](#)
- [Database Service Events capability for Exadata Database](#)
- [Exadata Database on Dedicated Infrastructure: 'Create database from backup' now available for databases using customer-managed encryption](#)
- [Support for DB Home Minor Version Selection \(N-3\)](#)
- [Oracle Cloud Infrastructure Operations Insights Support for Oracle Cloud Databases](#)
- [Specify the Same SID for Primary and Standby Databases in Data Guard Association](#)
- [Exadata Cloud Infrastructure: Pluggable database lifecycle support](#)
- [Exadata Cloud Infrastructure: Set DB_UNIQUE_NAME and Oracle SID prefix during database creation](#)
- [Elastic Expansion](#)
- [Oracle Database: Encryption key options updated for Exadata Cloud Infrastructure databases](#)
- [Performance Hub Exadata Tab](#)
- [Exadata Cloud Infrastructure: custom SCAN listener port for VM cluster](#)
- [Performance Hub & metrics available for databases running in Exadata Cloud Infrastructure](#)
- [Maintenance advisory contacts for Exadata infrastructure](#)
- [Data Guard protection mode enhancement for Exadata Cloud Infrastructure instances](#)
- [Exadata Cloud Infrastructure: Non-rolling infrastructure patching option now available](#)
- [Customer-managed encryption keys available with Oracle Data Guard-enabled databases in Exadata Cloud Infrastructure](#)
- [ExaDB-D OS/DomU Patching Project](#)
- [Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure](#)
- [Exadata Cloud Infrastructure: Oracle Database 19c upgrade feature available](#)

- [Create custom database software images for Exadata Cloud Infrastructure instances](#)
- [Exadata Cloud Infrastructure: grid infrastructure upgrade for cloud VM clusters](#)
- [Exadata Cloud Infrastructure: the flexible X8M shape now available](#)
- [Exadata Cloud Infrastructure: use an existing Database Home when setting up a Data Guard standby database](#)

Support for Private Service Access (PSA)

- **Service:** [Database](#)
- **Release Date:** January 15, 2026

This enhancement introduces Private Service Access (PSA), which enforces a private-endpoint model and changes how VCN resources connect to OCI services.

Key benefits include:

- Control which credentials can be used from your networks to help prevent cross-tenancy data exfiltration
- Control credential usage across specific endpoints within a Virtual Cloud Network (VCN)
- Optional ZPR Policy Enforcement for PSA: When ZPR security attributes are configured on a PSA, only network traffic explicitly permitted by ZPR rules is allowed. All other traffic is implicitly denied.

Example:

The following ZPR policy allows compute resources tagged with `app:backend` to connect only to PSA objects tagged with the attribute `myservices:secure`. Customers can assign this attribute either when creating a PSA or by updating an existing PSA.

```
In networks:web VCN allow app:backend endpoints to connect to
myservices:secure endpoints
```

For more information about Private Service Access Endpoints, see [Access to Oracle Services: Private Service Access Endpoints](#).

Related Topics

- [Private Service Access for OCI Service Access from VCN](#)
You can use Private Service Access (PSA) instead of a Service Gateway to provide access to the OCI services required to support your database service.

Enhanced Out-of-Place Restore Options for ExaDB-D

- **Service:** [Database](#)
- **Release Date:** January 14, 2026

Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) now expands its out-of-place restore capabilities when creating a new database.

Previously supported restore options included:

- Creating a database from a specific backup
- Creating a database from the most recent backup

With this enhancement, an additional restore option is introduced:

- Create a database from a **specified point-in-time (timestamp)**

These enhancements provide greater flexibility and precision for database recovery scenarios, enabling customers to meet stricter recovery objectives and simplify restore operations.

Related Topics

- [To create a database from the latest backup](#)

Support for Concurrent Data Guard Operations in Multiple Standby Environments

- **Service:** [Database](#)
- **Release Date:** November 20, 2025

In addition to the [Enhancements to support concurrent Data Guard, Container Database \(CDB\), and Pluggable Database \(PDB\) Operations](#), these enhancements enable you to perform concurrent operations on CDBs and PDBs alongside Data Guard migration operations in environments with multiple standby databases.

Related Topics

- [Support for Concurrent Data Guard Operations in Multiple Standby Environments](#)

Enhanced Data Guard Health Insights

- **Service:** [Database](#)
- **Release Date:** November 19, 2025

Enhanced Data Guard health status reporting now provides detailed metrics on overall health, switchover/failover readiness, data loss exposure, and disaster recovery, making it especially useful for managing multiple standby databases.

Related Topics

- [Enhanced Data Guard Health Status Reporting](#)
The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.
- [Data Guard Event Types](#)
Review the list of event types that Data Guard group and Data Guard Associations emit.

AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS

- **Service:** [Database](#)
- **Release Date:** November 18, 2025

Exadata Database Service on Oracle Database@AWS supports integration with AWS Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.

Related Topics

- [AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS](#)

Exadata Database Service on Oracle Database@AWS supports integration with AWS Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.

Convert a Physical Standby Database to a Snapshot Standby Database and vice versa

- **Service:** [Database](#)
- **Release Date:** October 07, 2025

With this enhancement, you can temporarily convert a physical standby database to an updatable (read-write) copy, then revert it to its original state.

Snapshot standby allows administrators to perform root cause analysis and test with production data in a very simple and risk-free manner, thereby drastically reducing the risk of errors and improving resiliency. Until now, customers were able to convert a standby database into snapshot mode outside of cloud automation. Customers will now be able to create a snapshot standby database via OCI Console/API/SDK, and Terraform.

The snapshot standby feature temporarily converts the standby database to read-write mode providing full access to update the database without any disruption to the business or risk of data loss. A snapshot standby database receives and archives redo logs from the primary database but does not apply it. The snapshot standby database can be converted back to a standby database at any point in time. When converting back, all the local updates made to the snapshot standby database will be discarded and the redo logs received from the primary database will be applied converting the standby database back into an exact physical replica of the primary database.

Related Topics

- [Convert a Physical Standby Database to Snapshot Standby Database](#)
A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.
- [Convert a Snapshot Standby Database to Physical Standby Database](#)
Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.

Exadata Exascale with Oracle Exadata Database Service on Dedicated Infrastructure

- **Service:** [Database](#)
- **Release Date:** September 24, 2025

Exadata Exascale with Oracle Exadata Database Service on Dedicated Infrastructure introduces a next-generation intelligent data architecture for the cloud, combining the strengths

of Exadata with the agility of the cloud. Exascale reimagines how compute and storage resources are managed on Exadata platforms by decoupling and simplifying storage management, paving the way for innovative capabilities. It ensures industry-leading database performance, availability, and security standards that organizations expect from Exadata.

With this release, customers can harness the power of Exascale's intelligent data architecture, effortlessly configuring it for Oracle databases on Oracle Exadata Database Service on Dedicated Infrastructure infrastructure. A new Exascale storage option is available during VM cluster provisioning, enabling customers to deploy VM clusters with Exascale alongside Automatic Storage Management (ASM), all on the same Oracle Exadata Database Service on Dedicated Infrastructure infrastructure without impacting existing workloads and at no additional cost.

Minimum requirements for configuring Exascale on Oracle Exadata Database Service on Dedicated Infrastructure:

- This feature is supported on Exadata Infrastructure Model X8M (or later).
- This feature requires Exadata DB Server Version 25.1.7 (or later) and Exadata Storage Server Version 25.1.8 (or later) or 25.2.2 (or later).
- This feature requires Oracle Grid Infrastructure version 26ai (23.6 or later) and supports Oracle database versions 26ai (23.6 or later).

For Oracle Database release and software support timelines, see [Release Schedule of Current Database Releases \(Doc ID 742060.1\)](#) in the My Oracle Support portal.

Exadata Exascale with Oracle Exadata Database Service on Dedicated Infrastructure offers the following benefits:

- **Efficient and scalable storage:** Exascale improves overall storage capacity utilization and efficiency by dynamically allocating storage capacity across all storage servers and sharing it with multiple VM clusters. Exascale also allows the scaling of database storage dynamically at any time without affecting overall performance.
- **Agile development with Exascale thin clones:** Exascale enables space-efficient thin clones from any read/write pluggable database, significantly boosting developer productivity. Each developer can get a thinly provisioned clone database with all Exadata benefits for sole use, ensuring a production-like environment for application development and testing while drastically reducing the overall storage requirement and associated costs.
- **Powerful:** Exascale seamlessly integrates with development, test, and deployment pipelines, accelerating production applications on Exascale while leveraging Exadata optimizations that deliver extreme performance, reliability, availability, and security. This strengthens Exadata's mission to be the ideal platform for running all Oracle database workloads.

For more information about Oracle Exadata Exascale, see the [Oracle® Exadata Exascale User's Guide](#).

Related Topics

- [Storage Configuration Requirements for Oracle Exadata Database Service on Dedicated Infrastructure](#)
With the introduction of Exascale technology in Oracle Exadata Database Service on Dedicated Infrastructure, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

- [To create an Exascale cloud VM cluster](#)
To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [Cloning an Exadata Pluggable Database](#)
You can create local, remote, and refreshable clones.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for Cloud Exadata Infrastructures](#)
- [Permissions and API operation details for Exascale DB Storage Vaults](#)
- [Oracle Exadata Database Service on Dedicated Infrastructure Event Types](#)
The events in this section are emitted by the cloud Exadata infrastructure resource
- [Exascale DB Storage Vaults Event Types](#)
The events in this section are emitted by the `exascale-db-storage-vaults` resource.

Exadata Database Service for Developers

- **Service:** [Database](#)
- **Release Date:** September 17, 2025

The newly introduced Exadata Database-Developer VM Cluster type allows developers to build applications on Oracle Cloud without incurring Oracle Database license fees. With this VM Type, developers pay only for the underlying infrastructure. Customers using Oracle Exadata Database Service on Dedicated Infrastructure can develop for no additional cost, as the license fee is waived and the infrastructure costs are already covered under their subscription.

You can allocate as many OCPUs/ECPU, memory, and storage to the Developer VM Cluster as you require. Within the VM cluster you can have multiple CDBs, and each CDB can have multiple PDBs.

However, there are restrictions on the databases that can be hosted with the Exadata Database-Developer VM Cluster type. These restrictions are designed to prevent running production workloads on this cluster type, which is intended for development use cases. Restrictions on resources and functionality will be enforced by the database and cloud automation. Key limitations include:

- **Versions:** Only Oracle Database versions 19.26 and 23.6 and later are supported
- **Single VM Cluster:** Exadata-Database-Developer VM clusters are limited to a single VM.
- **Threads per PDB:** 2 threads per PDB.
- **Memory per PDB:** 8 GB memory per PDB.
- **Database size per PDB:** 20 GB storage per PDB.
- **Sessions per PDB:** 30 sessions per PDB.
- **Data Guard:** Creating a Data Guard standby database is prohibited when using the Developer VM Clusters.

Additionally, Multi-VM RAC (Real Application Clusters) is not supported, and Developer VM clusters are limited to a single VM. Data Guard automation is blocked, and manual configuration of Data Guard is not allowed.

Note

You cannot change the VM cluster type after deploying the VM cluster. If you wish to change the VM cluster type, you must create a new VM cluster and migrate the database to the new cluster.

Related Topics

- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud

- **Service:** [Database](#)
- **Release Date:** September 02, 2025

Exadata Database Service on Oracle Database@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using GCP Customer Managed Encryption Keys (CMEKs).

Previously, Transparent Data Encryption (TDE) master encryption keys (MEKs) could only be stored in a file-based Oracle Wallet, Oracle Cloud Infrastructure (OCI) Vault, or Oracle Key Vault (OKV). With this update, users can now store and manage MEKs directly in GCP KMS, providing improved key lifecycle control and alignment with organization-specific security policies.

Related Topics

- [Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud](#)
Exadata Database Service on Oracle Database@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).

Cross-Service Data Guard Between ExaDB-D and ExaDB-XS

- **Service:** [Database](#)
- **Release Date:** August 20, 2025

We are pleased to announce cross-service Oracle Data Guard deployment support. In a cross-service deployment, you set up primary and secondary databases between two services: Exadata Database Service on Dedicated Infrastructure (ExaDB-D) and Exadata Database Service on Exascale Infrastructure (ExaDB-XS). The ability to deploy cross-service Oracle Data Guard provides enhanced availability.

For more information, see *Cross-Service Data Guard Between ExaDB-D and ExaDB-XS*.

Related Topics

- [Cross-Service Data Guard Between ExaDB-D and ExaDB-XS](#)
You can now create a cross-service Oracle Data Guard group across database services.

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To view Data Guard Group details of databases in a Cloud VM Cluster](#)
To view the role of each database in a Data Guard Group in an Cloud VM Cluster, follow this procedure.

Enhancements to Serial Console Functionality

- **Service:** [Database](#)
- **Release Date:** August 19, 2025

These new features include:

- Serial Console access via OCI Cloud Shell
- Console History

With this enhancement, you can easily connect to the serial console of your virtual machines through OCI Cloud Shell to perform corrective actions. Additionally, you can view serial console history to review and audit previous activities conducted through the serial console by all users.

Note

- You cannot concurrently connect to more than one DB node using Cloud Shell. As an example, if you have an open connection to *DBnode1* and want to connect to *DBnode2*, you must first exit the active Cloud Shell from *DBnode1* and then establish a connection to *DBnode2*.
- Cloud Shell access to the serial console require proper IAM permissions for Cloud Shell, see [OCI Cloud Shell](#) documentation for details.

Related Topics

- [Using Cloud Shell to Connect to the Serial Console](#)
- [Displaying the Console History for a Virtual Machine](#)
- [Permissions Required for Each API Operation](#)
- [Serial Console History Event Types](#)
Review the list of new event types that serial console history emits.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for DB Node Console History](#)

Change in Backup Destination for New Tenancies in Select OCI Regions (Effective August 06, 2025)

- **Service:** [Database](#)
- **Release Date:** August 06, 2024

Starting August 06, 2025, the Autonomous Recovery Service is now the exclusive backup destination during automatic backup configuration for newly created tenancies in the following OCI regions: Frankfurt (FRA), Phoenix (PHX), and Tokyo (NRT).

For more information, see [Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console](#).

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To configure automatic backups for a database](#)
- [To create an on-demand backup of a database](#)
- [Using the Console to restore a database](#)
You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.
- [To create a database from a backup](#)
- [To create a database from the latest backup](#)

Exadata System Software 25.1.0.0.0

- **Service:** [Database](#)
- **Release Date:** December, 2024

Exadata System Software Release 25.1 has been available since December 2024. New Exadata X11M Infrastructure deployments will include Exadata System Software 25.1. This release will be applied to the Exadata Cloud Infrastructure as part of Quarterly Maintenance, starting in May 2025. It builds upon the capabilities introduced in Exadata System Software 24ai and earlier releases. This release brings several key innovations, including:

- [AI Smart Scan Adaptive Top-K Pruning](#)
- [AI Smart Scan Vector Distance Projection on Storage Servers](#)
- [AI Smart Scan Offload for BINARY and INT8 vector dimension formats](#)

Exadata Software 25.1 for Guest OS is also available and it can be applied to the Guest VMs by customers.

To learn more, refer to the [Exadata System Software Release 25.1 documentation](#).

For update instructions, see [My Oracle Support Note 3021895.1](#) — Exadata System Software 25.1.0.0.0 Update.

Note

Some Exadata Software 25.1 features may not be available in the cloud service.

Expanded Tagging Support Across Database Workflows

- **Service:** [Database](#)
- **Release Date:** July 22, 2025

Tagging, previously supported during database creation, is now extended to additional workflows, including creating a standby database and creating a database from a backup.

Related Topics

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To create a database from a backup](#)

Amazon S3 Support for Oracle-Managed Backups on Exadata Database Service on Oracle Database@AWS

- **Service:** [Database](#)
- **Release Date:** July 09, 2025

This feature introduces Amazon S3 as an additional backup destination for Exadata Database Service on Oracle Database@AWS customers, enabling them to seamlessly schedule and manage Oracle-managed database backups directly from the OCI Exadata Database Service interface.

For more information, see [Automatic Backup](#).

Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure

- **Service:** [Database](#)
- **Release Date:** July 01, 2025

Exadata Database Service on Oracle Database@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault. This feature enables Exadata Database Service on Oracle Database@Azure users to utilize Azure Key Vault (AKV) Managed HSM, AKV Premium and AKV Standard for managing TDE MEKs. This integration allows applications, Azure services, and databases to use a centralized key management solution for enhanced security and simplified key lifecycle management.

Related Topics

- [Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure](#)
Exadata Database Service on Oracle Database@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.
- [Database Multicloud Integration for Oracle Database Cloud Services](#)

Oracle Data Guard Setup with Precheck Validation

- **Service:** [Database](#)
- **Release Date:** June 17, 2025

You are now able to perform a precheck before setting up Oracle Data Guard.

As part of the Data Guard (DG) configuration, the service performed an implicit precheck. With this enhancement, you can now run an explicit precheck to identify and address potential issues before proceeding with the Data Guard setup.

Related Topics

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.

Schedule VM Cloud Automation Updates

- **Service:** [Database](#)
- **Release Date:** June 17, 2025

We're pleased to announce the General Availability (GA) of scheduling VM Cloud Automation Updates for VM Clusters in Exadata Database Service on Dedicated Infrastructure (ExaDB-D).

Previously, Oracle applied these updates automatically in the background without disrupting guest VMs. With this new capability, you now have enhanced control—allowing you to define when updates are applied, prioritize which clusters receive them first, and set freeze periods aligned with your business policies to pause updates during critical timeframes.

While cloud automation software updates to Oracle-managed agents and tools on guest VMs are essential for accessing the latest Oracle cloud capabilities, many customers build scripts and workflows around current versions. To maintain operational continuity, you can now choose when updates are applied to your clusters. This release also gives you the flexibility to test updates on non-production clusters first, ensuring a smooth rollout with every new Oracle release. This enhancement allows customers to:

- **Schedule Updates:** Define specific times for the VM to check for and apply new updates.
- **Avoid Update During Freeze Period:** Prevent updates during periods of heavy database activity.
- **Reduce Risk:** Minimize the impact of updates on existing scripts and automation.

Related Topics

- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [VM Cloud Automation Software Update Management](#)

Classify and Control Access to Critical OCI REST APIs

- **Service:** [Database](#)
- **Release Date:** June 17, 2025

Oracle API Access Control enables you to manage and restrict access to REST APIs and Cloud Console operations exposed by various database cloud services. By designating specific APIs as privileged, you can enforce an approval workflow that requires authorization from a designated group within the tenancy before those APIs can be invoked.

Related Topics

- [Using the Console to Create API Access Control](#)
Use this procedure to create API Access Control for your ExaDB-D resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.

Oracle Key Vault (OKV) Integration with ExaDB-D to Manage Transparent Data Encryption (TDE) Keys

- **Service:** [Database](#)
- **Release Date:** May 06, 2025

Integrate your Oracle Key Vault (OKV) with Oracle Exadata Database Service on Dedicated Infrastructure and use customer-managed keys stored in Oracle Key Vault to encrypt your Oracle Transparent Data Encryption (TDE) Master Encryption Key (MEK).

You control the location of your OKV server, and it can be deployed in any location has network connectivity to your ExaDB-D instances.

Related Topics

- [Managing Encryption Keys on External Devices](#)
Learn how to store and manage database encryption keys.

Remove Provisioned Storage Servers

- **Services:** [Database](#)
- **Release Date:** March 19, 2025

Provisioned storage servers can be removed from a Cloud Exadata Infrastructure. A "provisioned" storage server refers to one that has had disk groups configured. Refer to the ASM required free disk group capacity and note that an ASM rebalancing operation will need to be completed after storage servers are removed.

Related Topics

- [Remove Storage Servers from a Multi-VM enabled Infrastructure](#)
Remove storage servers from an existing Multi-VM enabled Infrastructure

Enhancements to support concurrent Data Guard, Container Database (CDB), and Pluggable Database (PDB) Operations

With this enhancement, you can now perform concurrent operations on Container Databases (CDBs) and Pluggable Databases (PDBs) alongside Data Guard associations and actions. This improvement significantly enhances efficiency and flexibility in managing your Oracle databases. The supported concurrent operations include:

- Create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state.
- Creating or deleting a CDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.

- Creating or deleting a PDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Performing Data Guard actions (switchover, failover, and reinstate) while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Creating or deleting a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) on different databases within the same Oracle home, and vice versa.
- Creating or deleting a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) on different databases within the same Oracle home, and vice versa.
- Creating or deleting a CDB while concurrently creating or deleting a PDB of a different CDB within the same Oracle home, and vice versa.
- Creating or deleting a CDB concurrently on different databases within the same Oracle home.
- Creating or deleting a PDB concurrently on different databases within the same Oracle home.
- Performing Data Guard setup concurrently on different databases within the same Oracle home.
- Performing Data Guard actions (switchover, failover, and reinstate) concurrently on different databases within the same Oracle home.
- Creating or deleting a CDB/PDB, performing Data Guard setup, and performing Data Guard actions (switchover, failover, and reinstate) while simultaneously updating VM Cluster tags.

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To terminate a database](#)
- [To create pluggable database](#)
- [To delete a pluggable database](#)
- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To perform a database switchover](#)
You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- [To perform a database failover](#)
You can initiate a failover operation on a standby database that is a member of the Data Guard Group.
- [To reinstate a database](#)
After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary.
- [Concurrently Create or Delete Pluggable Databases \(PDBs\)](#)
You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while

a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.

- [Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel](#)

Dual Stack (IPv4 and IPv6) Network Support

- **Services:** [Database](#)
- **Release Date:** March 12, 2025

You can now provision VM clusters with IPv4/IPv6 dual-stack networking in Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D). This feature enables organizations to utilize both IPv4 and IPv6 concurrently, providing a cost-effective solution to manage IPv4 address scarcity, ensure compliance with regulatory requirements, and facilitate a smooth transition to IPv6 for future scalability and growth. The following scenarios are supported in this release:

- IPv4/IPv6 dual-stack support for new Exadata VM clusters on the client and backup subnets.
- Configure Data Guard to migrate Data from an IPv4-only Exadata VM cluster to a dual stack Exadata VM cluster.
- Configure an Application VIP with both IPv4 and IPv6 addresses.

Note

Minimum requirements for configuring a dual stack network:

- Exadata System Software 24.1.4
- Oracle Database and Oracle Grid Infrastructure:
 - 19c -> 19.26
 - 26ai -> 26ai (23.7)

Note

Oracle Exadata Database Service on Dedicated Infrastructure supports GUA, BYOIP and ULA IPv6 prefixes. While provisioning a VM Cluster a subnet should have only one IPv6 prefix. ExaDB-D does not support subnet with multiple IPv6 prefixes.

For more information, see [IPv6 Addresses](#).

For more information see, [Overview of VCNs and Subnets](#) and [Adding an IPv6 Prefix to a Subnet](#).

Related Topics

- [VCN and Subnets](#)
To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network and at least two subnets:
- [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#)
This section lists the security rules to use with Exadata Cloud Infrastructure.

- [General ingress rule 1: Allows SSH traffic from anywhere](#)
- [General ingress rule 2: Allows Path MTU Discovery fragmentation messages](#)
- [General ingress rule 3: Allows connectivity error messages within the VCN](#)
This rule enables the hosts in the VCN to receive connectivity error messages from each other.
- [General egress rule 1: Allows all egress traffic](#)
- [Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet](#)
The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.
- [Client ingress rule 2: Allows SQL*NET traffic from within the client subnet](#)
This rule is for SQL*NET traffic and is required in these cases:
- [Client egress rule 1: Allows all TCP traffic inside the client subnet](#)
This rule is for SQL*NET traffic as noted.
- [Client egress rule 2: Allows all egress traffic \(allows connections to the Oracle YUM repos\)](#)
Client egress rule 3 is important because it allows connections to the Oracle YUM repos.
- [About Application VIP](#)
Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.
- [To Attach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [To Detach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [Migrate from a Single Stack \(IPv4\) Exadata VM Cluster to a Dual Stack \(IPv4/IPv6\) Exadata VM Cluster with Data Guard Synchronization](#)

Long-Term Retention Backup (LTR)

- **Services:** [Database](#)
- **Release Date:** March 05, 2025

With Long-Term Retention Backup (LTR), you can store full backups for up to 10 years or a shorter duration, enabling you to search for and retrieve archived data to meet compliance, regulatory, or other business requirements. During this retention period, LTR backups can be restored to create a new database, a process referred to as "out-of-place restore."

Related Topics

- [Long-Term Retention Backup with Recovery Service](#)
Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.
- [Using the Console to restore a database](#)
You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.
- [To create an on-demand backup of a database](#)

- [To change the retention period of an LTR backup with Recovery Service](#)

Enhancements to Quarterly Exadata Infrastructure Maintenance Planning and Execution

- **Services:** [Database](#)
- **Release Date:** February 06, 2025

With this enhancement, you will have the flexibility to plan and apply quarterly infrastructure updates to fit smaller maintenance windows. You can choose to perform maintenance across all infrastructure components in a single window or split them into multiple smaller windows to align with your business needs. Based on your preferred time slots, Oracle automation will perform maintenance on specific infrastructure components across these maintenance windows to ensure all components have software updates applied to meet security and compliance guidelines.

Infrastructure components include:

- DB Servers
- Storage Servers

Related Topics

- [Configure Oracle-Managed Infrastructure Maintenance](#)
Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- [To create a Cloud Exadata infrastructure resource](#)
- [Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types](#)
The events in this section are emitted by the cloud Exadata infrastructure resource for Maintenance Events
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for Scheduling Policies](#)
- [Permissions and API operation details for Scheduling Windows](#)
- [Permissions and API operation details for Scheduling Plan](#)
- [Permissions and API operation details for Scheduled Action](#)
- [Permissions and API operation details for Execution Windows](#)
- [Permissions and API operation details for Execution Action](#)
- [Permissions Required for Each API Operation](#)

Granular Permissions for VM Cluster Update Operations

- **Services:** [Database](#)
- **Release Date:** January 23, 2025

This enhancement provides fine-grained control over VM cluster update operations.

You can now assign granular permissions for VM Cluster operations, such as enabling the DBA group to scale only memory or CPU, allowing the storage administrator group to manage

local/Exadata storage, or permitting the security administrator group to add SSH keys to a VM cluster.

For more information, see [Permissions and API operation details for VM Clusters](#).

Multiple Standby Databases

- **Services:** [Database](#)
- **Release Date:** January 22, 2025

This enhancement provides the ability to create and manage multiple local and remote standby databases linked to a primary database, providing flexibility for both data protection and disaster recovery. Local standby databases help minimize data loss, while remote standby databases safeguard against regional failures. This enhancement allows creation of up to 6 standby databases for a primary database.

In a typical Data Guard configuration, two standby databases are commonly used:

- **Local Standby:** A standby database in the same region as the production database is ideal for failover scenarios, offering zero data loss for local failures (such as database, cluster, or availability domain failures). Application failover impact is reduced in this case, as applications continue operating without the performance overhead of communicating with a remote region.
- **Remote (Cross-Region) Standby:** A remote standby database, located in a different region, is typically used for disaster recovery or to offload read-only query processing. A remote standby database setup ensures data protection against regional failures.

Some enterprise customers aim for symmetry after a site switch. For example, they may prefer to have both the primary and local standby in Region 1, and a remote standby with its own local standby in Region 2. In this configuration, there will be three standby databases. After a site switch, you will still have a primary database and a local standby readily available in the new primary region.

Additionally, customers can enhance their configurations by adding another standby database for testing purposes, leveraging our snapshot (read/write) standby capabilities.

Note

Creating a standby database associated with another standby database ("cascading standby") is not supported.

Related Topics

- [Use Oracle Data Guard with Exadata Cloud Infrastructure](#)
Learn to configure and manage Data Guard Groups in your VM cluster.
- [Using the API to manage Data Guard Group](#)
Use these API operations to manage a Data Guard Group on an Exadata Cloud Infrastructure instance:
- [Permissions and API operation details for Data Guard Group](#)
- [Data Guard Event Types](#)
Review the list of event types that Data Guard group and Data Guard Associations emit.

X11M System Support

- **Services:** [Database](#)
- **Release Date:** January 07, 2025

Note

This feature is rolled out only to the PHX, IAD, NRT, FRA, LHR, GRU, VCP, KIX, MAD, HYD, SIN, ORD, YYZ, CDG, JED, BOM, JNB, LIN, MXP, and TYO regions. It will be rolled out to other regions in phases.

Oracle Exadata Database Service on Dedicated Infrastructure has been extended to support Exadata Infrastructure X11M.

Related Topics

- [Scaling Options](#)
Introduction to the Scaling option on Exadata Cloud Infrastructure.
- [Scaling CPU cores within an Exadata Cloud Infrastructure instance](#)
If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores (ECPUs for X11M) symmetrically across all the nodes in the system as follows:
- [Scaling X6, X7 and X8 Exadata Cloud Infrastructure Instances Configurations](#)
Scaling an Exadata X6, X7, or X8 Exadata Cloud Infrastructure instance by moving to a shape with more capacity enables you meet the needs of your growing workload.
- [Exadata Shape Configuration](#)
This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.
- [Exadata X11M](#)
The values in the table that follows represent the specifications for an X11M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.
- [Exadata X9M](#)
The values in the table that follows represent the specifications for an X9M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.
- [Scaling Resources within an Exadata Infrastructure Instance](#)
If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.
- [Add Resources to a Multi-VM enabled Infrastructure](#)
Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure
- [Scaling CPU cores within an Exadata Cloud Infrastructure instance](#)
If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:
- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.

- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)
- [Estimating How Much Local Storage You Can Provision on Your VMs](#)
- [The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes](#)
- [To create a Cloud Exadata infrastructure resource](#)
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Microsoft Entra ID (MS-EI) Integration with Oracle Exadata Database Service on Dedicated Infrastructure

- **Services:** [Database](#)
- **Release Date:** November 15, 2024

Oracle Exadata Database Service on Dedicated Infrastructure now can accept Microsoft Entra ID (MS-EI) tokens to access the database. Azure users and applications can use the MS-EI token to access the database.

MS-EI integration will be available for databases patched to 19.17 and above. This feature is not available on Oracle Database release 21c.

For information on configuring MS-EI, configuring the database, and configuring the database client, see:

- [Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases](#) in the Oracle Database 19c Security Guide.
- [Authenticating and Authorizing Microsoft Azure Users for Oracle Databases](#) in the Oracle AI Database 26ai Security Guide.

Related Topics

- [Authenticating and Authorizing Microsoft Entra ID \(MS-EI\) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

Delegate Access Control for ExaDB-D

- **Services:** [Database](#)
- **Release Date:** November 01, 2024

Delegate Access Control service enables Oracle Exadata Database Service Dedicated customers to subscribe to VM and database maintenance and support services, delegate access to service providers, and control when those service providers can access VM and database resources. These service providers include Oracle global support, Oracle cloud support, and Oracle professional services.

Related Topics

- [Overview of Delegate Access Control](#)
- [Exadata Database Service Dedicated Security Controls](#)
- [Oracle Engineered System Deployment and Infrastructure Support](#)

Different RUs for Primary and Standby DB Homes in Data Guard Associations, Switchover, and Failover Operations

- **Services:** [Database](#)
- **Release Date:** October 05, 2024

In Oracle Data Guard configurations, it's common to have primary and standby databases in a synchronized state, including the same Release Updates (RUs) applied to both database homes. However, there are scenarios where you might need to allow different RUs between the primary and standby database homes, particularly during the patching cycle or for testing purposes.

Create Data Guard Associations:

- The primary and standby Oracle Homes must be of the same major database version.
- If the primary and standby Oracle Homes are running different RUs, a Data Guard association can be created only if the standby is on the same or higher RU than the primary database.
- The home for the standby can be a custom DSI or an Oracle image regardless of whether the primary is running on a custom DSI or Oracle image.

Switchover or Failover Database: The primary and standby Oracle Homes can be of any major database version or running different RUs.

Upgrade Database: The primary and standby Oracle Homes can be of different major database versions.

Patch Database: If the primary and standby Oracle Homes are running different RUs, the standby can be patched to a higher RU than the primary database.

Related Topics

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To perform a database switchover](#)
You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- [To perform a database failover](#)
You can initiate a failover operation on a standby database that is a member of the Data Guard Group.

Protect Sensitive Data With Oracle Cloud Infrastructure Zero Trust Packet Routing

- **Services:** [Database](#)

- **Release Date:** October 01, 2024

Protect your data against unauthorized access and exfiltration with Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR), a data-aware cloud security control.

Related Topics

- [Overview of Oracle Cloud Infrastructure Zero Trust Packet Routing \(ZPR\)](#)
Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) safeguards sensitive data from unauthorized access using intent-based security policies that you define for OCI resources assigned with security attributes.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To add security attributes to an Exadata VM Cluster](#)
To add security attributes to an Exadata VM Cluster, use this procedure.
- [To edit a security attribute](#)
To edit a security attribute of an Exadata VM Cluster, use this procedure.
- [To remove a security attribute](#)
To remove a security attribute of an Exadata VM Cluster, use this procedure.

Assign New Key Versions to KMS-Based Container Databases (CDBs) and Pluggable Databases (PDBs)

- **Services:** [Database](#)
- **Release Date:** September 30, 2024

Bring-Your-Own-Key (BYOK) allows users to import their own keys or key versions instead of having the vault service generate the keys internally. BYOK enables users to import keys into the OCI Vault service and utilize their imported keys for database encryption. Currently, when an ExaDB-D customer employs BYOK to create a new database (CDB), the customer's key is only used for the CDB, while a system-generated key version is assigned to the PDB. Users cannot assign specific key versions to the PDBs.

With this enhancement:

- Associate a key version with both CDB and PDB.
- Rotate keys independently at the CDB and PDB levels.

Note

In a Data Guard association, key rotation is only possible for the primary database, not the standby database.

Related Topics

- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or assign a new key version.

Enhancement to Backup and Restore from a Standby Database in a Data Guard Environment

- **Services:** [Database](#)
- **Release Date:** September 26, 2024

This enhancement allows you to use the Recovery Service or Object Storage to back up the standby databases.

Before getting started, note the following:

- You can schedule automatic backups and configure retention periods and backup schedules on the standby database.
- You can create a database in another availability domain (AD) within the same region or a different region from a backup of the standby database.
- Backups can be configured on the primary database, the standby database, or both. However, when configured, the primary and standby databases must share the same backup destination.
- For backups in the Recovery Service, the primary database can be restored or recovered from the backups of either the standby database or the primary database. Similarly, the standby database can be restored or recovered from the backups of either the primary database or the standby database.
- For backups in Object Storage, the primary and standby databases can only be restored or recovered using their respective backups.
- The backup destination of the primary database and standby database in a Data Guard association must be the same. For example, if the backup destination of the primary database is Recovery Service, then the backup destination of the standby database must also be Recovery Service. Similarly, if the backup destination of the standby database is Recovery Service, then the backup destination of the primary database must also be Recovery Service.
- The backup destination can be changed only after disabling the backup on either the primary or standby database in a Data Guard association. For example, if the backup destination of both the primary and standby databases is Object Storage and you want to change the backup destination of the primary database to Recovery Service, you must first disable the backup on the standby database.
- If automatic backups are configured on the primary database, upon switchover, the backups will continue on the new standby database.
- If automatic backups are configured on the standby database, upon failover, the backups will continue on the new primary database. However, the backups will be disabled on the new standby database.

Related Topics

- [To enable automatic backups on a standby database](#)
Learn to enable automatic backups on a standby database.

VM Cluster on a Single VM

- **Services:** [Database](#)

- **Release Date:** September 17, 2024

With this enhancement, you can deploy and run multiple databases in a VM cluster running on a single VM without requiring RAC licenses.

For more information, see [Overview of VM Cluster Node Subsetting](#).

Multicloud Oracle Database Backup Support

- **Services:** [Database](#)
- **Release Date:** August 29, 2024

Recovery Service supports multicloud Oracle Databases, such as Oracle Database@Azure, and provides the flexibility to store backups in the same cloud location where the source database resides.

Recovery Service creates protected databases and related backups in Oracle Cloud by default. You can optionally override this default behavior for your multicloud Oracle Databases such as Oracle Database@Azure.

If you enable the **Store backups in the same cloud provider as the database** option for a protection policy, then the policy-linked protected database and backups will be stored in the same cloud location where the Oracle Database is provisioned. For example, for Oracle Database@Azure, Recovery Service stores the associated protected database backups in Azure if you have selected the **Store backups in the same cloud provider as the database** option in the protection policy.

If you do not select the **Store backups in the same cloud provider as the database** for a protection policy, then the policy-linked protected database and backups will be stored in Oracle Cloud even if your Oracle Database is provisioned in a different cloud location.

Related Topics

- [About Configuring Protection Policies](#)
- [Creating a Protection Policy](#)
- [Enable Automatic Backups to Recovery Service](#)

Restore a Backup to Create a Database Across Regions

- **Services:** [Database](#)
- **Release Date:** August 22, 2024

With this enhancement, when the region is up, you can:

- use an existing backup and restore it to create a database (out-of-place restore) either within the same availability domain or in a different availability domain across regions, regardless of whether the backup was created with backup destination Object Storage or Autonomous Recovery Service
- restore a backup taken on either a database that was configured using host-based wallets (local wallet) or OCI Vault

Related Topics

- [Prerequisites for Oracle Database Autonomous Recovery Service Cross Region Restore \(Same Tenancy\)](#)

- [Prerequisites for Oracle Database, Object Storage Cross Region Restore \(Same Tenancy\)](#)
- [To create a database from a backup](#)
- [To create a database from the latest backup](#)

Cost and Usage Attribution for Pluggable Databases (PDBs)

- **Services:** [Database](#)
- **Release Date:** July 10, 2024

With this enhancement to the Cost Analysis feature of the OCI Cost Management Service, you can view the attributed usage and cost for all the PDBs in a VM Cluster. This data will be available on the cost analysis dashboard and the reports.

Related Topics

- [Cost and Usage Attribution for Pluggable Databases \(PDBs\)](#)

Use the Same Custom Software Image Across OCI Regions

- **Services:** [Database](#)
- **Release Date:** May 15, 2024

With this enhancement, you can use the software image created in one region in a different region while:

- updating a database software
- updating a Grid Infrastructure software
- provisioning a new Database Home
- provisioning a new database
- enabling a Data Guard association
- creating a database from a backup

Related Topics

- [To update database software using custom database software image](#)
Use the following instructions to update database software using a custom database software image.
- [To update Grid Infrastructure software using custom Grid Infrastructure software image](#)
Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To create a database from a backup](#)

Ability to Increase the Size of Guest VM Local File Systems

- **Services:** [Database](#)
- **Release Date:** May 09, 2024

Currently, you can only increase or decrease the size of the `/u02` file system in the Guest VM. Now, using the OCI Console or API, you can increase the size of additional local file systems such as `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home`.

Note

For X8M and later, a rolling restart is not required when expanding any of the Guest VM file systems. However, a rolling restart of each VM is required when the size of `/u02` is reduced.

Related Topics

- [Estimating How Much Local Storage You Can Provision on Your VMs](#)
- [The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes](#)
- [Scaling Local Storage](#)
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Create and Use Custom Software Images

- **Services:** [Database](#)
- **Release Date:** May 08, 2024

The ability to create a custom software image (Database and Grid Infrastructure) with all the required patches bundled together and certified in the customer environment will allow developers and database administrators to build an approved and reusable "gold image".

Related Topics

- [Manage Software Images](#)

Manage Serial Console Access to Oracle Exadata Database Service on Dedicated Infrastructure Systems

- **Services:** [Database](#)
- **Release Date:** May 07, 2024

Note

The serial console feature requires (at a minimum) Exadata System Software 23.1.13. Once the necessary software is installed via Quarterly Maintenance and a reboot of your VMs takes place, you will be able to use the new serial console feature.

You can create and delete serial console connections to your Oracle Exadata Database Service on Dedicated Infrastructure systems to diagnose and resolve VM guest operating system issues using an SSH connection in case standard SSH access to the VMs is not possible.

Requirements: Exadata System Software 23.1.13 is the minimum required version. Also, make sure to review all prerequisites stated below, including setting a password for either the `opc` or the `root` user. Failure to make necessary changes to meet these requirements in advance will result in the inability to urgently connect to the serial console when the need arises when the VM is not otherwise accessible.

Related Topics

- [Troubleshooting Virtual Machines Using Console Connections](#)
You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.
- [Create the Virtual Machine Serial Console Connection](#)
Before you can make a local connection to the serial console, you need to create the virtual machine console connection.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for DB Nodes](#)
- [Permissions and API operation details for DB Node Console Connection](#)
- [Serial Console Connection Event Types](#)
Review the list of event types that serial console connection emits.
- [Viewing Audit Log Events](#)
Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.
- [Permissions Required for Each API Operation](#)

Oracle AI Database 26ai on Exadata Database Service on Dedicated Infrastructure

- **Services:** [Database](#)
- **Release Date:** May 02, 2024

Oracle AI Database 26ai is a regular production release available on Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D). With this release, you can perform all the lifecycle operations on the 26ai databases.

Related Topics

- [To upgrade the Oracle Grid Infrastructure of a cloud VM cluster](#)
Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Enable Unified Auditing While Creating a Database Home

- **Services:** [Database](#)
- **Release Date:** April 30, 2024

With this enhancement, you can enable Unified Auditing during the creation of a database home, a feature available since Oracle Database version 12.1.

- **For Oracle Database versions lower than 12.1:** You cannot use the Unified Auditing framework and should instead use the Traditional Audit, the legacy Oracle Database audit framework.
- **For Oracle Database versions 12.1 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.1 or higher but lower than version 26ai, the **Unified Auditing** check box is not selected by default. However, it is selected by default for Oracle AI Database 26ai.

Note

You cannot disable Unified Auditing after provisioning the Database Home.

Related Topics

- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Provision a VM Cluster with Either an OL7 or OL8-Based Image

- **Services:** [Database](#)
- **Release Date:** February 28, 2024

With this enhancement, you can provision a VM cluster with either an OL7-based image or an OL8-based one if the infrastructure is X9 or prior.

Related Topics

- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Enhancement to the OCI Console to Remove Database and Storage Servers

- **Services:** [Database](#)
- **Release Date:** January 29, 2024

With this enhancement, you can:

- Scale down your Exadata Infrastructure resources by changing your server count to a lower value than the current assignment.
- Scaling down to a lower count is supported for both DB and Storage Servers.
- Database servers will be removed if there are no VMs running on them.

Note

You will not be able to choose the DB Server to remove. This functionality will automatically remove Database Servers in which there are no VMs.

- Storage server will be removed if the server has not been used to expand Exadata Infrastructure storage.
- Remove a VM from a provisioned VM cluster in a non multi-VM-enabled infrastructure. The procedure is similar to terminating a VM from a VM Cluster in a multi-VM-enabled infrastructure.

The 'Add Capacity' step runs as part of the storage server scale-up workflow, creates disk groups, and rebalances data across all storage servers. For more information, see *Scale VM Resources in Multi VM Enabled Infrastructure*.

Related Topics

- [Scale VM Resources in Multi VM Enabled Infrastructure](#)
Increase or decrease the OCPUs (ECPUs for X11M), memory, storage, or local disk size (/u02) storage available to a VM cluster

Enable Data Guard Across Different VCNs or Compartments in the Same OCI Region

- **Services:** [Database](#)
- **Release Date:** January 16, 2024

With this enhancement, you can enable Data Guard if the Exadata Cloud Infrastructure is in different VCNs or compartments in the same OCI region.

Enhancement to Pluggable Database (PDB) Management

- **Services:** [Database](#)
- **Release Date:** October 18, 2023

With this enhancement, you can restore, refresh, and relocate a Pluggable Database (PDB).

Related Topics

- [Create and Manage Exadata Pluggable Databases](#)
You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.
- [To create pluggable database](#)
- [To relocate a pluggable database](#)
- [Cloning an Exadata Pluggable Database](#)
You can create local, remote, and refreshable clones.
- [Restoring an Exadata Pluggable Database](#)
You can perform in-place and out of place restore of an Exadata pluggable database.
- [Permissions Required for Each API Operation](#)
- [Permissions and API operation details for Pluggable Databases \(PDBs\)](#)
- [Pluggable Database Event Types](#)
These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

Manage Administrator (SYS User) and TDE Wallet Passwords

- **Services:** [Database](#)
- **Release Date:** September 28, 2023

With this enhancement, you can manage the administrator and TDE wallet passwords.

Note

Changing a TDE wallet password for Oracle Key Vault (OKV) or OCI Vault Key management-enabled databases is currently not supported.

Related Topics

- [To manage SYS user and TDE Wallet passwords](#)
Learn to manage administrator (SYS user) and TDE wallet passwords.

Backup and Restore from a Standby Database with OCI Object Storage in a Data Guard Environment

- **Services:** [Database](#)
- **Release Date:** August 24, 2023

This enhancement:

- Enables customers to offload backups to the standby database with OCI Object Storage in a Data Guard environment thereby freeing up resources in the production database environment.

- Allows customers to schedule automatic backups on the standby database in a Data Guard environment and configure retention period and backup schedules.
- Enables customers to create a database in another Availability Domain (AD) within the same region from a backup of the standby database.
- Allows customers to restore and recover a standby database using a backup of the standby database.
- Provides the flexibility to take backups only on the primary database, only on the standby database, or both.
- Allows customers to create a manual full backup of a standby database.
- Enables customers to enable or disable backup on the standby database only if the backup destination of the primary database is Object Storage.

① Note

- You cannot change the backup destination of the primary database to Autonomous Recovery Service if the backup destination of the primary and standby databases is Object Storage. To change the backup destination of the primary database to Autonomous Recovery Service, first, disable backup on the standby database.
- You cannot use standby-side backups to perform restore/recover operations on the primary database.
- Switchover scenarios:
 - If automatic backups was configured on the primary with backup destination of Object Storage, upon switchover, the backups will continue on the new standby database
 - If automatic backups was configured on the primary with backup destination of Autonomous Recovery Service, upon switchover, backup and restore will be disabled on the new standby database
 - If automatic backups was configured on the standby with backup destination of Object Storage, upon switchover, the backups will continue on the new primary database
- Failover scenarios:
 - If automatic backups was configured on the primary with backup destination of Object Storage or Autonomous Recovery Service, upon failover the backups be disabled on the new *Disabled Standby* database
 - If automatic backups was configured on the standby with backup destination of Object Storage, upon failover the backups will continue on the new primary database

Related Topics

- [To enable automatic backups on a standby database](#)
Learn to enable automatic backups on a standby database.
- [To restore a database](#)
- [To create a database from the latest backup](#)

Cancel a Running Full or Incremental Backup

- **Services:** [Database](#)
- **Release Date:** August 21, 2023

You now have the ability to cancel an ongoing backup, allowing you to free up system resources. You will no longer have to call the operations team to have this backup job canceled.

As part of the Create Database workflow and independently (after the database has been created), you may enable Automatic Backup and select the desired backup destination. Depending on the backup destination selected, you may have one or more full backups and several incremental backups. Once any of these backups have started, you will not have the option to cancel that backup midway.

This feature allows you to cancel any running backup (automatic or standalone) from the OCI console or via OCI API.

You can also:

- Cancel a manual backup, which is triggered when you click the **Create backup** button
Note: All manual backups are full backups.
- Delete a canceled manual backup

Related Topics

- [Using the Console to Manage Backups](#)
- [To view backup status](#)
- [To cancel a backup](#)

Autonomous Recovery Service as the Default Backup Destination

- **Services:** [Database](#)
- **Release Date:** August 17, 2023

This Console enhancement sets Autonomous Recovery Service as the default backup destination for automatic backups in all regions and includes default limits automatically without having to request them.

For more information about Service Limits, Quotas, and Usage, see [Autonomous Recovery Service limits](#).

Related Topics

- [To configure automatic backups for a database](#)

Exadata Fleet Update

- **Services:** [Database](#)
- **Release Date:** August 02, 2023

Exadata Fleet Update simplifies, standardizes, and enhances the Oracle Database and Grid Infrastructure patching experience. Exadata Fleet Update achieves this by grouping components based on the customers' business needs into collections that can be patched as one entity within a given maintenance cycle.

Exadata Fleet Update brings this patching engine to OCI as a native cloud service, accessible from the OCI Console, OCI API, and via the OCI CLI.

Exadata Fleet Update is available free of charge on Oracle's Exadata Database Service including Cloud@Customer (ExaDB-C@C) and Exadata Database Service on Dedicated Infrastructure (ExaDB-D).

For more information, see:

- [Exadata Fleet Update Overview](#)
- [Exadata Fleet Update service AP](#)

Update Guest VM (domU) Operating System to Oracle Linux 8

- **Services:** [Database](#)
- **Release Date:** August 01, 2023

Update the Guest VM operating system to Oracle Linux 8 using the Console or API. This enhancement is limited to Exadata X7, X8M, and X9M systems.

Related Topics

- [Updating an Exadata Cloud VM Cluster Operating System](#)
Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- [Supported Software Versions and Update Restrictions](#)
Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- [Updating the Operating System using the Console](#)
- [Add a VM to a VM Cluster](#)
Add a Virtual Machine to a VM Cluster

Use a Backup to Create a Database Across Availability Domains within the Same Region

- **Services:** [Database](#)
- **Release Date:** July 26, 2023

With this enhancement, when AD is up, you can:

- use an existing backup and restore it to create a database (out-of-place restore) either within the same availability domain or in a different availability domain within the same region, regardless of whether the backup was created with backup destination Object Storage or Autonomous Recovery Service
- restore a backup taken on either a database that was configured using host-based wallets (local wallet) or OCI Vault

Related Topics

- [To create a database from a backup](#)
- [To create a database from the latest backup](#)

Interim Software Updates

- **Services:** [Database](#)
- **Release Date:** June 07, 2023

This feature enables cloud-only customers to download one-off patches from the OCI console and API. There is no option to apply the downloaded patch via console and API. To apply these patches, customers must log in to their VM and run the patch apply utility.

Note

To be able to download interim software update, you should at least have an ExaDB-D infrastructure provisioned.

Downloading one-off patches does not replace Database Software Image (DSI) creation. Customers must continue to use Database Software Images (DSI) to build and deploy their customized images.

Related Topics

- [Interim Software Updates](#)
For authorized environments, learn how to download interim software updates.
- [Permissions Required for Each API Operation](#)
- [Permissions and API operation details for Interim Software Updates](#)
- [Interim Software Updates Event Types](#)
These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.

Enhanced Controls to Configure Automatic Full (L0) and Incremental (L1) Backups

- **Services:** [Database](#)
- **Release Date:** May 17, 2023

Enabling Automatic Backup during the Create Database workflow or as a separate step afterward starts the first full backup ("initial L0") immediately.

Similarly, for subsequent full backups (future L0) and daily incremental backups (L1), you can specify a time window but cannot change the day of the week when these backups must begin.

Future L0 and L1 backups will begin during the 2-hour scheduling window that the user selects for the database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting at an even-numbered hour. For example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM. Backup jobs do not necessarily complete within the scheduled window. If you do not specify a window, the default

6-hour backup window of 00:00 to 06:00 is chosen. In this case, the time zone corresponds to the region of the Exadata Cloud infrastructure instance.

Here are the current defaults for the backup destinations, Object Storage Service, and Autonomous Recovery Service:

- **Initial full L0 backup:** Immediate
- **Subsequent full L0 backups:** Every Sunday
- **Daily incremental L1 backups:** Every Monday - Saturday

With these enhanced controls, you can:

1. Aside from configuring the initial L0 backup to start immediately, you can also specify whether you want the initial L0 backup to start immediately or according to the L0 schedule.
2. Choose a time window for the future full backups to start.
3. Choose a time window for the incremental backups to start, which can be different from the time window for the L0 backups.
The time windows will remain the same, 2-hour scheduling windows and the default 6-hour window.

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To configure automatic backups for a database](#)

Configure Oracle Database Autonomous Recovery Service as a Backup Destination

- **Services:** [Database](#)
- **Release Date:** April 12, 2023

Oracle Database Autonomous Recovery Service provides an optimized policy-driven automatic backup and recovery system for the Exadata Database on Dedicated Infrastructure. It also offers a real-time data protection feature that enables protected databases with zero data loss recovery in the event of a database failure. Since Real-time data protection is an extra cost option, you can choose to enable or disable it.

Related Topics

- [Overview of Oracle Database Autonomous Recovery Service](#)
- [Network Requirements for Oracle Database Autonomous Recovery Service](#)
Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).
- [Create a Service Gateway to Object Storage](#)
In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.
- [Creating Protection Policies](#)
Recovery Service uses protection policies to control database backup retention in Oracle Cloud.

- [Managing Exadata Database Backups](#)
Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.
- [Managed Backup Types and Usage Information](#)
There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.
- [Prerequisites for Backups on Exadata Cloud Infrastructure](#)
- [To configure automatic backups for a database](#)
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To view details of a Protected Database](#)
To view the details of a Protected Database, use this procedure.
- [To designate Autonomous Recovery Service as a Backup Destination for an Existing Database](#)
To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.
- [To terminate a database](#)
- [To Terminate a VM Cluster](#)
- [Recovering an Exadata Database from Backup Destination](#)
This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.
- [Using the Console to restore a database](#)
You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.
- [To restore a database](#)
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Application VIP Support

- **Services:** [Database](#)
- **Release Date:** April 12, 2023

The VM Cluster now supports attaching and detaching Application Virtual IP Addresses.

Related Topics

- [About Application VIP](#)
Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.
- [To Attach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [To Detach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for Application VIPs](#)
- [Permissions Required for Each API Operation](#)

- [Application VIP Event Types](#)
These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Monthly ExaDB-D Infrastructure Security Maintenance

- **Services:** [Database](#)
- **Release Date:** March 01, 2023

Security maintenance, performed alongside the quarterly maintenance, is executed once a month and includes fixes for vulnerabilities with CVSS scores greater than or equal to 7.

Related Topics

- [About Oracle-managed Exadata Cloud Infrastructure Maintenance](#)
Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.
- [Overview of Monthly Security Maintenance](#)
Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- [To view or edit quarterly maintenance preferences](#)
- [To view the maintenance history of an Exadata Cloud Infrastructure resource](#)
This task describes how to view the maintenance history for a cloud Exadata infrastructure resource.

Identity and Access Management (IAM) Integration with Oracle Exadata Database Service on Dedicated Infrastructure

- **Services:** [Database](#)
- **Release Date:** January 24, 2023

With the latest Release Update, you can now configure the database in a virtual machine cluster to use Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) authentication and authorization to allow IAM users to access the database with IAM credentials.

As of this release, IAM authentication and authorization are supported under the following conditions:

- **Supported Environments:** Available on newly provisioned databases and on existing databases patched to 19.17. This feature is not supported on Oracle Database 21c.
- **Unsupported Configurations:** IAM authentication and authorization cannot be used with databases configured with Data Guard.

Related Topics

- [Connect Identity and Access Management \(IAM\) Users to Oracle Exadata Database Service on Dedicated Infrastructure](#)
You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.

Exadata Cloud Infrastructure: Private DNS

- **Services:** [Database](#)
- **Release Date:** January 18, 2023

Allow users to choose the private view and private zone while provisioning a new VM cluster for ExaCS. All the underlying resources of the VCN, including those of ExaDB-D, should be created in the same private zone. The private zones can be associated with subnets inside the VCN. This configuration cannot be changed later.

Private DNS resolver is going to resolve queries in customer VCN and queries coming from the on-premise networks. Ability to provide the DNS address and seed that into DB resources using conditional forwarding, is provided by the private DNS feature. With a private resolver, customers can resolve the A-record across different VCNs (with VCN peering- local/remote).

Related Topics

- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To view details about private DNS configuration](#)

Enhanced Infrastructure Maintenance Controls

- **Services:** [Database](#)
- **Release Date:** January 18, 2023

The Exadata Cloud Infrastructure Oracle-managed infrastructure maintenance now allows greater control and visibility including:

- Choice of rolling and non-rolling maintenance methods.
- Ability to perform custom actions before maintenance on each database server by having the automated maintenance wait before shutting down VMs until the maintenance is resumed or the configured timeout is reached.
- Visibility into the database server update order.
- Granular tracking of the maintenance progress at a component level.

Related Topics

- [To create a Cloud Exadata infrastructure resource](#)
- [Configure Oracle-Managed Infrastructure Maintenance](#)
Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- [To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure](#)
Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.

- [Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls](#)
Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.
- [Oracle Exadata Database Service on Dedicated Infrastructure Event Types](#)
The events in this section are emitted by the cloud Exadata infrastructure resource

Database Management Support for Pluggable Databases in Oracle Exadata Database Service on Dedicated Infrastructure

- **Services:** [Database](#)
- **Release Date:** January 11, 2023

You can now enable Database Management for Pluggable Databases (PDBs) on Oracle Exadata Database Service on Dedicated Infrastructure, and use Database Management features for monitoring, performance management, and tuning.

Related Topics

- [Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases](#)
Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.
- [Permissions and API operation details for Databases \(CDBs\)](#)
- [Permissions and API operation details for Pluggable Databases \(PDBs\)](#)
- [Permissions Required for Each API Operation](#)

Microsoft Azure Active Directory Integration with Oracle Cloud Infrastructure Databases

- **Services:** [Database](#)
- **Release date:** January 10, 2023

Oracle Exadata Database on Dedicated Infrastructure now can accept Azure AD tokens to access the database. Azure AD users can access the database directly using their Azure AD token, and applications can use their service tokens to access the database.

Azure AD integration will be available for databases patched to 19.17 and above. This feature is not available on Oracle Database release 21c.

Related Topics

- [Authenticating and Authorizing Microsoft Entra ID \(MS-EI\) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

Create and Manage Multiple Virtual Machines per Exadata System (MultiVM) and VM Cluster Node Subsetting

- **Services:** [Database](#)
- **Release Date:** Starting November 9, 2022 (the release date varies by region)

NOT_SUPPORTED

Slice Exadata resources into multiple virtual machines. Define up to 8 multiple virtual machine (VM) clusters on an Oracle Exadata Database Service on Dedicated Infrastructure, and specify how the overall system resources are allocated to them.

VM Cluster Node Subsetting enables you to allocate a subset of database servers to new and existing VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

Note

For existing Exadata Infrastructure, MultiVM will be enabled as part of your next scheduled maintenance run after MultiVM migration on December 20, 2022. All newly provisioned Exadata Infrastructure after the release of MVM on November 15, 2022, will have MultiVM enabled.

Related Topics

- [To create an ASM cloud VM cluster](#)
- [Scale VM Resources in Multi VM Enabled Infrastructure](#)
- [Exadata Cloud Infrastructure VM Cluster Event Types](#)
- [Overview of VM Cluster Node Subsetting](#)
- [Add a VM to a VM Cluster](#)
- [To View a List of DB Servers on an Exadata Infrastructure](#)
- [To Terminate a VM Cluster](#)
- [Terminate or Remove a VM from a VM Cluster](#)
- [Permissions Required for Each API Operation](#)
- [Permissions and API operation details for DB Servers](#)
- [VM Node Subsetting Event Types](#)
- [Adding a VM to a VM Cluster Fails](#)
- [CPU Offline Scaling Fails](#)

VM Cluster and Database Health and Performance Metrics in the OCI Console

- **Services:** [Database](#)
- **Release Date:** October 7, 2022

With this release Oracle will provide health metrics for databases and VM clusters in the Oracle Cloud Infrastructure (OCI) console.

Note

When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.

Potentially one hour of metrics will be lost between network restore and the first metric posted.

Related Topics

- [Monitor Metrics for VM Cluster Resources](#)
- [Prerequisites for Using Metrics](#)
- [View Metrics for VM Cluster](#)
- [View Metrics for a Database](#)
- [View Metrics for VM Clusters in a Compartment](#)
- [View Metrics for Databases in a Compartment](#)
- [Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service](#)

This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.
- [Metrics for Exadata Cloud Infrastructure in the Database Management Service](#)

Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.

Oracle Standard Tagging for Resources on Oracle Exadata Database Service on Dedicated Infrastructure

- **Services:** [Database](#)
- **Release Date:** September 15, 2022

Exadata Cloud Infrastructure resources can now be tagged using Oracle Standard tags according to your organizational scheme. By tagging resources, you can group them, manage costs, and gain insight into how they are being used.

Related Topics

- [Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources](#)

Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.

Automatic Diagnostic Collection

- **Services:** [Database](#)
- **Release Date:** August 31, 2022

This feature extends the Database Service Events feature implementation that enables you to get notified about health issues with your Oracle Databases or other components on the Guest VM. With this enhancement, you can allow:

- Oracle to proactively collect detailed health metrics for diagnosis and issue resolution
- Oracle to reactively collect Incident logs and trace files on demand for a deeper diagnosis and issue resolution

Collecting Guest VM events, health metrics, incident logs, and trace files, will help Oracle to enhance service operations as well as provide proactive support by early detection and correlation.

Related Topics

- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- [Incident Logs and Trace Files](#)
This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- [Health Metrics](#)
Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- [Database Service Events](#)
The Database Service emits events, which are structured messages that indicate changes in resources.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Exadata Database on Dedicated Infrastructure: Key Management Service for Cross Region Data Guard

- **Services:** [Database](#)
- **Release Date:** Aug 8, 2022

You can now have the encryption keys used for the primary and standby databases to be available in the primary and standby regions respectively so that it provides protection against a single point of failure for the OCI Vault key. This is possible if the keys are in OCI Virtual Private Vault. So, cross-region Data Guard can be set up between two databases if their keys are residing in a virtual private vault (VPV) and are managed by the OCI Vault service.

Related Topics

- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.

- [Customer-Managed Keys in Exadata Cloud Infrastructure](#)
Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- [Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure](#)
An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.

Concurrently Create or Terminate Oracle Databases in a VM Cluster

- **Services:** [Database](#)
- **Release Date:** August 3, 2022

With this enhancement, you can now concurrently create or terminate Oracle databases even if the VM cluster is in the Updating state.

- The number of databases that can be created on a cluster depends on the available memory on the VMs. For each database, by default, 12.6 GB (7.6 GB for SGA and 5 GB for PGA) is allocated if the VM has greater than 60 GB of memory. If the VM has less than or equal to 60 GB, then 6.3 GB (3.8 GB for SGA and 2.5 GB for PGA) is allocated. Also, Grid Infrastructure and ASM consume some memory, approximately 2 to 4 GB.
- A database that is being created cannot be terminated. You can, however, terminate other databases in the VM Cluster.

VM Guest Exadata OS Image Major Version Update

- **Services:** [Database](#)
- **Release Date:** July 11, 2022

In addition to performing minor version updates to the Exadata VM Cluster images, you can update to a new major version if the currently installed version is 19.2 or higher. For example, if the VM cluster is on version 20, then you can update it to version 21.

Related Topics

- [Supported Software Versions and Update Restrictions](#)
Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):

Database Service Events capability for Exadata Database

- **Services:** [Database](#)
- **Release Date:** July 8, 2022

This feature allows customers to use OCI Console or API/CLI/SDK/Terraform to receive event notifications about health issues with your Oracle Databases or other components on the Guest VM.

Customers currently have basic lifecycle management events like backup begin, backup end, patching begin, etc. We are extending that capability to include a comprehensive set of Database Service events to help customers troubleshoot issues.

Database Service Events monitor guest VM operations and conditions and generate diagnostic notifications for customers by leveraging the existing OCI Events service and Notification mechanisms in their tenancy. Customers can then create topics and subscribe to these topics through email, functions, streams, etc. For more information about using the Events Notification Service, see *Notifications Overview*

Key Customer Benefits

- Ability to receive notifications for Guest VM operations via an opt-in mechanism.
- Allows customers to proactively address issues before they may become serious.

OCI Console Experience

Customers can navigate to the VM Cluster details page from the OCI console menu by selecting Oracle Database → **Oracle Exadata Database Service on Dedicated Infrastructure** → a specific VM Cluster to enable Diagnostics Notification for a VM Cluster.

Related Topics

- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.
- [Receive Notifications about Database Service Events](#)
Subscribe to the Database Service Events and get notified.
- [Database Service Event Types](#)
Review the list of event types that the Database Service emits.
- [Temporarily Restrict Automatic Diagnostic Collections for Specific Events](#)
Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [Notifications Overview](#)

Exadata Database on Dedicated Infrastructure: 'Create database from backup' now available for databases using customer-managed encryption

- **Services:** [Database](#)
- **Release Date:** June 29, 2022

Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D): Now allows you to create a database from backup, when the backup is of the database using customer-managed encryption. This is in addition to the existing ability to create a database from backup, when the backup is of the database using oracle-managed encryption

Related Topics

- [To create a database from a backup](#)

Support for DB Home Minor Version Selection (N-3)

- **Services:** [Database](#)
- **Release Date:** May 23, 2022

Provision a DB Home using a major version and RU version of your choice.

While provisioning, if you opt to use **Oracle Provided Database Software Images** as the image type, then you can use the **Display all available versions** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Oracle Cloud Infrastructure Operations Insights Support for Oracle Cloud Databases

- **Services:** [Database](#)
- **Release Date:** March 22, 2022

Operations Insights now allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases.

Related Topics

- [Enabling Database Cloud Service Databases](#)
- [Oracle Cloud Infrastructure Operations Insights](#)
Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Exadata Cloud Infrastructure.

Specify the Same SID for Primary and Standby Databases in Data Guard Association

- **Services:** [Database](#)
- **Release Date:** March 14, 2022

The same SID prefix used for the primary database can now also be used for the standby database when creating a Data Guard Association.

Related Topics

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.

Exadata Cloud Infrastructure: Pluggable database lifecycle support

- **Services:** [Database](#)
- **Release Date:** Jan. 12, 2022

You can now create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the OCI console and APIs. See [Create and Manage Exadata Pluggable Databases](#) for details.

Exadata Cloud Infrastructure: Set DB_UNIQUE_NAME and Oracle SID prefix during database creation

- **Services:** [Database](#)
- **Release Date:** January 12, 2022

You can now specify the `DB_UNIQUE_NAME` value and the Oracle SID prefix when creating a new Oracle Database in Exadata Cloud Infrastructure. You can also set these values when creating a standby database in an Oracle Data Guard association. See the following topics for instructions:

[To create a database in an existing VM Cluster](#)

[To Enable Data Guard on an Exadata Cloud Infrastructure System](#)

Elastic Expansion

- **Services:** [Database](#)
- **Release Date:** December 23, 2021

With elastic provisioning and expansion, you can dynamically increase your CPU and storage capacity to meet your growing workload requirements.

Expand the infrastructure capacity on-demand by scaling up the infrastructure with additional database or storage servers without being constrained by the standard supported shapes. You can allocate CPU and storage capacity available on X8M and X9M servers up to the system

limits when you provision new VM Clusters on the infrastructure, or to already deployed VM Clusters without disrupting the current running workloads.

Related Topics

- [Scaling CPU cores within an Exadata Cloud Infrastructure instance](#)
If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores (ECPUs for X11M) symmetrically across all the nodes in the system as follows:
- [Exadata X9M](#)
The values in the table that follows represent the specifications for an X9M Exadata Cloud Infrastructure with 2 database and 3 storage servers that has not been expanded.
- [Exadata X8M](#)
The values in the table that follows represent the specifications for an X8M cloud instance with 2 database and 3 storage servers that has not been expanded.
- [The Cloud Exadata Infrastructure Resource](#)
The infrastructure resource is the top-level (parent) resource.
- [Resources to Be Created](#)
- [To create a Cloud Exadata infrastructure resource](#)
- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)
- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.
- [Permissions Required for Each API Operation](#)
- [Permissions and API operation details for Cloud Exadata Infrastructures](#)

Oracle Database: Encryption key options updated for Exadata Cloud Infrastructure databases

- **Services:** [Database](#)
- **Release Date:** Nov. 17, 2021

When provisioning a new Oracle Database, the TDE wallet password parameter is now optional, and it is not used if you configure an Exadata Cloud Infrastructure database to use customer-managed keys with the OCI Vault service.

For more information on creating Oracle Databases, see [Creating and Managing Exadata Databases](#).

For information on using the Vault service to store and manage encryption keys and other secrets used by OCI resources, see [Vault service](#).

Performance Hub Exadata Tab

- **Services:** [Database](#)
- **Release Date:** Nov. 16, 2021

The Exadata tab provides a unified view of Oracle Exadata hard disk and flash performance statistics with deep insight into the health and performance of all components such as the Oracle databases, Oracle Exadata storage cells, and Automatic Storage Management (ASM).

It is available for Exadata Cloud deployments and external databases that use Exadata infrastructure. For more information, see [Using Performance Hub to Analyze Database Performance](#).

Exadata Cloud Infrastructure: custom SCAN listener port for VM cluster

- **Services:** [Database](#)
- **Release Date:** Nov. 3, 2021

You can now specify a custom SCAN listener port for your Exadata cloud VM cluster. See [SCAN Listener Port Setting](#) for more information.

Performance Hub & metrics available for databases running in Exadata Cloud Infrastructure

- **Services:** [Database](#)
- **Release Date:** Sept. 9, 2021

You can now use the Performance Hub tool and view metrics on cloud databases that run on the following systems: Exadata Cloud Infrastructure instances. This feature provides additional monitoring and management functions to these databases. For more information, see [Analyzing Exadata Cloud Service Database Performance](#).

Maintenance advisory contacts for Exadata infrastructure

- **Services:** [Database](#)
- **Release Date:** July 28, 2021

You can choose to specify up to 10 valid email addresses to which Oracle sends maintenance notifications when updates are made to an Exadata infrastructure. The email addresses you specify are used only for service-related operational issues. See [Oracle-Managed Infrastructure Maintenance Updates](#) for more information.

Data Guard protection mode enhancement for Exadata Cloud Infrastructure instances

- **Services:** [Database](#)
- **Release Date:** July 21, 2021

You can now specify the Data Guard protection mode for Exadata Cloud Infrastructure instances. For more information, see [Use Oracle Data Guard with Exadata Cloud Infrastructure](#).

Exadata Cloud Infrastructure: Non-rolling infrastructure patching option now available

- **Services:** [Database](#)
- **Release Date:** May 26, 2021

You can now configure Exadata infrastructure patching to take place in a non-rolling fashion across the database nodes. This option allows you to reduce the total time your system is undergoing quarterly maintenance, but does involve system downtime. See [Oracle-Managed Infrastructure Maintenance Updates](#) for more information.

Customer-managed encryption keys available with Oracle Data Guard-enabled databases in Exadata Cloud Infrastructure

- **Services:** [Database](#)
- **Release Date:** April 16, 2021

Customer-managed keys for Exadata Cloud Infrastructure is an encryption key management service that enables you to encrypt your data using encryption keys that you control. You can use customer-managed keys on databases you provision in Exadata Cloud Infrastructure that are enabled with Oracle Data Guard.

ExaDB-D OS/DomU Patching Project

- **Services:** [Database](#)
- **Release Date:** Feb. 5, 2021

DomU OS Patching is a feature which enables ExaDB-D customers to upgrade the Exadata OS image on their domU nodes in an automated manner from their OCI console and APIs. The following information explains about a recent change in the feature that could not be added to the docs in time for release, but will be added soon.

Rollback required if the patch fails.

On multi-node systems, if one of the nodes fails the patch, you must roll back all nodes to get them all on the same version. Then run precheck, fix any problems, and run the patch again.

Example: If you run precheck on Monday and all nodes pass, but do not apply the patch until Wednesday, it is possible that one or more of the nodes may fail the patch because of changes on the nodes or a maintenance conflict.

To prevent this from happening, Oracle recommends that you run precheck right before applying the patch.

For more information, see [Updating an Exadata Cloud Service VM Cluster Operating System](#).

Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure

- **Services:** [Database](#)
- **Release Date:** Dec. 18, 2020

Oracle Cloud Infrastructure Vault service integration with Exadata Cloud Infrastructure enables database encryption with customer-managed keys. For more information, see [Customer-Managed Keys in Exadata Cloud Service](#).

Exadata Cloud Infrastructure: Oracle Database 19c upgrade feature available

- **Services:** [Database](#)
- **Release Date:** Dec. 3, 2020

You can now upgrade Exadata Cloud Infrastructure databases to Oracle Database version 19c using the Oracle Cloud Infrastructure Console or API. For information and instructions, see [Upgrading Exadata Databases](#).

Create custom database software images for Exadata Cloud Infrastructure instances

- **Services:** [Database](#)
- **Release Date:** Nov. 10, 2020

You can now create custom Oracle Database software images to use for provisioning Database Homes and patching databases in Exadata Cloud Infrastructure instances. For more information, see [Oracle Database Software Images](#).

Exadata Cloud Infrastructure: grid infrastructure upgrade for cloud VM clusters

- **Services:** [Database](#)
- **Release Date:** Oct. 15, 2020

You can now upgrade the grid infrastructure (GI) of an Exadata Cloud Infrastructure VM cluster using the Console. For more information, see [Upgrading Exadata Grid Infrastructure](#).

Exadata Cloud Infrastructure: the flexible X8M shape now available

- **Services:** [Database](#)
- **Release Date:** Oct. 15, 2020

You can now provision an Exadata Cloud Infrastructure instance using the flexible X8M shape. This shape allows you to expand your system after provisioning, as your databases grow and you need more storage servers, compute servers, or both. For more information, see [Overview of X8M Scalable Exadata Infrastructure](#).

Exadata Cloud Infrastructure: use an existing Database Home when setting up a Data Guard standby database

- **Services:** [Database](#)
- **Release Date:** Oct. 15, 2020

You can now choose to use an existing Database Home when setting up a Data Guard standby database in your Exadata Cloud Infrastructure instance. See [Using Oracle Data Guard with Exadata Cloud Service Instances](#) for information on setting up Data Guard for your Exadata databases.

3

Preparing for Exadata Cloud Infrastructure

Review OCI as well as the site, network and storage requirements to prepare and deploy Exadata Cloud Infrastructure in your data center.

- [Oracle Cloud Infrastructure \(OCI\) Requirements for Oracle Exadata Database Service on Dedicated Infrastructure](#)
Learn the basic concepts to get started using Oracle Cloud Infrastructure.
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- [Creating Protection Policies](#)
Recovery Service uses protection policies to control database backup retention in Oracle Cloud.
- [Storage Configuration Requirements for Oracle Exadata Database Service on Dedicated Infrastructure](#)
With the introduction of Exascale technology in Oracle Exadata Database Service on Dedicated Infrastructure, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both.
- [Cross-Service Data Guard Between ExaDB-D and ExaDB-XS](#)
You can now create a cross-service Oracle Data Guard group across database services.

Oracle Cloud Infrastructure (OCI) Requirements for Oracle Exadata Database Service on Dedicated Infrastructure

Learn the basic concepts to get started using Oracle Cloud Infrastructure.

Exadata Cloud Infrastructure is managed by the Oracle Cloud Infrastructure (OCI) control plane. The Exadata Cloud Infrastructure resources are deployed in your OCI Tenancy.

Before you can provision Exadata Cloud Infrastructure infrastructure, your Oracle Cloud Infrastructure tenancy must be enabled to use Oracle Exadata Database Service on Dedicated Infrastructure. Review the information in this publication for further details.

The following tasks are common for all OCI deployments, refer to the links in the Related Topics to find the associated Oracle Cloud Infrastructure documentation.

- Getting Started with OCI.
If you are new to OCI, learn the basic concepts to get started by following the *OCI Getting Started Guide*.
- Setting Up Your Tenancy.
After Oracle creates your tenancy in OCI, an administrator at your company will need to perform some set up tasks and establish an organization plan for your cloud resources and users. The information in this topic will help you get started.
- Managing Regions
This topic describes the basics of managing your region subscriptions.
- Managing Compartments

This topic describes the basics of working with compartments.

- **Managing Users**
This topic describes the basics of working with users.
- **Managing Groups**
This topic describes the basics of working with groups.
- [Required IAM Policy for Exadata Cloud Infrastructure](#)
Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

Related Topics

- [OCI Getting Started Guide](#)
- [Setting Up Your Tenancy](#)
- [Managing Regions](#)
- [Managing Compartments](#)
- [Managing Users](#)
- [Managing Groups](#)

Required IAM Policy for Exadata Cloud Infrastructure

Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

A **policy** is an IAM document that specifies who has what type of access to your resources. It is used in different ways:

- An individual statement written in the policy language
- A collection of statements in a single, named "policy" document, which has an Oracle Cloud ID (OCID) assigned to it
- The overall body of policies your organization uses to control access to resources

A **compartment** is a collection of related resources that can be accessed only by certain groups that have been given permission by an administrator in your organization.

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

If you're new to policies, then see "Getting Started with Policies" and "Common Policies". If you want to dig deeper into writing policies for databases, then see "Details for the Database Service".

For more details on writing policies specific to Exadata Cloud@Customer resources see "Policy Details for Exadata Cloud Infrastructure".

Related Topics

- [Getting Started with Policies](#)
- [Common Policies](#)
- [Policy Details for the Database Services](#)

- [Policy Details for Exadata Cloud Infrastructure](#)
This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.

Network Setup for Exadata Cloud Infrastructure Instances

This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.

Before you set up an Exadata Cloud Infrastructure instance, you must set up a virtual cloud network (VCN) and other [Networking service components](#).

- [VCN and Subnets](#)
To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network and at least two subnets:
- [Node Access to Object Storage: Static Route](#)
- [Service Gateway for the VCN](#)
Ensure that your VCN can reach the Oracle Services Network—specifically Object Storage for backups, Oracle YUM repositories for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS integration).
- [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#)
This section lists the security rules to use with Exadata Cloud Infrastructure.
- [Ways to Implement the Security Rules](#)
Learn how to implement security rules within your VCN using the networking service.
- [Network Requirements for Oracle Database Autonomous Recovery Service](#)
Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).

VCN and Subnets

To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network and at least two subnets:

To launch an Exadata Cloud Infrastructure instance, you must have a Virtual Cloud Network, at least two subnets and select the type of DNS resolver you will use:

- A [VCN](#) in the region where you want the Exadata Cloud Infrastructure instance
- At least two subnets in the VCN. The two subnets are:
 - Client subnet
 - Backup subnet
- Choose which method of DNS name resolution you will use. See *Choices for DNS in Your VCN*

Note

For Exadata Cloud Infrastructure instances using [The New Exadata Cloud Infrastructure Resource Model](#), networking is configured on the cloud VM cluster resource.

The client subnet can be configured with either IPv4 or IPv4/IPv6 dual-stack networking.

In general, Oracle recommends using **regional subnets**, which span all **availability domains** in the region. For more information, see [Overview of VCNs and Subnets](#).

You will create custom route tables [route tables](#) for each subnet. You will also create security rules [security rules](#) to control traffic to and from the client network and backup network of the Exadata compute nodes (for [The Cloud VM Cluster Resource](#), nodes are called virtual machines). More information follows about those items.

- [Option 1: Public Client Subnet with Internet Gateway](#)
This option can be useful when doing a proof-of-concept or development work.
- [Option 2: Private Subnets](#)
Oracle recommends private subnets for a production system.
- [Requirements for IP Address Space](#)
IP addresses must not overlap, especially when Exadata Cloud Infrastructure VM Clusters (and thus VCNs) are in more than one region.
- [Configuring a Static Route for Accessing the Object Store](#)
- [Setting Up DNS for an Exadata Cloud Infrastructure Instance](#)
DNS lets you use host names instead of IP addresses to communicate with an Exadata Cloud Infrastructure instance.
- [DNS: Short Names for the VCN, Subnets, and Exadata Cloud Infrastructure instance](#)
- [DNS: Between On-Premises Network and VCN](#)
Oracle recommends using a private DNS resolver to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other.
- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.

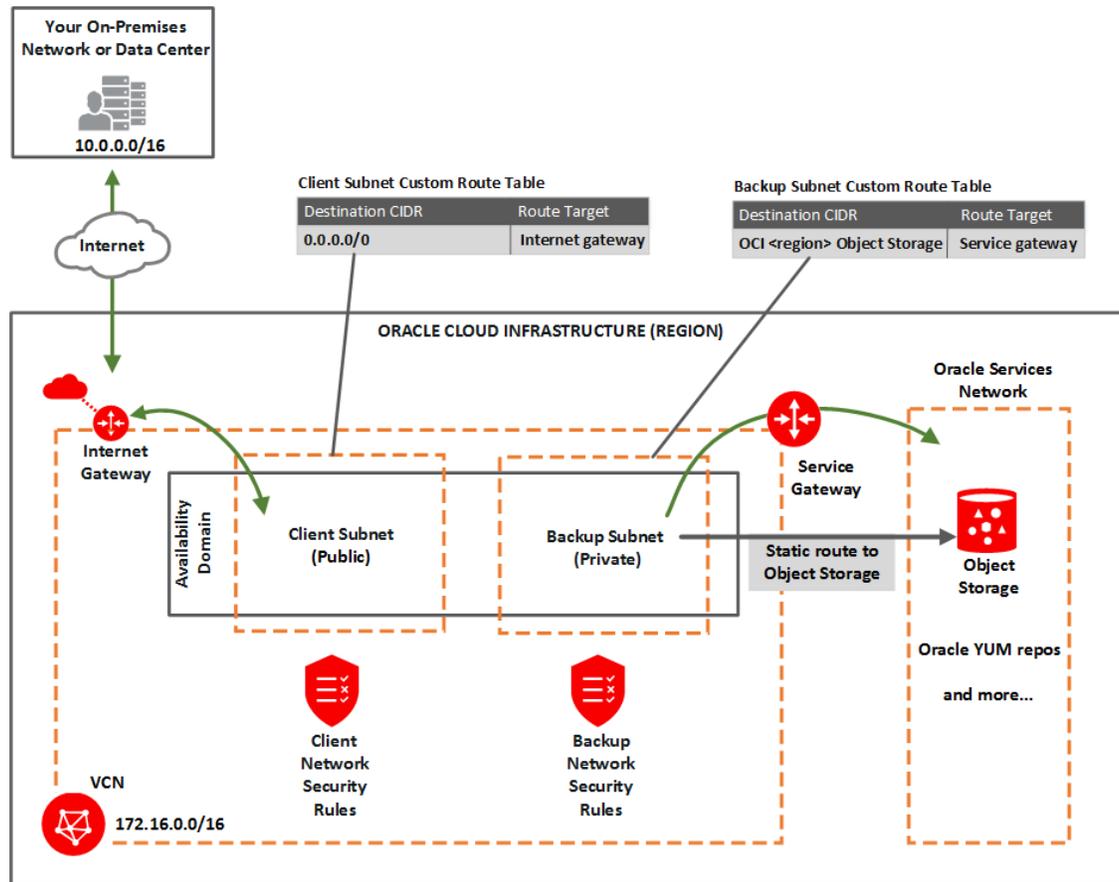
Related Topics

- [Choices for DNS in Your VCN](#)
- [Overview of VCNs and Subnets](#)
- [About Regions and Availability Domains](#)
- [Availability Domains and Your VCN](#)

Option 1: Public Client Subnet with Internet Gateway

This option can be useful when doing a proof-of-concept or development work.

You can use this setup in production if you want to use an **internet gateway** with the VCN, or if you have services that run only on a public network and need access to the database. See the following diagram and description.



You set up:

- [Subnets:](#)
 - *Public* client subnet (*public* means that the resources in the subnet can have public IP addresses at your discretion).
 - *Private* backup subnet (*private* means that the resources in the subnet cannot have public IP addresses and therefore cannot receive incoming connections from the internet).
- Gateways for the VCN:
 - [Internet gateway](#) (for use by the client subnet).
 - [Service gateway](#) (for use by the backup subnet). Also see [Option 1: Service Gateway Access to OCI Service for Backup Subnet](#).
- [Route tables:](#)
 - Custom route table for the public client subnet, with a route for 0.0.0.0/0, and target = the internet gateway.
 - Separate custom route table for the private backup subnet, with a route rule for the service CIDR labels (see about CIDR labels under [Overview of Service Gateways](#) and Available Service CIDR labels, and target = the service gateway. Also see [Option 1: Service Gateway Access to OCI Service for Backup Subnet](#)).
- [Security rules](#) to enable the desired traffic to and from the Exadata virtual machines compute nodes. See [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#).

- [Node Access to Object Storage: Static Route](#) on the Exadata Cloud Service instance's compute nodes (to enable access to OCI services by way of the backup subnet).

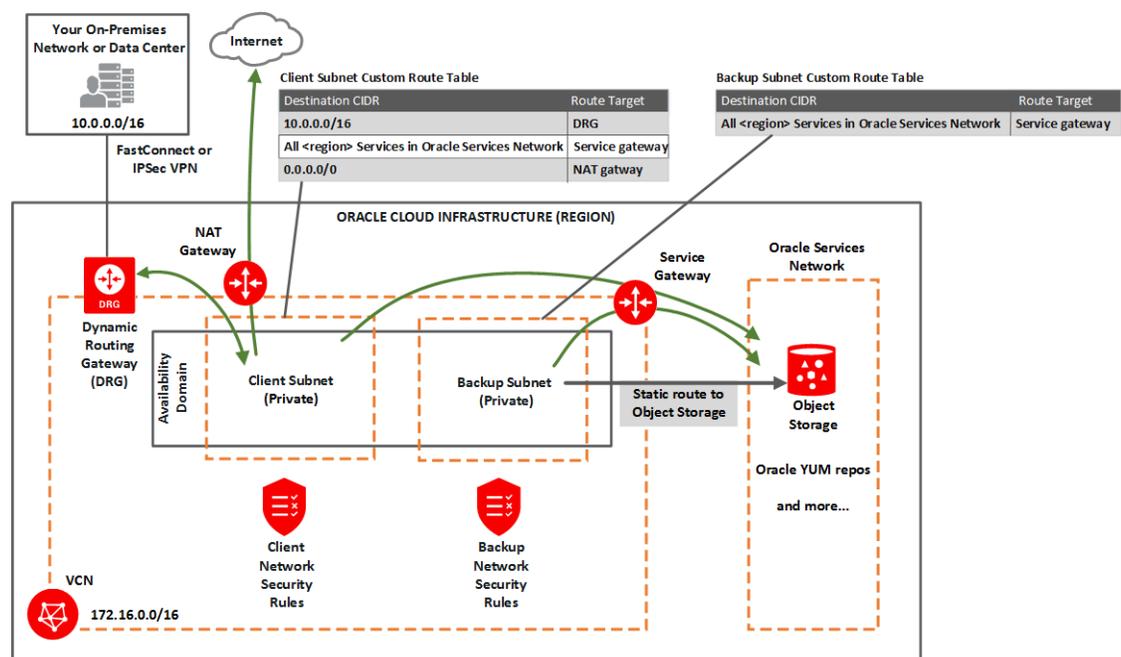
Note

Important See this [known issue](#) for information about configuring route rules with *service gateway* as the target on route tables associated with public subnets.

Option 2: Private Subnets

Oracle recommends private subnets for a production system.

Both subnets are private and cannot be reached from the internet. See the following diagram and description.



You set up:

- [Subnets](#):
 - Private client subnet.
 - Private backup subnet.
- Gateways for the VCN:
 - [Dynamic routing gateway \(DRG\)](#), with a [FastConnect](#) or [IPsec VPN](#) to your on-premises network (for use by the client subnet).
 - [Service gateway](#) For use by the backup and client subnets to reach OCI Services, such as Object Storage for backups, YUM for OS updates, IAM (Identity Access Management) and OCI Vault (KMS Integration) Also see [Option 2: Service Gateway Access to OCI Service for Both the Client and Backup Subnets](#).
 - [NAT gateway](#) (optional) For use by the client subnet to reach public endpoints not supported by the service gateway.
- [Route tables](#):

- Custom route table for the private client subnet, with the following rules:
 - * A rule for the on-premises network's CIDR, and `target = DRG`.
 - * A rule for the [service CIDR label](#) called `All <region> Services in Oracle Services Network`, and `target = the service gateway`. The *Oracle Services Network* is a conceptual network in Oracle Cloud Infrastructure that is reserved for Oracle services. The rule enables the client subnet to reach the regional Oracle YUM repository for OS updates. Also see [Option 2: Service Gateway Access to Both Object Storage and YUM Repos](#).
 - * Optionally, a rule for `0.0.0.0/0`, and `target = NAT gateway`.
- Separate custom route table for the private backup subnet, with one rule:
 - * The same rule as for the client subnet for the [service CIDR label](#) called `All <region> Services in Oracle Services Network`, and `target = the service gateway`. This rule enables the backup subnet to reach the regional Object Storage for backups.
- [Security rules](#) to enable the desired traffic to and from the Exadata nodes. See [Security Rules for the Exadata Cloud Service instance](#).
- Optionally add a [Static route](#) on the compute nodes to other OCI services (for VM clusters, the virtual machines) to enable access, if the services are only reachable on the backup subnet and not via the client subnet, e.g. when using a NAT Gateway.

Requirements for IP Address Space

IP addresses must not overlap, especially when Exadata Cloud Infrastructure VM Clusters (and thus VCNs) are in more than one region.

If you're setting up Exadata Cloud Infrastructure VM Clusters (and thus VCNs) in more than one region, make sure the IP address space of the VCNs does not overlap. This is important if you want to set up disaster recovery with Oracle Data Guard.

For Exadata X8 and lower, the two subnets you create for the Exadata Cloud Infrastructure VM Clusters must not overlap with `192.168.128.0/20`.

For Exadata X8M, X9M, and X11M, IP addresses `100.64.0.0/10` are reserved for the interconnect, and should not be used for the client or backup VCNs as well as database clients.

The following table lists the minimum required subnet sizes, depending on the Exadata VM Infrastructure for each VM Cluster size. For the client subnet, each node requires four IP addresses, and in addition, three addresses are reserved for Single Client Access Names ([SCANs](#)). For the backup subnet, each node requires three addresses.

Rack Size	Client Subnet: # Required IP Addresses	Client Subnet: Minimum Size Note	Backup Subnet: # Required IP Addresses	Backup Subnet: Minimum Size Note
Base System or Quarter Rack per VM Cluster	(4 addresses * 2 nodes) + 3 for SCANs = 11	/28 (16 IP addresses)	3 address * 2 nodes = 6	/28 (16 IP addresses)
Half Rack per VM Cluster	(4 * 4 nodes) + 3 = 19	/27 (32 IP addresses)	3 * 4 nodes = 12	/28 (16 IP addresses)
Full Rack per VM Cluster	(4 * 8 nodes) + 3 = 35	/26 (64 IP addresses)	3 * 8 nodes = 24	/27 (32 IP addresses)

Rack Size	Client Subnet: # Required IP Addresses	Client Subnet: Minimum Size Note	Backup Subnet: # Required IP Addresses	Backup Subnet: Minimum Size Note
Flexible infrastructure systems (X8M and higher) per VM Cluster	4 addresses per database node (minimum of 2 nodes) + 3 for SCANs	Minimum size determined by total number of IP addresses needed for database nodes	3 address per database node (minimum of 2 nodes)	Minimum size determined by total number of IP addresses needed for database nodes

① Note

The Networking service [reserves three IP addresses in each subnet](#). Allocating a larger space for the subnet than the minimum required (for example, at least /25 instead of /28) can reduce the relative impact of those reserved addresses on the subnet's available space.

Configuring a Static Route for Accessing the Object Store

All the traffic in an Exadata Cloud Infrastructure instance is, by default, routed through the data network. To route backup traffic to the backup interface (BONDETH1), you need to configure a static route on *each* of the compute nodes in the cluster.

For instructions, see [Node Access to Object Storage: Static Route](#).

Setting Up DNS for an Exadata Cloud Infrastructure Instance

DNS lets you use host names instead of IP addresses to communicate with an Exadata Cloud Infrastructure instance.

You can use the **Internet and VCN Resolver** (the DNS capability built into the VCN) as described in [DNS in Your Virtual Cloud Network](#). Oracle recommends using a VCN Resolver for DNS name resolution for the client subnet. It automatically resolves the Swift endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata instance.

DNS: Short Names for the VCN, Subnets, and Exadata Cloud Infrastructure instance

For the nodes to communicate, the VCN must use the [Internet and VCN Resolver](#). The Internet and VCN resolver enables hostname assignment to the nodes, and DNS resolution of those hostnames by resources in the VCN.

The Internet and VCN resolver enables round robin resolution of the database's [SCANs](#). It also enables resolution of important service endpoints required for backing up databases, patching, and updating the cloud tooling on an Exadata Cloud Infrastructure instance. The Internet and VCN Resolver is the VCN's default choice for DNS in the VCN. For more information, see [DNS in Your Virtual Cloud Network](#) and also [DHCP Options](#).

When you create the VCN, subnets, and Exadata, you must carefully set the following identifiers, which are related to DNS in the VCN:

- VCN domain label
- Subnet domain label
- Hostname prefix for the Exadata Cloud Infrastructure instance's cloud VM cluster resource

These values make up the node's fully qualified domain name (FQDN):

```
<hostname_prefix>-#####.<subnet_domain_label>.<vcn_domain_label>.oraclevcn.com
```

For example:

```
exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com
```

In this example, you assign `exacs` as the hostname prefix when you create the cloud VM cluster. The Database service automatically appends a hyphen and a five-letter string with the node number at the end. For example:

- Node 1: `exacs-abcde1.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 2: `exacs-abcde2.clientpvtad1.acmevcniad.oraclevcn.com`
- Node 3: `exacs-abcde3.clientpvtad1.acmevcniad.oraclevcn.com`
- And so on

Requirements for the hostname prefix:

- Recommended maximum: 12 characters. For more information, see the [example](#) under the following section, "Requirements for the VCN and subnet domain labels".
- Cannot be the string `localhost`

Requirements for the VCN and subnet domain labels:

- Recommended maximum: 14 characters each. The actual underlying requirement is a total of 28 characters *across both domain labels* (excluding the period between the labels). For example, both of these are acceptable: `subnetad1.verylongvcnphx` or `verylongsubnetad1.vcnphx`. For simplicity, the recommendation is 14 characters each.
- No hyphens or underscores.
- Recommended: include the region name in the VCN's domain label, and include the availability domain name in the subnet's domain label.
- In general, the FQDN has a maximum total limit of 63 characters. Here is a safe general rule:

```
<12_chars_max>-#####.<14_chars_max>.<14_chars_max>.oraclevcn.com
```

The preceding maximums are not enforced when you create the VCN and subnets. However, if the labels exceed the maximum, the Exadata deployment fails.

DNS: Between On-Premises Network and VCN

Oracle recommends using a private DNS resolver to enable the use of hostnames when on-premises hosts and VCN resources communicate with each other.

See [Private DNS resolvers](#) for information on creating and using private resolvers. For a reference architecture see [Use private DNS in your VCN](#) in the Oracle Architecture Center.

Configure Private DNS

Prerequisites needed to use Private DNS.

- Private view and private zone must be created before launching VM cluster provisioning. For details, see [Private DNS resolvers](#).

- Forwarding to another DNS server should be set up beforehand in the DNS console. This can be done by going to the VCN's resolver, and creating the endpoint and then the rules. For details, see [DNS in Your Virtual Cloud Network](#).
- Private zone's name cannot have more than 4 labels. For example, a.b.c.d is allowed while a.b.c.d.e is not.
- It is also required to add the private view to the resolver of the VCN. For details, see [Adding a Private View to a Resolver](#).
- When provisioning a Exadata VM Cluster using Private DNS feature, Exadata needs to create reverse DNS zones in the compartment of Exadata VM Cluster. If the compartment has defined tags or tag defaults, additional policies related to managing tags are needed. For details, see:
 - [Required Permissions for Working with Defined Tags](#)
 - [Required Permissions for Working with Tag Defaults](#)

Node Access to Object Storage: Static Route

To be able to back up databases, and patch and update cloud tools on an Exadata Cloud Infrastructure instance, you must configure access to Oracle Cloud Infrastructure Object Storage. Regardless of how you configure the VCN with that access (for example, with a service gateway), you may also need to configure a static route to Object Storage on each of the compute nodes in the cluster. This is only required if you are not using automatic backups. If you are using customized backups using the backup APIs, then you must route traffic destined for Object Storage through the backup interface (BONDETH1). This is not necessary if you are using the automatic backups created with the Console, APIs, or CLIs.

Caution

You must configure a static route for Object Storage access on each compute node in an Exadata Cloud Infrastructure instance if you *are not* creating automatic backups with the Console, APIs, or CLIs. Otherwise, attempts to back up databases, and patch or update tools on the system, can fail.

Note

When you enable the first automatic backup for a database the static route configuration will be automatically done on the service.

If you want to patch the service before creating a database, the manual static route is required to be able to patch the GI or DB Home.

The static route may also be required to access other services (IAM, KMS) if these are not reachable via client subnet and only the backup subnet uses the setting to access all services within a region.

- [Object Storage IP allocations](#)
- [To configure a static route for Object Storage access](#)

Object Storage IP allocations

Oracle Cloud Infrastructure Object Storage uses the CIDR block IP range 134.70.0.0/16 for all regions.

As of June 1, 2018, Object Storage no longer supports the following discontinued IP ranges. Oracle recommends that you remove these older IP addresses from your access-control lists, firewall rules, and other rules after you have adopted the new IP ranges.

The **discontinued** IP ranges are:

- Germany Central (Frankfurt): 130.61.0.0/16
- UK South (London): 132.145.0.0/16
- US East (Ashburn): 129.213.0.0/16
- US West (Phoenix): 129.146.0.0/16

To configure a static route for Object Storage access

1. SSH to a compute node in the Exadata Cloud Infrastructure instance.

```
ssh -i <private_key_path> opc@<node_ip_address>
```

2. Log in as opc and then sudo to the root user. Use `sudo su -` with a hyphen to invoke the root user's profile.

```
login as: opc  
[opc@dbsys ~]$ sudo su -
```

3. Identify the gateway configured for the BONDETH1 interface.

```
[root@dbsys ~]# grep GATEWAY /etc/sysconfig/network-scripts/ifcfg-bondeth1  
|awk -F"=" '{print $2}'
```

```
10.0.4.1
```

4. Add the following static rule for BONDETH1 to the `/etc/sysconfig/network-scripts/route-bondeth1` file:

```
10.0.X.0/XX dev bondeth1 table 211  
default via <gateway> dev bondeth1 table 211  
134.70.0.0/17 via <gateway_from_previous_step> dev bondeth1
```

5. Restart the interface.

```
[root@dbsys ~]# ifdown bondeth1; ifup bondeth1;
```

The file changes from the previous step take effect immediately after the `ifdown` and `ifup` commands run.

6. Repeat the preceding steps on *each* compute node in the Exadata Cloud Infrastructure instance.

Service Gateway for the VCN

Ensure that your VCN can reach the Oracle Services Network—specifically Object Storage for backups, Oracle YUM repositories for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS integration).

Note

If the VCN lacks connectivity to the Oracle Services Network, provisioning new VM Clusters may fail and the manageability of existing VM Clusters could also be affected.

Depending on whether you use *Option 1: Public Client Subnet with Internet Gateway* or *Option 2: Private Subnets*, you use the service gateway in different ways. See the next two sections.

- [Option 1: Service Gateway Access to OCI Service for Backup Subnet](#)
You configure the backup subnet to use the service gateway for access only to Object Storage.
- [Option 2: Service Gateway Access to OCI Service for Both the Client and Backup Subnets](#)
You configure both the client subnet and backup subnet to use the service gateway for access to the Oracle Services Network, which includes Object storage for backups, Oracle YUM repos for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS Integration).

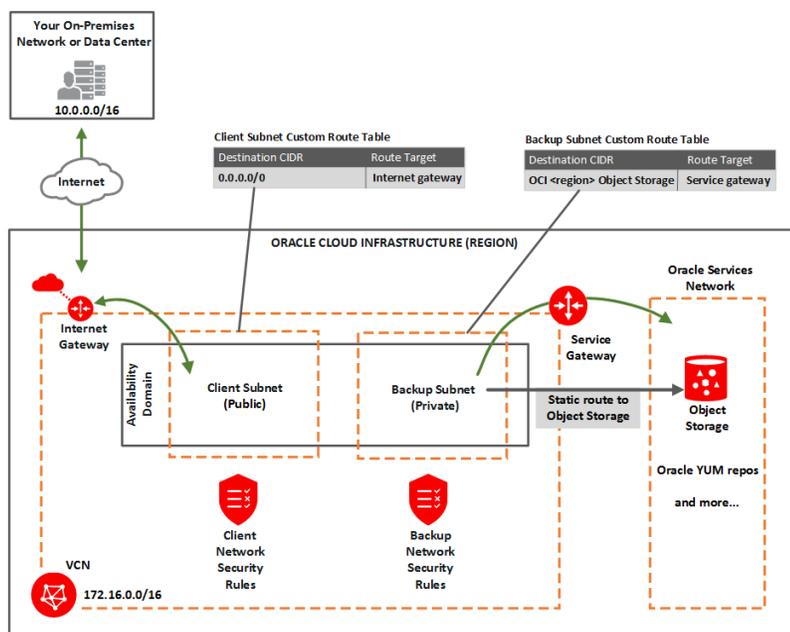
Related Topics

- [Option 1: Public Client Subnet with Internet Gateway](#)
This option can be useful when doing a proof-of-concept or development work.
- [Option 2: Private Subnets](#)
Oracle recommends private subnets for a production system.

Option 1: Service Gateway Access to OCI Service for Backup Subnet

You configure the backup subnet to use the service gateway for access only to Object Storage.

As a reminder, here's the diagram for option 1:



In general, you must:

- Perform the tasks for setting up a service gateway on a VCN, and specifically enable the service CIDR label called `OCI <region> Object Storage`.
- In the task for updating routing, add a route rule to the backup subnet's custom route table. For the destination service, use `OCI <region> Object Storage` and `target = the service gateway`.
- In the task for updating security rules in the subnet, perform the task on the backup network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to `OCI <region> Object Storage`. See *Rule Required Specifically for the Backup Network*.

Related Topics

- [Tasks for Setting Up a Service Gateway on a VCN in the Console](#)
- [Rule Required Specifically for the Backup Network](#)

The following security rule is important for the backup network because it enables the VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repositories if the client network doesn't have access to them).

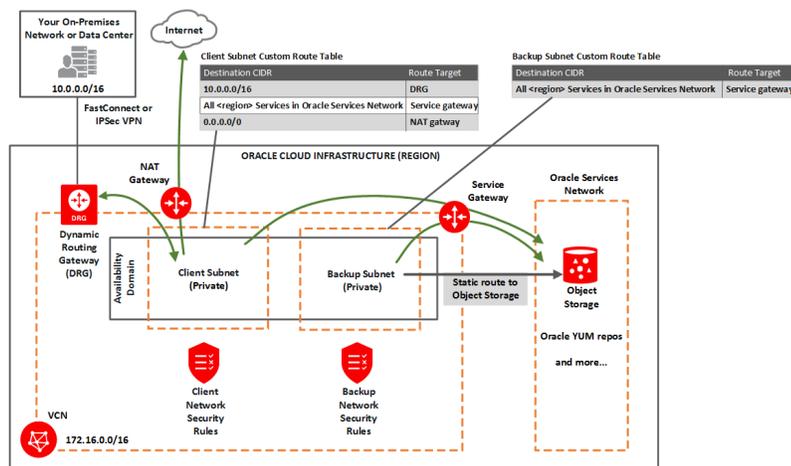
Option 2: Service Gateway Access to OCI Service for Both the Client and Backup Subnets

You configure both the client subnet and backup subnet to use the service gateway for access to the Oracle Services Network, which includes Object storage for backups, Oracle YUM repos for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS Integration).

Note

See this [known issues](#) for information about accessing Oracle YUM services through the service gateway.

As a reminder, here's the diagram for option 2:



In general, you must:

- Perform the tasks for setting up a service gateway on a VCN, and specifically enable the service CIDR label called All <region> Services in Oracle Services Network.
- In the task for updating routing in each subnet, add a rule to each subnet's custom route table. For the destination service, use All <region> Services in Oracle Services Network and target = the service gateway.
- In the task for updating security rules for the subnet, perform the task on the backup network's network security group (NSG) or custom security list. Set up a security rule with the destination service set to OCI <region> Object Storage. See *Security Rules for the Oracle Exadata Cloud Service* [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#). Note that the client subnet already has a broad egress rule that covers access to the YUM repos.

Here are a few additional details about using the service gateway for option 2:

- Both the client subnet and backup subnet use the service gateway, but to access different services. You cannot enable both the OCI <region> Object Storage service CIDR label and the All <region> Services in Oracle Services Network for the service gateway. To cover the needs of both subnets, you must enable All <region> Services in Oracle Services Network for the service gateway. The VCN can have only a single service gateway.
- Any route rule that targets a given service gateway must use an enabled service CIDR label and not a CIDR block as the destination for the rule. That means for option 2, the route tables for both subnets must use All <region> Services in Oracle Services Network for their service gateway rules.
- Unlike route rules, security rules can use either any service CIDR label (whether the VCN has a service gateway or not) or a CIDR block as the source or destination CIDR for the rule. Therefore, although the backup subnet has a route rule that uses All <region> Services in Oracle Services Network, the subnet can have a security rule that uses OCI <region> Object Storage. See *Security Rules for the Exadata Cloud Service instance*.

Related Topics

- [Oracle Service Gateway](#)
- [Tasks for Setting up a Service Gateway on a VCN](#)

- [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#)
This section lists the security rules to use with Exadata Cloud Infrastructure.

Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure

This section lists the security rules to use with Exadata Cloud Infrastructure.

Security rules control the types of traffic allowed for the client network and backup network of the Exadata's compute nodes. The rules are divided into three sections.

There are different ways to implement these rules. For more information, see [Ways to Implement the Security Rules](#).

Note

For X8M and X9M systems, Oracle recommends that all ports on the client subnet need to be open for ingress and egress traffic. This is a requirement for adding additional database servers to the system.

Note

If you plan to use Zero-trust Packet Routing to control the networks for your database services, and you plan to configure Data Guard peers within the same VCN, then all VM Clusters in the VCN and Data Guard configuration must have the same ZPR Security Attribute. Data Guard peers that are in a different VCN or different region must be specified by CIDR in the ZPR configuration.

Rules Required for Both the Client Network and Backup Network

This section has several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules, be aware that the rules that follow are included by default in the [default security list](#). Update or replace the list to meet your particular security needs. The two ICMP rules (general ingress rules 2 and 3) are required for proper functioning of network traffic within the Oracle Cloud Infrastructure environment. Adjust the general ingress rule 1 (the SSH rule) and the general egress rule 1 to allow traffic only to and from hosts that require communication with resources in your VCN.

General ingress rule 1: Allows SSH traffic from anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** SSH
- **Source Port Range:** All
- **Destination Port Range:** 22

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

This rule enables hosts in the VCN to receive Path MTU Discovery fragmentation messages. Without access to these messages, hosts in the VCN can have problems communicating with hosts outside the VCN.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** 4

General ingress rule 3: Allows connectivity error messages within the VCN

This rule enables the hosts in the VCN to receive connectivity error messages from each other.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: your VCN's IPv4 CIDR, IPv6: your VCN's IPv6 CIDR
- **IP Protocol:** ICMP
- **Type:** ALL
- **Code:** All

General egress rule 1: Allows all egress traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** All

Rules Required Specifically for the Client Network

The following security rules are important for the client network.

Important

- Client ingress rules 1 and 2 only cover connections initiated from within the client subnet. If you have a client that resides *outside the VCN*, Oracle recommends setting up two *additional* similar rules that instead have the **Source CIDR** set to the public IP address of the client.
- Client egress rules 1 and 2 in the client network configuration allow TCP and ICMP traffic, enabling secure node-to-node communication within the client network. These rules are critical, as they facilitate essential TCP connectivity across nodes. Should TCP connectivity fail between nodes, provisioning of the Exadata Cloud VM cluster resource will not complete successfully.

Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet

The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: client subnet's IPv4 CIDR, IPv6: client subnet's IPv6 CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 6200
- **Description:** An optional description of the rule.

Client ingress rule 2: Allows SQL*NET traffic from within the client subnet

This rule is for SQL*NET traffic and is required in these cases:

- If you need to enable client connections to the database
- If you plan to use Oracle Data Guard
- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: client subnet's IPv4 CIDR, IPv6: client subnet's IPv6 CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

Client egress rule 1: Allows SSH TCP traffic inside the client subnet

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22
- **Description:** An optional description of the rule.

Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)

Client egress rule 2 is important because it allows connections to the Oracle YUM repositories. It is redundant with the general egress rule in this topic (and in the [default security list](#)). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR

- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** All
- **Description:** An optional description of the rule.

Rule Required Specifically for the Backup Network

The following security rule is important for the backup network because it enables the VM Cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repos if the client network doesn't have access to them). It is redundant with the general egress rule in this topic (and in the [default security list](#)). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

Backup egress rule: Allows access to Object Storage

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
 - The service CIDR label called **OCI <region> Object Storage**
 - If the client network does not have access to the Oracle YUM repos, use the service CIDR label called **All <region> Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)
- **Description:** An optional description of the rule.
- [Rules Required for Both the Client Network and Backup Network](#)
This topic has several general rules that enable essential connectivity for hosts in the VCN.
- [Rules Required Specifically for the Client Network](#)
The following security rules are important for the client network.
- [Rule Required Specifically for the Backup Network](#)
The following security rule is important for the backup network because it enables the VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repositories if the client network doesn't have access to them).
- [Rules Required for Events Service](#)
The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Events service.
- [Rules Required for Monitoring Service](#)
The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

Rules Required for Both the Client Network and Backup Network

This topic has several general rules that enable essential connectivity for hosts in the VCN.

If you use security lists to implement your security rules, be aware that the rules that follow are included by default in the *default security list*. Update or replace the list to meet your particular security needs. The two ICMP rules (general ingress rules 2 and 3) are required for proper functioning of network traffic within the Oracle Cloud Infrastructure environment. Adjust the general ingress rule 1 (the SSH rule) and the general egress rule 1 to allow traffic only to and from hosts that require communication with resources in your VCN.

- [General ingress rule 1: Allows SSH traffic from anywhere](#)
- [General ingress rule 2: Allows Path MTU Discovery fragmentation messages](#)
- [General ingress rule 3: Allows connectivity error messages within the VCN](#)
This rule enables the hosts in the VCN to receive connectivity error messages from each other.
- [General egress rule 1: Allows all egress traffic](#)

Related Topics

- [default security list](#)

General ingress rule 1: Allows SSH traffic from anywhere

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** SSH
- **Source Port Range:** All
- **Destination Port Range:** 22

General ingress rule 2: Allows Path MTU Discovery fragmentation messages

This rule enables hosts in the VCN to receive Path MTU Discovery fragmentation messages. Without access to these messages, hosts in the VCN can have problems communicating with hosts outside the VCN.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** ICMP
- **Type:** 3
- **Code:** 4

General ingress rule 3: Allows connectivity error messages within the VCN

This rule enables the hosts in the VCN to receive connectivity error messages from each other.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: your VCN's IPv4 CIDR, IPv6: your VCN's IPv6 CIDR
- **IP Protocol:** ICMP
- **Type:** All
- **Code:** All

General egress rule 1: Allows all egress traffic

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR

- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** All

If the customer enables notification of Data Plane Guest VM Events, the default egress rule is sufficient.

Rules Required Specifically for the Client Network

The following security rules are important for the client network.

Note

- For X8M systems, Oracle recommends that all ports on the client subnet need to be open for ingress and egress traffic. This is a requirement for adding additional database servers to the system.
- Client ingress rules 1 and 2 only cover connections initiated from within the client subnet. If you have a client that resides outside the VCN, Oracle recommends setting up two additional similar rules that instead have the **Source CIDR** set to the public IP address of the client.
- Client ingress rules 3 and 4 and client egress rules 1 and 2 allow TCP and ICMP traffic inside the client network and enable the nodes to communicate with each other. If TCP connectivity fails across the nodes, the Exadata VM cluster resource fails to provision.

- [Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet](#)
The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.
- [Client ingress rule 2: Allows SQL*NET traffic from within the client subnet](#)
This rule is for SQL*NET traffic and is required in these cases:
- [Client Ingress Rule 3: Allows patching traffic from within the client subnet](#)
Allows TCP patching traffic on port 7085 from the client subnet to the Exadata Fleet Update private endpoint.
- [Client egress rule 1: Allows all TCP traffic inside the client subnet](#)
This rule is for SQL*NET traffic as noted.
- [Client egress rule 2: Allows all egress traffic \(allows connections to the Oracle YUM repos\)](#)
Client egress rule 3 is important because it allows connections to the Oracle YUM repos.

Client ingress rule 1: Allows ONS and FAN traffic from within the client subnet

The first rule is recommended and enables the Oracle Notification Services (ONS) to communicate about Fast Application Notification (FAN) events.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: client subnet's IPv4 CIDR, IPv6: client subnet's IPv6 CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 6200

- **Description:** An optional description of the rule.

Client ingress rule 2: Allows SQL*NET traffic from within the client subnet

This rule is for SQL*NET traffic and is required in these cases:

- If you need to enable client connections to the database
- If you plan to use Oracle Data Guard
- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** IPv4: client subnet's IPv4 CIDR, IPv6: client subnet's IPv6 CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 1521
- **Description:** An optional description of the rule.

Client Ingress Rule 3: Allows patching traffic from within the client subnet

Allows TCP patching traffic on port 7085 from the client subnet to the Exadata Fleet Update private endpoint.

- **Stateless:** No (all rules must be stateful)
- **Source Type:** CIDR
- **Source CIDR:** Client subnet's CIDR
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 7085
- **Description:** Optionally, add a meaningful description of the rule. For example: "Allow access to Exadata Fleet Update private endpoint within the subnet."

Client egress rule 1: Allows all TCP traffic inside the client subnet

This rule is for SQL*NET traffic as noted.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 22
- **Description:** An optional description of the rule.

Client egress rule 2: Allows all egress traffic (allows connections to the Oracle YUM repos)

Client egress rule 3 is important because it allows connections to the Oracle YUM repos.

It is redundant with the general egress rule 1: Allow all egress traffic (and in the *default security list*). It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0 (IPv4), ::/0 (IPv6)
- **IP Protocol:** All
- **Description:** An optional description of the rule.

Related Topics

- [default security list](#)

Rule Required Specifically for the Backup Network

The following security rule is important for the backup network because it enables the VM cluster to communicate with Object Storage through the service gateway (and optionally with the Oracle YUM repositories if the client network doesn't have access to them).

It is redundant with the *general egress rule 1: Allows all egress traffic* in. It is optional but recommended in case the general egress rule (or default security list) is inadvertently changed.

- [Backup egress rule: Allows access to Object Storage](#)

Related Topics

- [General egress rule 1: Allows all egress traffic](#)
- [default security list](#)

Backup egress rule: Allows access to Object Storage

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** Service
- **Destination Service:**
 - The service CIDR label called **OCI <region> Object Storage**
 - If the client network does not have access to the Oracle YUM repos, use the service CIDR label called **All <region> Services in Oracle Services Network**
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443 (HTTPS)
- **Description:** An optional description of the rule.

Rules Required for Events Service

The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Events service.

The default egress rules are sufficient to allow the compute instance to send compute instance metrics to the Events service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the

Events service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Events service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Events service.
- When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Access to Oracle Services: Service Gateway](#).

Rules Required for Monitoring Service

The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

The default egress rules are sufficient to allow the compute instance to send compute instance metrics to the Monitoring service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

- When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
- When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Access to Oracle Services: Service Gateway](#).

Ways to Implement the Security Rules

Learn how to implement security rules within your VCN using the networking service.

The Networking service offers two ways to implement security rules within your VCN:

- [Network security groups](#)
- [Security lists](#)

For a comparison of the two methods, see [Comparison of Security Lists and Network Security Groups](#).

- [If you use network security groups](#)
- [If you use security lists](#)

If you choose to use security lists, here is the recommended process:

If you use network security groups

If you choose to use network security groups (NSGs), here is the recommended process:

1. Create an NSG for the client network. Add the following security rules to that NSG:
 - The rules listed in [Rules Required for Both the Client Network and Backup Network](#)
 - The rules listed in [Rules Required Specifically for the Client Network](#)

2. Create a separate NSG for the backup network. Add the following security rules to that NSG:
 - The rules listed in [Rules Required for Both the Client Network and Backup Network](#)
 - The rules listed in [Rules Required Specifically for the Client Network](#)
3. When the database administrator is [Creating an Exadata Cloud Infrastructure Instance](#), they must choose several networking components (for example, which VCN and subnets to use):
 - When they choose the client subnet, they can also choose which NSG or NSGs to use. Make sure they choose the client network's NSG.
 - When they choose the backup subnet, they can also choose which NSG or NSGs to use. Make sure they choose the backup network's NSG.

You could instead create a separate NSG for the general rules. Then when the database administrator chooses which NSGs to use for the client network, make sure they choose both the general NSG and the client network NSG. Similarly for the backup network, they choose both the general NSG and the backup network NSG.

If you use security lists

If you choose to use security lists, here is the recommended process:

If you choose to use security lists, here is the recommended process:

1. Configure the client subnet to use the required security rules:
 - a. Create a custom security list for the client subnet and add the rules listed in [Rules Required Specifically for the Client Network](#).
 - b. Associate the following two security lists with the client subnet:
 - VCN's *default security list* with all its default rules. This automatically comes with the VCN. By default it contains the rules in [Rules Required for Both the Client Network and Backup Network](#).
 - The new custom security list you created for the client subnet.
2. Configure the backup subnet to use the required security rules:
 - a. Create a custom security list for the backup subnet and add the rules listed in [Rule Required Specifically for the Backup Network](#).
 - b. Associate the following two security lists with the backup subnet:
 - VCN's *default security list* with all its default rules. This automatically comes with the VCN. By default it contains the rules in [Rules Required for Both the Client Network and Backup Network](#).
 - The new custom security list you created for the backup subnet.

Later when the database administrator creates the Exadata Cloud Service instance, they must choose several networking components. When they select the client subnet and backup subnet that you've already created and configured, the security rules are automatically enforced for the nodes created in those subnets.

Warning

Do not remove the default egress rule from the default security list. If you do, make sure to instead include the following replacement egress rule in the client subnet's security list:

- **Stateless:** No (all rules must be stateful)
- **Destination Type:** CIDR
- **Destination CIDR:** 0.0.0.0/0
- **IP Protocol:** All

Network Requirements for Oracle Database Autonomous Recovery Service

Oracle Database Autonomous Recovery Service requires a registered Recovery Service subnet dedicated to backup and recovery operations in your database virtual cloud network (VCN).

To use Recovery Service for backups, follow the steps outlined in [Onboarding Oracle Database to Recovery Service](#).

- [Create a Service Gateway to Object Storage](#)
In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.

Create a Service Gateway to Object Storage

In the OCI Console, create a service gateway to Object Storage. The service gateway is required for automation updates and configuration metadata.

1. Open the navigation menu. Click **Networking**, and then click **Virtual Cloud Networks**.
2. Select the VCN where your database services to be backed up are located.
3. On the resulting Virtual Cloud Network Details page, under **Resources**, click **Service Gateways**.
4. Click **Create Service Gateway** and provide the following details.
 - a. **Name:** A descriptive name for the service gateway. It doesn't have to be unique. Avoid entering confidential information.
 - b. **Compartment:** The compartment where you want to create the service gateway, if different from the compartment you're currently working in.
 - c. **Services:** Select the service CIDR Label, All <region> Services in Oracle Services Network from the drop-down list.
 - d. **Tags:** (advanced option) 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.
5. Click **Create Service Gateway**.

Wait for the gateway to be created before proceeding to the next step.

6. Under **Resources**, click **Route Tables**.

Route Table Association: You can associate a specific VCN route table with this gateway. If you associate a route table, afterward the gateway must always have a route table associated with it. You can modify the rules in the current route table or replace them with another route table.

7. Click the **Route Table** name that is being used by the subnet for Recovery Service.
8. In the resulting Route Table Details page, click **Add Route Rules** in the **Route Rules** section.

When you configure a service gateway for a particular service CIDR label, you must also create a route rule that specifies that label as the destination and the target as the service gateway. You do this for each subnet that needs to access the gateway.

9. In the resulting Add Route Rules dialog, enter the following details:
 - a. **Target Type:** Service Gateway.
 - b. **Destination Service:** The service CIDR label that is enabled for the gateway. All `<region> Services in Oracle Services Network`
 - c. **Target Service Gateway:** Select the name that you provided in step 4.
 - d. **Description:** An optional description of the rule.
10. Click **Add Route Rules**.

Related Topics

- [Resource Tags](#)

Creating Protection Policies

Recovery Service uses protection policies to control database backup retention in Oracle Cloud.

Protection Policies provide automated retention management for protected databases, satisfying requirements for regulated environments. Each protected database must be associated with one protection policy.

A protection policy determines the maximum period (in days) allowed to retain backups created by Recovery Service. Based on your business requirements, you can assign separate policies for each protected database or use a single policy across all protected databases in a VCN. For more information, see *About Configuring Protection Policies*.

To use the Oracle Cloud Infrastructure (OCI) Console to configure and manage protection policies, follow the steps outlined in *Creating a Protection Policy*.

Related Topics

- [About Configuring Protection Policies](#)
- [Create Protection Policy](#)

Storage Configuration Requirements for Oracle Exadata Database Service on Dedicated Infrastructure

With the introduction of Exascale technology in Oracle Exadata Database Service on Dedicated Infrastructure, you can configure the Exadata infrastructure to use ASM, Exascale, or a combination of both.

When Exascale is configured to coexist with ASM, you must allocate storage capacity for Exascale from the total available Exadata storage on the infrastructure. The allocated storage capacity will be solely available for Exascale use.

Review the storage requirements for ASM, Exascale, and VM file systems to plan and optimize storage allocation based on your enterprise needs.

- [Configuring Exadata ASM Storage](#)
The storage space inside the Exadata storage servers is configured for use by Oracle Automatic Storage Management (ASM) When you launch an Exadata Cloud Infrastructure instance.
- [Configuring Exadata Exascale Storage](#)

Configuring Exadata ASM Storage

The storage space inside the Exadata storage servers is configured for use by Oracle Automatic Storage Management (ASM) When you launch an Exadata Cloud Infrastructure instance.

By default, the following ASM disk groups are created:

- The DATA disk group is intended for the storage of Oracle Database data files.
- The RECO disk group is primarily used for storing the Fast Recovery Area (FRA), which is an area of storage where Oracle Database can create and manage various files related to backup and recovery, such as RMAN backups and archived redo log files.
- The `/acfs` file systems contain system files that support various operations. You should not store custom files, Oracle Database data files, or backups inside the ACFS disk groups. Custom ACFS mounts can be created using the DATA ASM disk group for files that are not service-related.

The disk group names contain a short identifier string that is associated with your Exadata Database machine environment. For example, the identifier could be C2, in which case the DATA disk group would be named DATA_C2, the RECO disk group would be named RECO_C2, and so on.

In addition, you can create a SPARSE disk group. A SPARSE disk group is required to support Exadata snapshots. Exadata snapshots enable space-efficient clones of Oracle databases that can be created and destroyed very quickly and easily. Snapshot clones are often used for development, testing, or other purposes that require a transient database.

Note that you cannot change the disk group layout after service creation.

- [Impact of Configuration Settings on ASM Storage](#)
Your choices concerning database Backups or sparse disk groups profoundly affect how storage space in the Exadata storage servers is allocated to the ASM and sparse disk groups

Impact of Configuration Settings on ASM Storage

Your choices concerning database Backups or sparse disk groups profoundly affect how storage space in the Exadata storage servers is allocated to the ASM and sparse disk groups

If you choose to perform database backups to the Exadata storage, or to create a sparse disk group, or to do both, the storage space allocation in the Exadata storage servers will be affected.

The table that follows shows the approximate percentages of storage allocated for DATA, RECO, and SPARSE disk groups for each possible configuration.

Configuration Settings	DATA Disk Group	RECO Disk Group	SPARSE Disk Group
Database backups on Exadata storage: No Sparse disk group: No	80 %	20 %	0 %
Database backups on Exadata storage: Yes Sparse disk group: No	40 %	60 %	0 %
Database backups on Exadata storage: No Sparse disk group: Yes	60 %	20 %	20 %
Database backups on Exadata storage: Yes Sparse disk group: Yes	35 %	50 %	15 %

Configuring Exadata Exascale Storage

- [About Exascale Storage Configuration for Oracle Exadata Database Service on Dedicated Infrastructure](#)
To use Exascale storage during VM cluster provisioning, you must first configure it from the total Exadata storage available on your Exadata Infrastructure.
- [Using the Console to Configure Exascale Storage on ExaDB-D Infrastructure](#)
To configure Exascale storage on Oracle Exadata Database Service on Dedicated Infrastructure, be prepared to provide values for the infrastructure configuration.
- [Using the Console to Create an Exascale Storage Vault](#)
To create an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.
- [Using the Console to Scale an Exascale Storage Vault](#)
To scale an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.
- [Using the Console to Move an Exascale Storage Vault to Another Compartment](#)
To move an Exascale storage vault to another compartment, use this procedure.
- [Using the Console to Delete an Exascale Storage Vault](#)
To delete an Exascale storage vault, use this procedure.

About Exascale Storage Configuration for Oracle Exadata Database Service on Dedicated Infrastructure

To use Exascale storage during VM cluster provisioning, you must first configure it from the total Exadata storage available on your Exadata Infrastructure.

Note

During the initial Exascale storage configuration, all database servers on the Exadata Infrastructure will be restarted in a rolling reboot manner.

Specify the storage capacity you want to allocate for Exascale usage on the infrastructure. A minimum of 2TB of Exadata storage must be available to configure Exascale storage. Configuring Exascale storage requires a rolling reboot of the data nodes. After the initial configuration, you can scale the Exascale storage capacity online as needed. Note that scaling storage may trigger a storage rebalancing process.

If sufficient storage is not available on the Infrastructure to configure the Exascale storage or to scale the existing Exascale storage capacity, additional storage servers can be added to the Infrastructure using the Scale Infrastructure option as follows:

- Add new storage servers to the infrastructure.
- Make the new storage capacity available to all ASM and Exascale VM clusters.
- Adding new storage will proportionally increase the Exascale storage capacity.
- Expand Exascale storage vaults as needed to utilize the additional capacity.

For more information, see [Scaling Resources within an Exadata Infrastructure Instance](#).

Alternatively, customers can shrink the existing ASM or Exascale clusters to free up storage and allocate it to Exascale to meet the additional capacity demand.

Using the Console to Configure Exascale Storage on ExaDB-D Infrastructure

To configure Exascale storage on Oracle Exadata Database Service on Dedicated Infrastructure, be prepared to provide values for the infrastructure configuration.

Note

- All Exadata Infrastructure with DB Server Version 25.1.7 (or later) and Storage Server Version 25.1.8 (or later) or 25.2.2 (or later) will now see a banner with the message, "You can now configure Exascale storage on this infrastructure" to indicate that the infrastructure is ready for Exascale configuration. This banner will not be displayed for Exadata Infrastructures that are already configured to use Exascale.
- The minimum Exascale storage that can be configured on Exadata Infrastructure is 2 TB.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where you want to configure Exascale storage.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Actions**, and then select **Configure Exascale storage**.
6. On the resulting Configure Exascale storage window, enter the storage capacity you want to allocate for Exascale.
7. Click **Submit**.

The Exadata storage section on the Exadata Infrastructure Details page displays the storage details allocated for ASM and Exascale.

Using the Console to Create an Exascale Storage Vault

To create an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.

An Exascale vault is a logical storage container that uses the physical resources provided by Exascale storage pools. Each vault is associated with at least one storage pool.

For more information, see [1.2.5 Vaults](#).

Note

- The minimum configurable or resizable size for an Exascale Storage Vault is 2 TB.
- To successfully create a vault, ensure that the storage pool size exceeds the requested vault size. A portion of the storage pool is reserved for system overhead and is not available for vault allocation. For example, creating a 2 TB vault requires a storage pool larger than 2 TB (typically at least 2.2 TB).
- A vault can be deleted and moved to another compartment.
- A vault can be shared between two or more Exascale VM clusters.

Warning

Attempting to create a vault that is the same size as the total storage pool (for example, creating a 2 TB vault on a 2 TB storage pool) will fail due to insufficient available space. The system reserves a portion of the storage pool for overhead, so the full nominal capacity is not available for vault creation. Always allocate additional storage capacity beyond the requested vault size to avoid failures.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where you want to create an Exascale storage vault.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Exascale storage vaults** tab.
6. Click **Create Exascale storage vault**.
7. On the resulting Create Exascale storage vault window, enter the following:
 - **Compartment**: Select a compartment where you want this resource to be created.
 - **Name**: Enter a descriptive name for the vault.

- Storage capacity for the databases: Enter a reasonable storage capacity within the minimum and maximum values displayed on the screen.
8. Click **Create**.
The Exascale Storage Vault Details page provides key information, including the allocated storage and a list of VM Clusters associated with the vault.
 9. Alternatively, you can create a new vault on the fly using the **Create new storage vault** option while provisioning the Exascale VM cluster.

Using the Console to Scale an Exascale Storage Vault

To scale an Exascale storage vault, be prepared to provide values for the Exascale storage vault configuration.

Note

The minimum configurable or resizable size for an Exascale Storage Vault is 2 TB.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Exascale storage vaults** tab.
6. Click the name of the **Exascale storage vault** you want to scale.
7. On the resulting **Exascale storage vault Details** page, click **Scale storage vault**.
8. On the resulting Scale storage vault window, enter the storage capacity for the databases within the minimum and maximum values displayed on the screen.
9. Click **Scale**.

Using the Console to Move an Exascale Storage Vault to Another Compartment

To move an Exascale storage vault to another compartment, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Exascale storage vaults** tab.

6. Click the name of the **Exascale storage vault** you want to move.
7. On the resulting **Exascale storage vault Details** page, click **Actions**, and then select **Move resource**.
8. On the resulting panel, choose the new compartment for the Exascale storage vault, and click **Move Resource**.

Using the Console to Delete an Exascale Storage Vault

To delete an Exascale storage vault, use this procedure.

Note

An Exascale storage vault can only be deleted when it is not associated with any Exascale VM clusters. To delete an active vault, you must first terminate all VM clusters and the underlying databases that are using the vault.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure where the Exascale storage vault you want to scale reside.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click the **Exascale storage vaults** tab.
6. Click the name of the **Exascale storage vault** you want to delete.
7. On the resulting **Exascale storage vault Details** page, click **Actions**, and then select **Delete**.
8. On the resulting dialog, enter the name of the Exascale storage vault, and click **Delete** to confirm the action.

Cross-Service Data Guard Between ExaDB-D and ExaDB-XS

You can now create a cross-service Oracle Data Guard group across database services.

- Primary database on ExaDB-D with one or more standby databases on ExaDB-XS or ExaDB-D.
- Primary database on ExaDB-XS with one or more standby databases on ExaDB-D or ExaDB-XS.

Note

To ensure maximum availability, Oracle recommends that you place your peer VM cluster in a different Exadata infrastructure from the primary VM cluster.

Standby Database Configuration Options

When adding a standby database, you can specify the following details:

- **Peer VM Cluster Details:** If the target is ExaDB-D, you can select the Exadata Infrastructure.
- **Target Service Selection:** Choose either ExaDB-D or ExaDB-XS. By default, the service initiating the Data Guard group is selected. If a service is unavailable in the selected region, it is disabled with the message: 'Service is not available in this region'.
- **Data Guard Type:** The group can be configured with either Data Guard or Active Data Guard, and each standby database can have a different type.
- **Data Guard Group Limitation:** A primary database is limited to one Data Guard group.
- **Standby Database Creation:** Only one standby database can be added at a time. However, standby databases can be created in any of the following categories without restrictions on their number:
 - Within the same service
 - Across services
 - Within the same Availability Domain (AD)
 - Across Availability Domains within the same region
 - Across regions

Key Considerations for cross-service primary and standby databases

- Both the primary and standby databases must use the same key management solution.
- Data Guard can be configured in Max Performance or Max Availability protection mode with Sync or Async transport type, subject to the following rules:
 - For the first standby database:
 - * Defaults to Max Performance mode with Async transport.
 - * Protection mode and transport type cannot be changed.
 - For the second and subsequent standby databases:
 - * Inherits the protection mode from the first standby.
 - * Defaults to Async transport.
 - * Protection mode and transport type cannot be changed.

Viewing and Editing Data Guard Configuration

- View all databases that are part of the Data Guard group in the Data Guard Group table, irrespective of whether you are on the primary or standby database pages.
- Edit the Data Guard configuration to update:
 - **Data Guard type:** Can be different for each standby database. This can be changed only from a standby database.
 - **Database administrator password:** Can be edited from any database.
 - **Protection mode:** Can be switched between Max Performance and Max Availability from the primary or any standby database.

- **Transport type:** Can be switched between Async and Sync only from a standby database.

① Note

If the Protection Mode is set to Max Availability, then the system verifies that at least one standby database uses Sync transport. If not found, then an error message is displayed:

To achieve zero data loss, you require at least one standby with Sync transport. It is recommended to have a standby with Sync transport in the same service as the primary database.

Switchover and Failover

- **Switchover (on any standby database)**
 - Switchover is performed without enforcing a major version or Release Update (RU) check. However, a warning appears if the standby database has asymmetric configuration, such as differing node counts, memory, or service type.
- **Failover (on any standby database)**
 - Failover is performed without enforcing a major version or Release Update (RU) check. However, a warning appears if the standby database has an asymmetric configuration, such as differing node counts, memory, or service type.

Backup and Restore

- **Scheduled and Automatic Backups**
 - You can schedule automatic backups on the primary database, standby database, or both.
 - Both Object Storage and Recovery Service backups are supported.
 - If backups are configured on both primary and standby databases, they must use the same backup destination type.
- **In-Place Restore of the Same Database within ExaDB-XS**

Restore and recover a database using a backup taken from the same database in ExaDB-XS:

 - Restore primary using a backup taken on the primary database.
 - Restore standby using a backup taken on the standby database
- **In-Place Restore of a Peer Database within ExaDB-XS**

Restore and recover a peer database (which has no backups configured) using a backup from the source database with Recovery Service:

 - **Scenario 1:** Restore the primary using the standby backup.
 - * **Primary database:** ExaDB-XS (no backups configured)
 - * **Standby database:** ExaDB-XS (backups configured to Recovery Service)
 - **Scenario 2:** Restore the standby using the primary backup.
 - * **Primary database:** ExaDB-XS (backups configured to Recovery Service)
 - * **Standby database:** ExaDB-XS (no backups configured)

- **In-Place Restore of a Peer Database Across Services**
Restore and recover a database in ExaDB-XS or ExaDB-D using a backup from the source database in the opposite service (ExaDB-D or ExaDB-XS) with Recovery Service:
 - **Scenario 1:** Restore a peer database in ExaDB-XS using a backup from ExaDB-D
 - * **Use Case 1:** Restore primary using the standby backup.
 - * **Use Case 2:** Restore standby using the primary backup.
 - **Scenario 2:** Restore a peer database in ExaDB-D using a backup from ExaDB-XS
 - * **Use Case 1:** Restore primary using the standby backup.
 - * **Use Case 2:** Restore standby using the primary backup.

Out-of-Place Restore (Creating a New Database)

- **Within ExaDB-XS** You can create a new database in ExaDB-XS using a backup from the same service.
 - Restore within:
 - * The same Availability Domain (AD)
 - * A different AD within the same region
 - * A different region
 - Supports Object Storage or Recovery Service as the backup destination..
- **Across Services**
 - **Scenario 1:** Create a New Database in ExaDB-D Using a Backup from ExaDB-XS
 - * **Source:** ExaDB-XS (database and backups)
 - * **Target:** ExaDB-D (any region or AD)
 - * **Backup Destination:** Object Storage or Recovery Service
 - **Scenario 2:** Create a New Database in ExaDB-XS Using a Backup from ExaDB-D
 - * **Source:** ExaDB-D (database and backups)
 - * **Target:** ExaDB-XS (any region or AD)
 - * **Backup Destination:** Object Storage or Recovery Service

4

Getting Started with Exadata Cloud Infrastructure Deployment

After completing the preparation tasks in [Preparing for Exadata Cloud Infrastructure](#), get started with deploying your Exadata Cloud Infrastructure system following these procedures.

- [Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources](#)
Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.
- [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#)
Oracle Cloud Infrastructure scalable X8M, X9M, and X11M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- [Creating an Exadata Cloud Infrastructure Instance](#)
This topic explains how to create an Oracle Exadata Cloud Infrastructure instance. It also describes how to configure required access to the Oracle Cloud Infrastructure Object Storage service and set up DNS.
- [Connecting to an Exadata Cloud Infrastructure Instance](#)
This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- [Using the Console to Create API Access Control](#)
Use this procedure to create API Access Control for your ExaDB-D resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.
- [Best Practices for Exadata Cloud Infrastructure Instances](#)
Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Exadata Cloud Infrastructure instance:
- [Moving to Oracle Cloud Using Zero Downtime Migration](#)

Tagging Oracle Exadata Database Service on Dedicated Infrastructure Resources

Tagging is a powerful foundational service for Oracle Cloud Infrastructure (OCI) that enables users to search, control access, and do bulk actions on a set of resources based on the tag.

Importance of Tagging

Using the Oracle Cloud Infrastructure (OCI) tagging system, you can tag resources per your organizational scheme allowing you to group resources, manage costs, and give insights into usage. Tags also help you build a governance model around security and Maximum Availability Architecture (MAA). As your organization expands its cloud footprint, it can become challenging to keep track of the deployment architectures, security best practices, MAA, application tier, and so on. Using metadata tags to identify workload attributes can help keep up with the security and availability of your tenancy without cost overruns.

To enable customers to manage OCI resources securely and cost-effectively, Oracle provides a set of pre-defined tags in line with best practices for tagging resources. These tags are grouped into two namespaces, the `oracleStandard` namespace, and the `OracleApplicationName` namespace. You can think of a tag namespace as a container for your tag keys.

Consider a scenario where your organization has multiple cloud resources such as Exadata Infrastructure, VM Cluster, DB Home, Oracle Database and VM Cluster Networks across multiple compartments in your tenancy. Suppose you wish to track these cloud resources for specific purposes, report on them, or take bulk actions. In that case, you will need a system that lets you group these resources based on different criteria such as environment, criticality, target users, application, etc. You can achieve this by applying appropriate tags to these resources.

For example, you may tag all resources in your development stack with `Oracle-Standard.Environment=Dev` or for a business-critical application stack set `Oracle-Standard.Criticality=High` or `Extreme`. In the event of service disruptions due to various reasons, you would then be able to quickly identify all OCI resources associated with an application or business function or be able to separate critical and non-critical workloads.

Tagging can also help you deploy optimized configurations based on workload attributes identified via tags. For example, database deployments for the PeopleSoft application require a specific configuration. Setting the `ApplicationName` and `AppMajorVersion` tags while deploying an Oracle Database can ensure that the database is configured and ready for the particular application, for example, PeopleSoft out of the box.

Moreover, integration with the Cloud Advisor OCI service can provide you with direct, deep insight into how well your cloud services adhere to the corporate guidelines and help your management govern with a vision. See *Cloud Advisor Overview* for more details.

Adding Tags

You can tag resources using the Oracle Cloud Infrastructure (OCI) console, command-line interface, or SDK.

There are many cloud resources that can be tagged in an Oracle Exadata Database Service on Dedicated Infrastructure deployment. Exadata Infrastructure, VM Cluster, DB Home, Oracle Database, Autonomous Exadata VM Cluster, Autonomous Container Database, Autonomous Database, and VM Cluster Networks are some of them. Tags can either be applied while creating the resources or modified later. For example, you can apply tags to an Autonomous Container Database (ACD) while provisioning the ACD or add them later from its **Details** page.

See *How Tagging Works* for more details on using tags. Tagging integrates with Oracle Cloud Infrastructure authorization system. You can use IAM policy controls to enable delegation or restriction of tag manipulation. See *Authentication and Authorization* to learn about the permissions required to work with defined and free-form tags. (Required) Enter introductory text here, including the definition and purpose of the concept.

✓ Tip

For a "try it out" tutorial that demonstrates implementing tags in Oracle Autonomous Database, refer to *Lab 14: Oracle Standard Tags in Oracle Autonomous Database Dedicated for Fleet Administrators Workshop* on Oracle LiveLabs.

Your tenancies come with a library of standard tags that would apply to most resources. These tags are currently available as a set of Tag Namespaces that your governance administrators can deploy. OCI best practices recommend applying these tags to all resources a standard tag

can be applied to. Besides reporting and governance, OCI service automation can deliver workload-specific optimizations based on standard tag values.

For example, database deployments for the PeopleSoft application require a specific configuration. By setting the appropriate application tag key in the `Oracle-ApplicationName` tag namespace while deploying an Autonomous Database, can ensure that the database is configured ready for the particular application, for example, PeopleSoft out of the box.

Figure 4-1 Tagging Example

The screenshot shows the OCI console interface for tagging a resource. It features two tabs: 'Encryption Key' and 'Tags'. The 'Tags' tab is selected. Below the tabs, there is a heading: 'Optional tags to organize and track resources in your tenancy. [How do I use tags?](#)'. The main area contains three input fields: 'Tag Namespace' with a dropdown menu showing 'Oracle-ApplicationName', 'Tag Key' with a dropdown menu showing 'PeopleSoft', and 'Tag Value' with a dropdown menu showing '9.2'. A small 'X' icon is to the right of the 'Tag Value' field. Below the 'Tag Value' dropdown, a list of options is visible: '9.2' (highlighted) and '9.1'. A 'Tag' button is located to the right of the dropdown list.

Oracle Standard Tags

Your tenancy governance administrators can deploy the standard tags at the tenancy level and may also mark certain tags as required, thereby enforcing tags on resources in those compartments. The following are the standard tags defined in the namespace called `OracleStandard`. For more information about importing standard tags, see *To import standard tags* under the *Managing Tag Namespaces* section.

Table 4-1 Oracle Standard Tags

Tag Key	Tag Value Options	Description
<code>OracleStandard.Criticality</code>	<ul style="list-style-type: none"> • Extreme • High • Medium • Low 	<p>Enables tiering of resources in line with corporate application classification standards. Customer governance can use this tag for reporting and ensuring resources are configured as per the guideline for the tier they belong to.</p> <p>For example, a database resource with <code>OracleStandard.Criticality</code> set to Extreme or High may require the highest availability SLA and may need to be configured with Autonomous Data Guard.</p>

Table 4-1 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
OracleStandard.Environment	<ul style="list-style-type: none"> • Dev • Test • Prod • Pre-Prod • Staging • Trial • Sandbox • User Testing 	Denotes a resource lifecycle. In the case of databases, it helps determine consolidation density, database distribution across containers, set maintenance plans, and manage clones.
OracleStandard.Sensitivity	<ul style="list-style-type: none"> • Public • Internal • Sensitive • Highly Sensitive • Extremely Sensitive 	An application or database classification tag. OracleStandard.Sensitivity set to Highly Sensitive may indicate that an access control list or certain Network Security Group (NSG) enforcement is mandatory to restrict access.
OracleStandard.Regulation	Refer to <i>List of Compliance Regulations</i> for values.	Denotes one or more compliance regulations that a resource must adhere to. Tag administrators may add values to the list from the OCI Governance and Administration console. Refer to <i>Using Predefined Values</i> for more details.
OracleStandard.TargetUsers	<ul style="list-style-type: none"> • Public • Customers • Partners • Company • Division • Department • Workgroup 	Denotes the end users of a resource. Another form of resource classification that helps determine target users and allows governance teams to set corporate standards based on user or application type.
OracleStandard.EndUserCount	<ul style="list-style-type: none"> • 1 • 10 • 100 • 1000 • 10000 • 100000 • 1000000 • 10000000 	An approximate count of end-users. This tag helps determine the number of users impacted or the blast radius during an availability or security event. This also helps prioritize recovery efforts in the event of major outages affecting a large number of cloud resources.
OracleStandard.OwnerEmail	Free form tag. For example <i>john.smith@acme.com</i> or <i>app_support_grp@acme.com</i>	Denotes the email address of the resource owner.

Table 4-1 (Cont.) Oracle Standard Tags

Tag Key	Tag Value Options	Description
OracleStandard.Org	<ul style="list-style-type: none"> • HR • Finance • Marketing • Sales • Legal • R&D • Customer Support • Internal Support • Manufacturing 	Identifies the customer's line of business or department that owns or uses the resource. This may help with cost aggregation reports and determining usage across business units. Tag administrators may add relevant values to the list from the OCI Governance and Administration console. Refer to <i>Using Predefined Values</i> for more details.
OracleStandard.CostCenter	<ul style="list-style-type: none"> • 12345 • WebMarketing 	Freeform field for cost center.
OracleStandard.RecoveryTimeObjectiveMinutes	0-10080	Time in minutes. Denotes the maximum time within which the resource is required to recover from a failure.
OracleStandard.RecoveryPointObjectiveMinutes	0-1440	Time in minutes. Maximum data loss tolerance for a data store resource such as a database or a storage device.

List of Compliance Regulations

Table 4-2 List of Compliance Regulations

Regulation	Description
PCI DSS	Payment Card Industry Data Security Standard
HIPAA	Health Insurance Portability and Accountability Act
ISO	International Standards Organization
SOC1	System and Organization Controls 1
SOC 2	System and Organization Controls 2
FedRamp	Federal Risk and Authorization Management Program
GLBA	Gramm–Leach–Bliley Act
CCPA	California Consumer Privacy Act
SOX	Sarbanes Oxley
NIST	National Institute of Standards and Technology - Cyber Security
FISMA	Federal Information Security Management
HITECH	Health Information Technology for Economic and Clinical Health Act
FERPA	Family Educational Rights and Privacy Act (Student privacy)
FACTA	Fair and Accurate Credit Transaction Act

Table 4-2 (Cont.) List of Compliance Regulations

Regulation	Description
Texas HB300	Texas Medical Records Privacy Act
CIS	Center for Internet Security
CJIS	Criminal Justice Information Services Security Policy
C-TPAT	Customs-Trade Partnership Against Terrorism
COPPA	Children's Online Privacy Protection Act
PIPED Act, or PIPEDA	Personal Information Protection and Electronic Documents Act
GDPR	General Data Protection Regulation
PIPL	Personal Information Protection Law

Oracle Application Name Tags**Table 4-3 Oracle Application Name Tags**

Tag Key	Tag Value Options	Description
Hyperion	<ul style="list-style-type: none"> • 11.2 • 11.1 	Denotes the version of the Hyperion application.
JD Edwards	<ul style="list-style-type: none"> • 9.2 • 9.1 • 9.0 	Denotes the version of the JD Edwards application.
Oracle_E-Business_Suite	<ul style="list-style-type: none"> • 12.2 • 12.1 • 12.1 • 11i 	Denotes the version of the Oracle E-Business Suite application.
PeopleSoft	<ul style="list-style-type: none"> • 9.2 • 9.1 	Denotes the version of the PeopleSoft application.
Siebel	<ul style="list-style-type: none"> • 8.2 • 8.1 	Denotes the version of the Siebel application.
Other_Oracle_Application	Free form tag in string format.	Can be used to denote any application other than those listed above. You can enter the application name as a string value.

Related Topics

- [To Import standard tags](#)
- [Cloud Advisor Overview](#)
- [Oracle Autonomous Database Dedicated for Fleet Administrators Workshop](#)
- [How Tagging Works](#)
- [Authentication and Authorization](#)
- [Managing Tag Namespaces](#)

- [Using Predefined Values](#)

Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure

Oracle Cloud Infrastructure scalable X8M, X9M, and X11M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.

- [The Exadata Cloud Infrastructure Resource Model](#)
Exadata Cloud Infrastructure instances are provisioned with an infrastructure model that uses two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.
- [The Cloud Exadata Infrastructure Resource](#)
The infrastructure resource is the top-level (parent) resource.
- [The Cloud VM Cluster Resource](#)
The VM cluster is a child resource of the infrastructure resource.
- [Additional Exadata Cloud Infrastructure Instance Resources](#)
The new Exadata resource model includes databases, backups, Data Guard associations, work requests, database homes, and server nodes.
- [The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes](#)
- [The New Exadata Cloud Infrastructure Resource Model](#)

The Exadata Cloud Infrastructure Resource Model

Exadata Cloud Infrastructure instances are provisioned with an infrastructure model that uses two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.

The Cloud Exadata Infrastructure Resource

The infrastructure resource is the top-level (parent) resource.

At the infrastructure level, you control the number of database and storage servers. You also control Exadata system maintenance scheduling at the Exadata infrastructure level. This resource is created using the [CreateCloudExadataInfrastructure](#) API.

See *Scaling Exadata X8M and X9M Compute and Storage* for information on scaling the X8M or X9M cloud Exadata infrastructure resource.

Note

After adding storage or database servers to the infrastructure resource, you must then add them to the system VM clusters to utilize the new capacity.

Related Topics

- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.

The Cloud VM Cluster Resource

The VM cluster is a child resource of the infrastructure resource.

The VM cluster resource provides a link between your Exadata cloud infrastructure resource and Oracle Database. Networking, OCPU count, IORM, and Oracle Grid Infrastructure are configured and managed at the VM cluster level. This resource is created using the [CreateCloudVmCluster](#) API.

See [To add database server or storage server capacity to a cloud VM cluster](#) for information on adding available storage or database servers to the VM cluster. Note that you must add servers to the infrastructure resource before you can add capacity to the VM cluster.

Multi-VM enabled Infrastructure support multiple VM clusters in a single infrastructure

Exadata Cloud Infrastructure instances that are NOT Multi-VM enabled Infrastructure only support creating a single cloud VM cluster.

Related Topics

- [About IORM](#)
The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model.

Additional Exadata Cloud Infrastructure Instance Resources

The new Exadata resource model includes databases, backups, Data Guard associations, work requests, database homes, and server nodes.

Note

The database server file system for database server nodes (also known as "virtual machines") has changed with the X8M generation of hardware. See [The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes](#) for details on the X8M database server node file system.

The X8M, X9M, and X11M Virtual Machine File System Structure Important File System and Sizes

Table 4-4 The X8M and X9M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
/	15 GB	Max supported size 900 GB. File system size can only be increased.
/u01	20 GB	Max supported size 900 GB. File system size can only be increased.

Table 4-4 (Cont.) The X8M and X9M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
/u01/./grid	50 GB	File system is not resizable.
/u02	60 GB	Max supported size 900 GB. File system size can be increased or decreased.
/acfs01	100 GB	Contact Oracle Support to resize.
/boot	509 MB	File system is not resizable.
/crashfiles	20 GB	File system is not resizable.
/home	4 GB	Max supported size 900 GB. File system size can only be increased.
/var	5 GB	Max supported size 900 GB. File system size can only be increased.
/var/log	18 GB	Max supported size 900 GB. File system size can only be increased.
/var/log/audit	3 GB	Max supported size 900 GB. File system size can only be increased.
/tmp	3 GB	Max supported size 900 GB. File system size can only be increased.

Table 4-5 The X11M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
/	15 GB	Max supported size 900 GB. File system size can only be increased.
/u01	20 GB	Max supported size 900 GB. File system size can only be increased.
/u01/./grid	50 GB	File system is not resizable.
/u02	60 GB	Max supported size 900 GB. File system size can be increased or decreased.
/acfs01	100 GB	Contact Oracle Support to resize.
/boot	412 MB	File system is not resizable.
/crashfiles	20 GB	File system is not resizable.
/home	4 GB	Max supported size 900 GB. File system size can only be increased.

Table 4-5 (Cont.) The X11M Virtual Machine File System Structure Important File System and Default Sizes

Filesystem Mount	Size	Configuration
/var	5 GB	Max supported size 900 GB. File system size can only be increased.
/var/log	18 GB	Max supported size 900 GB. File system size can only be increased.
/var/log/audit	3 GB	Max supported size 900 GB. File system size can only be increased.
/tmp	3 GB	Max supported size 900 GB. File system size can only be increased.

The New Exadata Cloud Infrastructure Resource Model

Exadata Cloud Infrastructure instances can now only be provisioned with a new infrastructure resource model that replaced the DB system resource.

In the new model, there are two resources, the **cloud Exadata infrastructure** resource, and the **cloud VM cluster** resource.

The X8M and X9M system models are only compatible with the new resource model. For provisioning new X7 and X8 systems, Oracle recommends using the new resource model so that your instance will not have to be switched to the new resource model later.

Note

No new systems can be provisioned with the old DB system resource model/APIs after May 15th, 2021. Support for the old DB system resource model/APIs on existing systems will end on January 15th, 2021. After this date, old APIs will stop working and the only action available will be to list DB System details and perform the switch to the new API. Oracle recommends that you migrate your Exadata Cloud Infrastructure instances to the new resource model APIs as soon as possible. Converting to the new resource model does not involve any system downtime.

Creating an Exadata Cloud Infrastructure Instance

This topic explains how to create an Oracle Exadata Cloud Infrastructure instance. It also describes how to configure required access to the Oracle Cloud Infrastructure Object Storage service and set up DNS.

When you create an Exadata Cloud Infrastructure instance using the Console or the API, the system is provisioned to support Oracle databases. Along with the Infrastructure, a VM cluster, an initial Database Home and database are created. You can create additional Database Homes and databases at any time by using the Console or the Oracle Cloud Infrastructure

API. The service creates an initial database based on the options you provide and some default options described later in this topic.

- [Resources to Be Created](#)
- [Prerequisites for Creating an Cloud Exadata Infrastructure Instance](#)
You need a SSH key pair key and a Virtual Cloud Network (VCN) to create an infrastructure instance.
- [Default Options for the Initial Database](#)
Default option simplify launching an Exadata Cloud Infrastructure instance in the Console and when using the API.
- [Using the Console to Create Infrastructure Resources](#)
Console tasks required to create cloud resources

Resources to Be Created

You will provision Exadata Cloud Infrastructure infrastructure and VM cluster resources separately.

- **Cloud Exadata infrastructure** resource: The infrastructure resource is the top-level (parent) resource. At the infrastructure level, you control the number of database and storage servers. You also control Exadata system maintenance scheduling at the Exadata infrastructure level.
- **Cloud VM cluster** resource: The VM cluster is a child resource of the infrastructure resource, providing a link between your Exadata cloud infrastructure resource and Oracle Database. Networking, OCPU count, IORM (see [About IORM](#), and Oracle Grid Infrastructure are configured and managed at the VM cluster level. To create a cloud VM cluster, you must have an existing Cloud Exadata infrastructure resource to house the VM cluster.

Note

- Exadata Cloud Infrastructure only supports using the new resource model (consisting of separate Exadata infrastructure and VM cluster resources) to provision Exadata Cloud Infrastructure instances, regardless of the hardware shape family you are choosing (X7, X8, X8M, or X9M). The DB system resource model and APIs are deprecated for Exadata Cloud Infrastructure.
- Multi-VM enabled Infrastructure supports the creation of up to 8 VM clusters in an Infrastructure. >> Exadata Infrastructures with X8M and above generation of DB Servers can support a maximum of 8 VM clusters across all DB Servers. Maximum number of clusters across the infrastructure depends on the resources available per DB server and is subject to the per DB Server maximum VM limit. For more information, see [Overview of VM Cluster Node Subsetting](#).
- An Exadata Cloud Service Infrastructure instance that is NOT Multi-VM enabled supports only one cloud VM cluster

Prerequisites for Creating an Cloud Exadata Infrastructure Instance

You need a SSH key pair key and a Virtual Cloud Network (VCN) to create an infrastructure instance.

- The proper IAM policy is required to proceed. See [Required IAM Policy for Exadata Cloud Infrastructure](#)
- The public key, in OpenSSH format, from the key pair that you plan to use for connecting to the system via SSH. A sample public key, abbreviated for readability, is shown below.

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAA...lo/gKMLVM2xzclxJr/  
Hc26biw3TXWGEakrK10Q== rsa-key-20160304
```

For more information, see *Managing Key Pairs on Linux Instances* .

- A correctly configured virtual cloud network (VCN) to launch the system in. Its related networking resources (gateways, route tables, security lists, DNS, and so on) must also be configured as necessary for the system. For more information, see *Network Setup for Exadata Cloud Infrastructure Instances* .

Related Topics

- [Managing Key Pairs on Linux Instances](#)
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.

Default Options for the Initial Database

Default option simplify launching an Exadata Cloud Infrastructure instance in the Console and when using the API.

The following default options are used for the initial database:

- **Console Enabled:** False
- **Create Container Database:** False for version 11.2.0.4 databases. Otherwise, true.
- **Create Instance Only (for standby and migration):** False
- **Database Home ID:** Creates a database home
- **Database Language:** AMERICAN
- **Database Sizing Template:** odb2
- **Database Storage:** Automatic Storage Management (ASM)
- **Database Territory:** AMERICA
- **Database Unique Name:** The user-specified database name and a system-generated suffix, for example, dbtst_phx1cs.
- **PDB Admin Name:** pdbuser (Not applicable for version 11.2.0.4 databases.)

Using the Console to Create Infrastructure Resources

Console tasks required to create cloud resources

- [To create a Cloud Exadata infrastructure resource](#)
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [VM Cloud Automation Software Update Management](#)

- [Configuring Network Resources for Recovery Service](#)
Use an existing IP4-only subnet in the database VCN for Recovery Service operations. Define security rules to control the backup traffic between your database and Recovery Service. Finally, register the private subnet as a Recovery Service subnet.
- [Autonomous Recovery Service Checklist](#)

To create a Cloud Exadata infrastructure resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. Click **Create Exadata Cloud Infrastructure**.
4. **Compartment:** Select a compartment for the Exadata infrastructure.
5. **Display name:** Enter a display name for the Exadata infrastructure. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the Cloud Exadata infrastructure resource. Avoid entering confidential information.
6. **Select an availability domain:** The availability domain in which the Exadata infrastructure resides.
7. **Select the Exadata Cloud Infrastructure model:** Select either a fixed-shape system (quarter, half, or full rack X7-2 or X8-2 shapes), or a scalable system (X8M-2, X9M-2, or X11M).
 - X11M:** If you select the flexible X11M cloud infrastructure model, your initial Exadata Cloud Infrastructure instance can have a minimum of 2 database servers and 3 storage servers up to 32 database servers and 64 storage servers. You will also have to select the database server and storage server type. This will default to X11M for the database server type and X11M-HC for the storage server type. After provisioning, you can scale the service instance as needed by adding additional storage servers, compute servers, or both.
 - X9M-2:** If you select the flexible X9M-2 cloud infrastructure model, your initial Exadata Cloud Infrastructure instance can have a minimum of 2 database servers and 3 storage servers up to 32 database servers and 64 storage servers. After provisioning, you can scale the service instance as needed by adding additional storage servers, compute servers, or both.
 - X8M-2:** If you select the flexible X8M-2 cloud infrastructure model, your initial Exadata Cloud Infrastructure instance can have a minimum of 2 database servers and 3 storage servers (the equivalents of an X8 quarter rack shape) up to 32 database servers and 64 storage servers. After provisioning, you can scale the service instance as needed by adding additional storage servers, compute servers, or both.
 - X7 and X8:** If you select an X7 or X8 system, you are given the choice of provisioning a quarter, half, or full rack. See *Exadata Shape Configuration* for hardware and capacity details.
- Exadata Base:** The Exadata base shape comes in a single configuration, and provides an economical alternative to provisioning a quarter rack system. See *Exadata Shape Configuration*
8. If you selected a flexible shape (X8M-2, X9M-2, or X11M), specify the Compute and storage configuration. You can specify Database servers from minimum of 2 up to 32. You can specify Storage servers from minimum of 3 up to 64.
9. Select the time zone: Choose one of :

- **UTC**
 - **Select another timezone**
 - (Browser detected) timezone
10. Configure Maintenance.
Click **Edit Maintenance Preferences** to change the values.
- On the **Configure Maintenance** page, configure the following:
- **Maintenance scheduling preference:** Oracle managed schedule
 - **Choose a maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
 - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
 - * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- Click **Save Changes**.

Note

From the next maintenance run onwards, executions will occur according to Oracle's schedules.

- **Maintenance scheduling preference:** Customer managed schedule
 - **Maintenance schedule:** Define maintenance preferences for this infrastructure.

Note

Changes will take effect from the next maintenance run.

- * **Configure maintenance preference:** Define maintenance time preferences for each quarter. If more than one preference is defined for a quarter, Oracle automation will select one of them to perform maintenance on all components in your infrastructure.
Select at least one month every two quarters.
- * **Specify a schedule:** Choose your preferred week, weekday, start time, and lead time for infrastructure maintenance.
 - * Optional. Under **Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a week of the month, Oracle will run the maintenance update in a week to minimize disruption.
 - * Optional. Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.
 - * Optional. Under **Hour of the day**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
 - * Under **Notification Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly released maintenance update is scheduled to account for your required minimum period of advanced notification.
- * **Choose a maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
- * **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
- * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- * **Show advanced options:**
 - * Enable monthly security infrastructure maintenance: Select this check box to perform monthly security infrastructure maintenance.
 - **Maintenance schedule:** Use maintenance window preferences from a scheduling policy

During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers, into the maintenance windows from your policy based on duration. After provisioning the infrastructure, you can update the scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

 - * Click **Select policy**.
 - * In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.

You can also create a maintenance scheduling policy and use it. For more information, see [Create a Maintenance Scheduling Policy](#). Note that you can add additional maintenance windows to the policy after creating it. For more information, see [Add Additional Maintenance Windows to a Maintenance Scheduling Policy](#).
 - * Click **Save changes**.
- 11.** In the **Provide maintenance details** : **Provide up to 10 unique maintenance contact email addresses**. Click **Add Contact**.
In the **Contact Email** field, provide the email ID of a desired contact.

Note

At least one Contact is required.

Click **Add Contact** to add another contact.

- 12.** Click **Show Advanced Options** to specify advanced options for the initial database.

In the **Tags** tab, you can add tags to the database. 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 if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

13. Click **Create Exadata Infrastructure**. The Cloud Exadata infrastructure appears in the Exadata Infrastructure list with a status of Provisioning. The infrastructure's icon changes from yellow to green (or red to indicate errors).

WHAT NEXT?

After the Cloud Exadata infrastructure resource is successfully provisioned and in the Available status, you can create a cloud VM cluster as described in *To create a cloud VM cluster resource* on your infrastructure. You must provision both an infrastructure resource and a VM cluster before you can [create your first database](#) in the new Exadata Cloud Infrastructure instance.

Related Topics

- [Exadata Shape Configuration](#)
This topic describes the available Exadata Cloud Infrastructure instance shapes in Oracle Cloud Infrastructure.
- [Resource Tags](#)

To create an ASM cloud VM cluster

To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Note

To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first [created a Cloud Exadata infrastructure resource](#).

Note

Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature [Create and Manage Multiple Virtual Machines per Exadata System \(MultiVM\) and VM Cluster Node Subsetting](#) was released only support creating a single cloud VM cluster.

Note

When you provision an Exadata VM cluster in Exadata Database Service on Oracle Database@Google Cloud, an Identity Connector is automatically created and associated with the VM cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
4. **Compartment:** Select a compartment for the VM cluster resource.
5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the VM cluster. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

7. VM Cluster Type:**Note**

You cannot change the VM cluster type after deploying the VM cluster. If you wish to change the VM cluster type, you must create a new VM cluster and migrate the database to the new cluster.

- **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
 - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.
8. **Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Select DB Servers** to select from the available DB servers, and then click **Save**.

VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

In the **Resource allocation per VM** pane:

- Specify the number of OCPU/ECPU you want to allocate to each of the VM cluster's virtual machine compute nodes. For VM clusters created on X11M Exadata infrastructure specify ECPU's. For VM Clusters created on X10M and earlier Exadata infrastructure, specify OCPU's. The minimum is 2 OCPU per VM for X10M and earlier

infrastructure or 8 ECPUs per VM for VM clusters created on X11M Exadata infrastructure. The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU or ECPU cores you are allocating.

- Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
- Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.

Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

In addition to `/u02`, you can specify the size of additional local file systems.

For more information and instructions to specify the size for each individual VM, see [Introduction to Scale Up or Scale Down Operations](#).

- Click **Show additional local file systems configuration options**.
- Specify the size of `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

Note

- * You can only expand these file systems and cannot reduce the size once expanded.
- * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

- After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

9. Exadata storage:

- **Specify the usable Exadata storage TB.** Specify the storage in multiples of 1 TB. Minimum: 2 TB
- **Allocate storage for Exadata sparse snapshots:** Select this configuration option if you intend to use snapshot functionality within your VM cluster. If you select this option, the SPARSE disk group is created, which enables you to use VM cluster snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and snapshot functionality will not be available on any database deployments that are created in the environment.

Note

The storage configuration option for sparse snapshots cannot be changed after VM cluster creation.

- **Allocate storage for local backups:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata Cloud Infrastructure instance. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.

Note

The storage configuration option for local backups cannot be changed after VM cluster creation.

10. Version:

- **Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release (19c and 26ai) that you want to install on the VM cluster. The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

Note

Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 26ai:

- Exadata Guest VM running Exadata System Software 23.1.8
- Exadata Infrastructure running Exadata System Software 23.1.x

- **Exadata guest version:**
 - **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
 - * The **Change image** button is not enabled.
 - * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
 - * The Exadata guest version will be the same as that of the host OS.
 - **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
 - * The Exadata guest version defaults to the latest (23.1.3.0).
 - * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
 - * The **Change image** button is enabled.
 - * Click **Change image**.
The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).
The most recent release for each major version is indicated by "(latest)".
 - * Slide **Display all available versions**.
Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
 - * Choose a version.
 - * Click **Save Changes**.

11. SSH Keys: Add the public key portion of each key pair you want to use for SSH access to the VM cluster:

- **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.

- **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
- **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.

12. **Network settings:** Specify the following:

Note

IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect

You do not have the option to choose between IPv4 (single stack) and IPv4/IPv6 (dual stack) if both configurations exist. For more information, see [VCN and Subnet Management](#).

- **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
- **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment. Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
- **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable. Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in [Managing Exadata Database Backups](#).

Note

In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, [Configuring Network Resources for Recovery Service](#).

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress **security rules** to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified. For more information, see **Network Security Groups** and *Network Setup for Exadata Cloud Infrastructure Instances*.
Note that if you choose a subnet with a **security list**, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

To use network security groups:

- Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You

can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.

- Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.
- To use additional NSGs with the network, click **+;Another Network Security Group**.

Note

To provide your cloud VM Cluster resources with additional security, you can use Oracle Cloud Infrastructure Zero Trust Packet Routing to ensure that only resources identified with security attributes have network permissions to access your resources. Oracle provides Database policy templates that you can use to assist you with creating policies for common database security use cases. To configure it now, you must already have created security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. Click **Show Advanced Options** at the end of this procedure.

Be aware that when you provide security attributes for a cluster, as soon as it is applied, all resources require a Zero Trust Packet policy to access the cluster. If there is a security attribute on an endpoint, then it must satisfy both network security group (NSG) and Oracle Cloud Infrastructure Zero Trust Packet Routing policy (OCI ZPR) rules.

- **To use private DNS Service**

Note

A Private DNS must be configured before it can be selected. See *Configure Private DNS*

- Check the **Use private DNS Service** check box.
- Select a private view. Click **Change Compartment** to select a private view in a different compartment.
- Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.
- **Hostname prefix:** Your choice of hostname for the Exadata VM cluster. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata VM cluster is 12.

Caution

The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field

displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted.

If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.

- **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.
13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
- **License Included** means the cost of the cloud service includes a license for the Database service.
 - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the check boxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).

- **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).
15. Click **Show Advanced Options** to specify advanced options for the VM cluster:
- **Time zone:** This option is located in the **Management** tab. The default time zone for the VM cluster is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Note

If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

- **SCAN Listener Port:** This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521.

Note

Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

- **Zero Trust Packet Routing (ZPR):** This option is located in the **Security attributes** tab. Select a namespace, and provide the key and value for the security attribute. To complete this step during configuration, you must already have set up security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. You can also add security attributes after configuration, and add them later. For more information about adding Oracle Exadata Database Service on Dedicated Infrastructure specific policies, see [Policy Template Builder](#).
- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

Note

Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Note

- The freeze period can extend for a maximum of 45 days from the start date.
- Oracle automation will automatically apply updates with critical security fixes (CVSS \geq 9) even during a configured freeze period.

- **Tags:** 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.

16. Click **Create**.

WHAT NEXT?

After your VM cluster is successfully created and in the **Available** state.

- You can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can **create your first database** in the cluster by clicking **Create Database**
- The **SCAN IP address (IPv4)** and **SCAN IP address (IPv6)** fields in the **Network** section on the VM Cluster Details page displays the dual stack IP address details.
- The **Cloud Automation Update** field in the **Version** section on the VM Cluster Details page displays the freeze period you have set.

Related Topics

- [Network Security Groups](#)
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- [Security Lists](#)
- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.
- [Resource Tags](#)
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Getting Started with Events](#)
- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.

- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

VM Cloud Automation Software Update Management

VM Cloud Automation software updates to Oracle-managed agents and tools on guest VMs ensure access to the latest Oracle Cloud features and support ongoing operational efficiency. Environments may depend on specific versions for custom scripts and workflows, or require updates to align with business cycles and critical operations.

With these capabilities, customers benefit from greater control and flexibility over their update process, gaining the ability to minimize risk by proactively testing updates and ensuring operational stability by protecting critical business windows from disruption. This approach supports the stability of business operations and the efficient adoption of Oracle Cloud enhancements.

Customers can control the scheduling and application of updates to VM Cloud Automation software in several ways:

- **Daily update scheduling:** Configure specific periods during the day when guest VMs poll for and apply updates, aligning software maintenance with preferred maintenance windows and operational requirements.
- **Phased rollout and early testing:** Assign specific clusters to receive updates first, enabling phased rollouts and allowing updates to be evaluated in non-production environments before wider deployment.
- **Freeze periods:** Define windows during which updates are paused to ensure uninterrupted performance during business-critical operations.

All updates are applied online, maintaining continuity for running guest VMs and their databases.

- [What to Expect During VM Cloud Automation Software Updates](#)
- [Manage VM Cloud Automation Update Preferences](#)
- [Configure the freeze period](#)
- [Cancel Cloud Automation Update preference while the freeze period is in effect](#)
- [Re-enable Cloud Automation Update](#)

What to Expect During VM Cloud Automation Software Updates

Handling Background Activities

When an update is scheduled, Oracle ensures that updates are applied only when the VM Cluster is not engaged in background activities that could be incompatible with the update process. If the VM Cluster is performing tasks such as configuration changes, software updates, backup jobs, or other background workflows during the configured update timeslot, the automation update will be deferred. Updates will be retried on the next available day during your defined timeslot, provided no conflicting activities are occurring at that time.

Update Delivery and Early Access

VM Cloud Automation software updates are delivered to your VM Cluster according to your configured schedule once an update is published by Oracle and becomes generally available.

For customers who want to validate updates before deploying them more broadly, an option exists to enable early access for specific test clusters. To enable early access, set the early access flag in your cloud automation update configuration. A VM Cluster with early access enabled will receive and apply updates one week ahead of general availability, allowing you to thoroughly test and certify new updates in a non-production environment before production deployment.

This approach helps ensure your VM Clusters remain stable and reliable, while providing opportunities to proactively validate changes and minimize the risk of operational disruptions.

Manage VM Cloud Automation Update Preferences

Cloud Automation Update

Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster.

Set the start time for cloud automation updates.

Note

Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable Early Access for Cloud Tools Update

VM clusters designated for early access receive updates 1-2 weeks before they are available to other clusters. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period

Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Note

- The freeze period can extend for a maximum of 45 days from the start date.
- Oracle automation will automatically apply updates with critical security fixes (CVSS \geq 9) even during a configured freeze period.

Configure the freeze period

You can always update the freeze period to extend up to 45 days from the freeze period start date.

1. On the Exadata VM Cluster Details page, click **More actions** and then select **Manage Cloud Automation Update**.
2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period slider** to enable.

Note

The freeze period can be extended for a maximum of 45 days from the freeze period start date.

3. Click **Save**.

Cancel Cloud Automation Update preference while the freeze period is in effect

1. On the Exadata VM Cluster Details page, click **More actions** and then select **Manage Cloud Automation Update**.
2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period slider** to disable.
3. Click **Save**.
4. On the Cancel freeze period dialog, click **Cancel freeze period** to confirm
Oracle automation requires at least 7 days to apply any pending updates before a new freeze period can be configured. You can set up a new freeze period starting 7 days from the cancellation date of the previous freeze period.

Cloud Automation Update will be applied during the configured time period when available.

Re-enable Cloud Automation Update

1. On the Exadata VM Cluster Details page, click **More actions** and then select **Manage Cloud Automation Update**.
(or)
Click the Cloud Automation Update **Edit** link in the **Version** section.
2. On the resulting Manage Cloud Automation Update page, move the **Configure freeze period slider** to enable.

Note

Once the freeze period expires (the end date has been reached) or the freeze period is canceled, Oracle automation requires 7 days to apply any pending updates before the freeze period is re-enabled again on the cluster.

3. Click **Save**.

Configuring Network Resources for Recovery Service

Use an existing IP4-only subnet in the database VCN for Recovery Service operations. Define security rules to control the backup traffic between your database and Recovery Service. Finally, register the private subnet as a Recovery Service subnet.

- [About Using a Private Subnet for Recovery Service](#)
Recovery Service uses a private subnet inside a virtual cloud network (VCN) where your database resides. The private subnet defines the network path for backups between your database and Recovery Service.
- [Review Networking Service Permissions to Configure a Subnet](#)
Ensure that you have these Networking Service permissions required to create a subnet in the database VCN and to assign security rules for Recovery Service.
- [Review Subnet Size Requirements and Security Rules for Recovery Service Subnet](#)
The security rules are necessary to allow backup traffic between a database and Recovery Service.
- [Create a Recovery Service Subnet in the Database VCN](#)
Use the OCI Console to configure a private subnet for Recovery Service in your database virtual cloud network (VCN).
- [Register Recovery Service Subnet](#)
After you have created a private subnet for Recovery Service in your database VCN, use this procedure to register the subnet in Recovery Service.

About Using a Private Subnet for Recovery Service

Recovery Service uses a private subnet inside a virtual cloud network (VCN) where your database resides. The private subnet defines the network path for backups between your database and Recovery Service.

Oracle recommends that your database VCN must have a single private subnet dedicated for backups to Recovery Service. Your Oracle Cloud database can reside in the same private subnet used by Recovery Service, or in a different subnet within the same VCN.

You can either create a private subnet or use a preexisting subnet in your database VCN. Oracle recommends that you use a subnet size of /24 (256 IP addresses).

Note

Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet. See [Creating a Subnet](#) to learn more.

The database VCN requires security rules to allow backup traffic between your database and Recovery Service. Security rules must include stateful ingress rules to allow destination ports 8005 and 2484. You can use these Networking service features to implement security rules:

- [Security Lists](#)
A security list allows you to add security rules at the subnet level. In your database VCN, select the security list that is used for the Recovery Service subnet, and add the ingress rules to allow destination ports 8005 and 2484.
- [Network Security Groups \(NSG\)](#) Network security groups (NSG) enable granular control over security rules that apply to individual VNICs in a VCN. Recovery Service supports these options to configure security rules using NSGs:
 - To implement network isolation, create one NSG for the database VNIC (add egress rules to allow ports 2484 and 8005) and a separate NSG for Recovery Service (add ingress rules to allow ports 2484 and 8005).
 - Create and use a single NSG (with egress and ingress rules) for the database VNIC and Recovery Service.

Note

If you have configured a security list and an NSG within your database VCN, then the rules defined in the NSGs takes precedence over the rules defined in a security list.

See [Comparison of Security Lists and Network Security Groups](#) to learn more.

After you create a private subnet in the database VCN, assign the security rules and then register the subnet as a Recovery Service subnet in Recovery Service. If you have created NSGs to implement security rules, then you must also ensure to associate the Recovery Service NSG with the Recovery Service subnet.

Note

Oracle recommends using a private subnet for your backups. However, it is possible to use a public subnet.

Review Networking Service Permissions to Configure a Subnet

Ensure that you have these Networking Service permissions required to create a subnet in the database VCN and to assign security rules for Recovery Service.

Table 4-6 Networking Service Permissions Required to Create a Private subnet and Configure Security Rules for Recovery Service

Operation	Required IAM Policies
Configure a private subnet in a database VCN	<ul style="list-style-type: none"> • <code>use_vcns</code> for the compartment which the VCN is in • <code>use_subnets</code> for the compartment which the VCN is in • <code>manage_private-ips</code> for the compartment which the VCN is in • <code>manage_vnics</code> for the compartment which the VCN is in • <code>manage_vnics</code> for the compartment which the database is provisioned or is to be provisioned in

Alternatively, you can create a policy that allows a specified group with broader access to networking components.

For example, use this policy to allow a `NetworkAdmin` group to manage all networks in any compartment in a tenancy.

Example 4-1 Policy for Network Administrators

```
Allow group NetworkAdmin to manage virtual-network-family in tenancy
```

Review Subnet Size Requirements and Security Rules for Recovery Service Subnet

The security rules are necessary to allow backup traffic between a database and Recovery Service.

Note

Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet. See [Creating a Subnet](#) to learn more.

Table 4-7 Subnet size requirements and ingress rules for a private subnet used by Recovery Service

Item	Requirements
Minimum subnet size	/24 (256 IP addresses)
General ingress rule 1: Allow HTTPS traffic from Anywhere	<p>This rule allows backup traffic from your Oracle Cloud Infrastructure Database to Recovery Service.</p> <ul style="list-style-type: none"> • Stateless: No (all rules must be stateful) • Source Type: CIDR • Source CIDR: CIDR of the VCN where the database resides • IP Protocol: TCP • Source Port Range: All • Destination Port Range: 8005
General ingress rule 2: Allows SQLNet Traffic from Anywhere	<p>This rule allows recovery catalog connections and real-time data protection from your Oracle Cloud Infrastructure Database to Recovery Service.</p> <ul style="list-style-type: none"> • Stateless: No (all rules must be stateful) • Source Type: CIDR • Source CIDR: CIDR of the VCN where the database resides • IP Protocol: TCP • Source Port Range: All • Destination Port Range: 2484

Note

If you use network security groups (NSG) to implement security rules or if your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database NSG or subnet to the Recovery Service NSG or subnet that you create.

Create a Recovery Service Subnet in the Database VCN

Use the OCI Console to configure a private subnet for Recovery Service in your database virtual cloud network (VCN).

1. In the navigation menu, select **Networking**, and then select **Virtual cloud networks** to display the Virtual Cloud Networks page.
2. Select the VCN in which your database resides.
3. Use these steps to create a Recovery Service subnet with a security list. If you choose to use network security groups, then proceed to [step 4](#).
 - a. Under **Resources**, select **Security Lists**.

- b. Select the security list that is used for the VCN. You must add two ingress rules to allow destination ports 8005 and 2484.
- c. Click **Add Ingress Rules** and add these details to set up a stateful ingress rule that allows HTTPS traffic from anywhere:
 - **Source Type:** CIDR
 - **Source CIDR:** Specify the CIDR of the VCN where the database resides.
 - **IP Protocol:** TCP
 - **Source Port Range:** All
 - **Destination Port Range:** 8005
 - **Description:** Specify an optional description of the ingress rule to help manage the security rules.
- d. Click **Add Ingress Rule** and add these details to set up a stateful ingress rule that allows SQLNet traffic from anywhere:
 - **Source Type:** CIDR
 - **Source CIDR:** Specify the CIDR of the VCN where the database resides.
 - **IP Protocol:** TCP.
 - **Source Port Range:** All
 - **Destination Port Range:** 2484.
 - **Description:** Specify an optional description of the ingress rule to help manage the security rules.

Note

Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet. See [Creating a Subnet](#) to learn more. See: [Review Subnet Size Requirements and Security Rules for Recovery Service Subnet](#) for more information.

- e. In the Virtual Cloud Networks Details page, click **Create Subnet**.
- f. Create a private subnet or select a private subnet that already exists in the database VCN. Oracle recommends a subnet size of /24 (256 IP addresses) for the private subnet.
- g. In the Subnet Details page, under **Resources** select **Security Lists**. Add the security list that includes the ingress rules to allow destination ports 8005 and 2484.

Note

If your database VCN restricts network traffic between subnets, then ensure to add an egress rule for ports 2484 and 8005 from the database subnet to the Recovery Service subnet that you create.

4. Use these steps to create a Recovery Service subnet with network security groups (NSG).
 - a. Under **Resources**, select **Network Security Groups**.

- b. Click **Create Network Security Group**. Use one of these supported methods to configure security rules using NSGs:
 - To implement network isolation, create one NSG for the database VNIC (add egress rules to allow ports 2484 and 8005) and a separate NSG for Recovery Service (add ingress rules to allow ports 2484 and 8005).
 - Create and use a single NSG (with egress and ingress rules) for the database VNIC and Recovery Service.

The Network Security Group page lists the NSGs that you create.

Note

For additional configuration details, refer the relevant OCI Database Service documentation.

5. After you create the Recovery Service subnet in the database VCN, proceed to register the subnet as a Recovery Service subnet. Oracle recommends that you register a single Recovery Service subnet per VCN. If you have implemented security rules using NSGs, then you must also ensure to add the Recovery Service NSG to the Recovery Service subnet.

Register Recovery Service Subnet

After you have created a private subnet for Recovery Service in your database VCN, use this procedure to register the subnet in Recovery Service.

Multiple protected databases can use the same Recovery Service subnet. In order to ensure that the required number of IP addresses are available to support the Recovery Service private endpoints, you can assign multiple subnets to a Recovery Service subnet that is used by more than one protected database.

Note

Select an IPv4-only subnet for Recovery Service in your database VCN. Do not select an IPv6-enabled subnet as Recovery Service does not support using an IPv6-enabled subnet. Ensure that you have completed these prerequisite configuration tasks:

- [Assign Policies to Allow Access to Recovery Service and Related Resources](#)
 - [Create a Recovery Service Subnet in the Database VCN](#)
1. Log in to your OCI tenancy.
 2. In the navigation menu, click **Oracle AI Database**, and select **Database Backups** to display the Database Backups page.
 3. Click **Recovery Service Subnets**.
 4. In the **Compartment** field, select a compartment where you want to create the Recovery Service subnet.
 5. Click **Register Recovery Service subnet**, and specify the details.
 6. In the **Name** field, enter a name for the Recovery Service subnet.
 7. In the **Compartment** field, select the compartment where you want to create the Recovery Service subnet.

8. In the **Virtual cloud network** field, select the database VCN. Click **Change Compartment** to select a VCN belonging to a different compartment.
9. In the **Subnet** field, select a private subnet that you have configured for Recovery Service operations in your database VCN. Click **Change Compartment** to select a private subnet from a different compartment.
10. (Optional) Click **+Another Subnet** to assign an additional subnet to the Recovery Service subnet. If a single subnet does not contain enough IP addresses to support the Recovery Service private endpoints, then you can assign multiple subnets.
11. Click **Show advanced options** to configure these additional features.
 - Network security groups
 - Tags

If you have used a network security group (NSG) to implement security rules for Recovery Service in the database VCN, then you must add the Recovery Service NSG to the Recovery Service subnet. The Recovery Service NSG can reside in the same compartment or in a different compartment. However, the NSG must belong to the same VCN to which the specified subnet belongs.

- a. In the **Network security groups** section, select **Use network security groups to control traffic**.
- b. Select the Recovery Service NSG you have created in the database VCN.
- c. Click **+Another network security group** to associate multiple NSGs (maximum five).

(Optional) In the **Tag Namespace** field, consider adding a tag namespace, or tagging the control with an existing tag namespace.

12. Click **Register**.

You can replace a subnet or add more subnets to support the required number of private endpoints.
13. Use these steps to update an existing Recovery Service subnet:
 - a. In the Recovery Service subnet details page, under **Resources**, click **Subnets**.
 - b. Click **Add subnets** and select the subnets you want to add.
 - c. To replace an existing subnet, click the Action menu, and select **Remove subnet**. You can then add another subnet.

Note

A Recovery Service subnet must be associated with at least one subnet belonging to your database VCN.

14. Use these steps to manage the network security groups (NSGs) for an existing Recovery Service subnet:
 - a. In the **Network security groups** section, click **Add network security groups**.
 - b. Select and add the Recovery Service network security groups (maximum five).
 - c. To remove an NSG, select the resource and click **Remove**.

Autonomous Recovery Service Checklist

- [Ensure that the Recovery Service Subnet Can Communicate with Oracle Services](#)
The Recovery Service Subnet that you registered needs to communicate with the Recovery Service.
- [Ensure that Your Database Has TDE Fully Configured](#)
When using the Autonomous Recovery Service, you must have your database fully TDE encrypted.
- [Turn Off Any Manual Operational Backups](#)
In some cases, OCI users perform manual operational backups. These backups are run outside the standard tooling and support point-in-time recovery (non-KEEP backups).

Ensure that the Recovery Service Subnet Can Communicate with Oracle Services

The Recovery Service Subnet that you registered needs to communicate with the Recovery Service.

To access the service, the routing table for the private subnet needs to include "All IAD Services In Oracle Services Network".

For more information, see [Private Access to Oracle Services](#).

Ensure that Your Database Has TDE Fully Configured

When using the Autonomous Recovery Service, you must have your database fully TDE encrypted.

For new databases that are born in the cloud, this should already be done. But if you are creating a stub database in OCI and migrating a database to an Oracle Database Service from on-premises, or somewhere else, you might not meet all the criteria. For these databases you should verify that you meet the prerequisites for a backup up to the recovery service.

You must meet these 3 criteria:

- You need to have `WALLET_ROOT` configured in the database. If you are still using `sqlnet.ora`, you need to use `dbaascli` to properly set `WALLET_ROOT` for all databases that will be utilizing the Recovery Service.
To enable `wallet_root` SPFILE parameter for an existing database, run:

```
dbaascli tde enableWalletRoot
```

For more information, see [dbaascli tde enableWalletRoot](#).

Note

Setting `ENCRYPTION_WALLET_LOCATION` in `sqlnet.ora` is not supported and will be deprecated.

- You need to have an encryption key set for the CDB and all PDBs in your database.
- ALL tablespaces must be TDE encrypted before executing your first backup.

Turn Off Any Manual Operational Backups

In some cases, OCI users perform manual operational backups. These backups are run outside the standard tooling and support point-in-time recovery (non-KEEP backups).

If you are running any of these types of operational backups, it is critical that you turn them off at this point. Running operational backups to two different locations can cause issues with both backups.

Note

If you are using the tooling for object storage backups, and switch to the Recovery Service, the switchover will be automated by the tooling, and all of the previous backups will remain available.

Connecting to an Exadata Cloud Infrastructure Instance

This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.

How you connect depends on how your cloud network is set up. You can find information on various networking scenarios in [Networking Overview](#), but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

Note

Exadata Cloud Infrastructure servers cannot be joined to Active Directory domains, and the service does not support the use of Active Directory for user authentication and authorization.

- [Prerequisites](#)
List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.
- [SCAN Listener Port Setting](#)
When creating a cloud VM cluster, you can optionally designate a different SCAN listener port number.
- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [Using Oracle Net Services to Connect to a Database](#)
Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Prerequisites

List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.

You'll need the following:

- The full path to the file that contains the private key associated with the public key used when the system was launched.
- The public or private IP address of the Exadata Cloud Infrastructure instance.
Use the private IP address to connect to the system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located

on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Oracle Cloud InfrastructureConsole as follows:

- *Cloud VM clusters (new resource model)*: On the **Exadata VM Cluster Details** page, click Virtual Machines in the **Resources** list.

The values are displayed in the **Public IP Address** and **Private IP Address & DNS Name** columns of the table displaying the **Virtual Machines** or **Nodes** of the Exadata Cloud Infrastructure instance.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

SCAN Listener Port Setting

When creating a cloud VM cluster, you can optionally designate a different SCAN listener port number.

The default SCAN listener port for cloud VM clusters is 1521. When using the console [To create an ASM cloud VM cluster](#), you can optionally designate a different SCAN listener port number. In the OCI Console, this option appears under **Advanced Options** when creating the cluster.

Note

Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

Connecting to a Virtual Machine with SSH

You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Most Unix-style systems (including Linux, Oracle Solaris, and macOS) include an SSH client. For Microsoft Windows systems, you can download a free SSH client called PuTTY from the following site: "http://www.putty.org".

- [Connecting from a Unix-Style System](#)
To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.
- [Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY](#)
Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- [Accessing a Database After You Connect to the Virtual Machine](#)
After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

Related Topics

- <http://www.putty.org/>

Connecting from a Unix-Style System

To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.

- Enter the following SSH command to access the virtual machine:

```
ssh -i private-key user@node
```

In the preceding syntax:

- *private-key* is the full path and name of the file that contains the SSH private key that corresponds to a public key that is registered in the system.
- *user* is the operating system user that you want to use to connect:
 - * To perform operations as the Oracle Database software owner, connect as `opc` and `su oracle`. The `oracle` user does not have `root` user access to the virtual machine.
 - * To perform operations that require `root` access to the virtual machine, such as patching, connect as `opc`. The `opc` user can use the `sudo -s` command to gain `root` access to the virtual machine.
- *node* is the host name or IP address for the virtual machine that you want to access.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.
To download PuTTY, go to <http://www.putty.org/> and click the **You can download PuTTY here** link.
2. Run the PuTTY program (`putty.exe`).
The PuTTY Configuration window is displayed, showing the **Session** panel.
3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.
4. Confirm that the **Connection type** option is set to **SSH**.
5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.
The **Data** panel is displayed.
6. In the **Auto-login username** field, enter the operating system user you want to connect as:

- Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.
7. Confirm that the **When username is not specified** option is set to **Prompt**.
 8. In the **Category** tree, expand **SSH** and then click **Auth**.
The **Auth** panel is displayed.
 9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.
 10. In the **Category** tree, click **Session**.
The **Session** panel is displayed.
 11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
 12. Click **Open** to open the connection.
The PuTTY Configuration window closes and the PuTTY terminal window displays.
If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

Accessing a Database After You Connect to the Virtual Machine

After you connect to a virtual machine, you can use the following series of commands to identify a database and connect to it.

1. SSH in as the `opc` user.
2. `sudo su oracle`
3. Use the `srvctl` utility located under the Oracle Grid Infrastructure home directory to list the databases on the system. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl config database -v
nc122    /u02/app/oracle/product/12.2.0/dbhome_6 12.2.0.1.0
s12c     /u02/app/oracle/product/12.2.0/dbhome_2 12.2.0.1.0
```

4. Identify the database instances for the database that you want to access. For example:

```
/u01/app/12.2.0.1/grid/bin/srvctl status database -d s12c
Instance s12c1 is running on node node01
Instance s12c2 is running on node node02
```

5. Configure the environment settings for the database that you want to access. For example:

```
. oraenv
ORACLE_SID = [oracle] ? s12c
The Oracle base has been set to /u02/app/oracle
```

```
export ORACLE_SID=s12c1
```

6. You can use the `svrctl` command to display more detailed information about the database. For example:

```
svrctl config database -d s12c
Database unique name: s12c
Database name:
Oracle home: /u02/app/oracle/product/12.2.0/dbhome_2
Oracle user: oracle
Spfile: +DATAC4/s12c/spfiles12c.ora
Password file: +DATAC4/s12c/PASSWORD/passwd
Domain: example.com
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: DATAC4
Mount point paths:
Services:
Type: RAC
Start concurrency:
Stop concurrency:
OSDBA group: dba
OSOPER group: racoper
Database instances: s12c1,s12c2
Configured nodes: node01,node02
CSS critical: no
CPU count: 0
Memory target: 0
Maximum memory: 0
Default network number for database services:
Database is administrator managed
```

7. You can access the database by using SQL*Plus. For example:

```
sqlplus / as sysdba

SQL*Plus: Release 12.2.0.1.0 Production ...

Copyright (c) 1982, 2016, Oracle. All rights reserved.

Connected to:
Oracle Database 12c EE Extreme Perf Release 12.2.0.1.0 - 64bit Production
```

Using Oracle Net Services to Connect to a Database

Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Because Exadata Cloud Infrastructure uses Oracle Grid Infrastructure, you can make Oracle Net Services connections by using **Single Client Access Name (SCAN)** connections. SCAN is a feature that provides a consistent mechanism for clients to access the Oracle Database instances running in a cluster.

By default, the SCAN is associated with three virtual IP addresses (VIPs). Each SCAN VIP is also associated with a SCAN listener that provides a connection endpoint for Oracle Database connections using Oracle Net Services. To maximize availability, Oracle Grid Infrastructure distributes the SCAN VIPs and SCAN listeners across the available cluster nodes. In addition, if there is a node shutdown or failure, then the SCAN VIPs and SCAN listeners are automatically migrated to a surviving node. By using SCAN connections, you enhance the ability of Oracle Database clients to have a reliable set of connection endpoints that can service all of the databases running in the cluster.

The SCAN listeners are in addition to the Oracle Net Listeners that run on every node in the cluster, which are also known as the node listeners. When an Oracle Net Services connection comes through a SCAN connection, the SCAN listener routes the connection to one of the node listeners, and plays no further part in the connection. A combination of factors, including listener availability, database instance placement, and workload distribution, determines which node listener receives each connection.

Note

This documentation provides basic requirements for connecting to your Exadata Cloud Infrastructure databases by using Oracle Net Services.

- [Prerequisites for Connecting to a Database with Oracle Net Services](#)
Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.
- [Connecting to a Database with SQL Developer](#)
You can connect to a database with SQL Developer by using one of the following methods:
- [Connecting to a Database Using SCAN](#)
To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.
- [Connecting to a Database Using a Node Listener](#)
To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.

To connect to an Oracle Database on Exadata Cloud Infrastructure with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)

- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see [Updating the Security List](#).

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an Exadata Cloud Infrastructure instance using SCAN IP addresses or public IP addresses, depending on how your network is set up and where you are connecting from. You can find the IP addresses in the Console, in the **Database** details page.

- [To connect using SCAN IP addresses](#)
You can connect to the database using the SCAN IP addresses if your client is on-premises and you are connecting using a FastConnect or Site-to-Site VPN connection.
- [To connect using public IP addresses](#)
You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway.

To connect using SCAN IP addresses

You can connect to the database using the SCAN IP addresses if your client is on-premises and you are connecting using a FastConnect or Site-to-Site VPN connection.

You have the following options:

- Use the private SCAN IP addresses, as shown in the following `tnsnames.ora` example:

```
testdb=
  (DESCRIPTION =
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP)(HOST = <scanIP1>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = <scanIP2>)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

- Define an external SCAN name in your on-premises DNS server. Your application can resolve this external SCAN name to the VM cluster's private SCAN IP addresses, and then the application can use a connection string that includes the external SCAN name. In the following `tnsnames.ora` example, `extscanname.example.com` is defined in the on-premises DNS server.

```
testdb =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <extscanname.example.com>)(PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

To connect using public IP addresses

You can use the node's public IP address to connect to the database if the client and database are in different VCNs, or if the database is on a VCN that has an internet gateway.

However, there are important implications to consider:

- When the client uses the public IP address, the client bypasses the SCAN listener and reaches the node listener, so server side load balancing is not available.
- When the client uses the public IP address, it cannot take advantage of the VIP failover feature. If a node becomes unavailable, new connection attempts to the node will hang until a TCP/IP timeout occurs. You can set client side sqlnet parameters to limit the TCP/IP timeout.

The following `tnsnames.ora` example shows a connection string that includes the `CONNECT_TIMEOUT` parameter to avoid TCP/IP timeouts.

```
test=
  (DESCRIPTION =
    (CONNECT_TIMEOUT=60)
    (ADDRESS_LIST=
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP1>)(PORT = 1521))
      (ADDRESS = (PROTOCOL = TCP)(HOST = <publicIP2>)(PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <dbservice.subnetname.dbvcn.oraclevcn.com>)
    )
  )
```

Connecting to a Database Using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.

- [Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs](#)
You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.
- [Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name](#)
You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.

This approach requires you to supply all of the single client access name (SCAN) virtual IP (VIP) addresses, and enables Oracle Net Services to connect to an available SCAN listener.

- Use the following template to define a Net Services alias, which is typically used to provide a convenient name for the connect descriptor:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-1)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-2)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-3)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

SCAN-VIP-[1-3] are the IP addresses for the SCAN VIPs.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=`*sid-name*. For example: `SID=S12C1`.
- `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

Note

By default, Oracle Net Services randomly selects one of the addresses in the address list to balance the load between the SCAN listeners.

Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Using this approach, you define a custom single client access name (SCAN) name in your domain name server (DNS), which resolves to the three SCAN virtual IP addresses (VIPs).

- Use the following template to define a Net Services alias that references the custom SCAN name:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=scan-name)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

scan-name is the custom SCAN name.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=`*sid-name*. For example: `SID=S12C1`.
- `SERVICE_NAME=`*service-name*. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
scan-name:1521/sid-or-service-entry
```

For example:

```
exalscan.example.com:1521/S12C1
```

Or

```
exalscan.example.com:1521/PDB1.example.yourcloud.com
```

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

By using this method, you give up the high-availability and load-balancing provided by SCAN. However, this method may be desirable if you want to direct connections to a specific node or network interface. For example, you might want to ensure that connections from a program that performs bulk data loading use the backup network.

Using this approach, you direct your connection using the hostname or IP address of the node.

Example 4-2 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The `(CONNECT_TIMEOUT=timeout)` parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=sid-name`. For example, `SID=S12C1`.
- `SERVICE_NAME=service-name`. For example, `SERVICE_NAME=PDB1.example.oraclecloudatcust.com`.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
node:1521/sid-or-service-entry
```

For example:

```
exalnode01.example.com:1521/S12C1
```

Or

```
exalnode01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

Using the Console to Create API Access Control

Use this procedure to create API Access Control for your ExaDB-D resources—Exadata Infrastructure, VM Cluster, Database, and Pluggable Database.

1. Go to the details page of the respective resource.
2. In the **General Information** section, locate **API Access Control**, and click **Create**.
 - If API Access Control has not been created yet, this option will be available, and the status will be shown as **Disabled**.
 - After creation, the status changes to **Enabled**, and the **Create** option is replaced with the **OCID** of the newly created API Access Control.
3. Click the OCID to view the API Access Control details.
4. You can disable the API Access Control at any time, if needed.

Best Practices for Exadata Cloud Infrastructure Instances

Oracle recommends that you follow these best practice guidelines to ensure the manageability of your Exadata Cloud Infrastructure instance:

When followed best practice guidelines can prevent problems that could affect the manageability and performance of your Exadata Cloud Infrastructure instance:

- Wherever possible, use the Oracle-supplied cloud interfaces such as the Oracle Cloud Infrastructure Console, API, or CLI, or cloud-specific tools such as `dbaascli` to perform lifecycle management and administrative operations on your Exadata Cloud Infrastructure instance. For example, use the OCI console, API, CLI, or `dbaascli` to apply Oracle Database patches instead of manually running `opatch`. In addition, if an operation can be performed by using the Console as well as a command line utility, Oracle recommends that you use the Console. For example, use the Console instead of using `dbaascli` to create databases.
- Do not change the compute node OS users or manually manipulate SSH key settings associated with your Exadata VM cluster.
- Apply only patches that are available through the Database service. Do not apply patches from any other source unless you are directed to do so by Oracle Support.
- Apply the quarterly patches regularly, every quarter if possible.
- Do not change the ports for Oracle Net Listener.

Moving to Oracle Cloud Using Zero Downtime Migration

Oracle now offers the Zero Downtime Migration service, a quick and easy way to move on-premises databases to Oracle Cloud Infrastructure.

Zero Downtime Migration leverages Oracle Active Data Guard to create a standby instance of your database in an Oracle Cloud Infrastructure system. You switch over only when you are ready, and your source database remains available as a standby. Use the Zero Downtime Migration service to migrate databases individually or at the fleet level. See *Move to Oracle Cloud Using Zero Downtime Migration* for more information.

Related Topics

- [Move to Oracle Cloud Using Zero Downtime Migration](#)

5

How-to Guides

A collection of tasks and procedures for managing Exadata Database Service on Dedicated Infrastructure.

- [Manage Database Security with Oracle Data Safe](#)
- [Connecting to an Exadata Cloud Infrastructure Instance](#)
This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- [Manage Exadata Cloud Infrastructure](#)
Use the provided tools to manage the Infrastructure.
- [Configure Oracle-Managed Infrastructure Maintenance](#)
Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.
- [Manage VM Clusters](#)
Learn how to manage your VM clusters on Exadata Cloud Infrastructure.
- [Manage Software Images](#)
- [Create Oracle Database Homes on an Exadata Cloud Infrastructure System](#)
Learn to create Oracle Database Homes on Exadata Cloud Infrastructure.
- [Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance](#)
You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.
- [Manage Databases on Exadata Cloud Infrastructure](#)
- [Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure](#)
Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure.
- [Patch and Update an Exadata Cloud Infrastructure System](#)
- [Interim Software Updates](#)
For authorized environments, learn how to download interim software updates.
- [Use Oracle Data Guard with Exadata Cloud Infrastructure](#)
Learn to configure and manage Data Guard Groups in your VM cluster.
- [Configure Oracle Database Features for Exadata Cloud Infrastructure](#)
This topic describes how to configure Oracle Multitenant, tablespace encryption, and Huge Pages for use with your Exadata Cloud Infrastructure instance.
- [Managing Exadata Cloud Infrastructure I/O Resource Management \(IORM\)](#)
- [Managing Encryption Keys on External Devices](#)
Learn how to store and manage database encryption keys.
- [Migrate to Exadata Cloud Infrastructure](#)
For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Exadata Cloud Infrastructure see "Migrating Databases to the Cloud".

- [Connect Identity and Access Management \(IAM\) Users to Oracle Exadata Database Service on Dedicated Infrastructure](#)
You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.
- [Authenticating and Authorizing Microsoft Entra ID \(MS-EI\) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.
- [Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure](#)
Exadata Database Service on Oracle Database@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.
- [Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud](#)
Exadata Database Service on Oracle Database@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).
- [AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS](#)
Exadata Database Service on Oracle Database@AWS supports integration with AWS Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.
- [Database Multicloud Integration for Oracle Database Cloud Services](#)

Manage Database Security with Oracle Data Safe

- [About Oracle Data Safe](#)
- [Get Started](#)
- [Using Oracle Data Safe](#)

About Oracle Data Safe

Your corporate policy requires that you monitor your databases and retain audit records. Your developers are asking for copies of production data for that new application, and you're wondering what kinds of sensitive information it will contain. Meanwhile, you need to make sure that recent maintenance activities haven't left critical security configuration gaps on your production databases and that staff changes haven't left dormant user accounts on the databases. Oracle Data Safe assists you with these tasks and is included with your Exadata Database Service*.

Oracle Data Safe is a unified control center, that helps you to manage the day-to-day security and compliance requirements of Oracle Databases no matter if they are running in the Oracle Cloud Infrastructure, at Cloud@Customer, on-premises or in any other cloud.

Data Safe supports you to evaluate security controls, assess user security, monitor user activity, and address data security compliance requirements for your database by evaluating the sensitivity of your data as well as masking sensitive data for non-production databases.

Data Safe provides the following features:

- **Security Assessment:** Configuration errors and configuration drift are significant contributors to data breaches. Use security assessment to evaluate your database's configuration and compare it to Oracle and industry best practices. Security assessment reports on areas of risk and notifies you when configurations change.
- **User Assessment:** Many breaches start with a compromised user account. User Assessment helps you spot the riskiest database accounts - those accounts which, if compromised, could cause the most damage - and take proactive steps to secure them. User Assessment Baselines make it easy to know when new accounts are added, or an account's privileges are modified. Use OCI events to receive proactive notifications when a database deviates from its baseline.
- **Activity Auditing:** Understanding and reporting on user activity, data access, and changes to database structures supports regulatory compliance requirements and can aid in post-incident investigations. Activity auditing collects audit records from databases and helps you manage audit policies. Audit insights make it easy to identify inefficient audit policies, while alerts based on audit data proactively notify you of risky activity.
- **Sensitive Data Discovery:** Knowing what sensitive data is managed in your applications is critical for security and privacy. Data discovery scans your database for over 150 different types of sensitive data, helping you understand what types and how much sensitive data you are storing. Use these reports to formulate audit policies, develop data masking templates, and create effective access control policies.
- **Data Masking:** Minimizing the amount of sensitive data your organization maintains helps you meet compliance requirements and satisfy data privacy regulations. Data masking helps you remove risk from your non-production databases by replacing sensitive information with masked data. With reusable masking templates, over 50 included masking formats, and the ability to easily create custom formats for your organization's unique requirements, data masking can streamline your application development and testing operations.
- **SQL Firewall Management:** Protect against risks such as SQL injection attacks or compromised accounts. Oracle SQL Firewall is a new security capability built into the Oracle AI Database 26ai kernel and offers best-in-class protection against these risks. The SQL Firewall feature in Oracle Data Safe lets you centrally manage and monitor the SQL Firewall policies for your target databases. Data Safe lets you collect authorized SQL activities of a database user, generate and enable the policy with allowlists of approved SQL statements and database connection paths and provides a comprehensive view of any SQL Firewall violations across the fleet of your target databases.

**Includes 1 million audit records per database per month if using the audit collection for Activity Auditing*

Get Started

To get started you just need to register your database with Oracle Data Safe:

- Pre-requisite: Obtain the necessary Identity and Access Management (IAM) permissions to register your target database in Data Safe: [Permissions to Register an Oracle Cloud Database with Oracle Data Safe](#)
- Connecting your database to Data Safe
 - If your database is running in a private virtual cloud network (VCN), you can connect it to Data Safe via a **Data Safe private endpoint**.

The private endpoint essentially represents the Oracle Data Safe service in your VCN with a private IP address in a subnet of your choice.

You can create the private endpoint in the VCN of your database either before the registration or during the registration process. You can find more details on how to create the private endpoint under [Create an Oracle Data Safe Private Endpoint](#).

- [Register your database in Data Safe](#)

Using Oracle Data Safe

Once your database is registered in Data Safe, you can leverage all features.

Security Assessment

Security Assessments are automatically scheduled once a week in Data Safe and provide an overall picture of your database security posture. It analyzes your database configurations, users and user entitlements, as well as security policies to uncover security risks and improve the security posture of Oracle Databases within your organization. A security assessment provides findings with recommendations for remediation activities that follow best practices to reduce or mitigate risk.

Start by reviewing the security assessment report for your database: [View the latest assessment for a target database](#)

You can find more details on Security Assessment under [Security Assessment Overview](#).

User Assessment

User Assessments are automatically scheduled once a week in Data Safe and help you to identify highly privileged user accounts that could pose a threat if misused or compromised. User Assessment reviews information about your users in the data dictionaries on your target databases and then calculates a potential risk for each user, based on system privileges and role grants.

Start by reviewing the user assessment report for your database: [View the latest user assessment for a target database](#)

You can find more details on User Assessment under [User Assessment Overview](#).

Data Discovery

Data Discovery searches for sensitive columns in your database. It comes with over 150 pre-defined sensitive types and you can also create your own sensitive types. You tell Data Discovery if you want to scan your entire database or just certain schemas and what type of sensitive information to look for, and it finds the sensitive columns that meet your criteria and stores them in a sensitive data model (SDM).

Start by discovering sensitive data in your database: [Create Sensitive Data Models](#)

You can find more details on Data Discovery under [Data Discovery Overview](#).

Data Masking

Data masking, also known as static data masking helps you to replace sensitive or confidential information in your non-production databases with realistic and fully functional data with similar characteristics as the original data. Data Safe comes with pre-defined masking formats for each of the pre-defined sensitive types that can also be leveraged for your own sensitive types.

Once you know where sensitive data is stored in your database (for instance after running Data Discovery in Data Safe), you can start by creating a masking policy: [Create Masking Policies](#)

After you created a masking policy and copied your production database, you can mask your non-production copy: [Mask Sensitive Data on a Target Database](#)

You can find more details on Data Masking under [Data Masking Overview](#).

Activity Auditing

Activity Auditing in Oracle Data Safe helps to ensure accountability and improve regulatory compliance. With Activity Auditing, you can collect and retain audit records per industry and regulatory compliance requirements and monitor user activities on Oracle databases with pre-defined reports and alerts. For example, you can audit access to sensitive data, security-relevant events, administrator and user activities, activities recommended by compliance regulations like the Center for Internet Security (CIS), and activities defined by your own organization.

If you are using the audit collection in Data Safe, up to 1 million audit records per target database per month are included for your Cloud@Customer database.

To use activity auditing, start the audit trail for your target database in Data Safe: [Start an Audit Trail](#)

Once the audit trail is started, you can monitor and analyze your audit data with pre-defined audit reports: [View a Predefined or Custom Audit Report](#)

You can find more details on Activity Auditing under [Activity Auditing Overview](#).

SQL Firewall*

SQL Firewall in Oracle Data Safe lets you centrally manage the SQL Firewalls and provides a comprehensive view of SQL Firewall violations across the fleet of your target databases. Data Safe lets you collect authorized SQL activities of a database user you wish to protect, monitor the progress of the collection, generate and enable the policy with allowlists of approved SQL statements and database connection paths.

Start by enabling the SQL Firewall in your 26ai target database: [Enable SQL Firewall On Your Target Database](#).

Next, you need to generate and enable a SQL Firewall policy with allowlists for the database user you wish to protect: [Generate and Enforce SQL Firewall Policies](#).

Once you start enforcing the SQL Firewall policy, you can monitor and analyze the violations in the pre-defined violation reports: [View and Manage Violations Reports](#).

You can find more details on SQL Firewall under [SQL Firewall Overview](#).

*SQL Firewall is only available for Oracle AI Database 26ai.

Connecting to an Exadata Cloud Infrastructure Instance

This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.

How you connect depends on how your cloud network is set up. You can find information on various networking scenarios in [Networking Overview](#), but for specific recommendations on how you should connect to a database in the cloud, contact your network security administrator.

Note

Exadata Cloud Infrastructure servers cannot be joined to Active Directory domains, and the service does not support the use of Active Directory for user authentication and authorization.

- [Prerequisites](#)
List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.
- [About Connecting to a Compute Node with SSH](#)
You can connect to the compute nodes in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [Connect to the Exadata Cloud Infrastructure Service](#)
Learn how to connect to an Exadata Cloud Infrastructure system using SSH, and how to connect to an Exadata Cloud Infrastructure database using Oracle Net Services (SQL*Net).

Prerequisites

List of the requirements for SSH access to a compute node in an Exadata Cloud Infrastructure instance.

You'll need the following:

- The full path to the file that contains the private key associated with the public key used when the system was launched.
- The public or private IP address of the Exadata Cloud Infrastructure instance.
Use the private IP address to connect to the system from your on-premises network, or from within the virtual cloud network (VCN). This includes connecting from a host located on-premises connecting through a VPN or FastConnect to your VCN, or from another host in the same VCN. Use the public IP address to connect to the system from outside the cloud (with no VPN). You can find the IP addresses in the Oracle Cloud InfrastructureConsole as follows:
 - *Cloud VM clusters (new resource model)*: On the **Exadata VM Cluster Details** page, click Virtual Machines in the **Resources** list.

The values are displayed in the **Public IP Address** and **Private IP Address & DNS Name** columns of the table displaying the **Virtual Machines** or **Nodes** of the Exadata Cloud Infrastructure instance.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

About Connecting to a Compute Node with SSH

You can connect to the compute nodes in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Most Unix-style systems (including Linux, Oracle Solaris, and Apple MacOS) include an SSH client. For Microsoft Windows, you can download a free SSH client called PuTTY from the following address: <http://www.putty.org>

- [Connecting from a Unix-Style System](#)
To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.
- [Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY](#)
Learn to access a virtual machine from a Microsoft Windows system using PuTTY.
- [To access a database after you connect to the compute node](#)
To connect to the database, you set environment information for the database.

Connecting from a Unix-Style System

To access a virtual machine on an Oracle ExaDB-D system from a Unix-style system using SSH, use this procedure.

- Enter the following SSH command to access the virtual machine:

```
ssh -i private-key user@node
```

In the preceding syntax:

- *private-key* is the full path and name of the file that contains the SSH private key that corresponds to a public key that is registered in the system.
- *user* is the operating system user that you want to use to connect:
 - * To perform operations as the Oracle Database software owner, connect as `opc` and `su oracle`. The `oracle` user does not have `root` user access to the virtual machine.
 - * To perform operations that require `root` access to the virtual machine, such as patching, connect as `opc`. The `opc` user can use the `sudo -s` command to gain `root` access to the virtual machine.
- *node* is the host name or IP address for the virtual machine that you want to access.

Connecting to a Virtual Machine from a Microsoft Windows System Using PuTTY

Learn to access a virtual machine from a Microsoft Windows system using PuTTY.

Before you begin

Before you use the PuTTY program to connect to a virtual machine, you need the following:

- The IP address of the virtual machine
- The SSH private key file that matches the public key associated with the deployment. This private key file must be in the PuTTY `.ppk` format. If the private key file was originally created on the Linux platform, you can use the PuTTYgen program to convert it to the `.ppk` format.

To connect to a virtual machine using the PuTTY program on Windows:

1. Download and install PuTTY.

To download PuTTY, go to <http://www.putty.org/> and click the **You can download PuTTY here** link.

2. Run the PuTTY program (`putty.exe`).

The PuTTY Configuration window is displayed, showing the **Session** panel.

3. In the **Host Name (or IP address)** field, enter the host name or IP address of the virtual machine that you want to access.
4. Confirm that the **Connection type** option is set to **SSH**.
5. In the **Category** tree, expand **Connection** if necessary and then click **Data**.
The **Data** panel is displayed.
6. In the **Auto-login username** field, enter the operating system user you want to connect as:
 - Connect as the user `opc` to perform operations that require `root` or `oracle` access to the virtual machine, such as backing up or patching; this user can use the `sudo` command to gain `root` or `oracle` access to the VM.
7. Confirm that the **When username is not specified** option is set to **Prompt**.
8. In the **Category** tree, expand **SSH** and then click **Auth**.
The **Auth** panel is displayed.
9. Click the **Browse** button next to the **Private key file for authentication** field. Then, in the **Select private key file** window, navigate to and open the private key file that matches the public key that is associated with the deployment.
10. In the **Category** tree, click **Session**.
The **Session** panel is displayed.
11. In the **Saved Sessions** field, enter a name for the connection configuration. Then, click **Save**.
12. Click **Open** to open the connection.
The PuTTY Configuration window closes and the PuTTY terminal window displays.
If this is the first time you are connecting to the VM, the PuTTY Security Alert window is displayed, prompting you to confirm the public key. Click **Yes** to continue connecting.

To access a database after you connect to the compute node

To connect to the database, you set environment information for the database.

1. Log in as `opc` and then use `sudo` to connect as the `oracle` user.

```
login as: opc  
  
[opc@<host_name> ~]$ sudo su - oracle
```

2. Source the database's `.env` file to set the environment.

```
[oracle@<host_name>]# . <database_name>.env
```

In the following example, the host name is "ed1db01" and the database name is "cdb01".

```
[oracle@ed1db01]# . cdb01.env  
ORACLE_SID = [root] ? +ASM1  
The Oracle base has been set to /u01/app/grid
```

Connect to the Exadata Cloud Infrastructure Service

Learn how to connect to an Exadata Cloud Infrastructure system using SSH, and how to connect to an Exadata Cloud Infrastructure database using Oracle Net Services (SQL*Net).

- [Connecting to a Database with SQL Developer](#)
You can connect to a database with SQL Developer by using one of the following methods:
- [Connecting to a Database with Oracle Net Services](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system using Oracle Net Services.

Connecting to a Database with SQL Developer

You can connect to a database with SQL Developer by using one of the following methods:

- Create a temporary SSH tunnel from your computer to the database. This method provides access only for the duration of the tunnel. (When you are done using the database, be sure to close the SSH tunnel by exiting the SSH session.)
- Open the port used as the Oracle SCAN listener by updating the security list used for the cloud VM cluster resource in the Exadata Cloud Service instance. The default SCAN listener port is 1521. This method provides more durable access to the database. For more information, see [Updating the Security List](#).

After you've created an SSH tunnel or opened the SCAN listener port as described above, you can connect to an Exadata Cloud Infrastructure instance using SCAN IP addresses or public IP addresses, depending on how your network is set up and where you are connecting from. You can find the IP addresses in the Console, in the **Database** details page.

Connecting to a Database with Oracle Net Services

You can connect to the virtual machines in an Exadata Cloud Infrastructure system using Oracle Net Services.

- [Using Oracle Net Services to Connect to a Database](#)
Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.
- [Prerequisites for Connecting to a Database with Oracle Net Services](#)
Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.
- [Connecting to a Database Using SCAN](#)
To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.
- [Connecting to a Database Using a Node Listener](#)
To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

Using Oracle Net Services to Connect to a Database

Oracle Database Exadata Cloud Infrastructure supports remote database access by using Oracle Net Services.

Because Exadata Cloud Infrastructure uses Oracle Grid Infrastructure, you can make Oracle Net Services connections by using **Single Client Access Name (SCAN)** connections. SCAN is a feature that provides a consistent mechanism for clients to access the Oracle Database instances running in a cluster.

By default, the SCAN is associated with three virtual IP addresses (VIPs). Each SCAN VIP is also associated with a SCAN listener that provides a connection endpoint for Oracle Database connections using Oracle Net Services. To maximize availability, Oracle Grid Infrastructure distributes the SCAN VIPs and SCAN listeners across the available cluster nodes. In addition, if there is a node shutdown or failure, then the SCAN VIPs and SCAN listeners are automatically migrated to a surviving node. By using SCAN connections, you enhance the ability of Oracle Database clients to have a reliable set of connection endpoints that can service all of the databases running in the cluster.

The SCAN listeners are in addition to the Oracle Net Listeners that run on every node in the cluster, which are also known as the node listeners. When an Oracle Net Services connection comes through a SCAN connection, the SCAN listener routes the connection to one of the node listeners, and plays no further part in the connection. A combination of factors, including listener availability, database instance placement, and workload distribution, determines which node listener receives each connection.

Note

This documentation provides basic requirements for connecting to your Exadata Cloud Infrastructure databases by using Oracle Net Services.

Prerequisites for Connecting to a Database with Oracle Net Services

Review the prerequisites to connect to an Oracle Database instance on Oracle ExaDB-D using Oracle Net Services.

To connect to an Oracle Database on Exadata Cloud Infrastructure with Oracle Net Services, you need the following:

- The IP addresses for your SCAN VIPs, or the hostname or IP address for a virtual machine that hosts the database that you want to access.
- The database identifier: Either the database system identifier (SID), or a service name.

Connecting to a Database Using SCAN

To create an Oracle Net Services connection by using the SCAN listeners, you can choose between two approaches.

- [Identifying IP Addresses Using the SDK or CLI](#)
You can use the SDK or the OCI CLI to identify the IP addresses of Exadata Cloud Infrastructure compute nodes. You can then use the IP addresses to connect to your system.
- [Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs](#)
You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.

- [Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name](#)
You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Identifying IP Addresses Using the SDK or CLI

You can use the SDK or the OCI CLI to identify the IP addresses of Exadata Cloud Infrastructure compute nodes. You can then use the IP addresses to connect to your system.

NOT_SUPPORTED

1. Use the [GetDbNode](#) API to return the details of the Exadata Cloud Infrastructure [dbNode](#). Note the [OCIDs](#) returned for the `hostIpId` and `backupIpId` parameters of the `dbNode`.
2. With the OCIDs found in the `hostIpId` and `backupIpId` parameters, you can use the [GetPrivateIp](#) API to get the private IP addresses used by the client and backup subnets. For public subnet IP addresses, use the [GetPublicIpByPrivateIpId](#) API.

Connecting to a Database Using a Connect Descriptor that References All of the SCAN VIPs

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using multiple SCAN listeners.

This approach requires you to supply all of the single client access name (SCAN) virtual IP (VIP) addresses, and enables Oracle Net Services to connect to an available SCAN listener.

- Use the following template to define a Net Services alias, which is typically used to provide a convenient name for the connect descriptor:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-1)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-2)(PORT=1521))
    (ADDRESS=(PROTOCOL=tcp)(HOST=SCAN-VIP-3)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

SCAN-VIP-[1-3] are the IP addresses for the SCAN VIPs.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=sid-name`. For example: `SID=S12C1`.
- `SERVICE_NAME=service-name`. For example: `SERVICE_NAME=PDB1.example.yourcloud.com`.

Note

By default, Oracle Net Services randomly selects one of the addresses in the address list to balance the load between the SCAN listeners.

Connecting to a Database Use a Connect Descriptor that References a Custom SCAN Name

You can set up a connect descriptor for Oracle Exadata Database Service on Dedicated Infrastructure System using a custom SCAN name.

Using this approach, you define a custom single client access name (SCAN) name in your domain name server (DNS), which resolves to the three SCAN virtual IP addresses (VIPs).

- Use the following template to define a Net Services alias that references the custom SCAN name:

```
alias-name = (DESCRIPTION=
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=scan-name)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

scan-name is the custom SCAN name.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=sid-name`. For example: `SID=S12C1`.
- `SERVICE_NAME=service-name`. For example:
`SERVICE_NAME=PDB1.example.yourcloud.com`.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
scan-name:1521/sid-or-service-entry
```

For example:

```
exalscan.example.com:1521/S12C1
```

Or

```
exalscan.example.com:1521/PDB1.example.yourcloud.com
```

Connecting to a Database Using a Node Listener

To connect to an Oracle Database instance on Exadata Cloud Infrastructure with a connect descriptor that bypasses the SCAN listeners, use this procedure to route your connection directly to a node listener.

By using this method, you give up the high-availability and load-balancing provided by SCAN. However, this method may be desirable if you want to direct connections to a specific node or network interface. For example, you might want to ensure that connections from a program that performs bulk data loading use the backup network.

Using this approach, you direct your connection using the hostname or IP address of the node.

Example 5-1 Defining a Net Service Alias That Directly References the Node

```
alias-name = (DESCRIPTION=
  (CONNECT_TIMEOUT=timeout)
  (ADDRESS_LIST=(ADDRESS=(PROTOCOL=tcp)(HOST=node)(PORT=1521)))
  (CONNECT_DATA=(sid-or-service-entry)))
```

Where:

alias-name is the name you use to identify the alias.

timeout specifies a timeout period (in seconds), which enables you to terminate a connection attempt without having to wait for a TCP timeout. The `(CONNECT_TIMEOUT=timeout)` parameter is optional.

node is the hostname or IP address for the virtual machine that you want to use.

sid-or-service-entry identifies the database SID or service name using one of the following formats:

- `SID=sid-name`. For example, `SID=S12C1`.
- `SERVICE_NAME=service-name`. For example, `SERVICE_NAME=PDB1.example.oraclecloudatcust.com`.

Alternatively, you can use the easy connect method to specify a connect descriptor with the following format:

```
node:1521/sid-or-service-entry
```

For example:

```
exalnode01.example.com:1521/S12C1
```

Or

```
exalnode01.example.com:1521/PDB1.example.oraclecloudatcust.com
```

Manage Exadata Cloud Infrastructure

Use the provided tools to manage the Infrastructure.

- [Using the Console to Provision Exadata Cloud Infrastructure](#)
Learn how to provision an Exadata Cloud Infrastructure system.
- [Using the API to Create Infrastructure Components](#)
- [Using the API to Manage Exadata Cloud Infrastructure Instance](#)

Using the Console to Provision Exadata Cloud Infrastructure

Learn how to provision an Exadata Cloud Infrastructure system.

- [Lifecycle Management Operations](#)

- [Network Management Operations](#)
- [Management Tasks for the Oracle Cloud Infrastructure Platform](#)
- [Oracle Database License Management Tasks](#)
- [Scaling Resources within an Exadata Infrastructure Instance](#)
If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

Lifecycle Management Operations

- [To check the status of a Cloud Exadata infrastructure resource](#)
- [To change the infrastructure display name](#)
- [To check the status of a cloud VM cluster](#)
- [To start, stop, or reboot an Exadata Cloud Infrastructure cloud VM cluster](#)
- [To terminate Exadata Cloud Infrastructure infrastructure-level resources](#)
- [To View a List of DB Servers on an Exadata Infrastructure](#)
To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.

To check the status of a Cloud Exadata infrastructure resource

Note

This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Cloud Exadata infrastructure resources, click the name of the infrastructure you're interested in and check its icon. The icon text indicates the status of the system. The following lifecycle states apply to the Cloud Exadata infrastructure
 - **Provisioning:** Resources are being reserved for the Cloud Exadata infrastructure resource. Provisioning can take several minutes. The resource is not ready to be used.
 - **Available:** The Cloud Exadata infrastructure was successfully provisioned. You can create a cloud VM cluster on the resource to complete the infrastructure provisioning.
 - **Updating:** The Cloud Exadata infrastructure is being updated. The resource goes into the updating state during management tasks. For example, when moving the resource to another compartment, or creating a cloud VM cluster on the resource.
 - **Maintenance in Progress:** A maintenance update is currently being performed on the infrastructure resource. See [Maintaining an Exadata Cloud Service Instance](#) for details on infrastructure maintenance scheduling and impacts.
 - **Terminating:** The Cloud Exadata infrastructure is being deleted by the terminate action in the Console or API.

- **Terminated:**The Cloud Exadata infrastructure has been deleted and is no longer available.
- **Failed:** An error condition prevented the provisioning or continued operation of the Cloud Exadata infrastructure.

To change the infrastructure display name

Note

This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Service instance resource model.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Cloud Exadata infrastructure resources, click the name of the infrastructure you're interested in
5. On the **Infrastructure Details** page, click **Update Display Name** .
6. In the **Update Display Name** dialog, Enter the **New display name**, and the **current display name** as instructed.
7. Click **Update Display Name**.

To check the status of a cloud VM cluster

Note

This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of cloud VM clusters, find the cluster you're interested in and check its icon. The icon text indicates the status of the system. The following lifecycle states apply to the cloud VM cluster:
 - **Provisioning:** Resources are being reserved for the Cloud Exadata infrastructure resource. Provisioning can take several minutes. The resource is not ready to use yet.
 - **Available:** The Cloud Exadata infrastructure was successfully provisioned. You can create a cloud VM cluster on the resource to complete the infrastructure provisioning.

- **Updating:** The Cloud Exadata infrastructure is being updated. The resource goes into the updating state during management tasks. For example, when moving the resource to another compartment, or creating a cloud VM cluster on the resource.
- **Terminating:** The Cloud Exadata infrastructure is being deleted by the terminate action in the Console or API.
- **Terminated:** The Cloud Exadata infrastructure has been deleted and is no longer available.
- **Failed:** An error condition prevented the provisioning or continued operation of the Cloud Exadata infrastructure.

To view the status of a virtual machine (database node) in the cloud VM cluster, under Resources, click **Virtual Machines** to see the list of virtual machines. In addition to the states listed for a cloud VM cluster, a virtual machine's status can be one of the following:

- **Starting:** The database node is being powered on by the start or reboot action in the Console or API.
- **Stopping:** The database node is being powered off by the stop or reboot action in the Console or API.
- **Stopped:** The database node was powered off by the stop action in the Console or API.

To start, stop, or reboot an Exadata Cloud Infrastructure cloud VM cluster

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to start, stop, or reboot:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Virtual Machines** to display the compute nodes of the cloud service instance. Click Actions icon (three dots) for a node and then click one of the following actions:
 - **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
 - **Stop:** Shuts down the node. After the node is powered off, the **Start** action is enabled.
 - **Reboot:** Shuts down the node, and then restarts it.

Note

- For billing purposes, the **Stop** state has no effect on the resources you consume. Billing continues for virtual machines or nodes that you stop, and related resources continue to apply against any relevant quotas. You must **Terminate** a cloud VM cluster to remove its resources from billing and quotas.
- After you restart or reboot a node, the floating IP address might take several minutes to be updated and display in the Console.

To terminate Exadata Cloud Infrastructure infrastructure-level resources

This topic describes how to terminate a Cloud Exadata infrastructure or cloud VM cluster resource in an Exadata Cloud Infrastructure instance.

Note

The database data is local to the cloud VM cluster hosting it and is lost when the system is terminated. Oracle recommends that you back up any data in the cloud VM cluster before terminating it.

Terminating an Exadata Cloud Infrastructure resource permanently deletes it and any databases running on it. The data is deleted in compliance with the NIST SP-800-88r1 standard and implemented as an Exadata crypto erase using the hardware Instant Secure Erase (ISE) feature of the Exadata storage devices, as documented in the Exadata Database Machine Documentation at [Securely Erasing Exadata Database Machine](#)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Cloud Exadata infrastructure, cloud VM cluster you want to move:
Cloud Exadata infrastructure ([new resource model](#)): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
Cloud VM clusters ([new resource model](#)): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **More Actions**, then **Terminate** on the resource details page. Confirm when prompted.

The resource's icon indicates Terminating.

Note

If you are terminating a Cloud Exadata infrastructure resource that contains a cloud VM cluster, you must check the box labelled **Also delete the VM cluster associated with this infrastructure** to confirm that you intend to delete the VM cluster.

At this point, you cannot connect to the system and any open connections are terminated.

To View a List of DB Servers on an Exadata Infrastructure

To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. In the list of Exadata Infrastructures, click the display name of the infrastructure you wish to view details.
4. Under **Resources**, click **DB Servers**.
5. In the list of DB Servers, click the name of the DB Server that you wish to view details.
DB Server lists VMs from each cluster hosted on them along with resources allocated to them.

Network Management Operations

- [To edit the network security groups \(NSGs\) for your client or backup network](#)

To edit the network security groups (NSGs) for your client or backup network

Your client and backup networks can each use up to five network security groups (NSGs). Note that if you choose a subnet with a [security list](#), the security rules for the cloud VM cluster will be a union of the rules in the security list and the NSGs. For more information, see [Network Security Groups](#) and [Network Setup for Exadata Cloud Infrastructure Instances](#).

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to manage:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the **Network** details, click the **Edit** link to the right of the **Client Network Security Groups** or **Backup Network Security Groups** field.
5. In the **Edit Network Security Groups** dialog, click **+ Another Network Security Group** to add an NSG to the network.

To change an assigned NSG, click the drop-down menu displaying the NSG name, then select a different NSG.

To remove an NSG from the network, click the **X** icon to the right of the displayed NSG name.

6. Click **Save**.

Management Tasks for the Oracle Cloud Infrastructure Platform

- [To view a work request for your Exadata Cloud Infrastructure resources](#)
- [To move an Exadata Cloud Infrastructure resource to another compartment](#)
- [To manage tags for your Exadata Cloud Infrastructure resources](#)
- [Managing Infrastructure Maintenance Contacts](#)

Learn to manage your Exadata infrastructure maintenance contacts.

To view a work request for your Exadata Cloud Infrastructure resources

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, or database resource you're interested in, and click the name.
4. In the **Resources** section, click **Work Requests**. The status of all work requests appears on the page.
5. To see the log messages, error messages, and resources that are associated with a specific work request, click the operation name. Then, select an option in the **More information** section.
For associated resources, you can click the Actions icon (three dots) next to a resource to copy the resource's OCID.

Related Topics

- [Work Requests](#)

To move an Exadata Cloud Infrastructure resource to another compartment

Note

- To move resources between compartments, resource users must have sufficient access permissions on the compartment that the resource is being moved to, as well as the current compartment. For more information about permissions for Database resources, see *Details for the Database Service*.
- If your Exadata Cloud Infrastructure instance is in a security zone, the destination compartment must also be in a security zone. See the *Security Zone Policies* topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Cloud Exadata infrastructure, cloud VM cluster you want to move:
Cloud Exadata infrastructure (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Move Resource**.
5. Select the new compartment.
6. Click **Move Resource**.

Related Topics

- [Details for the Database Service](#)
- [Security Zone Policies](#)
- [The New Exadata Cloud Infrastructure Resource Model](#)

To manage tags for your Exadata Cloud Infrastructure resources

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Find the Cloud Exadata infrastructure, cloud VM cluster, or database resource you're interested in, and click the name.
4. Click the **Tags** tab to view or edit the existing tags. Or click **More Actions** and then **Apply Tags** to add new ones.

Related Topics

- [Resource Tags](#)

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- [To manage maintenance contacts in an Exadata Cloud Infrastructure](#)
Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To manage maintenance contacts in an Exadata Cloud Infrastructure

Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To prevent an Exadata infrastructure administrator from being overwhelmed by system update notifications, you can specify up to 10 email addresses of people to whom maintenance notifications are sent.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure** to display a list of Exadata infrastructures in the default compartment. You can select a different compartment from the **Compartment** drop-down located in the **List Scope** section.
3. In the list of Exadata infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
4. In the **Maintenance** section, click **Manage** in the **Customer Contacts** field to display the Manage Contacts dialog.
5. Click the **Add Contacts** button to display a field in which to enter a valid email address. You can have up to 10 maintenance contacts for each Exadata infrastructure.
6. To edit an email address, in the Manage Contacts dialog, select the box preceding the email address you want to edit and click the **Edit** button.
7. To remove an email address from the list, in the Manage Contacts dialog, select the box preceding the email address you want to remove and click the **Remove** button.

Oracle Database License Management Tasks

- [To manage your BYOL database licenses](#)
If you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.
- [To change the license type of a cloud VM cluster](#)

To manage your BYOL database licenses

If you want to control the number of database licenses that you run at any given time, you can scale up or down the number of OCPUs on the instance. These additional licenses are metered separately.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to scale:

Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

4. Click **Scale VM Cluster** and then specify a new number of CPU cores.
5. Click **Scale**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To change the license type of a cloud VM cluster

① Note

Updating the license type is not supported for systems running on the X6 shape. The feature is supported for X7 and higher shapes.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to manage:
 - a. *Cloud VM clusters (new resource model)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. On the resource details page, click **Update License Type**.
The dialog displays the options with your current license type selected.
5. Select the new license type.
6. Click **Save**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

Scaling Resources within an Exadata Infrastructure Instance

If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.

There are three distinct scaling paths:

- Scaling an X8M, X9M, or X11M based Multi-VM enabled Infrastructure by adding DB servers and Storage servers to an existing infrastructure. See *Add Resources To a Multi-VM Enabled Infrastructure*.
- Scaling an X8M or X9M based that is NOT a Multi-VM enabled Infrastructure. See *Scaling Exadata X8M and X9M Compute and Storage*.
- Scaling an X6, X7, and X8 Exadata infrastructure (fixed shape).
- [Add Resources to a Multi-VM enabled Infrastructure](#)
Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure
- [Remove DB Servers from a Multi-VM enabled Infrastructure](#)
Remove DB servers from an existing Multi-VM enabled Infrastructure
- [Remove Storage Servers from a Multi-VM enabled Infrastructure](#)
Remove storage servers from an existing Multi-VM enabled Infrastructure
- [Scaling CPU cores within an Exadata Cloud Infrastructure instance](#)
If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:

Related Topics

- [Add Resources to a Multi-VM enabled Infrastructure](#)
Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure

Add Resources to a Multi-VM enabled Infrastructure

Add DB servers or storage servers to an existing Multi-VM enabled Infrastructure

You can scale an X8M, X9M, or X11M Exadata Multi-VM enabled Infrastructure instance in the Console on the cloud Exadata infrastructure details page. After adding additional database or storage servers to your cloud Exadata infrastructure resource, you must add the increased capacity to the associated cloud VM cluster to utilize the newly-provisioned CPU or storage resources.

Adding DB servers, or Storage servers do not require any database downtime.

To add DB Servers to Multi-VM enabled Infrastructure

1. Navigate to **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Select **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**
3. Select the desired Infrastructure in the desired compartment.
4. On the **Infrastructure Details** page click **Scale Infrastructure**.

5. In the **Scale Infrastructure** page, set the **Database servers** to a value so that the total of DB servers is 8 or less.
6. Click **Scale**.

To add storage servers to Multi-VM enabled Infrastructure

1. Navigate to **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Select **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**
3. Select the desired Infrastructure in the desired compartment.
4. On the **Infrastructure Details** page click **Scale Infrastructure**.
5. In the **Scale Infrastructure** page, set the **Storage servers** to a value so that the total of storage servers is 12 or less.

Note

- This operation adds the storage servers to the infrastructure, but the storage capacity must be made available for VM Cluster consumption.
- You will be able to scale down a storage server if the server has not been used to expand Exadata infrastructure storage.

6. Click **Scale**.
7. On the **Infrastructure Details** page, a banner directs you to **Add Storage Capacity** to make the storage capacity available for VM Cluster consumption.
8. Click **Add Storage Capacity**.
9. In the **Add Storage Capacity** page, click **Add Storage Capacity**.

Remove DB Servers from a Multi-VM enabled Infrastructure

Remove DB servers from an existing Multi-VM enabled Infrastructure

Database servers will be removed if there are no VMs running on them.

Note

You will not be able to choose the DB Server to remove. This functionality will automatically remove Database Servers in which there are no VMs.

For more information about removing a VM, see [Terminate or Remove a VM from a VM Cluster](#).

To remove DB Servers from a Multi-VM enabled Infrastructure, follow these steps:

1. Navigate to the **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select Exadata Infrastructure under **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. Select the desired Infrastructure in the desired compartment.

4. On the Infrastructure Details page click **Scale Infrastructure**.
5. In the Scale Infrastructure page, set the database servers to the desired value by typing in the value of the target database servers.
6. Click **Scale**.

Remove Storage Servers from a Multi-VM enabled Infrastructure

Remove storage servers from an existing Multi-VM enabled Infrastructure

Before proceeding, ensure that the available storage on the Cloud Exadata Infrastructure is sufficient—it should equal the capacity of the storage server being deleted plus the additional storage required for ASM disk group rebalancing. Please refer to MOS note [1551288.1](#) for more details on the additional storage required.

To remove storage servers from a Multi-VM enabled Infrastructure, follow these steps:

1. Navigate to the **Oracle Cloud** menu and click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select Exadata Infrastructure under **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. Select the desired Infrastructure in the desired compartment.
4. On the Infrastructure Details page click **Scale Infrastructure**.
5. On the Scale Infrastructure page, set the storage servers to the desired value by typing in the value of the target database servers.
6. Click **Scale**.

During the storage server deletion process, the Exadata Infrastructure transitions to an UPDATING state. Once the deletion is successfully completed, the infrastructure returns to the AVAILABLE state. This workflow may take several hours to complete. After completion, the Exadata Resources section will reflect the updated Exadata Storage configuration.

Note

- The minimum number of provisioned storage servers cannot be less than three.
- You cannot remove provisioned storage to the extent that the total remaining available storage falls below the current amount of Exadata storage in use by the clusters.
- The existing storage (in TB) used by a cluster will remain unchanged after deleting the storage cell server.
- Only the total available Exadata storage will be updated, while the used storage will stay the same.

Scaling CPU cores within an Exadata Cloud Infrastructure instance

If an Exadata Cloud Infrastructure instance requires more compute node processing power, you can scale up the number of enabled CPU cores symmetrically across all the nodes in the system as follows:

The options for each of the shapes are:

X8M, X9M, or X11M flexible infrastructure systems: You can scale CPU cores in multiples of the number of database servers currently provisioned for the cloud VM cluster. For example, if you have 6 database servers provisioned, you can add CPU cores in multiples of 6. In the case of X11M database servers, you will need to add ECPU cores in multiples of 24 for 6 database servers. At the time of provisioning, X8M, X9M, and X11M systems have as few as 2 database servers or up to 32 storage servers. For more information on adding compute and storage resources to an X8M, X9M, or X11M system, see *Scaling Exadata X8M, X9M, and X11M Compute and Storage*.

Non-X8M fixed-shape systems: For a base system or an X7 or X8 quarter rack, you can scale in multiples of 2 across the 2 database compute nodes. For an X7 or X8 half rack, you can scale in multiples of 4 across the 4 database compute nodes. For an X7 or X8 full rack, you can scale in multiples of 8 across the 8 database compute nodes.

For non-metered service instances, you can temporarily modify the compute node processing power (bursting) or add compute node processing power on a more permanent basis. For a metered service instance, you can simply modify the number of enabled CPU cores.

You can provision an Exadata Cloud Infrastructure instance with zero CPU cores, or scale the service instance down to zero cores after you provision it. With zero cores, you are billed only for the infrastructure until you scale up the system. For detailed information about pricing, see [Exadata Cloud Service Pricing](#).

Note

OCPU scaling activities are done online with no downtime.

For information on CPU cores per configuration, see *Exadata Shape Configurations*. To learn how to scale a system, see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster*.

- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.

Related Topics

- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)

Scaling Exadata X8M, X9M, and X11M Compute and Storage

The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.

You can scale an X8M, X9M, and X11M Exadata cloud infrastructure instance in the Console on the cloud Exadata infrastructure details page. After adding additional database or storage servers to your cloud Exadata infrastructure resource, you must add the increased capacity to the associated cloud VM cluster to utilize the newly-provisioned CPU or storage resources. After adding additional database servers to a VM cluster, you can then allocate the new CPU cores as described in see *To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster*. After adding additional storage servers to your VM cluster, you do not need to take any further action to utilize the new storage.

Note

- Neither the Exadata X8M, X9M, nor the X11M shapes support removing storage from an existing cloud infrastructure instance. For more information, see [Remove DB Servers from a Multi-VM enabled Infrastructure](#).

- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)

Related Topics

- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)
- [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#)
Oracle Cloud Infrastructure scalable X8M, X9M, and X11M Exadata cloud infrastructure model allows you to add additional database and storage servers after provisioning and create a system that matches your capacity needs.
- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- [To add database server or storage server capacity to a cloud VM cluster](#)
This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

To add compute and storage resources to a flexible cloud Exadata infrastructure resource

This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Currently, only Exadata X8M, X9M, and X11M systems in Oracle Cloud Infrastructure have the ability to add database (compute) and storage servers to an existing service instance.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**.
3. In the list of cloud Exadata infrastructure resources, click the name of the resource you want to scale.
4. Click **Scale Infrastructure**.
5. Add either Database servers or Storage Servers by selecting the proper radio button
 - a. *Adding database servers:* To add compute servers to the infrastructure resource, select the **Database Servers** radio button, then enter the number of servers you want to add in the **Database servers** field.
 - b. *Adding storage servers:* To add storage servers to the infrastructure resource, select the **Storage Servers** radio button, then enter the number of servers you want to add in the **Storage servers** field.
6. Click **Scale**.

Note

After scaling your infrastructure, you must add the new capacity to the cloud VM cluster before you can use the additional CPU and storage resources in the Exadata Cloud Infrastructure instance.

Related Topics

- [To add database server or storage server capacity to a cloud VM cluster](#)
This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster

Note

For information on adding additional database (compute) and storage servers to X8M, X9M, or X11M cloud VM clusters, see [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#) and [To add database server or storage server capacity to a cloud VM cluster](#). Adding additional database servers to your X8M cloud VM cluster will increase the number of CPU cores available for scaling.

If an Exadata Cloud Infrastructure VM Cluster requires more compute node processing power or memory, you can scale up (increase) the number of enabled CPU cores (OCPU or ECPUS for X11M) or memory in the VM cluster.

You can also scale a cloud VM cluster down to zero (0) CPU cores to temporarily stop the system and be charged only for the hardware infrastructure. For more information about scaling down, see [Scaling Options](#). Oracle recommends that if you are not scaling down to a stopped system (0 cores), you scale to at least 2 cores per guest VM or ECPUS per X11M guest VM.

Note

The minimum number of cores is 1 for X8 and older, and 2 for X8M and newer. The minimum number of ECPUS for an X11M guest VM is 8.

CPU cores must be scaled symmetrically across all nodes in the cloud VM cluster. Use multiples of two CPUs per database server or multiples of 8 ECPUS per X11M database server. For example, if you have two database servers, a minimum of 2 CPU cores per server or a total of 4 CPU cores. In the case of an X11M database server, this is a minimum of 8 ECPUS per server or a total of 16 ECPUS. The total number of CPU cores must not exceed the maximum limit for that shape and/or resources.

Memory must be scaled symmetrically across all the nodes in the cloud VM cluster. The minimum memory per VM is 30 GB, but you can scale it up to the maximum available memory on the database server.

Since the database server CPU and memory are shared across the VM clusters you need to be careful when resizing memory and CPU.

Note

When you scale up or down the memory, the associated compute nodes are rebooted in a rolling manner one compute node at a time to minimize the impact on the VM cluster.

Tip

OCPU (ECPUs for X11M) scaling activities are done online with no downtime.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to scale:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Scale VM Cluster**, and then specify a new number of CPU cores (ECPUs for X11M) and memory.
5. Click **Scale**.

Note

If you scale a cloud VM cluster down to zero (0) CPU cores, the floating IP address of the nodes might take several minutes to update and display in the Console.

Using the API to Create Infrastructure Components

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

Use these API operations to create Exadata Cloud Infrastructure components.

APIs for the New Exadata Cloud Infrastructure Resource Model

The new Exadata resource model is compatible with all offered Exadata shape families (X7, X8, and X8M). See [The Exadata Cloud Infrastructure Resource Model](#) for more information.

Tip

As of November 15, 2021 new Exadata Cloud Infrastructure instances may only be provisioned using the new resource model.

Cloud Exadata infrastructure resource:

- [GetCloudExadataInfrastructure](#)

- [CreateCloudExadataInfrastructure](#)
- [ListCloudExadataInfrastructures](#)

Cloud VM cluster resource:

- [GetCloudVmCluster](#)
- [CreateCloudVmCluster](#)
- [ListCloudVmClusters](#)
- [ListCloudVmClusters](#)
- [CreateCloudVmCluster](#)
- [GetCloudVmClusterIormConfig](#)
- [UpdateCloudVmClusterIormConfig](#)

Databases

- [GetDatabase](#)
- [ListDatabases](#)

Shapes and Database Versions

- [ListDbSystemShapes](#)
- [ListDbVersions](#)

Database Homes

- [CreateDbHome](#)
- [GetDbHome](#)
- [ListDbHomes](#)

NOT_SUPPORTED

Note

The DB system APIs are deprecated for Exadata Cloud Infrastructure. Oracle recommends converting existing Exadata DB systems to the new resource model as soon as possible. Converting to the new resource model does not involve system downtime. [Learn more](#).

- [GetDbSystem](#)
- [LaunchDbSystem](#)
- [ListDbSystems](#)

Using the API to Manage Exadata Cloud Infrastructure Instance

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

Use these API operations to manage Exadata Cloud Infrastructure instance components.

Cloud Exadata infrastructure resource (new resource model):

- [ListCloudExadataInfrastructures](#)
- [GetCloudExadataInfrastructure](#)
- [ChangeCloudExadataInfrastructureCompartment](#)
- [UpdateCloudExadataInfrastructure](#)
- [DeleteCloudExadataInfrastructure](#)

Cloud VM cluster (new resource model)

- [ListCloudVmClusters](#)
- [GetCloudVmCluster](#)
- [ChangeCloudVmClusterCompartment](#)
- [UpdateCloudVmCluster](#)
- [DeleteCloudVmCluster](#)

DB systems (old resource model):

- [ListDbSystems](#)
- [GetDbSystem](#)
- [ChangeDbSystemCompartment](#)
- [UpdateDbSystem](#)
- [TerminateDbSystem](#)

Virtual machines nodes (all Exadata Cloud Infrastructure instances):

- [DbNodeAction](#)
- [ListDbNodes](#)
- [GetDbNode](#)

Configure Oracle-Managed Infrastructure Maintenance

Oracle performs the updates to all of the Oracle-managed infrastructure components on Exadata Cloud Infrastructure.

Oracle performs periodic maintenance on your Exadata Cloud Infrastructure infrastructure to ensure it is free of potential issues that could affect database availability, integrity, and security. Updating the software running on the infrastructure with the latest product and security fixes protects your data and the overall compliance of the Oracle Cloud. These updates are performed in an automated way and include all the best practices, relieving you of the need to invest any effort in maintaining your infrastructure. Oracle updates include the physical database server hosts, storage servers, network fabric and management switches (InfiniBand systems only), power distribution units (PDUs), integrated lights-out management (ILOM) interfaces, and control plane servers.

Oracle performs two types of infrastructure maintenance:

- Quarterly maintenance is applied every three months and can include product fixes, enhancements, and security fixes.
- Monthly maintenance only applies critical security fixes that can be applied online to ensure components are maintained at the highest security standards with any security vulnerabilities fixed as soon as possible.

Quarterly Maintenance

Oracle minimizes the impact of quarterly maintenance on your applications using rolling maintenance operations, preserving database availability throughout the update process. Rolling maintenance reboots each database server, one at a time, with at most one server offline at any time. Applications designed for high availability automatically and transparently migrate their database connections between available database instances without disruption, eliminating the need for scheduling downtime. Storage server updates are also applied in a rolling manner. Rebooting storage servers has no effect on the database service, and thus has no impact on your applications.

Oracle allows you to fully control quarterly maintenance schedules, so you can schedule maintenance during a period which will have the least impact on your business users. You have full control and visibility over when quarterly maintenance will be applied. You may also reschedule maintenance should unexpected business issues occur.

Monthly Security Maintenance

Monthly security maintenance is performed on the database servers online, with no reboot, and no impact to your applications. Monthly updates are applied to storage servers in a rolling manner, also with no impact to your applications.

Monthly security maintenance can also be scheduled at a specific time during the month, albeit in a single maintenance window. Oracle will publish a schedule for monthly maintenance at least one week prior to start of the maintenance period, and you can reschedule if required.

- [About Oracle-managed Exadata Cloud Infrastructure Maintenance](#)
Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.
- [Overview of the Quarterly Infrastructure Maintenance Process](#)
By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.
- [Overview of Monthly Security Maintenance](#)
Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.
- [Understanding Monthly and Quarterly Maintenance in the Same Month](#)
- [Maintenance Scheduling Policy](#)
Learn how to use the OCI Console to configure and manage maintenance scheduling policies.
- [Using the Console to Configure Oracle-Managed Infrastructure Updates](#)
Software updates are scheduled quarterly and monthly. You can use the console to schedule and plan for them.
- [Manage Quarterly Maintenance Run Created From Scheduling Plan](#)
- [Monitor Infrastructure Maintenance Using Lifecycle State Information](#)
The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.
- [Receive Notifications about Your Infrastructure Maintenance Updates](#)
There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.
- [Managing Infrastructure Maintenance Contacts](#)
Learn to manage your Exadata infrastructure maintenance contacts.

- [Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls](#)
Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.

About Oracle-managed Exadata Cloud Infrastructure Maintenance

Oracle performs patches and updates to all of the Oracle-managed system components on Exadata Cloud Infrastructure.

The frequency of the updates depends on the region type, as follows:

- Commercial regions: Oracle performs full quarterly infrastructure updates and monthly security infrastructure updates.
- Government regions: Oracle performs monthly full infrastructure maintenance updates.

In all but rare exceptional circumstances, you receive advance communication about these updates to help you plan for them. If there are corresponding recommended updates for your VM cluster virtual machines (VMs), then Oracle provides notifications about them.

Wherever possible, scheduled updates are performed in a manner that preserves service availability throughout the update process. However, there can be some noticeable impact on performance and throughput while individual system components are unavailable during the update process.

For example, database server patching typically requires a reboot. In such cases, wherever possible, the database servers are restarted in a rolling manner, one at a time, to ensure that the service remains available throughout the process. However, each database server is unavailable for a short time while it restarts, and the overall service capacity diminishes accordingly. If your applications cannot tolerate the restarts, then take mitigating action as needed. For example, shut down an application while database server patching occurs.

Note

Customers using Exadata Database on Dedicated Infrastructure in Oracle Cloud Infrastructure (OCI) US Government (OC2) and US Department of Defense (OC3) regions can use the OCI console to reschedule monthly and quarterly patching events.

At this time specifying a maintenance schedule, all so known as “Setting Patch Management Schedule for Exadata Cloud Infrastructure”, is still not available in the OCI US Government (OC2) and US DOD (OC3) realms for Exadata patch management. For more information on Exadata Database on Dedicated Infrastructure on Patch Management Rescheduling can be found [here](#).

Overview of the Quarterly Infrastructure Maintenance Process

By default, infrastructure maintenance updates the Exadata database server hosts in a rolling fashion, followed by updating the storage servers.

Rolling infrastructure maintenance begins with the Exadata database server hosts. For the rolling maintenance method, database servers are updated one at a time. Each of the database server host's VMs is shut down, the host is updated, restarted, and then the VMs are started, while other database servers remain operational. This rolling maintenance does not impact applications designed for high availability. Older applications not written to handle a rolling instance restart can be impacted. This process continues until all servers are updated.

After database server maintenance is complete, storage server maintenance begins. For the rolling maintenance method, storage servers are updated one at a time and this does not impact database or application availability. However, the rolling storage server maintenance can result in reduced IO performance as storage servers are taken offline one at a time (reducing available IO capacity) and resynced when brought back online (small overhead on database servers). Properly sizing the database and storage infrastructure to accommodate increased work distributed to database and storage servers that are not under maintenance will minimize (or eliminate) any performance impact.

You can also choose non-rolling maintenance to update database and storage servers. The non-rolling maintenance method first updates your storage servers at the same time, then your database servers at the same time. Although non-rolling maintenance minimizes maintenance time, it incurs full system downtime while the storage servers and database servers are being updated.

Note that while databases are expected to be available during the rolling maintenance process, the automated maintenance verifies Oracle Clusterware is running but does not verify that all database services and pluggable databases (PDBs) are available after a server is brought back online. The availability of database services and PDBs after maintenance can depend on the application service definition. For example, a database service, configured with certain preferred and available nodes, may be relocated during the maintenance and wouldn't automatically be relocated back to its original node after the maintenance completes. Oracle recommends reviewing the documentation on [Achieving Continuous Availability for Your Applications on Exadata Cloud Systems](#) to reduce the potential for impact to your applications. By following the documentation's guidelines, the impact of infrastructure maintenance will be only minor service degradation as database servers are sequentially updated.

Oracle recommends that you follow the Maximum Availability Architecture (MAA) best practices and use Data Guard to ensure the highest availability for your critical applications. For databases with Data Guard enabled, Oracle recommends that you separate the maintenance windows for the infrastructure instances running the primary and standby databases. You may also perform a switchover prior to the maintenance operations for the infrastructure instance hosting the primary database. This allows you to avoid any impact on your primary database during infrastructure maintenance.

Prechecks are performed on the Exadata Cloud Infrastructure components prior to the start of the maintenance window. The goal of the prechecks is to identify issues that may prevent the infrastructure maintenance from succeeding. The Exadata infrastructure and all components remain online during the prechecks. An initial precheck is run approximately 5 days prior to the maintenance start and another precheck is run approximately 24 hours prior to maintenance start. If the prechecks identify an issue that requires rescheduling the maintenance notification is sent to the maintenance contacts.

Review [MOS Note KB181723](#) for the Exadata Infrastructure maintenance timeframes.

Note

Do not perform major maintenance operations on your databases or applications during the patching window, as these operations could be impacted by the infrastructure maintenance operations

Related Topics

- [Achieving Continuous Availability For Your Applications](#)
- [Maximum Availability Architecture \(MAA\) Best Practices](#)
- <https://support.oracle.com/rs?type=doc&id=2333222.1>

Overview of Monthly Security Maintenance

Security maintenance, performed alongside the quarterly maintenance, is executed in months when important security updates are needed and includes fixes for vulnerabilities across all CVSS scores.

Note

For more information about the CVE release matrix, see Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1). To view the CVE release matrix specific to an Exadata Infrastructure version, click the Exadata version, for example, Exadata 23. Version-specific CVE release matrices are listed in the Notes column of the table.

Security maintenance, when needed, is scheduled to be applied during a 21-day window that begins between the 18th-21st of each month and will run till the 9th-12th of the next month. Customers will receive notification of the proposed schedule at least 7 days before the start of the monthly maintenance window and can reschedule monthly maintenance to another date in the window if desired. The monthly security maintenance process updates the physical database servers to fix critical security vulnerabilities and critical product issues. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues. No updates are applied to the customer-managed guest VMs. Monthly maintenance also updates storage servers to an Exadata Storage Software image that resolves known security vulnerabilities and product issues.

Updates to database servers are applied online via Ksplice technology, and have no impact to workloads running on the compute (database) servers, as database server security updates are applied online to the host server while your VM and all processes within the VM, including databases, remain up and running. Servers and VMs are not restarted. Updates to storage servers are applied in a rolling fashion. As with quarterly maintenance, the impact of rebooting storage servers should be minimal to applications.

While updating your services infrastructure, some operations including memory, and storage scaling, operating system and Grid Infrastructure patching (including prechecks), and elastic expansion of compute and storage servers may be blocked.

Note

CPU scaling and VM startup/shutdown are the only operations supported during monthly infrastructure maintenance.

Please plan to defer these operations until after the updates are complete. If you attempt an affected operation, the console will notify you of the ongoing security updates. No software is updated in the guest VMs.

Related Topics

- <https://support.oracle.com/rs?type=doc&id=888828.1>
- <https://support.oracle.com/rs?type=doc&id=2333222.1>

Understanding Monthly and Quarterly Maintenance in the Same Month

Special considerations are made when both quarterly and monthly security maintenance are scheduled to run in the same month. Quarterly maintenance will reapply any security fixes already applied by security maintenance, and neither quarterly nor monthly maintenance will apply a storage server update if the existing storage server version is the same or newer than the version contained in the update.

- The contents of the updates applied during quarterly maintenance are determined at the start of the maintenance quarter and use the latest Exadata release from the month prior to the start of the maintenance quarter. If any additional security fixes are available at that time, those updates are included in the quarterly maintenance. That image is then used throughout the quarter. For example, the January release is used for quarterly maintenance in Feb, March, and April.
- When quarterly maintenance is applied it is possible there are security updates previously installed on the database servers are not included in the quarterly maintenance release to be applied. In that case, the automation will apply the same security fixes to new release installed by the quarterly maintenance so there will not be any regression in security fixes. If the current image on the storage server is the same or newer than that to be applied by the quarterly or monthly security maintenance, that maintenance will be skipped for the storage servers.

If quarterly maintenance is scheduled within 24 hours of the time the monthly is scheduled, the scheduled monthly maintenance will be skipped and the monthly update will instead be applied immediately following the quarterly maintenance.

- When scheduled at the same time, the monthly update is executed immediately following the completion of the quarterly maintenance.
- If monthly maintenance is scheduled to begin 0-24 hours ahead of the quarterly maintenance, then the monthly maintenance will not execute as scheduled, but instead, wait and be executed immediately following the quarterly maintenance. If the quarterly maintenance is subsequently rescheduled, then the monthly security maintenance will begin immediately. Oracle, therefore, recommends scheduling quarterly and monthly maintenance at the same time. As a result, if you reschedule the quarterly at the last moment, the monthly maintenance will run at the scheduled time instead of immediately upon editing the schedule. You can also reschedule the monthly security maintenance when rescheduling the quarterly maintenance as long as you keep the monthly within the current maintenance window. Monthly maintenance can be rescheduled to another time in the maintenance window, but cannot be skipped.

Monthly Security Maintenance before Quarterly Maintenance

- To apply security maintenance before quarterly maintenance, reschedule the monthly security maintenance to occur more than 24 hours prior to the quarterly maintenance. The security maintenance will online apply security patches to the database servers with no impact to applications, and apply an update to the storage servers with minimal to no impact (may be slight performance degradation) on applications. The quarterly maintenance will follow as scheduled, and will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. As part of the quarterly maintenance, it will apply the same security updates to the database server that are already installed on the system (no security regression).
- If you are concerned about getting the latest security updates applied, schedule the monthly security maintenance to run after the new monthly maintenance window opens (usually on the 21st of the month).

- The impact of the monthly security maintenance rebooting the storage servers should be minimal, so impact to the applications during this month will only be due to the restart of the database servers during the quarterly maintenance. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows.

Quarterly Maintenance before Monthly Security Maintenance

- To run the quarterly maintenance before the monthly security maintenance, reschedule the security maintenance to run no earlier than 24 hours before the quarterly maintenance is scheduled to start. The security maintenance will be deferred until the quarterly maintenance is completed. The quarterly maintenance will perform rolling maintenance on the database servers, which will impact applications not written to handle a rolling reboot. The quarterly maintenance may or may not skip the storage server patching. That depends on if it is newer or older than the release currently installed. In most cases, the version installed should be newer than the version associated with the quarterly maintenance. Exceptions to this rule may occur if it is the first month of a maintenance quarter, or you skipped the security maintenance in one or more prior months. The security maintenance will run either immediately after the quarterly maintenance is completed, or when scheduled, whichever is later. It will apply online updates to the database servers (no application impact) and will likely update the storage servers in a rolling manner. In some corner cases, the quarterly maintenance may contain the same storage server release as the security maintenance and the security maintenance storage server updates will be skipped.
- The impact to end users of running the quarterly maintenance before the security maintenance should be roughly the same as running the security maintenance first. The quarterly maintenance will be a disruptive event, but the security maintenance rebooting the storage servers should cause minimal disruption, and the security maintenance is applied to the database servers online. However, if you must coordinate a maintenance window with your end users for the security maintenance, this will require two maintenance windows. You can schedule those two maintenance windows to run back-to-back, to appear as single maintenance window to end users. To do this, reschedule the security maintenance to start at the same time (or up to 24 hours prior) as the quarterly maintenance. The security maintenance will be deferred until the quarterly maintenance is completed. Assuming you have been regularly applying monthly security maintenance, the storage servers will be skipped by the quarterly maintenance and will be updated by the security maintenance immediately upon the completion of the quarterly maintenance.

Minimizing Maintenance Windows

- To minimize the number of maintenance windows (you have to negotiate those with end users), schedule the quarterly maintenance and monthly maintenance at the same time. The security maintenance will be blocked. The quarterly maintenance will update the database servers in a rolling manner and will most likely skip the storage server. The security maintenance will follow up immediately and update the database servers online and the storage servers in a rolling manner. The result is a single database and storage server restart in a single maintenance window.
- There are two exceptions to this. 1. If the quarterly and monthly maintenance contain the same storage server release, the quarterly maintenance will apply the storage server update, and the security maintenance will be skipped. From your perspective, this is still a single rolling reboot in a single maintenance window. 2. The currently installed release on the storage servers is older than that contained in the quarterly maintenance, which in turn is older than that in the security maintenance. That would cause the quarterly maintenance to update the storage, and then the security maintenance to do it as well. This can only happen if you skipped a prior month's security maintenance, because it requires the current image to be at least 2 months out of date. In such a scenario, you may want to

schedule the security maintenance first and then the quarterly maintenance. This would result in one storage server reboot, but two distinct maintenance windows — the first for the security maintenance, and then later the quarterly maintenance.

- To minimize the impact to your end users, always apply the monthly security updates, and in months where both are scheduled, schedule them at the same time.

① Note

If the Exadata Infrastructure is provisioned before Oracle schedules the security maintenance, then it will be eligible for security maintenance.

Any time before the scheduled monthly Exadata Infrastructure maintenance, you can reschedule it.

Maintenance Scheduling Policy

Learn how to use the OCI Console to configure and manage maintenance scheduling policies.

The Maintenance Scheduling Policy standardizes scheduling across the fleet, ensuring consistency and efficiency. By defining a single policy and applying it to multiple resources, it streamlines the scheduling process and aligns maintenance activities with business best practices.

Serving as a central repository, the policy documents and coordinates maintenance commitments with stakeholders, improving compliance and operational efficiency. Its centralized management ensures adherence to compliance requirements while allowing efficient coordination of changes from a single point of control. Additionally, the policy enhances communication about planned maintenance across environments in the fleet, fostering better coordination and awareness.

If a maintenance scheduling policy is used, all the scheduling preferences for the infrastructure maintenance are derived from the policy. Any preference defined within the infrastructure is not valid while the policy is in use.

- [View the List of Maintenance Scheduling Policy](#)
- [Create a Maintenance Scheduling Policy](#)
- [View Details of a Maintenance Scheduling Policy](#)
- [Edit a Maintenance Scheduling Policy](#)
- [Edit a Maintenance Scheduling Policy that's in Needs Attention State](#)
- [View a Maintenance Window of a Maintenance Scheduling Policy](#)
- [Edit a Maintenance Window of a Maintenance Scheduling Policy](#)
- [Add Additional Maintenance Windows to a Maintenance Scheduling Policy](#)
- [View Resources Associated with a Policy](#)
- [Delete a Maintenance Window of a Maintenance Scheduling Policy](#)
- [Move a Maintenance Scheduling Policy to a Different Compartment](#)
- [Add Tags to a Maintenance Scheduling Policy](#)
- [Delete a Maintenance Scheduling Policy](#)

View the List of Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Use the **Compartment** filter to view the list of policies created in a specific compartment.
4. Use the **State** filter the policies by their states.
The policy states include:
 - Creating
 - Needs attention
 - Available
 - Updating
 - Failed
 - Deleting
 - Deleted

Create a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click **Create maintenance scheduling policy**.
4. In the resulting Create maintenance scheduling policy window, enter the following:
 - a. **Name:** Enter a descriptive name for the policy.
 - b. **Compartment:** Choose a compartment where you want to create this resource.
 - c. **Cadence:** Choose a frequency.
Available options:
 - i. Every six months
 - ii. Every quarter

Note

To use the scheduling policy to plan and automate quarterly infrastructure maintenance, you must create a policy with a 'Quarterly' cadence and a scheduled start month value set to 'February.'

- iii. Every month
- d. **Schedule start month:** Choose a start month.
Oracle automation requires a quarterly maintenance schedule to start in February if this policy is used for infrastructure maintenance.

- e. **Maintenance window:** Oracle automation will perform maintenance to run scheduled actions as per your maintenance windows defined in the schedule.
 - i. **Month:** Choose a month. The month selection option depends on the selected policy cadence. For a six-month cadence, you can choose between the First, Second, Third, Fourth, Fifth, or Sixth months. For a quarterly cadence, you can choose between the First, Second, or Third months of every quarter.
 - ii. **Week:** You can choose between the month's First, Second, Third, or Fourth weeks.
 - iii. **Day:** You can choose between the week's days: Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, or Saturday.
 - iv. **Start time:** Specify the start time when you want to begin maintenance.
 - v. **Duration:** Plan your maintenance window based on average time estimates to complete Database Cloud service actions. Refer to [Overview of the Quarterly Infrastructure Maintenance Process](#) for time estimates.
 - vi. **Enforce window duration:** With this option enabled, any scheduled action that goes over the configured window duration will be paused and re-scheduled to resume in a future maintenance window.

Note

If an update or action is already underway and cannot be paused and resumed without causing disruption, we will continue and complete the action. Oracle automation will reschedule any action planned to start after the configured window duration to a future maintenance window.

- f. **Show Advanced Options:**
 - i. **Tags:** (Optional) You can choose to apply tags. If you have permission to create a resource, then you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, then skip this option (you can apply tags later) or ask your administrator.

Note

You can add additional maintenance windows to the scheduling policy after creation.

View Details of a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to view details.

Edit a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to edit.
4. In the resulting Maintenance scheduling policy details page, click **Edit scheduling policy**.
5. In the resulting Edit scheduling policy window, you can only edit the name of the policy.
6. Edit the name and then click **Edit scheduling policy**.

Edit a Maintenance Scheduling Policy that's in Needs Attention State

Note

You cannot use a schedule policy without maintenance windows to plan and automate maintenance activity across services; hence, it has a 'Needs Attention' life cycle state. While the policy has no windows defined, you can change all its properties, including cadence and schedule start month.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that's in **Needs attention** state.
4. In the resulting Maintenance scheduling policy details page, click **Edit scheduling policy**.
5. In the resulting Edit scheduling policy window, update the policy settings.
6. Click **Edit scheduling policy**.

View a Maintenance Window of a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to view Maintenance window details.
The **Maintenance windows** section in the resulting Maintenance scheduling policy details page lists the maintenance windows associated with the policy.

Edit a Maintenance Window of a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to edit Maintenance window details.

The **Maintenance windows** section in the resulting Maintenance scheduling policy details page lists the maintenance windows associated with the policy.

4. Click the Actions menu (three dots) of the Maintenance window you want to edit and then select **Edit maintenance window**.
5. In the resulting Edit maintenance window page, make changes and then click **Save**.

Add Additional Maintenance Windows to a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to add Maintenance windows.
The **Maintenance windows** section in the resulting Maintenance scheduling policy details page lists the maintenance windows associated with the policy.
4. Click **Add maintenance window**.

Note

For security maintenance, you can add the scheduled start time within the 21-day window.

5. In the resulting Add maintenance window page, add details and then click **Add**.

View Resources Associated with a Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to view the associated resources.
The **Associated resources** section in the resulting Maintenance scheduling policy details page lists the resources associated with the policy.

Delete a Maintenance Window of a Maintenance Scheduling Policy

Note

Only maintenance windows not used by any resources to plan and automate maintenance activity can be deleted from the policy. Any window already used by services to automate maintenance cannot be deleted.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.

3. Click the name of the policy that you want to delete Maintenance windows. The **Maintenance windows** section in the resulting Maintenance scheduling policy details page lists the maintenance windows associated with the policy.
4. Click the Actions menu of the Maintenance window you want to delete and then select **Delete**.
5. On the resulting Delete maintenance window dialog, enter the name of the Maintenance window, and then click **Delete**.

Move a Maintenance Scheduling Policy to a Different Compartment

Note

You cannot move a policy across compartments if any resource uses it to plan and automate maintenance activity.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**. The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to move to a different compartment.
4. On the resulting Maintenance scheduling policy details page, click the Actions menu and then select **Move resource**.
5. On the resulting Move resource to a different compartment dialog, choose a compartment, and then click **Move**.

Add Tags to a Maintenance Scheduling Policy

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**. The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to add tags to.
4. In the resulting Maintenance scheduling policy details page, click **Add tags**.
5. In the resulting Add tags dialog, add tags, and then click **Add tags**.

Delete a Maintenance Scheduling Policy

Note

You cannot delete a policy if any resource uses it to plan and automate maintenance activity.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Click the name of the policy that you want to delete.
4. On the resulting Maintenance scheduling policy details page, click the Actions menu and then select **Delete**.
5. On the resulting Delete maintenance policy dialog, enter the name of the policy, and then click **Delete**.

Note

You cannot delete a policy if it's used in a maintenance scheduling plan.

Using the Console to Configure Oracle-Managed Infrastructure Updates

Software updates are scheduled quarterly and monthly. You can use the console to schedule and plan for them.

Full Exadata Cloud Infrastructure software updates are scheduled on a quarterly basis for commercial regions, and monthly for government regions. In addition, important security updates are scheduled monthly. While you cannot opt-out of these infrastructure updates, Oracle alerts you in advance through the Cloud Notification Portal and allows scheduling flexibility to help you plan for them.

For quarterly infrastructure maintenance, you can set a maintenance window to determine when the maintenance will begin. You can also edit the maintenance method, enable custom action, view the scheduled maintenance runs and the maintenance history, and manage maintenance contacts in the in the Exadata Infrastructure Details page of the Oracle Cloud Infrastructure Console.

Note

You can use a scheduling policy for quarterly infrastructure maintenance updates. However, monthly security updates do not support scheduling policies for maintenance. Oracle automation automatically schedules monthly security updates for your infrastructure.

- [To view or edit quarterly maintenance preferences](#)
- [Manage Quarterly Maintenance Plan using Scheduling Policy](#)
- [To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure](#)
Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.
- [View and Edit Maintenance While Maintenance is In Progress](#)
While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

- [View and Edit Maintenance While Maintenance is Waiting for Custom Action](#)
While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.
- [To view or edit a scheduled security maintenance](#)
Learn how to view and edit the next scheduled security maintenance.
- [To view the maintenance history of an Exadata Cloud Infrastructure resource](#)
This task describes how to view the maintenance history for a cloud Exadata infrastructure resource.
- [To set the node patching order for a scheduled infrastructure maintenance run](#)
This task describes how to set the node patching order for a scheduled infrastructure maintenance run for a cloud Exadata infrastructure resource.

To view or edit quarterly maintenance preferences

This task describes how to set quarterly maintenance preferences. The changes you make will only apply to future maintenance runs, not those already scheduled.

Note

Specifying a maintenance schedule is not available in Government regions.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure you want to access:
In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

The infrastructure details page is displayed.

3. Click the Actions menu and then select **Edit Maintenance Preferences**.
Edit Maintenance Preferences page is displayed.

Note

Changes made to maintenance preferences apply only to future maintenance, not the maintenance that has already been scheduled.

4. On the Edit Maintenance Preferences page, configure the following:
 - **Maintenance scheduling preference: Oracle managed schedule**
 - **Choose a maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
 - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this

option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.

- * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

① Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- Click **Save Changes**.

① Note

From the next maintenance run onwards, executions will occur according to Oracle's schedules.

- **Maintenance scheduling preference: Customer managed schedule**
 - **Maintenance schedule:** Define maintenance preferences for this infrastructure. Changes will take effect from the next maintenance run.
 - * **Configure maintenance preference:** Define maintenance time preferences for each quarter. If more than one preference is defined for a quarter, Oracle automation will select one of them to perform maintenance on all components in your infrastructure.
Select at least one month every two quarters.
 - * **Specify a schedule:** Choose your preferred week, weekday, start time, and lead time for infrastructure maintenance.
 - * Optional. Under **Week of the month**, specify which week of the month, maintenance will take place. Weeks start on the 1st, 8th, 15th, and 22nd days of the month, and have a duration of 7 days. Weeks start and end based on calendar dates, not days of the week. Maintenance cannot be scheduled for the fifth week of months that contain more than 28 days. If you do not specify a week of the month, Oracle will run the maintenance update in a week to minimize disruption.
 - * Optional. Under **Day of the week**, specify the day of the week on which the maintenance will occur. If you do not specify a day of the week, Oracle will run the maintenance update on a weekend day to minimize disruption.

- * Optional. Under **Hour of the day**, specify the hour during which the maintenance run will begin. If you do not specify a start hour, Oracle will pick the least disruptive time to run the maintenance update.
- * Under **Notification Lead Time**, specify the minimum number of weeks ahead of the maintenance event you would like to receive a notification message. Your lead time ensures that a newly released maintenance update is scheduled to account for your required minimum period of advanced notification.
- * **Choose a maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
- * **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
- * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- * **Show advanced options:**
 - * Enable monthly security infrastructure maintenance: Select this check box to perform monthly security infrastructure maintenance.
- **Maintenance schedule:** Use maintenance window preferences from a scheduling policy
During infrastructure provisioning, after the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers, into the maintenance windows from your policy based on duration. After provisioning the infrastructure, you can update the

scheduling plan by editing the 'Maintenance Scheduling Plan' resource and customize the update to specific components to align with different windows in your scheduling policy.

- * Click **Select policy**.
- * In the resulting Select maintenance scheduling policy window, choose a compartment and a policy.
You can also create a maintenance scheduling policy and use it. For more information, see [Create a Maintenance Scheduling Policy](#). Note that you can add additional maintenance windows to the policy after creating it. For more information, see [Add Additional Maintenance Windows to a Maintenance Scheduling Policy](#).
- * Click **Save changes**.

Note

Changes will take effect from the next maintenance run.

You must confirm your choice by entering the currently used policy name in a confirmation dialog before making any changes that delete the associated maintenance plan created with the attached policy.

- * Changing from one scheduling policy to another scheduling policy after the recommended maintenance plan is created and saved for the infrastructure
- * Changing from using a scheduling policy to not using a policy and defining your maintenance preference in line with your infrastructure
- * Change from using the scheduling policy to not using the policy and apply updates as per the Oracle-managed schedule.

All of the above changes delete the scheduling plan for your infrastructure created with the current policy, and you will lose any customizations made to the Oracle recommended plan if you attach the same policy later.

5. Click Save Changes.

If you switch from rolling to non-rolling maintenance method, then Confirm Non-rolling Maintenance Method dialog is displayed.

- a. Enter the name of the infrastructure in the field provided to confirm the changes.
- b. Click **Save Changes**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

Manage Quarterly Maintenance Plan using Scheduling Policy

After the scheduling policy is selected, Oracle generates a recommended maintenance scheduling plan to apply updates to all the components in your infrastructure. The recommended plan schedules all DB Servers, followed by Storage Servers, into the maintenance windows from your policy based on duration. You can update the maintenance scheduling plan and customize the update to specific components to align with different windows in your scheduling policy.

- [View Quarterly Maintenance Scheduling Policy](#)
To view the quarterly maintenance scheduling policy for your infrastructure, use this procedure.
- [Change the Quarterly Maintenance Scheduling Policy](#)
To change the quarterly maintenance scheduling policy for your infrastructure, use this procedure.
- [View Maintenance Scheduling Plan](#)
To view the maintenance scheduling plan for your infrastructure, use this procedure.
- [Edit Maintenance Plan Scheduled Actions](#)
To edit scheduled actions of an Exadata Cloud@Customer Infrastructure maintenance scheduling plan, use this procedure.

View Quarterly Maintenance Scheduling Policy

To view the quarterly maintenance scheduling policy for your infrastructure, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view the quarterly maintenance scheduling policy.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure. In the Maintenance section, find the Quarterly Maintenance Schedule policy.
5. Click the link to view details of the policy.

Change the Quarterly Maintenance Scheduling Policy

To change the quarterly maintenance scheduling policy for your infrastructure, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view the quarterly maintenance scheduling policy.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure. In the Maintenance section, find the Quarterly Maintenance Schedule policy.
5. Click the Actions menu and then select **Edit Maintenance Preferences**.
The Customer Managed Schedule and Use maintenance window preferences from a scheduling policy are selected as part your policy selection.
6. Click **Select Policy**.
7. Select an existing policy or create a policy and select it.
8. Click **Save Changes**.

Note

Changes will take effect from the next maintenance run.

- To change from one scheduling policy to another scheduling policy after the recommended maintenance plan is created and saved for the infrastructure, you must confirm your choice by entering the currently used policy name in a confirmation dialog before.
- Policy change deletes the scheduling plan created with the current policy for your infrastructure, and you will lose any customizations made to the Oracle recommended plan if you attach the same policy later.

View Maintenance Scheduling Plan

To view the maintenance scheduling plan for your infrastructure, use this procedure.

A maintenance scheduling plan for this infrastructure is created according to the associated scheduling policy to update infrastructure components across maintenance windows.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance scheduling plan.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Maintenance scheduling plan**.
The Maintenance scheduling Plan section displays the maintenance scheduling plan associated with the infrastructure. The details include maintenance windows associated with the plan, lifecycle state, duration in hours, scheduled actions, and the estimated time.

Note

- If all infrastructure components are not scheduled for an update in the maintenance scheduling plan, you will see a banner in the console indicating the components missing from the plan, and the plan will reflect a 'Needs Attention' lifecycle state.
- If the missing components are not included in the plan before Oracle automation schedules the next quarterly maintenance run, they will be automatically added to an 'Unplanned' maintenance window for that quarter's maintenance run.

Edit Maintenance Plan Scheduled Actions

To edit scheduled actions of an Exadata Cloud@Customer Infrastructure maintenance scheduling plan, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance scheduling plan.
The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.
5. Click **Maintenance scheduling plan**.
6. In the Maintenance scheduling plan section, click the Actions menu (three dots) of the Maintenance window you want to edit action, and select **Edit scheduled actions**.
7. In the resulting Edit scheduled actions window, do the following:
 - **Add a scheduled action:**
 - **Create new action:** When you are creating a new maintenance action, you can choose to add components already scheduled to update in different maintenance windows to a new window.
 - * **Select action type:**
 - * **DB Server Exadata full software update**
 - * **Configure maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
 - * **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
 - * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- * **Add DB Servers:**

- * **Select DB Servers:** Updates to selected DB Servers will be moved from their currently scheduled window to the window you are adding the maintenance action.

- * **Storage Server Exadata full software update**

- * **Configure maintenance method:**

- * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
- * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.

Note

All storage servers in the infrastructure must be scheduled to update in a single maintenance action to apply non-rolling storage updates. While these updates are applied, your database workloads will incur complete downtime.

- * **Select storage server from:** Select a window from where you want to add storage servers from.
- * **Select maintenance action to add from:** Select the action from where you want to add storage servers from.
- * **Select number of storage servers to add:** Select the number of storage servers to add to this action.

- **Move action from another window:** When you are moving an action, you can choose to move all components scheduled to update in a specific window to a new window.

- * **Select the window to move action from:** Choose a window from the maintenance run.

- * **Select the action to move:** Choose a specific action from the maintenance window to move.
- **Remove a scheduled action:**
 - Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

Note

- * Any action can be removed as long as there are no components scheduled for update in that action.
- * Any window can be removed as long as there are no actions scheduled for update in that window.
- **Edit a scheduled action:**
 - Click the Actions menu (three dots) of the maintenance action type **Db server full software update** and **Storage server full software update**, and then select **Edit scheduled action**.

To view or edit the properties of the next scheduled quarterly maintenance for Exadata Cloud Infrastructure

Review and change the properties of the Exadata Cloud Infrastructure scheduled quarterly maintenance.

NOT_SUPPORTED

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure you want to access:
Cloud Exadata infrastructure (new resource model): Under **Exadata at Oracle Cloud**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

The Infrastructure Details page displays information about the selected Oracle Exadata infrastructure.

Note

An information block is displayed 6 hours before the start of a maintenance run, regardless of whether you've chosen rolling or non-rolling maintenance method. When the maintenance begins, it is automatically removed.

3. On the resource details page, under **Maintenance**, click the **View** link in the **Next Quarterly Maintenance** field.
The Exadata Infrastructure Maintenance page is displayed.

Note

If you've opted to use a Scheduling Policy, link to the policy will appear. Click the link to view the Maintenance scheduling policy details. To edit a maintenance window, refer to the topics listed under *Manage Quarterly Maintenance Run created from Scheduling Plan*.

4. On the **Exadata Infrastructure Maintenance** page, scheduled maintenance details are listed.

Target DB Server Version and **Target Storage Server Version**: These fields display the Exadata software version to be applied by the scheduled maintenance. The version applied will be the most recent certified update for Exadata infrastructures in the cloud. If the next quarterly update is not yet certified when the maintenance is scheduled, then the versions may show "LATEST" until the new quarterly update becomes available. Once the update becomes available the new version will be displayed.

To find information on the Database Server Exadata software version or the Storage Server Exadata software version, see My Oracle Support note *Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 888828.1)*.

For each scheduled Exadata Infrastructure resource maintenance event, the Maintenance page lists the following details:

- The status of the event
- The OCID of the event
- The scheduled start time and date of the event
- Click **Patch Now** to start the maintenance event immediately. When prompted, click **Run Maintenance** to confirm that you want to start the event now.

If a maintenance event is already in progress on one or more of the VM Clusters hosted by an Exadata Infrastructure resource when a maintenance event on that resource is to start, the Exadata Infrastructure resource maintenance event is queued and begins immediately after all VM Cluster maintenance events complete.

5. To change the next scheduled maintenance settings, click **Edit Maintenance Run**.
6. On the **Edit Maintenance** page, do the following:
 - Select a maintenance method, **Rolling** or **Non-rolling**.

Note

If you select the **Non-rolling** option, components will be updated simultaneously, resulting in full system downtime.

- **Enable custom action before performing maintenance on DB servers**: Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
 - **Custom action timeout (in minutes)**: Maximum timeout available to perform custom action before starting maintenance on the DB Servers.

Default: 30 minutes

Minimum: 15 minutes

Maximum: 120 minutes

- To reschedule the next maintenance run, enter a date and time in the **Scheduled Start time** field.

The following restrictions apply:

- Oracle expects to be able to perform infrastructure maintenance at least once per quarter. You should not defer maintenance beyond the end of a maintenance quarter unless unexpected issues prevent your accommodating it before the next maintenance quarter.
- In the event unexpected issues prevent your accommodating the scheduled infrastructure maintenance run, you can reschedule the infrastructure maintenance to another date no more than 180 days from the prior infrastructure maintenance. Since normal maintenance should be performed quarterly, this provides approximately 90 additional days for you to reschedule the infrastructure maintenance. Oracle strongly recommends you not schedule maintenance at or close to the 180 day limit, as you will have no flexibility to reschedule further if additional unexpected issues arise.
- If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date.
- You can reschedule your maintenance to take place earlier than it is currently scheduled.
- Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.

- Click **Save Changes**.

7. To view estimated maintenance time details for various components, click the **View** link is displayed in the Total Estimated Maintenance Time field. The **View** link is displayed in the **Total Estimated Maintenance Time** field only if the Maintenance Method is **Rolling**.

The **Estimated Maintenance Time Details** page is displayed with details that include:

- Total Estimated Maintenance Time
 - Database Servers Estimated Maintenance Time
 - Storage Servers Estimated Maintenance Time
 - Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed
 - a. To view the number of VMs that will be restarted as part of Database Server maintenance, click the **Show details** link. The **VM Location** dialog is displayed.
 - b. In the **VM Cluster Name** field, you can find out what VM cluster a particular VM belongs to.
 - c. Click **Close**.
8. Click **Close** to close the **Estimated Maintenance Time Details** page.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

- [Exadata Database Machine and Exadata Storage Server Supported Versions \(Doc ID 888828.1\)](#)

View and Edit Maintenance While Maintenance is In Progress

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

5. Click **Maintenance**.

The **Exadata Infrastructure Maintenance** page is displayed.

6. Click **Edit Maintenance Run**.

Edit Maintenance page is displayed.

Note

You can only make edits to the custom action configuration, not the maintenance method or scheduled start time. Enabling or disabling the custom action or modifying the custom action timeout while maintenance is in progress will apply to all database servers that have yet to be updated.

7. On the **Edit Maintenance** page, do the following:
 - **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
 - **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.
Default: 30 minutes
Maximum: 120 minutes
8. Click **Save Changes**.

If you have configured the rolling maintenance method, then the **View** link is displayed in the **Total Estimated Maintenance Time** field.

 - a. Click **View**.
Estimated Maintenance Time Details page is displayed with details that include:

- Total Estimated Maintenance Time
- Database Servers Estimated Maintenance Time
- Storage Servers Estimated Maintenance Time
- Network Switches Estimated Maintenance Time (InfiniBand systems only)
- Order in which components are updated. In rolling maintenance, components are updated in the sequence displayed.

b. Click **Close**.

View and Edit Maintenance While Maintenance is Waiting for Custom Action

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit.

The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

5. Click **Maintenance**.

Exadata Infrastructure Maintenance page is displayed.

Note

- Editing a maintenance run is not available while waiting for custom action.
- While maintenance is waiting for custom action, an information block is displayed. The information block is removed after the maintenance resumes.

6. On the information block, do one of the following:
 - a. Click **Resume Maintenance Now** to resume the maintenance, proceeding to the next database server.
Resume Maintenance dialog is displayed. Click **Resume Maintenance Now**.
 - b. Click **Extend Custom Action Timeout**.
You can extend timeout multiple times within the maximum allowable time of 2 hours. If you try extending beyond the maximum limit, then the system displays the Cannot Extend Custom Action Timeout dialog indicating that the custom action timeout has already been extended to the maximum allowable 2 hours and you cannot extend it further.

To view or edit a scheduled security maintenance

Learn how to view and edit the next scheduled security maintenance.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to view maintenance details.
The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.
Quarterly and monthly maintenance details are displayed under the **Maintenance** section. Monthly security maintenance is labeled as **Next Security Maintenance**.
5. On the **Infrastructure Details** page, click **Maintenance**.
The Exadata Infrastructure Maintenance page is displayed. The Exadata Infrastructure Maintenance page includes details such as **Type: Monthly Security Maintenance**, **Scheduled Start Time**, and so on.
6. To reschedule monthly security maintenance, click **Edit Maintenance Run** and pick a new date within the 21-day cycle.

Note

Certain black-out dates are not available for security maintenance and are grayed out in the rescheduling calendar.

7. Select a date and then click **Save Changes**.
8. To view the maintenance history, click **Maintenance History**.
The Maintenance History page displays details including the type of maintenance, **Monthly** or **Quarterly**.
When a monthly security maintenance is in progress, the Infrastructure resource state will be **Available**.

To view the maintenance history of an Exadata Cloud Infrastructure resource

This task describes how to view the maintenance history for a cloud Exadata infrastructure resource.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Navigate to the Cloud Exadata infrastructure resource you want to access.
Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
3. On the Infrastructure Details page, click **Maintenance History** to see a list of past maintenance events including details on their completion state and the target database and storage server versions.
The Maintenance jobs, State and type of patching are displayed.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)
- [My Oracle Support note Exadata Database Machine and Exadata Storage Server Supported Versions \(Doc ID 888828.1\)](#)

To set the node patching order for a scheduled infrastructure maintenance run

This task describes how to set the node patching order for a scheduled infrastructure maintenance run for a cloud Exadata infrastructure resource.

Note

By default, all scheduled maintenance runs are initially set to use rolling patching. To use non-rolling patching, you must change this setting for each maintenance run after it is scheduled.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Exadata on Oracle Public Cloud**.
2. Navigate to the cloud Exadata infrastructure you want to access:
 - *Cloud Exadata infrastructure (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata Infrastructure**. In the list of infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.
3. On the resource details page, click **Maintenance**.
4. On the Maintenance page, click **Edit maintenance run**.
5. Change the maintenance method to either **Rolling** or **Non-rolling** as needed.
6. Click **Save**.

Manage Quarterly Maintenance Run Created From Scheduling Plan

- [View Maintenance Windows Associated with a Maintenance Run](#)
- [Edit Maintenance Window Associated with a Maintenance Run](#)
- [View Maintenance Actions Associated with a Maintenance Run](#)
- [Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run](#)
- [View and Edit Maintenance While Maintenance is In Progress](#)
- [View and Edit Maintenance While Maintenance is Waiting for Custom Action](#)
- [Cancel a Maintenance Run While In Progress](#)
- [View Maintenance Activity in a Compartment](#)
- [View the Details of a Maintenance Activity](#)
- [Review and Respond to Unplanned Maintenance Activity](#)

View Maintenance Windows Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.

3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance updates scheduled to run in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
The Maintenance Windows section in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
 - **Start time (UTC):** Window start time in UTC For example, Sun, Jun 23, 2024, 18:30:58 UTC.
The following restrictions apply:
 - Oracle expects to be able to perform infrastructure maintenance at least once per quarter. You should not defer maintenance beyond the end of a maintenance quarter unless unexpected issues prevent your accommodating it before the next maintenance quarter.
 - In the event unexpected issues prevent your accommodating the scheduled infrastructure maintenance run, you can reschedule the infrastructure maintenance to another date no more than 180 days from the prior infrastructure maintenance. Since normal maintenance should be performed quarterly, this provides approximately 90 additional days for you to reschedule the infrastructure maintenance. Oracle strongly recommends you not schedule maintenance at or close to the 180 day limit, as you will have no flexibility to reschedule further if additional unexpected issues arise.
 - If a new maintenance release is announced prior to your rescheduled maintenance run, the newer release will be applied on your specified date.
 - You can reschedule your maintenance to take place earlier than it is currently scheduled.
 - Oracle reserves certain dates each quarter for internal maintenance operations, and you cannot schedule your maintenance on these dates.
 - **Type:** Planned vs Unplanned. All windows created from the infrastructure maintenance scheduling plan or added by you to this maintenance run are 'Planned' windows. All other windows Oracle automation creates to address failures, duration enforcement, or unforeseen events are defined as 'Unplanned' windows. Always review activities scheduled to run in an 'Unplanned' window.
 - **Maintenance action:** The summary of actions scheduled to update in a given window. The server name identifies updates scheduled for DB servers. For example, Apply full update to DB servers dbServer-1 and dbServer-2. The storage server updates are identified as count since all storage servers have identical storage layouts. For example, Apply full update to 2 Storage Servers.
 - **Estimated time:** The estimated time for Oracle automation to complete maintenance actions scheduled to apply updates to all infrastructure components across all windows in the maintenance run.

Edit Maintenance Window Associated with a Maintenance Run

Note

You can update the window configuration, like window schedule start time, duration, and duration enforcement, while the window is still in the 'Scheduled' life cycle state. Once the window is in progress, you cannot make changes to the configuration. You can choose to cancel a running maintenance for a window. Details covered in the *Cancel Maintenance Window Associated with a Maintenance Run* section.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
The Maintenance Windows section in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window you want to edit, and then select **Edit maintenance window**.
7. In the resulting Edit maintenance window dialogs, update the **Maintenance window start time**, **Duration in hours**, and **Enforce window duration** fields
8. Click **Edit maintenance window**.

View Maintenance Actions Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance actions.
6. Click **Maintenance actions**.
7. To add actions:
 - a. Click **Add actions**.
 - b. In the resulting Add maintenance actions window, **Select action type**, and then click **Add maintenance action**.
8. To delete an action:
 - a. Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

Edit Maintenance Actions of a Maintenance Window Associated with a Maintenance Run

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
The Maintenance Windows section in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window you want to edit, and then select **Edit maintenance actions**. The resulting Edit maintenance action page displays the list of actions. You can either add more actions or delete the existing ones.
7. To add actions:
 - a. Click **Add actions**.
 - b. Do the following in the resulting Add maintenance action window:
 - **Create new action:** When you are creating a new maintenance action, you can choose to add components already scheduled to update in different maintenance windows to a new window.
 - **Select action type:**
 - * **DB Server Exadata full software update**
 - * **Configure maintenance method:**
 - * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
 - * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.
 - * **Enable custom action before performing maintenance on DB servers:** Enable custom action only if you want to perform additional actions outside of Oracle's purview. For maintenance configured with a rolling software update, enabling this option will force the maintenance run to wait for a custom action with a configured timeout before starting maintenance on each DB server. For maintenance configured with non-rolling software updates, the maintenance run will wait for a custom action with a configured timeout before starting maintenance across all DB servers. The maintenance run, while waiting for the custom action, may also be resumed prior to the timeout.
 - * **Custom action timeout (in minutes):** Timeout available to perform custom action before starting maintenance on the DB Servers.

Note

Custom action timeout applies only to DB servers. Customer can specify a minimum 15 minutes and a maximum of 120 minutes of custom action time-out before DB server patching starts. Within this time, they can perform whatever actions they have planned. In case, they want to extend the custom action, they can extend the same by going to "edit maintenance window" option. If custom action is in progress, customer get 2 options - either extend Custom action timeout or resume maintenance window.

Default: 15 minutes

Maximum: 120 minutes

- * **Add DB Servers:**

- * **Select DB Servers:** Updates to selected DB Servers will be moved from their currently scheduled window to the window you are adding the maintenance action.

- * **Storage Server Exadata full software update**

- * **Configure maintenance method:**

- * **Rolling:** By default, Exadata Infrastructure is updated in a rolling fashion, one server at a time with no downtime.
- * **Non-rolling:** Update database and storage servers at the same time. The non-rolling maintenance method minimizes maintenance time but incurs full system downtime.

Note

All storage servers in the infrastructure must be scheduled to update in a single maintenance action to apply non-rolling storage updates. While these updates are applied, your database workloads will incur complete downtime.

- * **Select storage server from:** Select a window from where you want to add storage servers from.
 - * **Select maintenance action to add from:** Select the action from where you want to add storage servers from.
 - * **Select number of storage servers to add:** Select the number of storage servers to add to this action.
- **Move action from another window:** When you are moving an action, you can choose to move all components scheduled to update in a specific window to a new window.
 - **Select the window to move action from:** Choose a window from the maintenance run.

- **Select the action to move:** Choose a specific action from the maintenance window to move.
8. To delete an action:
 - a. Click the Actions menu (three dots) of the maintenance action, and then select **Remove**.

Note

- Any action can be removed as long as there are no components scheduled for update in that action.
- Any window can be removed as long as there are no actions scheduled for update in that window.

View and Edit Maintenance While Maintenance is In Progress

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit. The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

Note

Maintenance In Progress status is displayed in the **Next Quarterly Maintenance** field.

5. Click the **View** link in the **Next Quarterly Maintenance** field. You will be on the maintenance run details page that is in progress.
6. Click the Actions menu (three dots) and select **View details**.
7. Click **Maintenance windows** on the Maintenance Run details page.
8. Identify the maintenance window that's in progress
9. Click the Actions menu (three dots) and select **Edit maintenance actions**.
10. Click the Actions menu (three dots) and select **Edit custom action configuration**.
11. In the resulting **Edit custom action configuration** page, enter **Custom action in minutes**.

Note

While maintenance is in progress you can only change the custom action time for DB Server action type. You cannot change the custom action time for support this option for any other action type.

12. Click **Save changes**.

View and Edit Maintenance While Maintenance is Waiting for Custom Action

While maintenance is in progress, you can enable or disable custom action and change the custom action timeout. While maintenance is waiting for a custom action, you can resume the maintenance prior to the timeout or extend the timeout.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select **Region** and **Compartment**, and provide the region and the compartment where the Oracle Exadata infrastructure you want to edit is located.
3. Click **Exadata Infrastructure**.
4. Click the name of the Exadata infrastructure that you want to edit. The **Infrastructure Details** page displays information about the selected Oracle Exadata infrastructure.

Note

Maintenance In Progress status is displayed in the **Next Quarterly Maintenance** field.

5. Click the **View** link in the **Next Quarterly Maintenance** field. You will be on the maintenance run details page that is in progress.
6. In the resulting **Maintenance Run** details page, click **Maintenance actions**. While maintenance is waiting for custom action, an information block is displayed. The information block is removed after the maintenance resumes.
7. On the information block, do the following:
 - a. Click **Resume Maintenance Now** to resume the maintenance, proceeding to the next database server. Resume Maintenance dialog is displayed. Click **Resume Maintenance Now**.
 - b. Click **Extend Custom Action Timeout**. You can extend timeout multiple times within the maximum allowable time of 2 hours. If you try extending beyond the maximum limit, then the system displays the Cannot Extend Custom Action Timeout dialog indicating that the custom action timeout has already been extended to the maximum allowable 2 hours and you cannot extend it further.

Cancel a Maintenance Run While In Progress

To cancel a Maintenance Window Associated with a Maintenance Run, follow these steps:

Note

You can cancel a running maintenance while the scheduled updates for a window are in progress. Canceling the maintenance while the updates are in progress allows you to reschedule all actions that have not yet started to a future maintenance window of your choice. You can choose a new start time and duration to finish all the actions rescheduled from the maintenance window you decided to cancel.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the activity you want to view associated maintenance windows.
The Maintenance Windows section in the resulting Maintenance run page lists the maintenance windows associated with the chosen activity.
6. Click the Actions menu (three dots) of the Maintenance window you want to cancel, and then select **Cancel maintenance window**.

To cancel a Maintenance Run While In Progress, follow these steps:

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Maintenance**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.
5. Click the name of the **'In Progress'** activity you want to cancel.
6. Click **Maintenance windows** on the Maintenance Run details page.
7. Identify the maintenance window that's in progress.
8. Click the Actions menu (three dots) and select **Cancel maintenance window**.
9. On the resulting Cancel maintenance run window, reconfigure **Maintenance window start time**.
10. Select the **Enforce window duration** checkbox to pause and re-schedule any scheduled action that goes over the configured window duration to resume in a future maintenance window.
11. Click **Reschedule maintenance run**.

Note

The maintenance run will complete the current operation. All remaining actions scheduled for this window will be rescheduled to a new maintenance window.

View Maintenance Activity in a Compartment

Note

Maintenance activity lists all the maintenance updates scheduled to run for all infrastructure resources in a given compartment for the selected Exadata cloud service.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance activities in the chosen compartment.

View the Details of a Maintenance Activity

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Scheduling policy**.
The resulting Scheduling Policy page displays the list of policies.
3. Choose a compartment from the **Compartment** filter.
4. Under **Maintenance**, click **Activity**.
The resulting Activity page lists maintenance updates scheduled to run in the chosen compartment.
5. Click the Actions menu (three dots) of the maintenance activity you want to view details.
The resulting Maintenance run page displays the details of the chosen maintenance activity.

Review and Respond to Unplanned Maintenance Activity

Unplanned maintenance activity when infrastructure is scaled after maintenance is planned

- After scaling your infrastructure by adding DB or storage servers, you may need to update your maintenance scheduling plan to include these new components. If any infrastructure component is missing from the maintenance plan, a warning will appear in the infrastructure maintenance plan details section.
- When Oracle automation creates the maintenance run for the quarter, any components not included in the maintenance plan will be automatically added to an 'Unplanned' maintenance window. This ensures that all components have the correct system software applied each quarter, maintaining OCI software compliance.
- You can edit the scheduled start time of the 'Unplanned' window or move the updates for the missing components to an existing planned window as needed.

Unplanned maintenance activity when a scheduled update fails to apply

- If a scheduled update fails, the Oracle operations team will engage, evaluate the failure, and reschedule the failed update along with any unfinished updates to a future maintenance window. Oracle automation will mark this rescheduled window as 'Unplanned' and notify you to review the rescheduled maintenance activity.
- You can edit the 'Unplanned' window's scheduled start time or move the failed and unfinished updates to an existing planned window as needed.

Unplanned maintenance when schedule activity exceeds enforced window duration

- For maintenance windows configured with duration enforcement, Oracle automation will check if the estimated time to execute and apply the scheduled update is sufficient within the remaining window duration. If not, Oracle automation will automatically reschedule all unfinished updates to a future 'Unplanned' window, mark the current window as 'Duration Exceeded,' and notify you to review the rescheduled maintenance activity.
- Any updates already in progress will continue past the enforced window duration to ensure a consistent state of the underlying resources.

Monitor Infrastructure Maintenance Using Lifecycle State Information

The lifecycle state of your Exadata Infrastructure resource enables you to monitor when the maintenance of your infrastructure resource begins and ends.

In the Oracle Cloud Infrastructure Console, you can see lifecycle state details messages on the **Exadata Infrastructure Details** page when a tooltip is displayed beside the **Status** field. You can also access these messages using the `ListCloudExadataInfrastructures` API, and using tools based on the API, including SDKs and the `OCI CLI`.

During infrastructure maintenance operations, you can expect the following:

- If you specify a maintenance window, then patching begins at your specified start time. The infrastructure resource's lifecycle state changes from **Available** to **Maintenance in Progress**.

Note

The prechecks are now done prior to the start of the maintenance.

- When Exadata database server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (dbnodes) is being updated**.
- When storage server maintenance starts, the infrastructure resource's lifecycle state is **Maintenance in Progress**, and the associated lifecycle state message is, **The underlying infrastructure of this system (cell storage) is being updated and this will not impact Database availability**.
- After storage server maintenance is complete, the networking switches are updated one at a time, in a rolling fashion (InfiniBand systems only).
- When maintenance is complete, the infrastructure resource's lifecycle state is **Available**, and the Console and API-based tools do not provide a lifecycle state message.

Related Topics

- [ListCloudExadataInfrastructures](#)
- [Software Development Kits and Command Line Interface](#)
- [Command Line Interface \(CLI\)](#)

Receive Notifications about Your Infrastructure Maintenance Updates

There are two ways to receive notifications. One is through email to infrastructure maintenance contacts and the other one is to subscribe to the maintenance events and get notified.

Oracle schedules maintenance run of your infrastructure based on your scheduling preferences and sends email notifications to all your infrastructure maintenance contacts. You can login to the console and view details of the schedule maintenance run. Appropriate maintenance related events will be generated as Oracle prepares for your scheduled maintenance run, for example, schedule reminder, patching started, patching end, and so on. For more information about all maintenance related events, see *Oracle Cloud Exadata Infrastructure Events*. In case, if there are any failures, then Oracle reschedules your maintenance run, generates related notification, and notifies your infrastructure maintenance contacts.

For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. To receive additional notifications other than the ones sent to infrastructure maintenance contacts, you can subscribe to infrastructure maintenance events and get notified using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- [Oracle Exadata Database Service on Dedicated Infrastructure Events](#)
Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- [Overview of Events](#)
- [Notifications Overview](#)
- [Managing Infrastructure Maintenance Contacts](#)
Learn to manage your Exadata infrastructure maintenance contacts.

Managing Infrastructure Maintenance Contacts

Learn to manage your Exadata infrastructure maintenance contacts.

- [To manage maintenance contacts in an Exadata Cloud Infrastructure](#)
Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To manage maintenance contacts in an Exadata Cloud Infrastructure

Manage contacts for Exadata infrastructure maintenance notifications using the Console.

To prevent an Exadata infrastructure administrator from being overwhelmed by system update notifications, you can specify up to 10 email addresses of people to whom maintenance notifications are sent.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the **Oracle Exadata Database Service on Dedicated Infrastructure** section, click **Exadata Infrastructure** to display a list of Exadata infrastructures in the default compartment. You can select a different compartment from the **Compartment** drop-down located in the **List Scope** section.
3. In the list of Exadata infrastructure resources, find the infrastructure you want to access and click its highlighted name to view its details page.

4. In the **Maintenance** section, click **Manage** in the **Customer Contacts** field to display the Manage Contacts dialog.
5. Click the **Add Contacts** button to display a field in which to enter a valid email address. You can have up to 10 maintenance contacts for each Exadata infrastructure.
6. To edit an email address, in the Manage Contacts dialog, select the box preceding the email address you want to edit and click the **Edit** button.
7. To remove an email address from the list, in the Manage Contacts dialog, select the box preceding the email address you want to remove and click the **Remove** button.

Using the API to Manage Exadata Cloud Infrastructure Maintenance Controls

Use these API operations to manage Exadata Cloud Infrastructure maintenance controls and resources.

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

Use these API operations to manage Exadata Cloud Infrastructure maintenance controls.

Cloud Exadata infrastructure resource (new resource model):

- [ListCloudExadataInfrastructures](#)
- [GetCloudExadataInfrastructure](#)
- [UpdateCloudExadataInfrastructure](#)
- [GetMaintenanceRun](#)
- [ListMaintenanceRuns](#)
- [UpdateMaintenanceRun](#)
- [ListSchedulingPolicies](#)
- [CreateSchedulingPolicy](#)
- [GetSchedulingPolicy](#)
- [UpdateSchedulingPolicy](#)
- [DeleteSchedulingPolicy](#)
- [ChangeSchedulingPolicyCompartment](#)
- [ListRecommendedScheduledActions](#)
- [ListSchedulingWindows](#)
- [CreateSchedulingWindow](#)
- [GetSchedulingWindow](#)
- [UpdateSchedulingWindow](#)
- [DeleteSchedulingWindow](#)
- [ListSchedulingPlans](#)
- [CreateSchedulingPlan](#)
- [GetSchedulingPlan](#)

- [DeleteSchedulingPlan](#)
- [ChangeSchedulingPlanCompartment](#)
- [ReorderScheduledActions](#)
- [ListScheduledActions](#)
- [CreateScheduledAction](#)
- [GetScheduledAction](#)
- [UpdateScheduledAction](#)
- [DeleteScheduledAction](#)
- [ListParamsForActionType](#)
- [ListExecutionWindows](#)
- [CreateExecutionWindow](#)
- [GetExecutionWindow](#)
- [UpdateExecutionWindow](#)
- [DeleteExecutionWindow](#)
- [ReorderExecutionActions](#)
- [ListExecutionActions](#)
- [CreateExecutionAction](#)
- [GetExecutionAction](#)
- [UpdateExecutionAction](#)
- [DeleteExecutionAction](#)
- [MoveExecutionActionMember](#)

Manage VM Clusters

Learn how to manage your VM clusters on Exadata Cloud Infrastructure.

- [Introduction to Scale Up or Scale Down Operations](#)
With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.
- [Overview of VM Cluster Node Subsetting](#)
VM Cluster Node Subsetting enables you to allocate a subset of database servers to new VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.
- [Overview of Oracle Cloud Infrastructure Zero Trust Packet Routing \(ZPR\)](#)
Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) safeguards sensitive data from unauthorized access using intent-based security policies that you define for OCI resources assigned with security attributes.
- [Private Service Access for OCI Service Access from VCN](#)
You can use Private Service Access (PSA) instead of a Service Gateway to provide access to the OCI services required to support your database service.
- [About Application VIP](#)
Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.

- [Using the Console to Manage VM Clusters on Exadata Cloud Infrastructure](#)
Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Dedicated Infrastructure.
- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- [Incident Logs and Trace Files](#)
This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- [Health Metrics](#)
Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- [Using the API to Manage Exadata Cloud Infrastructure Instance](#)
- [Troubleshooting Virtual Machines Using Console Connections](#)
You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

Related Topics

- [Application Checklist for Continuous Service for MAA Solutions](#)

Introduction to Scale Up or Scale Down Operations

With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

- [Scale VM Resources in Multi VM Enabled Infrastructure](#)
Increase or decrease the OCPUs (ECPUs for X11M), memory, storage, or local disk size (/u02) storage available to a VM cluster
- [Resizing Memory and Large Pages](#)
You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect.
- [Calculating the ASM Storage](#)
- [Estimating How Much Local Storage You Can Provision on Your VMs](#)
- [Scaling Local Storage](#)

Scale VM Resources in Multi VM Enabled Infrastructure

Increase or decrease the OCPUs (ECPUs for X11M), memory, storage, or local disk size (/u02) storage available to a VM cluster

Note

Oracle doesn't stop billing when a VM or VM Cluster is stopped. To stop billing for a VM Cluster, lower the OCPU (ECPUs for X11M) count to zero.

Scaling up or down of these resources requires thorough auditing of existing usage and capacity management by the customer's DB administrator. Review the existing usage to avoid failures during or after a scale down operation. While scaling up, consider how much of these

resources are left for the next VM cluster you are planning to create. Oracle Exadata Database Service on Dedicated Infrastructure tooling calculates the current usage of memory, local disk, and ASM storage in the VM cluster, adds headroom to it, and arrives at a "minimum" value below which you cannot scale down, and expects that you specify the value above this minimum value.

Note

- When creating or scaling a VM Cluster, setting the number of OCPUs (ECPUs for X11M) to zero will shut down the VM Cluster and eliminate any billing for that VM Cluster, but the hypervisor will still reserve the minimum 2 OCPUs (8 ECPUs for X11M) for each VM. These reserved OCPUs (ECPUs for X11M) cannot be allocated to any other VMs, even though the VM to which they are allocated is shut down. The Control Plane does not account for reserved OCPUs (ECPUs for X11M) when showing maximum available OCPUs (ECPUs for X11M), so you should account for these reserved OCPUs (ECPUs for X11M) when performing any subsequent scaling operations to ensure the operation can acquire enough OCPUs (ECPUs for X11M) to successfully complete the operation.
- For memory and /u02 scale up or scale down operations, if the difference between the current value and the new value is less than 2%, then no change will be made to that VM. This is because memory change involves rebooting the VM, and /u02 change involves bringing down the Oracle Grid Infrastructure stack and unmounting /u02. Production customers will not resize for such a small increase or decrease, and hence, such requests are a no-op.
- You can scale down the following resources in any combinations:
 - OCPU (ECPU for X11M)
 - Memory
 - Local storage
 - Exadata storage

Each scaling operation can take several minutes to complete. The time for each operation will vary based on activity in the system, but as a general rule, most operations should complete within 15 minutes for a quarter rack, 20 minutes for a half rack, and 30 minutes for a full or larger rack. Performing multiple OCPU (ECPU for X11M) scaling operations over a short period of time can lengthen the time for completion. Although online, OCPU (ECPU for X11M) scaling is not implemented on all VMs in parallel, so as to detect and protect from any anomalies before they affect the entire system. Memory and Local Storage scaling require a VM reboot, and are performed one VM at a time in a rolling manner.

If you run multiple scale-down operations, then each operation is performed serially. For example, if you scale memory and local storage from the Console, then the system will first scale memory, and when that operation completes, it will scale storage. The time to complete all operations will be the sum of the time to complete individual operations.

- Storage servers added to the infrastructure but yet to run the 'Add Capacity' step will not have any disk groups created on them.

Resizing Memory and Large Pages

You can scale the database server memory up and down in a VM Cluster. Scaling memory requires a rolling restart of the database servers to take effect.

Changing the memory in a VM Cluster will affect the large pages (HugePages) settings for the VMs in that cluster. When a VM is initially created, each VM's operating system is configured with 50% of the memory allocated to the VM for large pages, and databases are configured to use that memory for their SGA. Oracle recommends that you not modify the large pages configuration unless you understand the implication of any changes you make. Improper configurations can prevent all databases from starting, and even prevent the VM from starting up.

Although not recommended, you are allowed to modify the large pages configuration. Any changes you make may be modified by automation should you subsequently resize the memory available to the VM. After a memory resize operation, the cloud automation will attempt to maintain the same amount of large pages memory as a percentage of the total memory, with a cap of 60%. If you configure large pages to be greater than 60% of total memory, then the cloud automation will resize it to 60% of total memory. This automatic resize is to ensure sufficient conventional memory for the virtual machine to start. The automation will perform a precheck to determine the actual large pages memory in use by the running database instances, and ensure after the resize that there is enough large pages memory available to support those same databases. If there will not be sufficient memory available after the resize, then the precheck will fail and the resize will not continue.

Calculating the ASM Storage

Each ASM VM Cluster will have its own pool of Exadata Storage. That pool will be split between the following disk groups, with each disk group getting the percent specified in the table. The ratios vary depending on whether the VM Cluster is configured to support local backups, or for Exadata snapshots. The disk groups sizes are fixed and cannot be changed post deployment.

Table 5-1 Disk Group sizes

Disk Group	Local backups = No Snapshots =No	Local backups = Yes Snapshots = No	Local backups = No Snapshots= Yes	Local backups = Yes Snapshots = Yes
+DATA	80%	40%	60%	35%
+RECO	20%	60%	20%	50%
+SPARSE	0%	0%	20%	15%

Use the following formula to calculate the minimum required ASM storage:

- For each disk group, for example, DATA, RECO, note the total size and free size by running the `asmcmd lsdg` command on any Guest VM of the VM cluster.
- Calculate the used size as (Total size - Free size) / 3 for each disk group. The /3 is used because the disk groups are triple mirrored.
- DATA:RECO ratio is:
 - 80:20 if **Local Backups** option was NOT selected in the user interface.
 - 40:60 if **Local Backups** option was selected in the user interface.

- Ensure that the new total size as given in the user interface passes the following conditions:
Used size for DATA * 1.15 <= (New Total size * DATA %)
Used size for RECO * 1.15 <= (New Total size * RECO %)

Example 5-2 Calculating the ASM Storage

1. Run the `asmcmd lsdg` command in the Guest VM:

- Without SPARSE:

```
/u01/app/19.0.0.0/grid/bin/asmcmd lsdg
ASMCMD>
State   Type Rebal Sector Logical_Sector Block AU      Total_MB
Free_MB Req_mir_free_MB Usable_file_MB Offline_disks
Voting_files Name
MOUNTED HIGH N      512    512      4096 4194304 12591936
10426224 1399104      3009040      0
Y      DATA5/
MOUNTED HIGH N      512    512      4096 4194304 3135456
3036336  348384      895984      0
N      RECO5/
ASMCMD>
```

- With SPARSE:

```
/u01/app/19.0.0.0/grid/bin/asmcmd lsdg
ASMCMD>
State   Type Rebal Sector Logical_Sector Block AU      Total_MB
Free_MB Req_mir_free_MB Usable_file_MB Offline_disks
Voting_files Name
MOUNTED HIGH N      512    512      4096 4194304 12591936
10426224 1399104      3009040      0
Y      DATA5/
MOUNTED HIGH N      512    512      4096 4194304 3135456
3036336  348384      895984      0
N      RECO5/
MOUNTED HIGH N      512    512      4096 4194304 31354560
31354500 3483840      8959840      0
N      SPRC5/
ASMCMD>
```

Note

The listed values of all attributes for SPARSE diskgroup (SPRC5) present the virtual size. In Exadata Cloud Infrastructure, we use the ratio of 1:10 for `physicalSize:virtualSize`. Hence, for all purposes of our calculation we must use 1/10th of the values displayed above in case of SPARSE for those attributes.

2. Used size for a disk group = $(\text{Total_MB} - \text{Free_MB}) / 3$

- Without SPARSE:
Used size for DATA5 = $(12591936 - 10426224) / 3 = 704.98$ GB
Used size for RECO5 = $(3135456 - 3036336) / 3 = 32.26$ GB

- With SPARSE:
Used size for DATA5 = $(12591936 - 10426224) / 3 \approx 704.98$ GB
Used size for RECO5 = $(3135456 - 3036336) / 3 \approx 32.26$ GB
Used size for SPC5 = $(1/10 * (31354560 - 31354500)) / 3 \approx 0$ GB
3. Storage distribution among diskgroups
- Without SPARSE:
DATA:RECO ratio is 80:20 in this example.
 - With SPARSE:
DATA RECO: SPARSE ratio is 60:20:20 in this example.
4. New requested size should pass the following conditions:
- Without SPARSE: (For example, 5 TB in user interface.)
5 TB = 5120 GB ; $5120 * .8 = 4096$ GB; $5120 * .2 = 1024$ GB
For DATA: $(704.98 * 1.15) \leq 4096$ GB
For RECO: $(32.36 * 1.15) \leq 1024$ GB
 - With SPARSE: (For example, 8 TB in the user interface.)
8 TB = 8192 GB; $8192 * .6 = 4915$ GB; $8192 * .2 = 1638$ GB; $8192 * .2 = 1638$ GB
For DATA: $(704.98 * 1.15) \leq 4915$ GB
For RECO: $(32.36 * 1.15) \leq 1638$ GB
For SPARSE: $(0 * 1.15) \leq 1638$ GB

Above resize will go through. If above conditions are not met by the new size, then resize will fail the precheck.

Estimating How Much Local Storage You Can Provision on Your VMs

Note

The following does not apply to X6, X7, X8, and Base Systems as they do not support multiple VMs. The Base System has 200 GB available for `/u02`.

VM Images include the files necessary to boot and run the VM and its operating system, as well as space for Oracle Homes stored in `/u02`. To estimate how much additional local storage space beyond the minimum can be allocated to any file system associated with a VM, subtract the size of the VM images for all VMs on a server from the total available space. If you have not modified the default VM Image size by expanding any file systems, use the VM Image size (default and minimum) below. If you have or plan to modify your VM Image size, you must use the OCI console and "Scale VM Cluster" action to check the allocated and available for an existing VM Cluster as expanding some non-`/u02` file systems will consume more incremental storage than was added to the file system. This information is also available in the "Configure VM Cluster" action while creating a new VM Cluster.

X8M-2 Systems

- Total space available for VM images (X8M): 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

X9M-2 Systems

- Total Available for VM Images: 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

X11 Systems

- Total Available for VM Images: 2243 GB
- VM Image size (default and minimum) including `/u02`: 244 GB
- Default (minimum) `/u02`: 60 GB

Example: If you have an X9M Elastic System with 2 VMs per physical server, and have not made any changes to any of the file systems, you will have 2243 GB available for all VMs, and each will consume 244 GB (488 total), leaving 1755 GB to expand any VM Images. The default VM image will include 60GB of `/u02` per VM to store Oracle homes. The 1755 GB of available space can be used to expand `/u02`, or can be used to expand other file systems in the VM Image. Every GB used to expand `/u02` will consume a GB of available space. Every GB used to expand other file systems in the VM image may consume more than a GB of space. Refer to the information in the console when expanding non-`/u02` file systems to see the actual available space impact of expanding these file systems.

Scaling Local Storage

Scale Local Space Operation Guidelines

You can scale local storage by modifying the size of many of the individual file systems in a VM. By default, the file systems are created at their minimum size. You can increase the size of the file systems as required. However, note that you can only shrink `/u02`. Other file systems can only be increased in size. The maximum supported size of any file system is 900 GB.

The storage consumed by all file systems is greater than the sum of the file system sizes. Refer to the calculations displayed in the OCI console to see the effects on free local storage when resizing a file system.

Using the OCI Console or API, you can increase or decrease the size of the following local file systems:

- `/u02`

Using the OCI Console or API, you can increase the size of following local file systems:

- `/`
- `/u01`
- `/tmp`
- `/var`
- `/var/log`
- `/var/log/audit`
- `/home`

However, you cannot resize the following local file systems:

- `/crashfiles`
- `/boot`

- `/acfs01`
- `/u01/app/19.0.0.0/grid`

Note

- With the exception of `/u02`, you can only expand the file systems and cannot reduce their size once they have been expanded.
- For X8M and later models, scaling up or scaling down the `/u02` file system, as well as expanding other Guest VM file systems, are online operations that do not require a rolling restart of the Guest VM. However, for X8 systems that use the Xen-based hypervisor, file system resize operations still require a rolling restart of the Guest VM.
- Each file system can only be expanded to a maximum of 900 GB
- Ability to increase the size of additional local file systems is only supported on X8M and later systems.

For more information about resizing these file systems, see *Estimating How Much Local Storage You Can Provision to Your VMs*.

Resource Limit Based On Current Utilization

- Any scale-down operation must leave 15% buffer on top of highest local space utilization across all nodes in the cluster.
- The lowest local space per node allowed is higher of the above two limits.
- Run the `df -kh` command on each node to find out the node with the highest local storage.
- You can also use the utility like `cssh` to issue the same command from all hosts in a cluster by typing it just once.
- Lowest value of local storage each node can be scaled down to would be = 1.15x (highest value of local space used among all nodes).

ACFS File Systems

If requested by support, you can also resize the `/acfs01` file system. This file system is used by the system to stage software. It uses Exadata storage and is not subject to the limits described above for `/u02`. It is a shared file system visible from all nodes in the cluster, and can be online resized from the command line of any VM.

- **Default size:** The default size of `/acfs01` is 100 GB.
- **Scaling `/acfs01`:** You can scale `acfs01` as user `grid` from any VM via the `/sbin/acfsutil` command. No reboot is required. The resize operation will not affect the availability of the database service running in the VM cluster. The following command issued by the `grid` user will increase the size of `/acfs01` by 100 GB: `/sbin/acfsutil size +100 GB /acfs01`.
- You can create additional ACFS file systems if required. These will also consume storage from the Exadata Storage diskgroups and may be shared across all VMs in the cluster. Refer to the ACFS documentation for more information.

Overview of VM Cluster Node Subsetting

VM Cluster Node Subsetting enables you to allocate a subset of database servers to new VM clusters to enable maximum flexibility in the allocation of compute (CPU, memory, local storage) resources.

With VM Cluster Node Subsetting, you can:

- Create a smaller VM cluster to host databases that have low resource and scalability requirements or to host a smaller number of databases that require isolation from the rest of the workload.
- Expand or shrink an existing VM cluster by adding and removing nodes to ensure optimal utilization of available resources.

Consider reviewing the points below that will assist you in subsetting VM cluster nodes.

- VM Cluster Node Subsetting capability is available for new VM clusters in Oracle Exadata Database Service on Dedicated Infrastructure service.
- All VMs across a VM cluster will have the same resource allocation per VM irrespective of whether the VM was created during cluster provisioning or added later by extending an existing VM cluster.
- VM Clusters only need a minimum of 1 VM with node subsetting. However, Oracle recommends a minimum of 2 VMs per VM Cluster to provide high availability.
- You can host a maximum of 8 VMs per DB Server on X8M and above generations.
- Exadata Infrastructures with X8M and above generation of DB Servers can support a maximum of 8 VM clusters across all DB Servers.
- Maximum number of clusters across the infrastructure depends on the resources available per DB server and is subject to the per DB Server maximum VM limit.

With the release of Multi-VM, the add and remove virtual machine API for cloud VM clusters will not be supported via terraform.

You can perform these operations through UI, SDK, OCI CLI, OCI Ansible or similar tools. Terraform states should be managed similar to other operations which happen outside of terraform but need to be managed in terraform.

For more information, see *Detecting and Managing Drift with Terraform*.

- [Add a VM to a VM Cluster](#)
Add a Virtual Machine to a VM Cluster
- [Terminate or Remove a VM from a VM Cluster](#)
To terminate or remove a virtual machine from a provisioned cluster, use this procedure.

Related Topics

- [Scaling Resources within an Exadata Infrastructure Instance](#)
If an Exadata Cloud Infrastructure instance requires more resources, you can scale up the number of DB servers, or storage servers.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [Add a VM to a VM Cluster](#)
Add a Virtual Machine to a VM Cluster

- [To View a List of DB Servers on an Exadata Infrastructure](#)
To view a list of database server hosts on an Oracle Exadata Database Service on Dedicated Infrastructure system, use this procedure.
- [Terminate or Remove a VM from a VM Cluster](#)
To terminate or remove a virtual machine from a provisioned cluster, use this procedure.
- [Detecting and Managing Drift with Terraform](#)

Add a VM to a VM Cluster

Add a Virtual Machine to a VM Cluster

Note

Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.

Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2023 update cycle.

Note

- This operation is only available with Multi-VM enabled Infrastructure.
- To add a VM to a VM Cluster requires that all TCP ports to be open for the client subnet CIDR for ingress and egress.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **VM Clusters**.
4. Click the name of the VM cluster to which you want to add a virtual machine.
5. Under Resources, select **Virtual Machines**, and click the **Add Virtual Machines** button.
6. In the Add Virtual Machines window, select the DB server where you want the new VM to reside.

Note

The VM that is added will have the same resources as the other VMs in the cluster.

7. Click **Add**.

Note

Add a VM to a VM Cluster is NOT supported using Terraform.

Related Topics

- [Adding a VM to a VM Cluster Fails](#)

Terminate or Remove a VM from a VM Cluster

To terminate or remove a virtual machine from a provisioned cluster, use this procedure.

Note

To remove a VM from a provisioned VM cluster in a non-multi-VM-enabled infrastructure, follow a procedure similar to terminating a VM from a VM Cluster in a multi-VM-enabled infrastructure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to scale the CPU resources.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to remove a virtual machine.
5. On the Exadata VM Cluster Details page, in the Virtual Machines section, select the Virtual Machine that will be removed, click the more commands symbol (three dots) and click **Terminate**

Note

Remove a VM from a VM Cluster is NOT supported using Terraform at this time.

Overview of Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR)

Oracle Cloud Infrastructure Zero Trust Packet Routing (ZPR) safeguards sensitive data from unauthorized access using intent-based security policies that you define for OCI resources assigned with security attributes.

These security attributes act as labels, enabling ZPR to identify and organize OCI resources. ZPR enforces these policies at the network level every time access is requested, regardless of any changes or misconfigurations in the network architecture.

ZPR is layered on top of existing network security group (NSG) and security control list (SCL) rules. For a packet to successfully reach its target, it must pass through all applicable NSG, SCL, and ZPR policies. If any rule or policy blocks the traffic, the request is denied.

You can secure your networks using Zero Trust Packet Routing (ZPR) in three steps:

1. Creating and managing security attribute namespaces and security attributes

For more information, see [Managing Security Attributes](#).

2. Writing policies using security attributes to control access to resources

For more information, see [Understanding Zero Trust Packet Routing Policy](#) and [Policy Template Builder](#)

3. Applying security attributes to specified resources

For more information, see [Adding a Resource to Zero Trust Packet Routing](#).

 **Caution**

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

Private Service Access for OCI Service Access from VCN

You can use Private Service Access (PSA) instead of a Service Gateway to provide access to the OCI services required to support your database service.

Create the necessary PSA endpoints in your Virtual Cloud Network (VCN) to enable private access to the OCI services used with your Exadata Database Service on Dedicated Infrastructure VM, such as:

- Database Service: Used for resource principal authentication and access to your service metadata
- Identity and Access Management Data Plane API: Used for service authentication
- Object Storage Service API: Used for software updates and custom image downloads
- Functions Service Invocation: Used for service monitoring
- OCI Monitoring Ingestion: Used for service monitoring
- Logging Ingestion API: Used for service monitoring

You must also create a network security rule to allow inbound access to your subnet from the network that you configured your PSA on.

1. On the **Virtual Cloud Networks** list page, select the VCN or subnet that you want to create a PSA endpoint in.
If you need help finding the list page for the VCN, see [Listing VCNs](#) or [Listing Subnets](#).
2. On the details page, go to the **Private Service Access** tab and select **Create**.
3. Enter a friendly name for the PSA endpoint.
It doesn't have to be unique. Avoid entering confidential information.
4. Verify the compartment that you want to create the PSA endpoint in.
Select another compartment if needed.
5. (Optional) In the **Tags** section, add one or more tags. 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're not sure whether to apply tags, skip this option or ask an administrator. You can apply tags later.
6. In the **Service** section, select the OCI Service for which you want to enable access via the PSA endpoint.
You can select only one service per endpoint.

7. In the **Network** section, select the compartment for the PSA endpoint's subnet, and the subnet itself.
8. Select whether the PSA endpoint's private IPv4 address is automatically or manually assigned.
If you select **Automatically assign a private IPv4 address**, you can decide between ephemeral (dynamically allocated from the available IP addresses in the subnet by Oracle) or persistent IPv4 addresses (selected from existing [reserved private IPs](#)).

If you select **Manually assign a private IPv4 address**, you can either **Provide private IPv4 address** and enter the address in the field, or **Select existing reserved IPv4 address** from the list of available addresses.
9. (Optional) Decide whether to add ZPR security attributes on the PSA endpoint.
If you select this option, you can add up to three security attributes to restrict access to resources. If you have permissions to create a resource, then you might also have permissions to apply security attributes to that resource. To apply a security attribute, you must have permissions to use the security attribute namespace. For more information about security attributes and security attribute namespaces, see [Zero Trust Packet Routing](#). If you're not sure whether to apply security attributes, skip this option or ask an administrator. You can apply security attributes later.
10. (Optional) Decide whether to add the PSA endpoint to an [NSG](#).
If you select this option, select the compartment that contains the NSG and then select the NSG you want.
11. Select **Create private service access endpoint**.

About Application VIP

Oracle Exadata Database Service on Dedicated Infrastructure fully supports creating additional Virtual IP Addresses on an Exadata VM Cluster.

These application VIPs are required to protect additional applications such as Oracle GoldenGate installed on an Oracle Exadata Database Service on Dedicated Infrastructure system or other services such as XA-Agents, and to provide High Availability to these additional applications. For more information, see [Oracle Grid Infrastructure Standalone Agents for Oracle Clusterware](#) and [Making Applications Highly Available Using Oracle Clusterware](#).

Within the Oracle Cloud Infrastructure, adding Virtual IP Addresses on the cluster stack alone is not sufficient as these additional (secondary) IP addresses also have to be added to the VCN layer as "**Floating IP**" addresses so that the VCN layer knows where these IP addresses are running, and in case of failover by the Clusterware to change the VNIC the floating IP address is attached to. For more information, see [Creating an Application VIP Managed by Oracle Clusterware](#) and [Overview of IP Addresses](#).

Adding an Application VIP to an Exadata VM Cluster consists of the following steps:

1. Add the Virtual IP address to the Clusterware layer within the Exadata DomU, by following the standard [Oracle Clusterware documentation](#), or the guide provided by the application, for example, by using

```
IPv4: appvipcfg create -network=1 -ip=10.10.10.10 -vipname=applicationvip
```

```
IPv4/IPv6 dual-stack: appvipcfg create -network=1 -ip="10.10.10.10  
2607:9b80:9a0a:9401:1801:2645:c0a2:4283" -vipname=applicationvip
```

2. Attach the Application Floating IP address object on the Exadata VM Cluster to add the knowledge of the floating IP to the VCN layer. Ensure that you choose the same subnet as you created the backend application VIP, which normally is the client subnet. The private IP address needs to be the same as the one specified in the `appvipcfg` command above. The **Virtual IP Address Hostname** is the name under which the IP address is reachable via DNS and does not have to be the same as the `vipname`.

If you have already started the VIP in the backend, ensure that the **Virtual Machine Name** reflects the host on which the VIP was started in the backend.

3. Test the relocation of the VIP. The VIP should stay available (test this via `ping`), and the user interface should display after a short while that the floating IP also has moved to another host. If you did choose the wrong host while creating the VCN attachment, then simply relocate the VIP within the cluster. The VCN layer will detect the change and the user interface should get updated after a short while.

Note

A single Virtual VM Cluster has a limitation of 8 additional Application VIPs. The limitation exists because a single VNIC can only have 31 additional secondary IP addresses. For more information, see [Overview of IP Addresses](#). If all VIPs are started on the same node, then the application VIPs cannot be reached.

If more application VIPs are required, then raise an SR to have this limit increased. However, there are a few additional steps required then, to ensure that under no scenario more than 31 secondary IP addresses are attached to a single Exadata VM Cluster node. One way to accomplish this would be to ensure application VIPs are bound by the Clusterware to certain nodes so that this scenario is prevented.

A setup with 32 additional application VIPs would look as follows:

Floating IP	Node 1	Node 2	Node 3	Node 4
Private Hostname	1	1	1	1
VIP Hostname	4	4	4	4
SCAN	3	3	3	3
Appvip 1-8	8	8	-	-
Appvip 9-16	-	8	8	-
Appvip 17-24	-	-	8	8
Appvip 25-32	8	-	-	8
Max Possible VIP if all Floating IPs failover	24	24	24	24

Related Topics

- [To Attach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [To Detach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Permissions and API operation details for Application VIPs](#)

- [Permissions Required for Each API Operation](#)
- [Application VIP Event Types](#)
These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Using the Console to Manage VM Clusters on Exadata Cloud Infrastructure

Learn how to use the console to create, edit, and manage your VM Clusters on Oracle Exadata Database Service on Dedicated Infrastructure.

- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To create an Exascale cloud VM cluster](#)
To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To add security attributes to an Exadata VM Cluster](#)
To add security attributes to an Exadata VM Cluster, use this procedure.
- [To edit a security attribute](#)
To edit a security attribute of an Exadata VM Cluster, use this procedure.
- [To remove a security attribute](#)
To remove a security attribute of an Exadata VM Cluster, use this procedure.
- [To add database server or storage server capacity to a cloud VM cluster](#)
This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.
- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- [To Update the License Type on a VM Cluster](#)
To modify licensing, be prepared to provide values for the fields required for modifying the licensing information.
- [To add SSH keys to a VM cluster](#)
The VM cluster exists, and you wish to add a another user which requires another SSH key.
- [To Add SSH Keys After Creating a VM Cluster](#)
- [To Stop, Start, or Reboot a VM Cluster Virtual Machine](#)
Use the console to stop, start, or reboot a virtual machine.
- [To Check the Status of a VM Cluster Virtual Machine](#)
Review the health status of a VM cluster virtual machine.
- [To Move a VM Cluster to Another Compartment](#)
To change the compartment that contains your VM cluster on Exadata Cloud Infrastructure, use this procedure.
- [To change the VM cluster display name](#)
- [To Terminate a VM Cluster](#)
- [To view details about private DNS configuration](#)

- [To Attach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [To Detach a Virtual IP Address](#)
Attach a Virtual IP address from a VM cluster using this procedure.
- [Migrate from a Single Stack \(IPv4\) Exadata VM Cluster to a Dual Stack \(IPv4/IPv6\) Exadata VM Cluster with Data Guard Synchronization](#)

To create an ASM cloud VM cluster

To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Note

To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first [created a Cloud Exadata infrastructure resource](#).

Note

Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature [Create and Manage Multiple Virtual Machines per Exadata System \(MultiVM\) and VM Cluster Node Subsetting](#) was released only support creating a single cloud VM cluster.

Note

When you provision an Exadata VM cluster in Exadata Database Service on Oracle Database@Google Cloud, an Identity Connector is automatically created and associated with the VM cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
4. **Compartment:** Select a compartment for the VM cluster resource.

5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the VM cluster. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

7. VM Cluster Type:

Note

You cannot change the VM cluster type after deploying the VM cluster. If you wish to change the VM cluster type, you must create a new VM cluster and migrate the database to the new cluster.

- **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
 - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.
8. **Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Select DB Servers** to select from the available DB servers, and then click **Save**.

VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

In the **Resource allocation per VM** pane:

- Specify the number of OCPU/ECPU you want to allocate to each of the VM cluster's virtual machine compute nodes. For VM clusters created on X11M Exadata infrastructure specify ECPUs. For VM Clusters created on X10M and earlier Exadata infrastructure, specify OCPUs. The minimum is 2 OCPU per VM for X10M and earlier infrastructure or 8 ECPUs per VM for VM clusters created on X11M Exadata infrastructure. The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU or ECPU cores you are allocating.
- Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
- Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.
Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

In addition to /u02, you can specify the size of additional local file systems.

For more information and instructions to specify the size for each individual VM, see [Introduction to Scale Up or Scale Down Operations](#).

- Click **Show additional local file systems configuration options**.
- Specify the size of `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

Note

- * You can only expand these file systems and cannot reduce the size once expanded.
- * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

- After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

9. Exadata storage:

- **Specify the usable Exadata storage TB.** Specify the storage in multiples of 1 TB. Minimum: 2 TB
- **Allocate storage for Exadata sparse snapshots:** Select this configuration option if you intend to use snapshot functionality within your VM cluster. If you select this option, the SPARSE disk group is created, which enables you to use VM cluster snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and snapshot functionality will not be available on any database deployments that are created in the environment.

Note

The storage configuration option for sparse snapshots cannot be changed after VM cluster creation.

- **Allocate storage for local backups:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata Cloud Infrastructure instance. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.

Note

The storage configuration option for local backups cannot be changed after VM cluster creation.

10. Version:

- **Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release (19c and 26ai) that you want to install on the VM cluster.

The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

Note

Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 26ai:

- Exadata Guest VM running Exadata System Software 23.1.8
- Exadata Infrastructure running Exadata System Software 23.1.x

- **Exadata guest version:**

- **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**

- * The **Change image** button is not enabled.
- * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
- * The Exadata guest version will be the same as that of the host OS.

- **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**

- * The Exadata guest version defaults to the latest (23.1.3.0).
- * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
- * The **Change image** button is enabled.
- * Click **Change image**.
The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).
The most recent release for each major version is indicated by "(latest)".
- * Slide **Display all available versions**.
Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
- * Choose a version.
- * Click **Save Changes**.

11. SSH Keys: Add the public key portion of each key pair you want to use for SSH access to the VM cluster:

- **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.
- **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
- **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.

12. Network settings: Specify the following:

Note

IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect

You do not have the option to choose between IPv4 (single stack) and IPv4/IPv6 (dual stack) if both configurations exist. For more information, see [VCN and Subnet Management](#).

- **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
- **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment. Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
- **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable. Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in [Managing Exadata Database Backups](#).

Note

In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, [Configuring Network Resources for Recovery Service](#).

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress **security rules** to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified. For more information, see **Network Security Groups** and *Network Setup for Exadata Cloud Infrastructure Instances*. Note that if you choose a subnet with a **security list**, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

To use network security groups:

- Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.
- Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.

- To use additional NSGs with the network, click **+;Another Network Security Group**.

Note

To provide your cloud VM Cluster resources with additional security, you can use Oracle Cloud Infrastructure Zero Trust Packet Routing to ensure that only resources identified with security attributes have network permissions to access your resources. Oracle provides Database policy templates that you can use to assist you with creating policies for common database security use cases. To configure it now, you must already have created security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. Click **Show Advanced Options** at the end of this procedure.

Be aware that when you provide security attributes for a cluster, as soon as it is applied, all resources require a Zero Trust Packet policy to access the cluster. If there is a security attribute on an endpoint, then it must satisfy both network security group (NSG) and Oracle Cloud Infrastructure Zero Trust Packet Routing policy (OCI ZPR) rules.

- **To use private DNS Service**

Note

A Private DNS must be configured before it can be selected. See *Configure Private DNS*

- Check the **Use private DNS Service** check box.
- Select a private view. Click **Change Compartment** to select a private view in a different compartment.
- Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.
- **Hostname prefix:** Your choice of hostname for the Exadata VM cluster. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata VM cluster is 12.

Caution

The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted. If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.

- **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.
13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
- **License Included** means the cost of the cloud service includes a license for the Database service.
 - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the check boxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
 - **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
 - **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).
15. Click **Show Advanced Options** to specify advanced options for the VM cluster:
- **Time zone:** This option is located in the **Management** tab. The default time zone for the VM cluster is UTC, but you can specify a different time zone. The time zone

options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Note

If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

- **SCAN Listener Port:** This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521.

Note

Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

- **Zero Trust Packet Routing (ZPR):** This option is located in the **Security attributes** tab. Select a namespace, and provide the key and value for the security attribute. To complete this step during configuration, you must already have set up security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. You can also add security attributes after configuration, and add them later. For more information about adding Oracle Exadata Database Service on Dedicated Infrastructure specific policies, see [Policy Template Builder](#).
- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

Note

Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Note

- The freeze period can extend for a maximum of 45 days from the start date.
- Oracle automation will automatically apply updates with critical security fixes (CVSS \geq 9) even during a configured freeze period.

- **Tags:** 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.

16. Click Create.**WHAT NEXT?**

After your VM cluster is successfully created and in the **Available** state.

- You can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can **create your first database** in the cluster by clicking **Create Database**
- The **SCAN IP address (IPv4)** and **SCAN IP address (IPv6)** fields in the **Network** section on the VM Cluster Details page displays the dual stack IP address details.
- The **Cloud Automation Update** field in the **Version** section on the VM Cluster Details page displays the freeze period you have set.

Related Topics

- [Network Security Groups](#)
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- [Security Lists](#)
- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.
- [Resource Tags](#)
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Getting Started with Events](#)
- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.
- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

To create an Exascale cloud VM cluster

To create your Exascale VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Note

To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first [created a Cloud Exadata infrastructure resource](#).

Note

Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature [Create and Manage Multiple Virtual Machines per Exadata System \(MultiVM\) and VM Cluster Node Subsetting](#) was released only support creating a single cloud VM cluster.

Note

When you provision an Exadata VM cluster in Exadata Database Service on Oracle Database@Google Cloud, an Identity Connector is automatically created and associated with the VM cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
4. **Compartment:** Select a compartment for the VM cluster resource.
5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the VM cluster. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

7. VM Cluster Type:**Note**

You cannot change the VM cluster type after deploying the VM cluster. If you wish to change the VM cluster type, you must create a new VM cluster and migrate the database to the new cluster.

- **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
 - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.
- 8. Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Select DB Servers** to select from the available DB servers, and then click **Save**.

VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.

VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

In the **Resource allocation per VM** pane:

- Specify the number of OCPU/ECPU you want to allocate to each of the VM cluster's virtual machine compute nodes. For VM clusters created on X11M Exadata infrastructure specify ECPU. For VM Clusters created on X10M and earlier Exadata infrastructure, specify OCPUs. The minimum is 2 OCPU per VM for X10M and earlier infrastructure or 8 ECPU per VM for VM clusters created on X11M Exadata infrastructure. The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU or ECPU cores you are allocating.
- Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
- Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.
Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

In addition to /u02, you can specify the size of additional local file systems.

For more information and instructions to specify the size for each individual VM, see [Introduction to Scale Up or Scale Down Operations](#).

- Click **Show additional local file systems configuration options**.
- Specify the size of /, /u01, /tmp, /var, /var/log, /var/log/audit, and /home file systems as needed.

Note

- * You can only expand these file systems and cannot reduce the size once expanded.
- * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

- After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

9. Exadata storage:

The following settings define how the Exadata storage is configured for use with the VM cluster. The storage type once selected cannot be changed later on once the VM cluster is provisioned with the desired storage type. You have two options to choose: Automatic storage type (ASM) and Exascale. For more information about ASM storage type, see *To create an ASM cloud VM cluster resource*.

Note

Minimum requirement to configure Exascale storage.

- This feature is supported on Exadata Infrastructure Model X8M and later.
- This feature is available on Exadata system software release 24.1 and later.
- This feature requires Oracle Grid Infrastructure version 26ai (24.3) and supports Oracle database versions 26ai (23.4) and later. Exascale option will be disabled if the minimum requirement is not met.

Exascale database storage vault: Choose this option to create a new Exascale database storage vault during VM cluster provisioning.

- **Create new storage vault:**
 - **Storage vault name:** Enter a descriptive name for the vault. Click the **change compartment** link and choose a compartment if you want to create this vault in a different compartment.
 - **Storage capacity for databases:** Enter the storage capacity for the databases within the minimum and maximum values displayed on the screen.

Note

If additional space is needed beyond the maximum shown, the Exascale capacity must be increased. For more information, see *Using the Console to Scale an Exascale Storage Vault*.

- **Select existing storage vault:** Select a vault that resides in the compartment of your choice.

10. Version:

- **Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release 26ai to install on the VM cluster. The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

Note

Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 26ai:

- Exadata Guest VM running Exadata System Software 23.1.8
- Exadata Infrastructure running Exadata System Software 23.1.x

- **Exadata guest version:**
 - Exadata infrastructure with Oracle Linux 8 and Exadata image version 25.1.7 (or later) for DB server and 25.1,8 (or later) or 25.2.2 (or later) for storage server.
 - * The Exadata guest version defaults to the latest available Exadata Image.
 - * Click **Change image**.
The resulting Change image panel displays the list of available major versions of Exadata images.
The most recent release for each major version is indicated by "(latest)".
 - * Slide **Display all available versions**.
Six past versions including the latest versions of each Exadata Major release are displayed.
 - * Choose the desired version.
 - * Click **Save Changes**.
- 11. SSH Keys:** Add the public key portion of each key pair you want to use for SSH access to the VM cluster:
- **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.
 - **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
 - **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.
- 12. Network settings:** Specify the following:

Note

IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect.

You do not have the option to choose between IPv4 (single stack) and IPv4/IPv6 (dual stack) if both configurations exist. For more information, see [VCN and Subnet Management](#).

- **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
- **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment. Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
- **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable. Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in [Managing Exadata Database Backups](#).

Note

In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, [Configuring Network Resources for Recovery Service](#).

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress **security rules** to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified. For more information, see **Network Security Groups** and *Network Setup for Exadata Cloud Infrastructure Instances*. Note that if you choose a subnet with a **security list**, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

To use network security groups:

- Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.
- Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.
- To use additional NSGs with the network, click **+;Another Network Security Group**.

Note

To provide your cloud VM Cluster resources with additional security, you can use Oracle Cloud Infrastructure Zero Trust Packet Routing to ensure that only resources identified with security attributes have network permissions to access your resources. Oracle provides Database policy templates that you can use to assist you with creating policies for common database security use cases. To configure it now, you must already have created security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. Click **Show Advanced Options** at the end of this procedure.

Be aware that when you provide security attributes for a cluster, as soon as it is applied, all resources require a Zero Trust Packet policy to access the cluster. If there is a security attribute on an endpoint, then it must satisfy both network security group (NSG) and Oracle Cloud Infrastructure Zero Trust Packet Routing policy (OCI ZPR) rules.

- **To use private DNS Service**

Note

A Private DNS must be configured before it can be selected. See *Configure Private DNS*

- Check the **Use private DNS Service** check box.
- Select a private view. Click **Change Compartment** to select a private view in a different compartment.
- Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.
- **Hostname prefix:** Your choice of hostname for the Exadata VM cluster. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata VM cluster is 12.

Caution

The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted. If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.
- **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.

13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
 - **License Included** means the cost of the cloud service includes a license for the Database service.
 - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time.

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

Note

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the check boxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
 - **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
 - **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).
15. Click **Show Advanced Options** to specify advanced options for the VM cluster:
 - **Time zone:** This option is located in the **Management** tab. The default time zone for the VM cluster is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Note

If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

- **SCAN Listener Port:** This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521.

Note

Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

- **Zero Trust Packet Routing (ZPR):** This option is located in the **Security attributes** tab. Select a namespace, and provide the key and value for the security attribute. To complete this step during configuration, you must already have set up security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. You can also add security attributes after configuration, and add them later. For more information about adding Oracle Exadata Database Service on Dedicated Infrastructure specific policies, see [Policy Template Builder](#).
- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

Note

Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Note

- The freeze period can extend for a maximum of 45 days from the start date.
- Oracle automation will automatically apply updates with critical security fixes (CVSS \geq 9) even during a configured freeze period.

- **Tags:** 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.

16. Click Create.**WHAT NEXT?**

After your VM cluster is successfully created and in the **Available** state.

- You can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can **create your first database** in the cluster by clicking **Create Database**
- The **SCAN IP address (IPv4)** and **SCAN IP address (IPv6)** fields in the **Network** section on the VM Cluster Details page displays the dual stack IP address details.
- The **Cloud Automation Update** field in the **Version** section on the VM Cluster Details page displays the freeze period you have set.

Related Topics

- [Network Security Groups](#)
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- [Security Lists](#)
- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.
- [DB System Time Zone](#)
- [Resource Tags](#)
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Getting Started with Events](#)
- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.
- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

To add security attributes to an Exadata VM Cluster

To add security attributes to an Exadata VM Cluster, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. In the list of cloud VM clusters, click the name of the cluster to which you want to add security attributes.
4. On the VM Cluster Details page, click **More actions**, and then select **Add security attributes**.
or
On the VM Cluster Details page, click the **Security attributes** tab.
5. Click **Add security attributes**.
6. Select the **Namespace** in which the required security attribute is available.
7. Select the **Key** and **Value** of the required security attribute.
8. To use additional security attributes, click **Add security attribute**.

Note

A maximum of 3 security attributes can be specified for an Exadata VM Cluster.

9. Click **Add security attributes**.

To edit a security attribute

To edit a security attribute of an Exadata VM Cluster, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of cloud VM clusters, click the name of the cluster for which you want to edit a security attribute.
3. Click the **Security** tab.
A list of existing security attributes are displayed.
4. Click the icon beside the name of the security attribute to view its details.
5. Choose the required value of the security attribute.
6. Click **Save**.

To remove a security attribute

To remove a security attribute of an Exadata VM Cluster, use this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of cloud VM clusters, click the name of the cluster for which you want to edit a security attribute.

3. Click the **Security** tab.
A list of existing security attributes are displayed.
4. Click the icon beside the name of the security attribute to view its details.
5. Click **Remove security attribute**.

To add database server or storage server capacity to a cloud VM cluster

This topic describes how to use the Oracle Cloud Infrastructure (OCI) Console to add the new capacity to your cloud VM cluster.

Note

This procedure does not apply to Multi-VM enabled Infrastructure

If you have used the task *To add compute and storage resources to a flexible cloud Exadata infrastructure resource* by adding additional database (compute) or storage servers to the service instance, you must add the additional capacity to the cloud VM cluster to utilize the additional resources.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. In the list of cloud VM clusters, click the name of the cluster to which you want to add capacity.
4. On the VM Cluster Details page, click **Scale VM Cluster**.
5. If you have additional capacity available as a result of scaling the cloud Exadata infrastructure resource, a banner at the top of the **Scale VM Cluster** panel provides a message telling you the type and amount of additional capacity available to the VM cluster. Check the **Add Capacity** box.
6. Select either the **Add Database Server** or the **Add Storage** radio button, depending on which type of capacity you want to add to the cloud VM cluster.
7. Click **Update**. The cloud VM cluster goes into the Updating state. When the capacity has been successfully added, the cluster returns to the Available state.

Note

If you have added additional database servers to the cluster, you can allocate the new CPU cores once the cluster is in the Available state by clicking the **Scale VM Cluster** button again. See *To scale CPU cores in an Exadata Cloud Service cloud VM cluster* for more information on adding CPU cores to your cloud VM cluster.

Related Topics

- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.
- [To scale CPU cores in an Exadata Cloud Infrastructure cloud VM cluster](#)

To Enable, Partially Enable, or Disable Diagnostics Collection

You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Note

- You are opting in with the understanding that the list of events, metrics, and log files collected can change in the future. You can opt-out of this feature at any time.
- Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.
- If you have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Click **VM Clusters**.
4. Click the name of the VM cluster you want to enable or disable diagnostic data collection.
5. On the VM Cluster Details page, under **General Information**, enable, partially enable, or disable **Diagnostics Collection** beside **Diagnostics Collection**.
6. In the **Edit Diagnostics Collection Settings** dialog, enable or disable any of the Diagnostics Collections. By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
 - **Enable Diagnostics Events** Allow Oracle to collect and publish critical, warning, error, and information events to me. For more information, see *Overview of Database Service Events*
 - **Enable Health Monitoring** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
 - **Enable Incident logs and trace collection.** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.
Note: You had previously opted in for incident log and trace file collection and decide to opt-out when Oracle Cloud operations run a log collection job, the job will run its course and will not cancel. Future log collections will not run until you opt-in again to the incident logs and trace file collection option.
7. Select or clear the checkboxes and then click **Save Changes**.

Related Topics

- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.

To Update the License Type on a VM Cluster

To modify licensing, be prepared to provide values for the fields required for modifying the licensing information.

1. Open the navigation menu. Under **Oracle AI Database**, click **Exadata Cloud Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster for which you want to update the license type.
3. Click **VM Clusters**.
4. Click the name of the VM cluster for which you want to update the license type.
The VM Cluster Details page displays information about the selected VM cluster.
5. Click **Update License Type**.
6. In the dialog box, choose one of the following license types and then click **Save**.
 - **Bring Your Own License (BYOL)**: Select this option if your organization already owns Oracle Database software licenses that you want to use on the VM cluster.
 - **License Included**: Select this option to subscribe to Oracle Database software licenses as part of Exadata Cloud Infrastructure.

Updating the license type does not change the functionality or interrupt the operation of the VM cluster. Customers are permitted to change the license type for a VM Cluster at most once per month.

To add SSH keys to a VM cluster

The VM cluster exists, and you wish to add a another user which requires another SSH key.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of VM clusters, find the cluster you want to manage and click its highlighted name.
5. Click **Add SSH Keys**.
6. Select one of the following options:
 - **Generate SSH key pair**: Use this option to create a new SSH key pair. Click both **Save Private Key** and **Save Public Key** when using this option. The private key is downloaded to your local machine, and should be stored in a safe location. You cannot download another copy of the private key generated during this operation after completing the operation.
 - **Upload SSH key files**: Select this option to browse or drag and drop .pub files.
 - **Paste SSH keys**: Select this option to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.

7. Click **Save**.

To Add SSH Keys After Creating a VM Cluster

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click **VM Clusters**.
3. Click the name of the VM cluster that you want to add SSH key(s).
4. In the VM Cluster Details page, click **Add SSH Keys**.
5. In the ADD SSH Keys dialog, choose any one of the methods:
 - **Generate SSH key pair:** Select this option if you want the Control Plane to generate public/private key pairs for you. Click **Save Private Key** and **Save Public Key** to download and save SSH Key pair.
 - **Upload SSH key files:** Select this option to upload the file that contains SSH Key pair.
 - **Paste SSH keys:** Select this option to paste the SSH key string. To provide multiple keys, click **Another SSH Key**. For pasted keys, ensure that each key is on a single, continuous line. The length of the combined keys cannot exceed 10,000 characters.
6. Click **Save**.

Related Topics

- [Managing Key Pairs on Linux Instances](#)

To Stop, Start, or Reboot a VM Cluster Virtual Machine

Use the console to stop, start, or reboot a virtual machine.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the virtual machine that you want to stop, start, or reboot.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the virtual machine that you want to stop, start, or reboot.

The VM Cluster Details page displays information about the selected VM cluster.
5. In the **Resources** list, click **Virtual Machines**.

The list of virtual machines is displayed.
6. In the list of nodes, click the **Actions** icon (three dots) for a node, and then click one of the following actions:
 - a. **Start:** Restarts a stopped node. After the node is restarted, the **Stop** action is enabled.
 - b. **Stop:** Shuts down the node. After the node is stopped, the **Start** action is enabled.
 - c. **Reboot:** Shuts down the node, and then restarts it.

To Check the Status of a VM Cluster Virtual Machine

Review the health status of a VM cluster virtual machine.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that is associated with the VM cluster that contains the virtual machine that you are interested in.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that contains the virtual machine that you are interested in.

The VM Cluster Details page displays information about the selected VM cluster.

5. In the **Resources** list, click **Virtual Machines**.

The list of virtual machines displays. For each virtual machine in the VM cluster, the name, state, and client IP address are displayed.

6. In the node list, find the virtual machine that you are interested in and check its state.

The color of the icon and the associated text it indicates its status.

- **Available:** Green icon. The node is operational.
- **Starting:** Yellow icon. The node is starting because of a start or reboot action in the Console or API.
- **Stopping:** Yellow icon. The node is stopping because of a stop or reboot action in the Console or API.
- **Stopped:** Yellow icon. The node is stopped.
- **Failed:** Red icon. An error condition prevents the continued operation of the virtual machine.

To Move a VM Cluster to Another Compartment

To change the compartment that contains your VM cluster on Exadata Cloud Infrastructure, use this procedure.

When you move a VM cluster, the compartment change is also applied to the virtual machines and databases that are associated with the VM cluster. However, the compartment change does not affect any other associated resources, such as the Exadata infrastructure, which remains in its current compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to move.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that you want to move.

The VM Cluster Details page displays information about the selected VM cluster.

5. Click **More actions**, then click **Move resource**.
6. In the resulting dialog, choose the new compartment for the VM cluster, and click **Move**.

To change the VM cluster display name

Note

This topic only applies to Exadata Cloud Infrastructure instances using the new Exadata Cloud Infrastructure instance resource model.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of Exadata VM Clusters resources, click the name of the VM Cluster you're interested in.
5. On the **Infrastructure Details** page, click **More Actions** and **Update Display Name**.
6. In the **Update Display Name** dialog, Enter the **New display name**, and the **current display name** as instructed.
7. Click **Update**.

To Terminate a VM Cluster

Before terminating a VM cluster, ensure the following:

- If any databases have Data Guard configured, deconfigure them first.
- If any databases have a backup in progress, wait for the backup to complete.

Terminating a VM cluster removes it from the Cloud Control Plane. In the process, the virtual machines and their contents are destroyed.

Note

You cannot terminate a VM cluster from an infrastructure with less than 5 storage servers.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** and **Compartment** that contains the VM cluster that you want to terminate.
3. Click **VM Clusters**.
4. Click the name of the VM cluster that you want to terminate.
The VM Cluster Details page displays information about the selected VM cluster.
5. Click **More Actions**, and then click **Terminate**.
6. In the resulting dialog:
 - Review the message about the backup retention policy.

- Enter the name of the VM cluster
- Click **Terminate** to confirm the action.

Note

The database stays in a terminated state with backups listed until all backups are expired.

To view details about private DNS configuration

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **VM Clusters**.
5. Click the name of the VM cluster that is configured with a private DNS you want to view.
6. Under the Network section, Private DNS and Private Zone are displayed, if a private DNS is configured.
7. Click the **Private View** name to edit the configuration.

Related Topics

- [Using the Console to manage private DNS](#)

To Attach a Virtual IP Address

Attach a Virtual IP address from a VM cluster using this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **Exadata VM Clusters**.
5. Under the **Resources**, click **Virtual IP Address**.
6. Click **Attach Virtual IP Address**.
7. In the Attach Virtual IP Address dialog:
 - a. Select a subnet from the **Subnet** drop-down list.
 - b. Enter a hostname for the Virtual IP Address in the Virtual IP Address Hostname field.
 - c. You can choose to either Automatically assign IPv4/IPv6 addresses from the subnet or Manually assign IPv4/IPv6 addresses. If you opt for manual assignment, enter the desired IP address in the **Virtual IP address** field.
 - d. (Optional) You can enter a Virtual Machine name to be the default attachment in the **Virtual Machine** field.
 - e. Click **Attach**.

To Detach a Virtual IP Address

Attach a Virtual IP address from a VM cluster using this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Region** that contains your Exadata infrastructure.
3. Choose the **Compartment** that contains your Exadata infrastructure.
4. Click **Exadata VM Clusters**.
5. Under the **Resources**, click **Virtual IP Address**.
6. Click the Actions icon (three dots) to the right of the Virtual IP Address you wish to detach.
7. Click **Detach**.
8. In the Detach Virtual IP Address dialog, confirm by entering the VIP Address that you wish to detach and click **Detach**.

Migrate from a Single Stack (IPv4) Exadata VM Cluster to a Dual Stack (IPv4/IPv6) Exadata VM Cluster with Data Guard Synchronization

1. Create a dual stack Exadata VM cluster.
2. Create a Data Guard association between the existing single stack (IPv4) Exadata VM cluster that hosts the primary database and the new dual stack (IPv4/IPv6) Exadata VM cluster that hosts the standby database.
3. Perform a database switchover.

Ensure that both databases are synchronized before performing switchover. The primary and standby databases should be synchronized, meaning there should be no backlog of redo logs on the standby database. You can check synchronization by ensuring that no redo logs are pending to be applied on the standby. Use the `DGMGRL> SHOW CONFIGURATION;` command to check the synchronization status.

Overview of Automatic Diagnostic Collection

By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

- **Enable Diagnostic Events**
Allow Oracle to collect and publish critical, warning, error, and information events to you. For more information, see *Database Service Events*.
- **Enable Health Monitoring**
Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events. For more information, see *Health Metrics*.
- **Enable Incident Logs and Trace Collection**
Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution. For more information, see *Incident Logs and Trace Files*.

Diagnostics Collection is:

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

Disabling diagnostic events and health monitoring will only stop the collection and notification of data/events from the time you uncheck the checkboxes tied to the options. However, historical data will not be purged from Oracle Cloud Operations data repositories.

Related Topics

- [Database Service Events](#)
The Database Service emits events, which are structured messages that indicate changes in resources.
- [Incident Logs and Trace Files](#)
This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- [Health Metrics](#)
Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Incident Logs and Trace Files

This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.

Note

- Oracle will create a service request (SR) against the infrastructure Customer Support Identifier (CSI) when an issue is detected and needs customer interaction to resolve.
- The customer's Oracle CCloud Infrastructure tenancy admin email will be used as the CSI contact to create SR and attach logs to it. Ensure tenancy admin is added as a CSI contact in My Oracle Support (MOS).

Oracle Trace File Analyze (TFA) Component Driven Logs Collections

The directories are generally assigned to a component and that component can then be used to guide TFA to the files it needs to collect, for example, requesting the CRS component would

tell TFA to look at directories mapped to the CRS component and find files that match the required collection time frame.

Note

If have previously opted in for incident log and trace file collection and decide to opt out when Oracle Cloud operations run a log collection job, then the job will run its course and will not cancel. Future log collections won't happen until you opt-in again to the incident logs and trace file collection option.

TFA is shipped with scripts that run when a particular component is requested, for example, for CRS component, `crscollect.pl` will run a number of `crsctl` commands and gather the input. By default, TFA does not redact collected logs.

Table 5-2 Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
OS: Operating system logs	<code>oscollect.pl</code>	<ul style="list-style-type: none"> • <code>/var/log/messages</code> • OSWatcher archive • Exadata Only: ExaWatcher archive <code>/opt/oracle.ExaWatcher/archive/</code>

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
CRS: Grid Infrastructure and cluster logs	crscollect.pl	<ul style="list-style-type: none"> • /etc/oracle • GIHOME/crf/db/HOSTNAME1 • GIHOME/crs/log • GIHOME/css/log • GIHOME/cv/log • GIHOME/evm/admin/log • GIHOME/evm/admin/logger • GIHOME/evm/log • GIHOME/log/-/client • GIHOME/log/HOSTNAME1 • GIHOME/log/HOSTNAME1/admin • GIHOME/log/HOSTNAME1/client • GIHOME/log/HOSTNAME1/crflgd • GIHOME/log/HOSTNAME1/crfmond • GIHOME/log/HOSTNAME1/crsd • GIHOME/log/HOSTNAME1/cssd • GIHOME/log/HOSTNAME1/ctssd • GIHOME/log/HOSTNAME1/diskmon • GIHOME/log/HOSTNAME1/evmd • GIHOME/log/HOSTNAME1/gipcd • GIHOME/log/HOSTNAME1/gnsd • GIHOME/log/HOSTNAME1/gpnpd • GIHOME/log/HOSTNAME1/mdnsd • GIHOME/log/HOSTNAME1/ohasd • GIHOME/log/HOSTNAME1/racg • GIHOME/log/HOSTNAME1/srvm • GIHOME/log/HOSTNAME1/xag • GIHOME/log/diag/asmtool

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
		<ul style="list-style-type: none"> • GIHOME/log/diag/clients • GIHOME/log/procwatcher/PRW_SYS_HOSTNAME1 • GIHOME/network/log • GIHOME/opmn/logs • GIHOME/racg/log • GIHOME/scheduler/log • GIHOME/srvm/log • GRIDBASE/crsdata/@global/cvu • GRIDBASE/crsdata/HOSTNAME1/core • GRIDBASE/crsdata/HOSTNAME1/crsconfig • GRIDBASE/crsdata/HOSTNAME1/crsdiag • GRIDBASE/crsdata/HOSTNAME1/cvu • GRIDBASE/crsdata/HOSTNAME1/evm • GRIDBASE/crsdata/HOSTNAME1/output • GRIDBASE/crsdata/HOSTNAME1/ovmmwallets • GRIDBASE/crsdata/HOSTNAME1/scripts • GRIDBASE/crsdata/HOSTNAME1/trace • GRIDBASE/diag/crs/-/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/cdump • GRIDBASE/diag/crs/HOSTNAME1/crs/incident • GRIDBASE/diag/crs/HOSTNAME1/crs/trace

Table 5-2 (Cont.) Oracle Trace File Analyze (TFA) Component Driven Logs Collections

Component	Script	Files/Directories
Database: Oracle Database logs	No DB Specific Script - runs opatch lsinventory for the ORACLE_HOME the DB runs from TFA will run ipspack based on the time range for certain DB incidents.	<ul style="list-style-type: none"> • ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/cdump • ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/trace • ORACLE_BASE/diag/rdbms/<dbname>/<instance_name>/incident

Cloud Tool Logs

- **Creg files:** /var/opt/oracle/creg/*.ini files with masked sensitive info
- **Cstate file:** /var/opt/oracle/cstate.xml
- **Database related tooling logs:**
If dbName specified, /var/opt/oracle/log/<dbName>, else collect logs for all databases /var/opt/oracle/log/
If dbName specified, /var/opt/oracle/dbaas_acfs/log/<dbName>, else collect logs for all databases /var/opt/oracle/log/<dbName>
- **Database env files:** If dbName specified, /home/oracle/<dbName>.env, else collect logs for all databases /home/oracle/*.env
- **Pilot logs:** /home/opc/.pilotBase/logs
- **List of log directories:**
 - /var/opt/oracle/log
 - /var/opt/oracle/dbaas_acfs/log
 - /var/opt/oracle/dbaas_acfs/dbsystem_details
 - /var/opt/oracle/dbaas_acfs/job_manager
 - /opt/oracle/dcs/log

DCS Agent Logs

- /opt/oracle/dcs/log/

Tooling-Related Grid Infrastructure/Database Logs

- **Grid Infrastructure:** GI_HOME/cfgtoollogs
- **Database alertlog:** /u02/app/oracle/diag/rdbms/*/*/alert*.log

Related Topics

- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- [Health Metrics](#)
Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Health Metrics

Review the list of database and non-database health metrics collected by Oracle Trace File Analyzer.

① Note

Oracle may add more metrics in the future, but if you have already chosen to collect metrics, you need not update your opt-in value. It will remain enabled/disabled based on your current preference.

① Note

In addition to the metrics listed below, Oracle analyzes additional metrics to provide highest level of service operations and support for ensuring high availability of services.

Guest VM Health Metrics List - Database Metrics

Table 5-3 Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
CpuUtilization	CPU Utilization	Percentage	Mean	One minute	Five minutes	The CPU utilization is expressed as a percentage, which is aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.
StorageUtilization	Storage Utilization	Percentage	Mean	One hour	One hour	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.
BlockChanges	DB Block Changes	Changes per second	Mean	One minute	Five minutes	The Average number of blocks changed per second.
ExecuteCount	Execute Count	Count	Sum	One minute	Five minutes	The number of user and recursive calls that executed SQL statements during the selected interval.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
CurrentLogons	Current Logons	Count	Sum	One minute	Five minutes	The number of successful logons during the selected interval.
TransactionCount	Transaction Count	Count	Sum	One minute	Five minutes	The combined number of user commits and user rollbacks during the selected interval.
UserCalls	User Calls	Count	Sum	One minute	Five minutes	The combined number of logons, parses, and execute calls during the selected interval.
ParseCount	Parse Count	Count	Sum	One minute	Five minutes	The number of hard and soft parses during the selected interval.
StorageUsed	Storage Space Used	GB	Max	One hour	One hour	Total amount of storage space used by the database at the collection time.
StorageAllocated	Storage Space Allocated	GB	Max	One hour	One hour	Total amount of storage space allocated to the database at the collection time.

Table 5-3 (Cont.) Guest VM Health Metrics List - Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Interval	Collection Frequency	Description
StorageUsedByTablespace	Storage Space Used By Tablespace	GB	Max	One hour	One hour	Total amount of storage space used by tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageAllocatedByTablespace	Allocated Storage Space By Tablespace	GB	Max	One hour	One hour	Total amount of storage space allocated to the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.
StorageUtilizationByTablespace	Storage Space Utilization By Tablespace	Percentage	Mean	One hour	One hour	This indicates the percentage of storage space utilized by the tablespace at the collection time. In the case of container databases, this metric provides root container tablespaces.

Guest VM Health Metrics List - Non-Database Metrics

Table 5-4 Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
ASMDiskgroup Utilization	ASM Diskgroup Utilization	Percentage	Max	10 minutes	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.
FilesystemUtilization	Filesystem Utilization	Percentage	Max	One minute	Percent utilization of provisioned filesystem.
CpuUtilization	CPU Utilization	Percentage	Mean	One minute	Percent CPU utilization.
MemoryUtilization	Memory Utilization	Percentage	Mean	One minute	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: <code>cat /proc/meminfo.</code>
SwapUtilization	Swap Utilization	Percentage	Mean	One minute	Percent utilization of total swap space.
LoadAverage	Load Average	Number	Mean	One minute	System load average over 5 minutes.
NodeStatus	Node Status	Integer	Mean	One minute	Indicates whether the host is reachable.

Table 5-4 (Cont.) Guest VM Health Metrics List - Non-Database Metrics

Metric Name	Metric Display Name	Unit	Aggregation	Collection Frequency	Description
OcpusAllocated	OCPU Allocated	Integer	Max	One minute	The number of OCPUs allocated.

Related Topics

- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.
- [Incident Logs and Trace Files](#)
This section lists all of the files that can be collected by Oracle Support if you opt-in for incident logs and trace collection.
- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.

Using the API to Manage Exadata Cloud Infrastructure Instance

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

Use these API operations to manage Exadata Cloud Infrastructure instance components.

Cloud Exadata infrastructure resource (new resource model):

- [ListCloudExadataInfrastructures](#)
- [GetCloudExadataInfrastructure](#)
- [ChangeCloudExadataInfrastructureCompartment](#)
- [UpdateCloudExadataInfrastructure](#)
- [DeleteCloudExadataInfrastructure](#)

Cloud VM cluster (new resource model)

- [ListCloudVmClusters](#)
- [GetCloudVmCluster](#)
- [ChangeCloudVmClusterCompartment](#)
- [UpdateCloudVmCluster](#)
- [DeleteCloudVmCluster](#)

DB systems (old resource model):

- [ListDbSystems](#)
- [GetDbSystem](#)
- [ChangeDbSystemCompartment](#)
- [UpdateDbSystem](#)
- [TerminateDbSystem](#)

Virtual machines nodes (all Exadata Cloud Infrastructure instances):

- [DbNodeAction](#)
- [ListDbNodes](#)
- [GetDbNode](#)

Troubleshooting Virtual Machines Using Console Connections

You can troubleshoot malfunctioning virtual machines using console connections. For example, a previously working Guest VM stops responding.

Note

Exadata System Software 23.1.13 is the minimum required version. Also, make sure to review all prerequisites stated below, including setting a password for either the `opc` or the `root` user. Failure to make necessary changes to meet these requirements in advance will result in the inability to urgently connect to the serial console when the need arises when the VM is not otherwise accessible.

To connect to a running instance for administration and general use, use a Secure Shell (SSH). For more information, see [Connecting to a Virtual Machine with SSH](#).

To make an SSH connection to the serial console, follow these configuration steps.

1. Ensure that you have the correct permissions.
2. Complete the prerequisites, including creating your SSH key pair (in case you don't have one yet).
3. Create the Virtual Machine Serial Console.
4. Connect to the serial console via SSH.

To check the DB server version installed, follow these steps:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Click **Exadata Infrastructure** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. Click the name of the infrastructure that you are interested in.
5. In the resulting Infrastructure Details page, go to the Version section to find the DB Server version installed.

- [Required IAM Policies](#)
An administrator must grant you secure access to the virtual machine console on the Oracle Exadata Database Service on Dedicated Infrastructure system through an IAM policy.
- [Prerequisites](#)
You must install an SSH client and create SSH key pairs.
- [Create the Virtual Machine Serial Console Connection](#)
Before you can make a local connection to the serial console, you need to create the virtual machine console connection.
- [Make an SSH Connection to the Serial Console](#)
- [Using Cloud Shell to Connect to the Serial Console](#)
- [Displaying the Console History for a Virtual Machine](#)
- [Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems](#)
- [Exiting the Virtual Machine Serial Console Connection](#)

Required IAM Policies

An administrator must grant you secure access to the virtual machine console on the Oracle Exadata Database Service on Dedicated Infrastructure system through an IAM policy.

This access is required whether you're using the Console or the REST API with an SDK, CLI, or other tools. If you get a message that you don't have permission or are unauthorized, verify with your administrator what type of access you have and which compartment to work in.

To create virtual machine console connections, an administrator needs to grant user access to read and manage virtual machine console connections through an IAM policy. The resource name for virtual machine console connections is `dbnode-console-connection`. The resource name for the virtual machine is `db-nodes`. The following policies grant users the ability to create virtual machine console connections:

```
Allow group <group_name> to manage dbnode-console-connection in tenancy
Allow group <group_name> to read db-nodes in tenancy
```

Prerequisites

You must install an SSH client and create SSH key pairs.

Install an SSH Client and a Command-line Shell (Microsoft Windows)

Microsoft Windows does not include an SSH client by default. If you are connecting from a Windows client, you need to install an SSH client. You can use [PuTTY](#) `plink.exe` with [Windows PowerShell](#) or software that includes a version of OpenSSH such as:

- [Git for Windows](#)
- [Windows Subsystem for Linux](#)

The instructions in this topic frequently use PuTTY and Windows PowerShell.

If you want to make the console connection from Windows with [Windows PowerShell](#), PowerShell might already be installed on your Windows operating system. If not, follow the steps at the link. If you are connecting to the instance from a Windows client using PowerShell, `plink.exe` is required. `plink.exe` is the command link connection tool included with PuTTY. You

can install PuTTY or install plink.exe separately. For installation information, see <http://www.putty.org>.

Create SSH Key Pairs

To create the secure console connection, you need an SSH key pair. The method to use for creating key pairs depends on your operating system. When connecting to the serial console, you must use an RSA key. The instructions in this section show how to create an RSA SSH key pair.

Create the SSH key Pair for Linux

If you're using a UNIX-style system, you probably already have the `ssh-keygen` utility installed. To determine whether the utility is installed, type `ssh-keygen` on the command line. If the utility isn't installed, you can download OpenSSH for UNIX from <http://www.openssh.com/portable.html> and install it.

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.

Alternatively, you can type a complete `ssh-keygen` command, for example:

```
ssh-keygen -t rsa -N "" -b 2048 -C "<key_name>" -f <path/root_name>
```

Argument	Description
<code>-t rsa</code>	Use the RSA algorithm.
<code>-N "<passphrase>"</code>	A passphrase to protect the use of the key (like a password). If you don't want to set a passphrase, don't enter anything between the quotes. A passphrase is not required. You can specify one as a security measure to protect the private key from unauthorized use. If you specify a passphrase, when you connect to the instance you must provide the passphrase, which typically makes it harder to automate connecting to an instance.
<code>-b 2048</code>	Generate a 2048-bit key. You don't have to set this if 2048 is acceptable, as 2048 is the default. A minimum of 2048 bits is recommended for SSH-2 RSA.
<code>-C "<key_name>"</code>	A name to identify the key.
<code>-f <path/root_name></code>	The location where the key pair will be saved and the root name for the files.

Create the SSH Key Pair for Windows Using PuTTY

If you are using a Windows client to connect to the instance console connection, use an SSH key pair generated by PuTTY.

Note

Ensure that you are using the latest version of PuTTY, see <http://www.putty.org>.

1. Find `puttygen.exe` in the PuTTY folder 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 not already set.
3. Click **Generate**.
4. To generate random data in the key, move your mouse around the blank area in the PuTTY window.
When the key is generated, it appears under **Public key for pasting into OpenSSH authorized_keys file**.
5. A **Key comment** is generated for you, including the date and timestamp. You can keep the default comment or replace it with your own more descriptive comment.
6. Leave the **Key passphrase** field blank.
7. Click **Save private key**, and then click **Yes** in the prompt about saving the key without a passphrase.
The key pair is saved in the PuTTY Private Key (PPK) format, which is a proprietary format that works only with the PuTTY tool set.

You can name the key anything you want but use the `ppk` file extension. For example, `mykey.ppk`.
8. Select all of the generated keys that appear 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.

Note

Do not use the **Save public key** option 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. Write down the names and locations of your public and private key files. You need the public key when creating an instance console connection. You need the private key to connect to the instance console connection using PuTTY. For example: `$HOME\Documents\mykey.ppk`.

To create a connection using the SSH key pair generated using PuTTY

For more information about generating SSH key pairs, see [Create the SSH Key Pair for Windows Using PuTTY](#)

Do the following on the Create serial console access window:

1. Paste the SSH Key generated from OpenSSH format or choose **Upload SSH key file** and provide the path of the public key saved at step 8 in [Create the SSH Key Pair for Windows Using PuTTY](#).
2. Once the connection is **Active**, click **Copy serial console connection for Windows**.
3. Paste the connection string copied from the previous step into a text file.

4. In the text file, replace `<PATH_FILE_PUTTY_PRIVATE.ppk>` to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Documents\mykey.ppk`.
5. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

Sign in to a Virtual Machine From the Serial Console

If you want to sign in to a virtual machine using a virtual machine console connection, you can use Secure Shell (SSH) connection to sign in. If you want to sign in with a username and password, you need a user account with a password. Oracle Exadata Cloud does not set a default password for the `opc` or `root` users. Therefore, if you want to sign in as the `opc` or `root` user, you need to create a password for the `opc` or `root` user. Otherwise, add a different user with a password and sign in as that user. This should be completed in advance, before a potential situation that might require you to log in to the serial console.

Connect Through Firewalls

If the client you will use to access the serial console is behind a firewall, you must ensure that this client can reach the required endpoint to access the serial console of the virtual machine. The client system connecting to the serial console must be able to reach the serial console server (for example, `vm-console-ad1.exacs.us-ashburn-1.oci.oraclecloud.com`) over SSH using port 443, directly or through a proxy.

Create the Virtual Machine Serial Console Connection

Before you can make a local connection to the serial console, you need to create the virtual machine console connection.

Virtual machine 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 with the following message:

```
channel 0: open failed: administratively prohibited: console access is
limited to one connection at a time
```

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connections** is selected by default.
4. Click **Create console connection**.
5. In the resulting Create serial console access window, you have three options for adding the SSH key
 - **Generate a key pair for me:** You can have Oracle Cloud Infrastructure generate an SSH key pair to use. If you are using PowerShell or PuTTY to connect to the instance from a Windows client, you cannot use the generated SSH key pair without first converting it to a `.ppk` file.

- **Upload public key file:** Browse to a public key file on your computer. If you followed the steps in *Creating SSH Key Pairs* in the *Prerequisites* section to create a key pair, use this option to navigate to the `.pub` file.
 - **Paste public key:** Paste the content of your public key file into the text box.
6. Click **Create console connection**.
When the console connection has been created and is available, the state changes to **Active**.

Make an SSH Connection to the Serial Console

After you create the console connection for the virtual machine, you can connect to the serial console using a Secure Shell (SSH) connection. When making an SSH connection to the serial console, you must use an RSA key. You can use the same SSH key for the serial console that was used when you launched the instance, or you can use a different SSH key.

When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must reauthenticate to connect again.

Validate Server Host Keys

When you first connect to the serial console, you're prompted to validate the fingerprint of the server host key. The fingerprint of the server host key is the SHA256 hash of the server host's public SSH key. The server SSH handshake response is signed with the associated private key. Validating the server host key's fingerprint protects against potential attacks.

When you make a manual connection to the serial console, the fingerprint of the server host key is not automatically validated. To manually validate the fingerprint, compare the fingerprint value displayed in the Oracle Cloud Infrastructure Console to the value of the RSA key fingerprint that appears in the terminal when you connect.

To find the fingerprint of the server host key in the Console, on the Virtual Machine details page, under **Resources**, click **Console connection**. The table displays the fingerprint of the server host key. The fingerprint in the Console should match the value of the **RSA key fingerprint** shown in the terminal when you connect to the serial console.

The server host keys are periodically rotated for security purposes. Key rotation reduces the risk posed when keys are compromised by limiting the amount of data encrypted or signed by one key version. When your key is rotated and you try to connect to the serial console, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete that line in your `.ssh/known_hosts` file and then reconnect to the serial console. You are then prompted to accept a new server host key fingerprint.

Connect from Mac OS X and Linux Operating Systems

Use an SSH client to connect to the serial console. Mac OS X and most Linux and UNIX-like operating systems include the SSH client OpenSSH by default.

To connect to the serial console using OpenSSH on Mac OS X or Linux:

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the VM Cluster that you're interested in.

3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
4. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
5. Click the Actions menu (three dots), and then click **Copy serial console connection for Linux/Mac**.
6. Paste the connection string into a terminal window on a Mac OS X or Linux system, and then press Enter to connect to the console.
If you are not using the default SSH key or ssh-agent, modify the serial console connection string to include the identity file flag, `-i` to specify the private key portion for the SSH key to use, for example, `id_rsa`. Specify this flag for both the SSH connection and the SSH Proxy Command, as shown in the following line:

```
ssh -i /<path>/<ssh_key> -o ProxyCommand='ssh -i /<path>/<ssh_key> -W
%h:%p -p 443...
```

7. If prompted, validate and accept the fingerprint of the server host key.
If you have previously accepted a fingerprint for the server host key but the key has been rotated, a warning appears indicating a potential attack. The warning includes an `Host key verification failed` error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
8. Press **Enter** again to activate the console.
If the connection is active, a message appears in the console:

```
=====
IMPORTANT: You are now connected to the serial console for this VM. This should be
used in emergency situations only.
See product documentation for more details and alternative connectivity options for
normal operations
=====
```

9. Reboot your virtual machine.
You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

For more troubleshooting options, see *Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems*.
 - a. Go to the ExaDB-D VM Cluster Details page.
 - b. Under **Resources**, click **Virtual Machines**.
 - c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

Connect from Windows Operating Systems

The steps to connect to the serial console from Microsoft Windows PowerShell are different from the steps for OpenSSH. The following steps do not work in the Windows terminal.

If you are connecting to the instance from a Windows client using PowerShell, `plink.exe` is required. `plink.exe` is the command link connection tool included with PuTTY. You can install PuTTY or install `plink.exe` separately. For more information, see *Installing an SSH Client and a Command-line Shell (Windows)*.

To connect to the serial console on Microsoft Windows:

1. On the Virtual Machine details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
2. Click the Actions menu (three dots).
Depending on which SSH client you are using, do one of the following:
 - If you are using Windows PowerShell, click **Copy serial console connection for Windows**.
 - If you are using OpenSSH, click **Copy serial console connection for Linux/Mac**.

Note

The copied connection string for Windows contains the parameter `-i` specifying the location of the private key file. The default value for this parameter in the connection string references an environment variable that might not be configured on your Windows client, or it might not represent the location where the private key file is saved. Verify the value specified for the `-i` parameter and make any required changes before proceeding to the next step.

3. Paste the connection string copied from the previous step into a text file so that you can add the file path to the private key file.
4. In the text file, replace `$env:homedrive$env:homepath\oci\console.ppk` with the file path to the `.ppk` file on your computer. This file path appears twice in the string. Replace it in both locations.

Note

For PuTTY versions 0.82 and above, add parameter `-legacy-stdio-prompts` to the first call to `plink`.

5. Paste the modified connection string into the PowerShell window or your OpenSSH client, and then press **Enter** to connect to the console.
6. If prompted, validate and accept the fingerprint of the server host key.
If you have previously accepted a fingerprint for the server host key, but the key has been rotated, a warning appears indicating a potential attack. The warning includes a Host key verification failed error and a line number in your `.ssh/known_hosts` file. Delete the specified line in your `.ssh/known_hosts` file and then reconnect to the serial console. Validate and accept the new server host key fingerprint.
7. Press **Enter** again to activate the console.
8. Reboot your virtual machine.
You do not need to enter a user name or password. If the Virtual Machine is functional and the connection is active, the serial output appears in your console. If the serial output does not appear in the console, the Guest VM operating system is not booting.

For more troubleshooting options, see [Troubleshooting Virtual Machines from Guest VM Console Connections](#).

- a. Go to the ExaDB-D VM Cluster Details page.
- b. Under **Resources**, click **Virtual Machines**.
- c. Select **Reboot** from the Actions menu (three dots) for the virtual machine that you want to reboot.

To create a connection using the SSH key pair generated using the OCI Console:

Do the following on the Create serial console access window:

1. Click **Generate a key pair for me**.
2. Click **Save Private Key**.
3. Click **Create console connection**.

Note

Ensure that you are using the latest version of PuTTY, see <http://www.putty.org>.

4. Find `puttygen.exe` in the PuTTY folder on your computer, for example, `C:\Program Files (x86)\PuTTY`. Double-click `puttygen.exe` to open it.
5. On the PuTTY Key Generator, click the **Conversions** menu and then click **Import**.
6. On the Windows Explorer, select OCI Console generated SSH key (step 1) and then click **Open**.
PuTTY imports the key and displays information about the key on the PuTTY Key Generator window.
7. Click **Save private key**.
8. Click **Yes** when prompted about saving the key without a passphrase.
The key pair is saved in the PuTTY Private Key (PPK) format, which is the proprietary format that works only with the PuTTY tool set.

You can name the key anything you want but use the `.ppk` file extension. For example, `$HOME\Desktop\key-vm-console.ppk`.
9. Use a text editor to change the command to point to your PuTTY Private Key (PPK) path. Replace `<PATH_FILE_PUTTY_PRIVATE.ppk>` to point to your PuTTY Private Key (PPK) file path on your computer. For example, if you have saved `.ppk` file at `$HOME\Desktop\key-vm-console.ppk`.
10. Paste the modified connection string into the PowerShell window, and then press **Enter** to connect to the console.

To convert a generated .key private key file:

1. Open PuTTYgen.
2. Click **Load**, and select the private key generated when you created the instance.
The extension for the key file is `.key`.
3. Click **Save private key**.
4. Specify a name for the key.
The extension for the new private key is `.ppk`.
5. Click **Save**.

Using Cloud Shell to Connect to the Serial Console

You can connect to the serial console quickly and easily using the Cloud Shell integration. Cloud Shell is a web browser-based terminal accessible from the Console. The Cloud Shell integration automatically creates the instance console connection and a temporary SSH key. The only prerequisite for connecting to the serial console from Cloud Shell is granting users the correct permissions. For an introductory walkthrough of using Cloud Shell, see Using Cloud Shell.

Note

- By default, Cloud Shell limits network access to OCI internal resources in your tenancy home region only unless you have enabled the [Cloud Shell managed Public Network](#). Your administrator must configure an Identity policy to enable Cloud Shell Public Network. For more information, see [Cloud Shell Networking](#).
- You cannot concurrently connect to more than one DB node using Cloud Shell. As an example, if you have an open connection to *DBnode1* and want to connect to *DBnode2*, you must first exit the active Cloud Shell from *DBnode1* and then establish a connection to *DBnode2*.

When you are finished with the serial console and have terminated the SSH connection, you should delete the serial console connection. If you do not disconnect from the session, Oracle Cloud Infrastructure terminates the serial console session after 24 hours and you must re-authenticate to connect again.

- [To connect to the serial console using Cloud Shell](#)

Related Topics

- [Cloud Shell](#)
- [Using Cloud Shell](#)
- [Required IAM Policies](#)

To connect to the serial console using Cloud Shell

1. Sign in to the Console.
2. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. On the instance details page in the Oracle Cloud Infrastructure Console, under **Resources**, click **Console connection**.
4. Click **Launch Cloud Shell connection**.
This action displays the Cloud Shell in a "drawer" at the bottom of the Console.
5. If a console connection already exists, you are asked if you want to delete the existing resource. Press **y**, and then press **Enter**.
6. When you are done, exit the instance console connection.

Related Topics

- [Exiting the Virtual Machine Serial Console Connection](#)

Displaying the Console History for a Virtual Machine

You can capture and display recent serial console data for a Virtual Machine. The data includes configuration messages that occur when the Virtual Machine boots, such as kernel and BIOS messages, and is useful for checking the status of the Virtual Machine or diagnosing and troubleshooting problems.

The console history captures up to a megabyte of the most recent serial console data for the specified Virtual Machine. Note that the raw console data, including multi-byte characters, is captured.

The console history is a point-in-time record. To troubleshoot a malfunctioning Virtual Machine using an interactive console connection, use a serial console connection.

- [Managing Console History Data](#)

Managing Console History Data

You can use the Console or API to manage console history captures. Console history lets you see serial output from your Virtual Machine without having to connect to the instance remotely. The console history can be used to audit previous access and actions taken with the serial console.

On the instance details page in the Console, you can capture and download console histories, view and edit metadata details, and delete console history captures.

- [Using the Console to Capture the Console History](#)
- [Using the Console to Download Console History Captures](#)
- [Using the Console to View Console History Captures](#)
- [Using the Console to View and Edit the Metadata Details of a Console History Capture](#)
- [Using the Console to Delete Console History Captures](#)
- [Using the API to Manage the Console History Data](#)

Using the Console to Capture the Console History

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the resulting window, click **Download** to download a copy of the console history.
7. Click **Save and close** to save the history and close the window.

Using the Console to Download Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the resulting window, click **Download** to download a copy of the console history.

Using the Console to View Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. Click the name of the history that you're interested in.
6. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **View details**.

Using the Console to View and Edit the Metadata Details of a Console History Capture

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **View details**.
6. Optionally, edit the name for the console history. Avoid entering confidential information.
7. To view or edit tags, click **Show tagging options**.
8. To edit or remove tags, click the edit icon next to the tag. To edit a tag, in the **Edit Tag** dialog, make any changes, and then click **Save**. To remove a tag, click **Remove Tag**.
9. Click **Save and close**.

Using the Console to Delete Console History Captures

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the **VM Cluster** that you're interested in.
3. On the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click **Console history**.
5. On the console history list, for the console history capture that you want to view, click the **Actions** menu, and then click **Delete**.
6. On the confirmation dialog, click **Delete console history**.

Using the API to Manage the Console History Data

Review the list of API calls to manage console history data.

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

For the complete list of APIs, see [Database Service API](#).

Use the following API operations to manage the console history data.

- To capture the console history, use the **createDbNodeConsoleHistory** method.
- To get details of console history metadata, use the **getDbNodeConsoleHistory** method.
- To get the details of console history content, use the **getDbNodeConsoleHistoryContent** method.
- To edit console history metadata, use the **updateDbNodeConsoleHistory** method.
- To list console history captures, use the **listDbNodeConsoleHistories** method.
- To delete console history captures, use the **deleteDbNodeConsoleHistory** method.

Troubleshooting Virtual Machines from Guest VM Console Connections on Linux Operating Systems

After you are connected with an instance console connection, you can perform various tasks, such as:

- Edit system configuration files.
- Add or reset the SSH keys for the `opc` user.
- Reset the password for the `opc` user.

These tasks require you to boot into a Bash shell in maintenance mode.

To boot into maintenance mode

Note

Default user and password:

- **Account:** Grub boot loader
- **Username:** root
- **Default Password:** sos1Exadata
- **Account Type:** Operating system user

For more information, see *Default User Accounts for Oracle Exadata*.

1. Reboot the VM from the VM Cluster.
2. For virtual machines running Oracle Linux 7.x or Oracle Linux 8.x, when the reboot process starts, switch back to the terminal window, and you see Console messages start to appear in the window. As soon as the **GRUB boot menu** appears, use the **up/down arrow key** to stop the automatic boot process, enabling you to use the boot menu.
3. In the boot menu, highlight the top item in the menu, and press **e** to edit the boot entry.
4. In edit mode, use the **down arrow key** to scroll down through the entries until you reach the line that starts with **linux16**.

5. At the end of that line, add the following:

```
init=/bin/bash
```

6. Reboot the instance from the terminal window by entering the keyboard shortcut **CTRL+X**. When the instance has rebooted, you see the Bash shell command-line prompt, and you can proceed with the following procedures.

To edit the system configuration files

1. From the Bash shell, run the following command to load the SELinux policies to preserve the context of the files you are modifying:

```
/usr/sbin/load_policy -i
```

2. Run the following command to remount the root partition with read/write permissions:

```
/bin/mount -o remount, rw /
```

3. Edit the configuration files as needed to try to recover the instance.
4. After you have finished editing the configuration files, to start the instance from the existing shell, run the following command:

```
exec /usr/lib/systemd/systemd
```

Alternatively, to reboot the instance, run the following command:

```
/usr/sbin/reboot -f
```

To add or reset the SSH key for the opc user

1. From the Bash shell, run the following command to load the SELinux policies to preserve the context of the files you are modifying:

```
/usr/sbin/load_policy -i
```

2. Run the following command to remount the root partition with read/write permissions:

```
/bin/mount -o remount, rw /
```

3. From the Bash shell, run the following command to change to the SSH key directory for the `opc` user:
`cd ~opc/.ssh`
4. Include your public key entry to the `authorized_keys` file.

Note

You can edit the file and remove your previous key if you want to. However, make sure to keep the cloud automation keys to prevent cloud automation from breaking.

```
echo '<contents of public key file>' >> authorized_keys
```

5. Restart the instance by running the following command:

```
/usr/sbin/reboot -f
```

To reset the password for the `opc` user

1. From the Bash shell, run the following command to load the SELinux policies to preserve the context of the files you are modifying.
This step is necessary to sign in to your instance using SSH and the Console.

```
/usr/sbin/load_policy -i
```

2. Run the following command to remount the root partition with read/write permissions:

```
/bin/mount -o remount, rw /
```

3. Run the following command to reset the password for the `opc` user:

```
sudo passwd opc
```

4. Restart the instance by running the following command:

```
sudo reboot -f
```

Note

Setting a `root` password would be an acceptable alternative to setting an `opc` password.

Exiting the Virtual Machine Serial Console Connection

To exit the serial console connection

When using SSH, the `~` character at the beginning of a new line is used as an escape character.

1. To exit the serial console, enter:

```
~.
```

2. To suspend the SSH session, enter:

```
~^z
```

The `^` character represents the **CTRL** key.

3. To see all the SSH escape commands, enter:

```
~?
```

To delete the serial console connection for a Virtual Machine

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Click the VM Cluster that you're interested in.
3. In the resulting VM Cluster Details page, click the name of the virtual machine that you're interested in.
Under **Resources**, **Console connection** is selected by default.
4. Click the Actions menu, and then click **Delete**. Confirm when prompted.

Manage Software Images

- [Using Software Images in Oracle Cloud Infrastructure](#)
- [Using a Software Image with an Exadata Cloud Infrastructure Instance](#)
Create, save, and reuse a Software Image.
- [Using the Console for Software Images](#)
- [Using the API to manage database software images](#)
Use these API operations to manage database software images:

Using Software Images in Oracle Cloud Infrastructure

- [Creation and Storage of Software Images](#)
Software images are resources within your tenancy that you create before provisioning or updating a VM cluster, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.
- [Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home](#)
OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/opatch` directory.

Creation and Storage of Software Images

Software images are resources within your tenancy that you create before provisioning or updating a VM cluster, Exadata Cloud Infrastructure instance, Database Home, database, or Grid Infrastructure.

There are two types of software image resources:

- **Grid Infrastructure software image:** Grid Infrastructure software images are resources containing Oracle Grid Infrastructure software used to update Oracle Grid Infrastructure. Grid Infrastructure software images are either Oracle-published software releases or custom software images created by the customer that include the desired Grid Infrastructure release updates (GIRU) and additional one-off (interim) patches.
- **Database software image:** Database software images are resources containing Oracle Database software used to provision and update Oracle Databases and Oracle Database Homes. Database software images are either Oracle-published software releases or custom software images created by the customer that include the desired Database release updates (DBRU) and additional one-off (interim) patches.

There is no limit on the number of software images you can create in your tenancy, and you can create your images with any Oracle Database software or Oracle Grid Infrastructure version and update supported in Oracle Cloud Infrastructure.

Software images are automatically stored in Oracle-managed Object Storage and can be viewed and managed in the Oracle Cloud Infrastructure Console. Software images are regional-level resources but they can be accessed from any region within your tenancy.

Note: The software images incur Object Storage usage costs.

Using the OPatch Lsinventory Command to Verify the Patches Applied to an Oracle Home

OPatch utility enables you to apply the interim patches to Oracle Database Home or Oracle Grid Infrastructure Home. You can find the `opatch` utility in the `$ORACLE_HOME/OPatch` directory.

Using the `lsinventory` command provided by OPatch, you can create a file that lists the interim patches applied to an Oracle Database Home or Oracle Grid Infrastructure Home. This file can then be uploaded to the OCI Console during the creation of a custom software image to add the exact set of patches used by the source Oracle Database Home or Oracle Grid Infrastructure Home to the list of patches included in the software image. You can find the `opatch` utility in the `$ORACLE_HOME/OPatch` directory. The following example shows how to use the `lsinventory` command to create the `lsinventory` file:

1. Run the `opatch lsinventory` command to get the list of interim patches applied.

```
$ORACLE_HOME/OPatch/opatch lsinventory
Oracle Interim Patch Installer version 12.2.0.1.21
Copyright (c) 2021, Oracle Corporation. All rights reserved.

Oracle Home      : /u02/app/oracle/product/19.0.0.0/dbhome_2
Central Inventory : /u01/app/oraInventory
from             : /u02/app/oracle/product/19.0.0.0/dbhome_2/oraInst.loc
OPatch version   : 12.2.0.1.21
OUI version      : 12.2.0.7.0
Log file location : /u02/app/oracle/product/19.0.0.0/dbhome_2/cfgtoollogs/
opatch/opatch2021-01-21_09-22-45AM_1.log

Lsinventory Output file location : /u02/app/oracle/product/19.0.0.0/
dbhome_2/cfgtoollogs/opatch/lsinv/lsinventory2021-01-21_09-22-45AM.txt

Oracle Interim Patch Installer version 12.2.0.1.41
Copyright (c) 2024, Oracle Corporation. All rights reserved.

Oracle Home      : /u01/app/oracle/product/19.0.0.0/gridhome_1
Central Inventory : /u01/app/oraInventory
from             : /u01/app/oracle/product/19.0.0.0/gridhome_1/oraInst.loc
OPatch version   : 12.2.0.1.41
OUI version      : 12.2.0.7.0
Log file location : /u01/app/oracle/product/19.0.0.0/gridhome_1/
cfgtoollogs/opatch/opatch2024-04-19_19-24-22PM_1.log
```

```
Lsinventory Output file location : /u01/app/oracle/product/19.0.0.0/  
gridhome_1/cfgtoollogs/opath/lsinv/lsinventory2024-04-19_19-24-22PM.txt
```

2. Use the `lsinventory` output file to extract the additional interim patches applied to a specific Oracle Database Home or Oracle Grid Infrastructure Home.

Using a Software Image with an Exadata Cloud Infrastructure Instance

Create, save, and reuse a Software Image.

Creating a Software Image enables you to:

- Create custom Database and Grid Infrastructure images based on Software Images, RU, and one-off (interim) patches.
- Save a custom image automatically to Object Storage as a resource.
- Provision an Oracle Database home or Oracle Database with the desired RU and one-off (interim) patches.
- Update the Database Home and Grid Infrastructure using the Software Image.
- Clone Software Image to another service in the Data Guard creation process.

Note

The Software Images are created and managed by the customer and they are available for use until explicitly deleted.

Using the Console for Software Images

- [To view the list of software images](#)
- [To create a database software image](#)
- [To create a Grid Infrastructure software image](#)
- [To create a database software image from a Database Home](#)
- [To view the image details of a software image](#)
- [To move a software image to a different compartment](#)
- [To update database software using custom database software image](#)
Use the following instructions to update database software using a custom database software image.
- [To update Grid Infrastructure software using custom Grid Infrastructure software image](#)
Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.
- [To delete a software image](#)
Use the following instructions to delete a software image.

To view the list of software images

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Under **Resources**, click **Software images**.

The resulting Software images page displays the list of custom software images, which includes details such as Image type (Database, Grid infrastructure), and version.

To create a database software image

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Database Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Database release**.
8. Choose the **Database version** for your image. You can create a database software image using any supported Oracle Database release update (RU).
9. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
10. Optionally, you can upload an Oracle Home inventory file from an existing Oracle Database. See *Using the OPatch Isinventory Command to Verify the Patches Applied to an Oracle Home* for instructions on creating an inventory file using OPatch.
11. Click **Show Advanced Options** to add tags to your database software image. 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 if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
12. Click **Create software image**.

To create a Grid Infrastructure software image

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. Click **Create Software Image**.
4. In the resulting Create software image page, click **Grid Infrastructure Software image**.
5. In the **Display name** field, provide a display name for your image. Avoid entering confidential information.
6. Choose your **Compartment**.
7. Choose a **Grid Infrastructure release**.
8. Choose the **Grid Infrastructure** version for your image. You can create a Grid Infrastructure software image using any supported Oracle Grid Infrastructure release update (RU).
9. Optionally, you can enter a comma-separated list of one-off (interim) patch numbers.
10. Optionally, you can upload a Grid Infrastructure Home inventory file from an existing Oracle Grid Infrastructure. See *Using the OPatch Isinventory Command to Verify the*

Patches Applied to an Oracle Home for instructions on creating an inventory file using OPatch.

11. Click **Show Advanced Options** to add tags to your database software image. 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 if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
12. Click **Create software image**.

To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Database Homes** under **Resources**.
5. Find the Database Home you want to use to create the database software image in the list of Database Homes. Click the name of the Database Home to display details about it.
6. Click **Create Image from Database Home**.
7. In the **Create Database Software Image** panel, enter a **Display name** and select a compartment for the software image.
8. Click **Create**.

To view the image details of a software image

Use this procedure to view the image details such as Image type, Oracle release, and version used for creating the software image and one-off (interim) patches included (if any) in a software image.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to view and click on the display name of the image.
4. On the Software Image details page for your selected image, details about the image are displayed:
 - The **General Information** section includes details such as Image type.
 - The **Patch Information** section includes details such as release, version, and available interim patches.
 - The **One-Off Patches** field displays the number of one-off patches included (if any) in a software image. The count includes all patches specified when creating the image (including patches listed in `lsinventory`).
 - To view the included patches (if any are included),
 - * click the **Show** link to view the list in the One-Off Patches overlay window.

- * click the **Copy** link and paste the list of included patches into a text editor. The copied list of patch numbers is comma-separated and can be used to create additional database software images.

To move a software image to a different compartment

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software images**.
3. In the list of software images, find the image you want to move and click the Actions icon (three dots) at the end of the row.
4. Click **Move resource**.
5. On the resulting Move resource to a different compartment dialog, choose a target compartment.
6. Click **Move resource**.

To update database software using custom database software image

Use the following instructions to update database software using a custom database software image.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster that you want to update the database software image.
4. Click **Database homes**.
5. Click the name of the Database Home that you want to update.
6. Click **Updates**.
7. Choose a **Compartment**.
8. Click **Type**, select **Custom**, and then click **Apply Filter**.
9. Choose a **Region**.
Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
10. Click the Actions button (three dots) for the update you're interested in, and select **Precheck**.
11. On the resulting Confirm precheck dialog, click **Precheck** to continue.
12. After running the precheck successfully, select **Apply Database Home update** from the Actions button (three dots).
13. On the resulting Confirm apply dialog, click **Apply** to continue.

To update Grid Infrastructure software using custom Grid Infrastructure software image

Use the following instructions to update Grid Infrastructure software using a custom Grid Infrastructure software image.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster that you want to update the Grid Infrastructure software image.
4. On the resulting VM cluster details page, click **Updates (GI)**.
5. Choose a **Compartment**.
6. Click **Type**, select **Custom**, and then click **Apply Filter**.
7. Choose a **Region**.
Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
8. Click the Actions button (three dots) for the update you're interested in, and select **Precheck**.
9. On the resulting Confirm precheck dialog, click **Precheck** to continue.
10. After running the precheck successfully, select **Apply Grid Infrastructure update** from the Actions button (three dots).
11. On the resulting Apply Cloud VM Cluster Patch dialog, click **Apply** to continue.

To delete a software image

Use the following instructions to delete a software image.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Software Images**.
3. In the list of software images, find the image you want to delete and click the Actions icon (three dots) at the end of the row.
4. Click **Delete**.
5. In the resulting Delete software image dialog, enter the name of the software image to confirm your action.
6. Click **Delete**.

Using the API to manage database software images

Use these API operations to manage database software images:

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

- [CreateDatabaseSoftwareImage](#)
- [ListDatabaseSoftwareImages](#)
- [GetDatabaseSoftwareImage](#)
- [DeleteDatabaseSoftwareImage](#)
- [ChangeDatabaseSoftwareImageCompartment](#)

Create Oracle Database Homes on an Exadata Cloud Infrastructure System

Learn to create Oracle Database Homes on Exadata Cloud Infrastructure.

- [About Creating Oracle Database Homes on an Exadata Cloud Infrastructure System](#)
You can add Oracle Database homes (referred to as **Database Homes** in Oracle Cloud Infrastructure) to an existing VM cluster by using the Oracle Cloud Infrastructure Console, the API, or the CLI.
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- [To create a database software image from a Database Home](#)
- [Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure](#)
To create an Oracle Database home, review the list of API calls.

About Creating Oracle Database Homes on an Exadata Cloud Infrastructure System

You can add Oracle Database homes (referred to as **Database Homes** in Oracle Cloud Infrastructure) to an existing VM cluster by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

A Database Home is a directory location on the Exadata database virtual machines that contains Oracle Database software binary files.

Note

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

You can also add and remove Database homes, and perform other management tasks on a Database home by using the `dbaascli` utility.

Related Topics

- [Using the dbaascli Utility on Exadata Cloud Infrastructure](#)
Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.

To create a new Database Home in an existing Exadata Cloud Infrastructure instance

To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.

Note

Minimum requirements for provisioning a Database 26ai home:

- Grid Infrastructure 26ai
- Exadata Guest VM running Exadata System Software 23.1.8

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to create the new Database Home on:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Database Homes**.
A list of Database Homes is displayed.
5. Click **Create Database Home**.
6. In the **Create Database Home** dialog, enter the following:
 - **Database Home display name**: The display name for the Database Home. Avoid entering confidential information.
 - **Database image**: Determines what Oracle Database version is used for the database. You can have databases with different minor versions the same database home. The major versions must remain the same. By default, the latest Oracle-published database software image is selected.

Click **Change Database Image** to use a desired Oracle-published image or a custom [database software image](#) that you have created in advance, then select an **Image Type**:

 - **Oracle Provided Database Software Images**: These images contain generally available versions of Oracle Database software.
 - **Custom Database Software Images**: These images are [created by your organization](#) and contain customized configurations of software updates and patches. Use the **Select a compartment**, **Select a region**, and **Select a Database version** selectors to limit the list of custom database software images to a specific compartment, region, or Oracle Database software major release version.

Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.

Note

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

Note

The custom database software image must be based on an Oracle Database release that meets the following criteria:

- * The release is currently supported by Oracle Cloud Infrastructure.
- * The release is supported by the hardware model on which you are creating the Database Home.

After choosing a software image, click **Select** to return to the Create Database dialog.

- **Unified Auditing:** Select this check box to enable Unified Auditing framework.

Note

You cannot disable Unified Auditing after provisioning the Database Home.

- **For Oracle Database versions lower than 12.1:** You cannot use the Unified Auditing framework and should instead use the Traditional Audit, the legacy Oracle Database audit framework.
- **For Oracle Database versions 12.1 or higher:** You can enable Unified Auditing from the OCI Console. For Oracle Database versions 12.1 or higher but lower than version 23, the **Unified Auditing** check box is not selected by default. However, it is selected by default for Oracle Database version 23.

Unified Auditing field in the **General Information** section on the **Database Home Details** page displays if Unified Auditing is **Enabled** or **Disabled**.

- Click **Show Advanced Options** to specify advanced options for the Database Home.
 - **Tags:** 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.

7. Click **Create**.

When the Database Home creation is complete, the status changes from Provisioning to Available.

To create a database software image from a Database Home

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the Database Home: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Database Homes** under **Resources**.
5. Find the Database Home you want to use to create the database software image in the list of Database Homes. Click the name of the Database Home to display details about it.
6. Click **Create Image from Database Home**.
7. In the **Create Database Software Image** panel, enter a **Display name** and select a compartment for the software image.
8. Click **Create**.

Using the API to Create Oracle Database Home on Exadata Cloud Infrastructure

To create an Oracle Database home, review the list of API calls.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

To create Database Homes in Exadata Cloud Infrastructure, use the API operation `CreateDbHome`.

For the complete list of APIs, see "Database Service API".

Related Topics

- [REST APIs](#)
- [Security Credentials](#)
- [Software Development Kits and Command Line Interface](#)
- [CreateDbHome](#)
- [Database Service API](#)

Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance

You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

For information on how to perform these tasks manually, see [About Using the dbaascli Utility on Exadata Cloud Infrastructure](#).

- [Manage Database Home Using the Console](#)
Use the OCI console to manage the various operations needed on a Database Home.

- [Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure](#)
Review the list of API calls to manage Oracle Database home.

Manage Database Home Using the Console

Use the OCI console to manage the various operations needed on a Database Home.

- [To view information about a Database Home](#)
- [To delete a database home](#)
You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See [To terminate a database](#) to learn how to terminate a database.
- [To manage tags for your Database Home](#)
- [To move a database to another Database Home](#)
Learn to move a database to another Database Home.

To view information about a Database Home

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the Database Home:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster
4. On the VM Cluster Details page, under Resources, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you are interested in, and then click its name to display details about it.

To delete a database home

You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See [To terminate a database](#) to learn how to terminate a database.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the Database Home you want to delete:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. On the VM Cluster Details page, under Resources, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you want to delete, and then click its name to display details about it.

6. On the Database Home Details page, click **Delete**.

If the Database Home contains databases, you will not be able to proceed. You must cancel the deletion, empty the Database Home as applicable, and then retry the deletion.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To manage tags for your Database Home

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the Database Home:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Database Homes**.
5. In the list of Database Homes, find the Database Home you want to administer.
6. Click the the Actions icon (three dots) on the row listing the Database Home, and then click **Add Tags**.

To move a database to another Database Home

Learn to move a database to another Database Home.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to move.
3. Click **Exadata VM Clusters** in the left hand navigation.
4. Click the name of the VM cluster that contains the database that you want to move.
5. In the Resources list of the VM Cluster Details page, click **Databases**.
6. Click the name of the database that you want to move.

The Database Details page displays information about the selected database.
7. Click **More Actions** and **Move To Another Home**.
8. In the resulting dialog, select the target Database Home.

Note

Oracle recommends using Database Homes, which are running the latest (N) to 3 versions from the latest (N-3) RU versions when updating the software version of the database by moving them to a target DB Home. Only DB Homes provisioned with database versions, which meet this best practice criterion are available as target homes to move your database.

9. Click **Move Database**.

The database will be stopped in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Database**. When the operation completes, Database Home is updated with the current home. If the operation is unsuccessful, the status of the database displays as **Failed**, and the Database Home field provides information about the reason for the failure.

Using the API to Manage Oracle Database Home on Exadata Cloud Infrastructure

Review the list of API calls to manage Oracle Database home.

For information about using the API and signing requests, see "REST APIs" and "Security Credentials". For information about SDKs, see "Software Development Kits and Command Line Interface".

Use these API operations to manage Database Homes:

- [ListDbHomes](#)
- [GetDbHome](#)
- [DeleteDbHome](#)

For the complete list of APIs, see "Database Service API".

Related Topics

- [REST APIs](#)
- [Security Credentials](#)
- [Software Development Kits and Command Line Interface](#)
- [ListDbHomes](#)
- [GetDbHome](#)
- [DeleteDbHome](#)
- [Database Service API](#)

Manage Databases on Exadata Cloud Infrastructure

- [Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure.
- [Prerequisites for Oracle Database Autonomous Recovery Service Cross Region Restore \(Same Tenancy\)](#)
- [Prerequisites for Oracle Database, Object Storage Cross Region Restore \(Same Tenancy\)](#)
- [Oracle Database Releases Supported by Oracle Exadata Database Service on Dedicated Infrastructure](#)
Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.
- [Provisioning and Managing Exadata Databases](#)
This topic describes creating and managing Oracle Databases on an Exadata Cloud Infrastructure instance.

- [Using the API to manage Databases](#)
- [Create and Manage Exadata Pluggable Databases](#)
You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.
- [Restoring an Exadata Pluggable Database](#)
You can perform in-place and out of place restore of an Exadata pluggable database.
- [Cost and Usage Attribution for Pluggable Databases \(PDBs\)](#)
- [Changing the Database Passwords](#)
To change the SYS password, or to change the TDE wallet password, use this procedure.

Prerequisites and Limitations for Creating and Managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure

Review the prerequisites for creating and managing Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure.

Before you can create and use an Oracle Database on Exadata Cloud Infrastructure, you must:

- Provision Exadata Cloud Infrastructure infrastructure
- Configure a VM cluster
- Create any required backup destinations

You can create one or more databases on each Oracle Exadata Database Service on Dedicated Infrastructure system. Other than the storage and processing limits of your Oracle Exadata system, there is no maximum for the number of databases that you can create. By default, databases on Exadata Cloud Infrastructure use Oracle Database Enterprise Edition - Extreme Performance. This edition provides all the features of Oracle Database Enterprise Edition, plus all of the database enterprise management packs, and all of the Enterprise Edition options, such as Oracle Database In-Memory, and Oracle Real Application Clusters (Oracle RAC). If you use your own Oracle Database licenses, then your ability to use various features is limited by your license holdings. TDE Encryption is required for all cloud databases. All new tablespaces will automatically be enabled for encryption.

Prerequisites for Oracle Database Autonomous Recovery Service Cross Region Restore (Same Tenancy)

1. **VCN peering:** Both the VCNs in local and remote regions must be peered across regions. For more information, see [Access to Other VCNs: Peering](#).
2. Add security rules on the source and target VCNs.
 - a. Add Ingress rules on the source.
 - i. Click **Add Ingress Rule**, and add these details to set up a rule that allows HTTPS traffic from anywhere:
Source Type: CIDR
Source CIDR: Specify the CIDR of the VCN where the database resides.
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 8005

- Description:** Specify an optional description of the ingress rule to help manage the security rules.
- ii. Click **Add Ingress Rule**, and add these details to set up a rule that allows SQL*Net traffic from anywhere:
Source Type: CIDR
Source CIDR: Specify the CIDR of the VCN where the database resides.
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 2484
Description: Specify an optional description of the ingress rule to help manage the security rules.
 - iii. Click **Add Ingress Rule**, and add these details to set up a rule that allows HTTPS traffic from anywhere:
Source Type: CIDR
Source CIDR: Specify the CIDR of the target VCN
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 8005
Description: Specify an optional description of the ingress rule to help manage the security rules.
 - iv. Click **Add Ingress Rule**, and add these details to set up a rule that allows SQL*Net traffic from anywhere:
Source Type: CIDR
Source CIDR: Specify the CIDR of the target VCN
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 2484.
Description: Specify an optional description of the ingress rule to help manage the security rules.
- b. Add Egress rules on the target.
These are optional if the egress traffic is opened for all IPs and ports.
- i. Click **Add Egress Rule**, and add these details to set up a rule that allows HTTPS traffic from anywhere:
Source Type: CIDR
Source CIDR: Specify the CIDR of the source VCN
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 8005
Description: Specify an optional description of the ingress rule to help manage the security rules.
 - ii. Click **Add Egress Rule**, and add these details to set up a rule that allows SQL*Net traffic from anywhere:
Source Type: CIDR

Source CIDR: Specify the CIDR of the source VCN

IP Protocol: TCP

Source Port Range: All

Destination Port Range: 2484

Description: Specify an optional description of the ingress rule to help manage the security rules.

Note: Ensure that recovery service subnets (RSS) are present in both regions and are attached to the peer VCNs, namely, source RSS attached to source VCN and target RSS attached to target VCN. For more information, see [Creating a Recovery Service Subnet in the Database VCN](#).

3. Perform DNS peering between local and remote VCNs. For more information, see [Private DNS Implementation](#).

Note: Ensure that the customer adds the `oci.oraclecloud.com` domain while creating forwarding rules inside target/remote VCN.

Also, ensure that the following requirements are met for DNS peering between source and target.

- a. Listening endpoint at the source VCN
- b. Forwarding endpoint at the target VCN
- c. Forwarding rule at the target VCN with the destination being set as the listening endpoint
- d. Ingress and egress rules as stated in the aforementioned link

Prerequisites for Oracle Database, Object Storage Cross Region Restore (Same Tenancy)

The VCNs in region A, where the new database will be located, and region B, where the backups are stored, should be remote peered using a DRG. For more information, see [Remote VCN Peering through an Upgraded DRG](#).

Once the remote peer is established, the DRG in the region with Object Storage should be configured to advertise Object Storage routes towards region A. Go to [Private Access to Oracle Services](#) and follow the steps outlined under *For routing directly between gateways*.

Note

In the "Transit routing directly through gateways", the "on-premises network" will be Region A. Specifically the IP addresses of the "on-premises network" will be the Backup Subnet CIDR of Region A's VCN.

"For routing directly between gateways" steps:

- If you have a VCN and SGW in the region with Object Storage, skip Tasks 1 and 2
- Skip Task 3
- In Task 4, instead of selecting the "All OSN services" route, select the "Object Storage" route.

You'll also need to confirm security lists, and that the VCN route table applied to the backup subnet in Region A, has a route rule to the DRG for Region B's Object Storage CIDRs.

- You can obtain the Object Storage CIDRs for Region B by viewing the JSON file located at [Public IP Addresses for VCNs and the Oracle Services Network](#) under the *Downloading the JSON File* section.
- Within the JSON, locate the region attribute corresponding to Region B. Within the region, next locate the CIDR ranges for the Object Storage, the corresponding CIDR will have "tags" 0 and 1 of "OSN" and "OBJECT_STORAGE".
- Note, some regions will have multiple CIDRs for "OSN" and "OBJECT_STORAGE", create a route rule for each in the route table.

Once completed, confirm remote access to Region's B Object Storage from Region A.

This provides network connectivity to Object Storage. The network cannot permit or prohibit specific Object Storage operations. For that, look to use IAM policies.

Oracle Database Releases Supported by Oracle Exadata Database Service on Dedicated Infrastructure

Exadata Cloud Infrastructure databases require Enterprise Edition - Extreme Performance subscriptions or you can bring your own Oracle Enterprise Edition software licenses.

The Enterprise Edition - Extreme Performance provides all the features of Oracle Database Enterprise Edition, plus all the database enterprise management packs and all the Enterprise Edition options, such as Oracle Database In-Memory and Oracle Real Application Clusters (Oracle RAC).

Exadata Cloud Infrastructure supports the following database versions:

- Oracle AI Database 26ai
- Oracle Database 19c
- Oracle Database 12c Release 2 (12.2) (**Upgrade Support Required**)
- Oracle Database 12c Release 1 (12.1) (**Upgrade Support Required**)
- Oracle Database 11g Release 2 (11.2) (**Upgrade Support Required**)

Note

- Earlier database versions are supported on a 19c cloud VM cluster and can be created at anytime. Cloud VM clusters created with earlier Oracle Database versions will not automatically support Oracle Database 19c.
- For information on upgrading an existing database, see [Upgrading Exadata Databases](#).
- To use Autonomous Recovery Service as a backup destination, your target database must have a minimum compatibility level of 19.0 (the `COMPATIBLE` initialization parameter must be set to 19.0.0 or higher).

For Oracle Database release and software support timelines, see [Release Schedule of Current Database Releases \(Doc ID 742060.1\)](#) in the My Oracle Support portal.

Provisioning and Managing Exadata Databases

This topic describes creating and managing Oracle Databases on an Exadata Cloud Infrastructure instance.

In this documentation, "database" refers to a container database (CDB). When you provision a database in an Exadata cloud VM cluster, the database includes an initial pluggable database (PDB). For more information on these resource types, see [Multitenant Architecture](#) in the Oracle Database documentation. See [Exadata Pluggable Database Operations](#) for more information on pluggable databases in Exadata Cloud Infrastructure.

Note

You can only create CDB databases from the console and OCI APIs. To create nonCDB databases (release 19.x and earlier), you must use the `dbaascli` command line tool. See [dbaascli Command Reference](#) for more information.

You can create Database Homes, databases, and pluggable databases at any time by using the Console or the Database APIs.

When you add a database to a VM cluster on an Exadata instance, the database versions you can select from depend on the current patch level of that resource. You may have to patch your VM cluster to add later database versions.

After you provision a database, you can move it to another Database Home. Consolidating databases under the same home can facilitate management of these resources. All databases in a given Database Home share the Oracle Database binaries and therefore, have the same database version. The Oracle-recommended way to patch a database to a version that is different from the current version is to move the database to a home running the target version. For information about patching, see [Patching an Exadata Cloud Service Instance](#).

Note

When provisioning databases, make sure your VM cluster has enough OCPUs enabled to support the total number of database instances on the system. Oracle recommends the following general rule: for each database, enable 1 OCPU per node. See [To scale CPU cores in an Exadata Cloud Service cloud VM cluster](#) for information on scaling your OCPU count up or down.

When you create an Exadata database, you can choose to encrypt the database using your own encryption keys that you manage. You can rotate encryption keys, periodically, to maintain security compliance and, in cases of personnel changes, to disable access to a database.

Note

- The encryption key you use must be AES-256.
- To ensure that your Exadata database uses the most current versions of the Vault encryption key, rotate the key from the Database Details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service's Console pages to rotate your Database keys.

If you want to use your own encryption keys to encrypt a database that you create, then you must create a dynamic group and assign specific policies to the group for customer-managed encryption keys. See [Managing Dynamic Groups](#) and [Let security admins manage vaults, keys, and secrets](#). Additionally, see [To integrate customer-managed key management into Exadata Cloud Service](#) if you need to update customer-managed encryption libraries for the Vault service.

You can also add and remove databases, and perform other management tasks on a database by using command line utilities. For information and instructions on how to use these utilities, see [Creating and Managing Exadata Databases Manually](#).

- [Database Memory Initialization Parameters](#)
- [Customer-Managed Keys in Exadata Cloud Infrastructure](#)
Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- [Using the Console to Manage Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
To create or terminate a database, complete procedures using the Oracle Exadata console.
- [Known Issues in Exadata Cloud Infrastructure](#)
`rac stopdb failed`

Database Memory Initialization Parameters

- When creating a container database, the initialization parameter, `SGA_TARGET` is set by the automation. This will automatically size the SGA memory pools. The setting will vary depending on the size of the database VM total memory. If the VM has less than or equal to 60 GB of system memory, `SGA_TARGET` is set to 3800 MB. If the VM has 60 GB or more system memory, `SGA_TARGET` is set to 7600 MB.
- The database initialization parameter `USE_LARGE_PAGES` is set to `ONLY` upon database creation, which will require the use of large pages for SGA memory. If the VM is configured with insufficient large pages, the instance will fail to start.

Customer-Managed Keys in Exadata Cloud Infrastructure

Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.

The OCI Vault service provides you with centralized key management capabilities that are highly available and durable. This key-management solution also offers secure key storage using isolated partitions (and a lower-cost shared partition option) in FIPS 140-2 Level 3-

certified hardware security modules, and integration with select Oracle Cloud Infrastructure services. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Enable customer-managed keys when you create databases in Exadata Cloud Infrastructure
- Switch from Oracle-managed keys to customer-managed keys
- Rotate your keys to maintain security compliance

Requirements

To enable management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so, similar to the following: `allow dynamic-group dynamic_group_name to manage keys in tenancy`.

Another policy is needed if the Vault being used by the customer is replicated (<https://docs.oracle.com/en-us/iaas/Content/KeyManagement/Tasks/replicatingvaults.htm>). For vaults that are replicated, this policy is needed: `allow dynamic-group dynamic_group_name to read vaults in tenancy`

Limitations

To enable Data Guard on Exadata Cloud Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same [realm](#).

Task 1. Create a Vault and a Master Encryption Key

Create a vault in the Vault service by following the instructions in [To create a new vault](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, Oracle recommends that you create the vault in a compartment created specifically to contain the vaults containing customer-managed keys, as described in [Before You Begin: Compartment Hierarchy Best Practice](#).

After creating the vault, create at least one master encryption key in the vault by following the instructions in [To create a new master encryption key](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, make these choices:

- **Create in Compartment:** Oracle recommends that you create the master encryption key in the same compartment as its vault; that is, the compartment created specifically to contain the vaults containing customer-managed keys.
- **Protection Mode:** Choose an appropriate value from the drop-down list:
 - **HSM** to create a master encryption key that is stored and processed on a hardware security module (HSM).
 - **Software** to create a master encryption key that is stored in a software file system in the Vault service. Software-protected keys are protected at rest using an HSM-based root key. You may export software keys to other key management devices or to a different OCI cloud region. Unlike HSM keys, software-protected keys are free of cost.
- **Key Shape Algorithm:** AES
- **Key Shape Length:** 256 bits

Oracle strongly recommends that you create a separate master encryption key for each of your container databases (CDBs). Doing so makes management of key rotation over time much simpler.

Task 2. Create a Service Gateway, a Route Rule, and an Egress Security Rule

Create a service gateway in the VCN (Virtual Cloud Network) where your Oracle Exadata Database Service on Dedicated Infrastructure resources reside by following the instructions in [Task 1: Create the service gateway](#) in Oracle Cloud Infrastructure Documentation.

After creating the service gateway, add a route rule and an egress security rule **to each subnet** (in the VCN) where Oracle Exadata Database Service on Dedicated Infrastructure resources reside so that these resources can use the gateway to access the Vault service:

1. Go to the **Subnet Details** page for the subnet.
2. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
3. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:
 - **Destination:** All IAD Services In Oracle Services Network
 - **Target Type:** Service Gateway
 - **Target:** The name of the service gateway you just created in the VCN

If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.

4. Return to the **Subnet Details** page for the subnet.
5. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.
6. In the side menu, under **Resources**, click **Egress Rules**.
7. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:
 - **Stateless:** No
 - **Destination:** All IAD Services In Oracle Services Network
 - **IP Protocol:** TCP
 - **Source Port Range:** All
 - **Destination Port Range:** 443

If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Task 3. Create a Dynamic Group and a Policy Statement

To grant your Oracle Exadata Database Service on Dedicated Infrastructure resources permission to access customer-managed keys, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the master encryption keys you created in the Vault service.

When defining the dynamic group, you identify your Oracle Exadata Database Service on Dedicated Infrastructure resources by specifying the OCID of the compartment containing your Exadata Infrastructure resource.

1. Copy the OCID of the compartment containing your Exadata Infrastructure resource. You can find this OCID on the **Compartment Details** page of the compartment.

2. Create a dynamic group by following the instructions in [To create a dynamic group](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

```
ALL {resource.compartment.id = '<compartment-ocid>'}
```

where *<compartment-ocid>* is the OCID of the compartment containing your Exadata Infrastructure resource.

After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and keys. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group-name>
to manage keys
in compartment <vaults-and-keys-compartment>
where all {
target.key.id='<key-ocid>',
request.permission!='KEY_DELETE',
request.permission!='KEY_MOVE',
request.permission!='KEY_IMPORT',
request.permission!='KEY_BACKUP'
}
```

If you are using a replicated virtual private vault for the Oracle Data Guard deployment, add an additional policy statement in this format:

```
allow dynamic-group <dynamic-group>
to read vaults
in tenancy | compartment <vaults-and-keys-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-keys-compartment>* is the name of the compartment in which you created your vaults and master encryption keys.

- [To integrate customer-managed key management into Exadata Cloud Infrastructure](#)
If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.
- [Known Issues for Exadata Cloud Infrastructure and Data Guard](#)
Possible TDE key replication issue, and MRP and DG LCM operation failures.
- [To integrate customer-managed key management into Exadata Cloud Infrastructure](#)
If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using

Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

To integrate customer-managed key management into Exadata Cloud Infrastructure

If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using Red Hat Package Manager) to enable DBAASTOOLS to interact with the APIs that customer-managed key management uses.

KMS TDE CLI

To update the KMS TDE CLI package, you must complete the following task on all nodes in the Exadata Cloud Infrastructure instance:

1. Deinstall current KMS TDE CLI package, as follows:

```
rpm -ev kmstdecli
```

2. Install the updated KMS TDE CLI package, as follows:

```
rpm -ivh kms_tde_cli
```

LIBKMS

LIBKMS is a library package necessary to synchronize a database with customer-managed key management through PKCS11. When a new version of LIBKMS is installed, any databases converted to customer-managed key management continue to use the previous LIBKMS version, until the database is stopped and restarted.

To update the LIBKMS package, you must complete the following task on all nodes in the Exadata Cloud Infrastructure instance:

1. Confirm that the LIBKMS package is already installed, as follows:

```
rpm -qa --last | grep libkmsdepkcs11
```

2. Install a new version of LIBKMS, as follows:

```
rpm -ivh libkms
```

3. Use SQL*Plus to stop and restart all databases converted to customer-managed key management, as follows:

```
shutdown immediate;  
startup;
```

4. Ensure that all converted databases are using the new LIBKMS version, as follows:

```
for pid in $(ps aux | grep "<dbname>" | awk '{print $2;}'); do echo $pid;  
sudo lsof -p $pid | grep kms | grep "pkcs11_[0-9A-Za-z.]*" | sort -u; done  
| grep pkcs11
```

5. Deinstall LIBKMS packages that are no longer being used by any database, as follows:

```
rpm -ev libkms
```

Using the Console to Manage Databases on Oracle Exadata Database Service on Dedicated Infrastructure

To create or terminate a database, complete procedures using the Oracle Exadata console.

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To manage SYS user and TDE Wallet passwords](#)
Learn to manage administrator (SYS user) and TDE wallet passwords.
- [To view details of a Protected Database](#)
To view the details of a Protected Database, use this procedure.
- [To create a database from a backup](#)
- [To create a database from the latest backup](#)
- [To move a database to another Database Home](#)
This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.
- [To terminate a database](#)
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.

To create a database in an existing VM Cluster

This topic covers creating your first or subsequent databases.

Note

If IORM is enabled on the Exadata Cloud Infrastructure instance, then the default directive will apply to the new database and system performance might be impacted. Oracle recommends that you review the IORM settings and make applicable adjustments to the configuration after the new database is provisioned.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster you want to create the database in: **Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model)**: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Create Database**.
5. In the **Create Database** dialog, enter the following:

Note

You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- **Database name:** The name for the database. The database name must meet the requirements:
 - Maximum of 8 characters
 - Contain only alphanumeric characters
 - Begin with an alphabetic character
 - Cannot be part of the first 8 characters of a `DB_UNIQUE_NAME` on the VM cluster
 - DO NOT use the following reserved names: `grid`, `ASM`
- **Database unique name suffix:**
Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. The value is case insensitive.
The unique name must meet the requirements:
 - Maximum of 30 characters
 - Contain only alphanumeric or underscore (`_`) characters
 - Begin with an alphabetic character
 - Unique across the VM cluster. Recommended to be unique across the tenancy.
 If not specified, the system automatically generates a unique name value, as follows:


```
<db_name>_<3_chars_unique_string>_<region-name>
```
- **Database version:** The version of the database. You can mix database versions on the Exadata VM cluster.
- **PDB name:** *(Optional)* For Oracle Database 12c (12.1.0.2) and later, you can specify the name of the pluggable database. The PDB name must begin with an alphabetic character, and can contain a maximum of eight alphanumeric characters. The only special character permitted is the underscore (`_`).
To avoid potential service name collisions when using Oracle Net Services to connect to the PDB, ensure that the PDB name is unique across the entire VM cluster. If you do not provide the name of the first PDB, then a system-generated name is used.
- **Database Home:** The Oracle Database Home for the database. Choose the applicable option:
 - **Select an existing Database Home:** The Database Home display name field allows you to choose the Database Home from the existing homes for the database version you specified. If no Database Home with that version exists, you must create a new one.
 - **Create a new Database Home:** Use this option to provision a new Database Home for your Data Guard peer database.
Click **Change Database Image** to use a desired Oracle-published image or a custom *database software image* that you have created in advance, then select an **Image Type**:
 - * **Oracle Provided Database Software Images:**

You can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

Note

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

- * **Custom Database Software Images:** These images are *created by your organization* and contain customized configurations of software updates and patches. Use the **Select a compartment**, **Select a region**, and **Select a Database version** selectors to limit the list of custom database software images to a specific compartment, region, or Oracle Database software major release version.
Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
- **Create administrator credentials:** (*Read only*) A database administrator `SYS` user will be created with the password you supply.
 - **Username:** `SYS`
 - **Password:** Supply the password for this user. The password must meet the following criteria:
A strong password for `SYS`, `SYSTEM`, TDE wallet, and PDB Admin. The password must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the username (`SYS`, `SYSTEM`, and so on) or the word "**oracle**" either in forward or reversed order and regardless of casing.
 - **Confirm password:** Re-enter the `SYS` password you specified.
 - Using a **TDE wallet password** is optional. If you are using customer-managed encryption keys stored in a [vault](#) in your tenancy, the TDE wallet password is not applicable to your VM Cluster. Use **Show Advanced Options** at the end of the Create Database dialog to configure customer-managed keys.
If you are using customer-managed keys, or if you want to specify a different TDE wallet password, uncheck the **Use the administrator password for the TDE wallet box**. If you are using customer-managed keys, leave the TDE password fields blank. To set the TDE wallet password manually, enter a password in the **Enter TDE wallet password** field, and then confirm by entering it into the **Confirm TDE wallet password** field.
- **Configure database backups:** Specify the settings for backing up the database to Autonomous Recovery Service or Object Storage:

Note

To use Autonomous Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

- **Enable automatic backup:** Check the check box to enable automatic incremental backups for this database. If you are creating a database in a security zone compartment, you must enable automatic backups.
- **Backup Destination:** Your choices are **Autonomous Recovery Service** or **Object Storage**.
- **Backup Scheduling:**
 - * **Object Storage (L0):**
 - * **Full backup scheduling day:** Choose a day of the week for the initial and future L0 backups to start.
 - * **Full backup scheduling time (UTC):** Specify the time window when the full backups start when the automatic backup capability is selected.
 - * **Take the first backup immediately:** A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
 - * **Object Storage (L1):**
 - * **Incremental backup scheduling time (UTC):** Specify the time window when the incremental backups start when the automatic backup capability is selected.
 - * **Autonomous Recovery Service (L0):**
 - * **Scheduled day for initial backup:** Choose a day of the week for the initial backup.
 - * **Scheduled time for initial backup (UTC):** Select the time window for the initial backup.
 - * **Take the first backup immediately:** A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
 - * **Autonomous Recovery Service (L1):**

- * **Scheduled time for daily backup (UTC):** Specify the time window when the incremental backups start when the automatic backup capability is selected.
 - **Deletion options after database termination:** Options that you can use to retain protected database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
 - * **Retain backups for the period specified in your protection policy or backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Object Storage Backup retention period or Autonomous Recovery Service protection policy after the database is terminated.
 - * **Retain backups for 72 hours, then delete:** Select this option to retain backups for a period of 72 hours after you terminate the database.
 - **Backup Retention Period/Protection Policy:** If you choose to enable automatic backups, you can choose a policy with one of the following preset retention periods, or a Custom policy.

Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

Autonomous Recovery Service protection policy:

 - * **Bronze:** 14 days
 - * **Silver:** 35 days
 - * **Gold:** 65 days
 - * **Platinum:** 95 days
 - * Custom defined by you
 - * **Default:** Silver - 35 days
 - **Enable Real-Time Data Protection:** Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
6. Click **Show Advanced Options** to specify advanced options for the database:
- **Management:**

Oracle SID prefix: The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. The `INSTANCE_NAME` parameter is also known as the `SID`. The `SID` is unique across the cloud VM Cluster. If not specified, `SID` prefix defaults to the `db_name`.

Note

Entering an `SID` prefix is only available for Oracle 12.1 databases and above.

The `SID` prefix must meet the requirements:

- Maximum of 12 characters
- Contain only alphanumeric characters. You can, however, use underscore (`_`), which is the only special character that is not restricted by this naming convention.

- Begin with an alphabetic character
- Unique in the VM cluster
- DO NOT use the following reserved names: `grid`, `ASM`
- **Character set:** The character set for the database. The default is `AL32UTF8`.
- **National character set:** The national character set for the database. The default is `AL16UTF16`.
- **Encryption:**
Choose the key store for the TDE keys. Oracle Wallet is the default for managing TDE keys. The other available options are OCI Vault and Oracle Key Vault (OKV).

To configure the database with encryption based on encryption keys you manage:

- OCI Vault:
 - a. You must have a valid key in Oracle Cloud Infrastructure Vault service. See [Let security admins manage vaults, keys, and secrets](#).

Note

You must use AES-256 keys for your database.

- b. Choose a **Vault**.
- c. Select a **key**.
- d. To specify a key version other than the latest version of the selected key, check **Choose the key version** and enter the OCID of the key you want to use in the **Key version OCID** field.

Note

The Key version will only be assigned to the container database (CDB), and not to its pluggable database (PDB). PDB will be assigned an automatically generated new key version.

- Oracle Key Vault: Choose a compartment and select a key store from the chosen compartment.
 - **Tags:** 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.
7. Click **Create Database**.

Note

You can now:

- Create or delete a CDB while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Create or delete a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.
- Create or delete a CDB while concurrently creating or deleting a PDB within the same Oracle home, and vice versa.
- Create or delete a CDB concurrently within the same Oracle home.
- Create or delete a CDB while simultaneously updating VM Cluster tags.

After database creation is complete, the status changes from **Provisioning** to **Available**, and on the database details page for the new database, the **Encryption** section displays the encryption key name and the encryption key OCID.

Warning

Do not delete the encryption key from the vault. This causes any database protected by the key to become unavailable.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)
- [security zone compartment](#)
- [Resource Tags](#)
- [Let security admins manage vaults, keys, and secrets](#)

To manage SYS user and TDE Wallet passwords

Learn to manage administrator (SYS user) and TDE wallet passwords.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment** that contains the VM cluster that hosts the database that you want to change passwords.
3. Click the name of the VM cluster that contains the database that you want to change passwords.
4. In the **Resources** list of the VM Cluster Details page, click **Databases**.
5. Click the name of the database that you want to change passwords. The Database Details page displays information about the selected database.
6. On the Database Details page, click More actions, and then click **Manage passwords**.
7. In the resulting **Manage passwords** dialog, click **Update Administrator Password** or **Update TDE Wallet Password**. Depending on the option you select, the system displays the fields to edit.

- **Update Administrator Password:** Enter the new password in both the New administrator password and Confirm administrator password fields.

Note

The **Update Administrator Password** option will change the sys user password only. Passwords for other administrator accounts such as system, pdbadmin, and TDE wallet will not be changed.

- **Update TDE Wallet Password:** Enter the current wallet password in the **Enter existing TDE wallet password** field, and then enter the new password in both the **New TDE wallet password** and **Confirm TDE wallet password** fields.
8. Click **Apply** to update your chosen password.

To view details of a Protected Database

To view the details of a Protected Database, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:
Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Click **Exadata VM Clusters**.

In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM Cluster Details page, in the Databases table, click the name of the database to display the **Database Details** page. The Backup section displays the state of the automatic backups. If the Autonomous Recovery Service is the destination, a link will be available which includes additional details. You can also check if Real-time Data Protection is enabled or disabled. Click the **Autonomous Recovery Service** link to be taken to the page containing the Protected Database details. For more information about Protected Databases, see *Viewing Protected Database Details*.

Related Topics

- [Viewing Protected Database Details](#)

To create a database from a backup

Before you begin, note the following:

- When you create a database from a backup, the availability domain can be the same as where the backup is hosted or a different one across regions.
- The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
- If you are creating a database from an automatic backup, then you can choose any level 0 weekly backup, or a level 1 incremental backup created after the most recent level 0 backup. For more information on automatic backups, see [To configure automatic backups for a database](#)

- If the backup being used to create a database is in a security zone compartment, the database cannot be created in a compartment that is not in a security zone. See the [Security Zone Policies](#) topic for a full list of policies that affect Database service resources.
- 1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
- 2. Choose your **Compartment**.
- 3. Navigate to a backup.
 - *Standalone backups*: Click **Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
 - *Automatic backups*: Navigate to the Database Details page of the database associated with the backup:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

Click the name of the database associated with the backup that you will use to create the new database. Locate the backup in the list of backups on the Database Details page.
- 4. Click the Actions icon (three dots) for the backup you chose.
- 5. Click **Create Database**. On the **Create Database from Backup** page, configure the database as follows.
- 6. In the **Provide basic information for the Exadata infrastructure** section:
 - **Select a region**: The target region where you want to create the database.
 - **Select an availability domain**: It could be the same as the availability domain that hosts the backup or a different one within the same region
 - **Select Exadata infrastructure**: Select an Exadata infrastructure from the chosen compartment. Click the **Change Compartment** hyperlink to choose a different compartment.
- 7. In the **Configure your VM Cluster** section:
 - *Backups created in cloud VM clusters*: Choose a cloud VM cluster to run the database from the **Select a VM cluster** drop-down list.
- 8. In the **Configure Database Home** section:
 - **Select an existing Database Home**: If you choose this option, make a selection from the **Select a Database Home** drop-down list.

Note

You can not create a database from backup in the same Database Home where the source database exists.

- **Create a new Database home**: If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
 - a. Select the compartment containing the database software image you want to use to create the new Database Home.

- b. Select the region containing the database software image you want to use to create the new Database Home. Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
- c. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.

Note

Database restore operations for Databases of 12.2.0.1 and earlier are not allowed at this time.

- d. Click **Select**.
9. In the **Configure database** section:

Note

You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- In the **Database name** field, name the database or accept the default name. The database name must meet the requirements:
 - Maximum of 8 characters
 - Contain only alphanumeric characters
 - Begin with an alphabetic character
 - Cannot be part of first 8 characters of a different database's `db_unique_name` on the VM cluster
 - DO NOT use the following reserved names: `grid`, `ASM`
- **Database unique name:** Specify a value for the `DB_UNIQUE_NAME` database parameter. The unique name must meet the requirements:
 - Maximum of 30 characters
 - Contain only alphanumeric or underscore (`_`) characters
 - Begin with an alphabetic character
 - Unique across the VM cluster. Recommended to be unique across the tenancy.

If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

- **Administrator username:** This read-only field displays the username for the administrator, "sys".
- In the **Password** and **Confirm password** fields, enter and re-enter a password. A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the user name

(SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.

10. In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.
11. Click **Show Advanced Options** to specify advanced options for the database:
 - **Management**
Oracle SID prefix: This option is in the **Management** tab. The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. If not provided, then the SID prefix defaults to the first twelve characters of the `db_name`.

Note

Entering an SID prefix is only available for Oracle 12.1 databases and above.

The SID prefix must meet the requirements:

- Maximum of 12 characters
- Contain only alphanumeric characters
- Begin with an alphabetic character
- Unique in the VM cluster
- DO NOT use the following reserved names: grid, ASM
- In the **Tags** tab, you can add tags to the database. 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 if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

12. Click **Create Database**.

NOT_SUPPORTED

1. Click the Exadata cloud VM cluster name that contains the specific database to display the details page.
2. From the list of databases, click the database name associated with the backup you want to use to display a list of backups on the database details page. You can also access the list of backups for a database by clicking **Backups** in the **Resources** section.

NOT_SUPPORTED

1. Click **Standalone Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of standalone backups, find the backup you want to use to create the database.
- [To navigate to the list of standalone backups for your current compartment](#)

To navigate to the list of standalone backups for your current compartment

1. Click **Backups** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. In the list of standalone backups, find the backup you want to use to create the database.

To create a database from the latest backup

Before you begin, note the following:

- When you create a database from a backup, the availability domain can be the same as where the backup is hosted or a different one across regions.
 - The Oracle Database software version you specify must be the same or later version as that of the backed-up database.
 - If the backup being used to create a database is in a security zone compartment, the database cannot be created in a compartment that is not in a security zone. See the [Security Zone Policies](#) topic for a full list of policies that affect Database service resources.
1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
 2. Choose your **Compartment**.
 3. Navigate to the cloud VM cluster that contains the source database you are using to create the new database:
 - *Cloud VM clusters* (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
 4. Under **Databases**, click the name of the database you are using as the source for the new database.
 5. On the Database Details page, click the **Actions** menu, and then select **Create database from backup**.
 6. On the Create database from backup panel, select **Create database from last backup** or **Create database from specified timestamp**, and then click **Create**.
 7. On the Create Database from Backup page, configure the database as follows:
 - **PDB:** You can include all PDBs in the backup or choose specific PDBs to restore.
 - **All PDBs:** Select this option to include all PDBs in the backup.
 - **Choose the PDBs to restore:** Specify the PDBs to restore. You can select one or more PDBs. If selecting multiple PDBs, provide a comma-delimited list.
 - **Region:** The target region where you want to create the database.
 - **Availability domain:** It could be the same as the availability domain that hosts the backup or a different one within the same region.
 - **Select a service:** Select **Exadata Database Service on Dedicated Infrastructure**.
 - **Exadata infrastructure:** Select an Exadata infrastructure from the chosen compartment.
 - **VM Cluster:**
 - **VM Cluster:** Choose a cloud VM cluster to run the database from the chosen compartment.
 - **Database Home:**
 - * **Select an existing Database Home:** If you choose this option, make a selection from the **Select a Database Home** drop-down list.

- * **Create a new Database home:** If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
 - a. Select the compartment containing the database software image you want to use to create the new Database Home.
 - b. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.
 - c. Click **Select**.
- **Database Home:**

Note

You cannot modify the `db_name`, `db_unique_name`, and SID prefix after creating the database.

- **Database name:** The name for the database. The database name must meet the requirements:
 - * Maximum of 8 characters
 - * Contain only alphanumeric characters
 - * Begin with an alphabetic character
 - * Cannot be part of first 8 characters of a `DB_UNIQUE_NAME` on the VM cluster
 - * DO NOT use the following reserved names: grid, ASM
- **Database unique name:** Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. The value is case insensitive. The unique name must meet the requirements:
 - * Maximum of 30 characters
 - * Contain only alphanumeric or underscore (`_`) characters
 - * Begin with an alphabetic character
 - * Unique across the VM cluster. Recommended to be unique across the tenancy.

If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

- **Administrator username:** This read-only field displays the username for the administrator, "sys".
- In the **Password** and **Confirm password** fields, enter and re-enter a password. A strong password for SYS administrator must be 9 to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. The password must not contain the user name (SYS, SYSTEM, and so on) or the word "oracle" either in forward or reverse order and regardless of casing.

- In the **Enter the source database's TDE wallet or RMAN password** field, enter a password that matches either the Transparent Data Encryption (TDE) wallet password or RMAN password for the source database.
8. Click **Show Advanced Options** to specify advanced options for the database.
 - **Management:**
Oracle SID prefix: The Oracle Database instance number is automatically added to the `SID` prefix to create the `INSTANCE_NAME` database parameter. The `INSTANCE_NAME` parameter is also known as the `SID`. The `SID` is unique across the cloud VM cluster. If not specified, `SID` prefix defaults to the first 12 characters of the `db_name`.
- Note**

Entering an `SID` prefix is only available for Oracle 12.1 databases and above. The `SID` prefix must meet the requirements:

 - Maximum of 12 characters
 - Contain only alphanumeric characters
 - Begin with an alphabetic character
 - Unique in the VM cluster
 - DO NOT use the following reserved names: `grid`, `ASM`
- **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
9. Click **Create**.

To move a database to another Database Home

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

You can move a database to any Database Home that meets at either of the following criteria:

- The target Database Home uses the same Oracle Database software version (including patch updates) as the source Database Home
- The target Database Home is based on either the latest version of the Oracle Database software release used by the database, or one of the three prior versions of the release

Moving a database to a new Database Home brings the database up to the patch level of the target Database Home. For information on patching Database Homes, see [Database Home Patching](#).

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.
Cloud VM clusters ([The New Exadata Cloud Infrastructure Resource Model](#)): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM**

Clusters. In the list of VM clusters, click the name of the VM cluster that contains the database you want to move.

4. Click **More Actions**, then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move**.
7. Confirm the move operation.

The database is moved in a rolling fashion. The database instance will be stopped, node by node, in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Database**. When the operation completes, Database Home is updated with the current home. Datapatch is executed automatically, as part of the database move, to complete post-patch SQL actions for all patches, including one-offs, on the new Database Home. If the database move operation is unsuccessful, then the status of the database displays as `Failed`, and the Database Home field provides information about the reason for the failure.

To terminate a database

You'll get the chance to back up the database prior to terminating it. This creates a standalone backup that can be used to create a database later. We recommend that you create this final backup for any production (non-test) database.

Note

Terminating a database removes all automatic incremental backups of the database from Oracle Cloud Infrastructure Object Storage. However, all full backups that were created on demand, including your final backup, will persist as standalone backups.

You cannot terminate a database that is assuming the primary role in a Data Guard association. To terminate it, you can switch it over to the standby role.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database:
Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the Databases table, click the name of the database to display the Database Details page.

4. Click **More Actions**, and then click **Terminate**.
For databases using Oracle Cloud Infrastructure Object Storage or the Autonomous Recovery Service:

In the confirmation dialog, perform the following steps:

- Review the message about the backup retention policy.
- Configure automatic backups, if needed.
- Type the name of the database to confirm the termination.

5. Click **Terminate**.
The database's status indicates Terminating.

Note

The database stays in a terminated state with backups listed until all backups are expired.

Note

You can now:

- Terminate a CDB when Data Guard setup is running on another database in the same Oracle home, and vice versa.
- Create or delete a CDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To administer Vault encryption keys

Use this procedure to rotate the Vault encryption key or change the encryption management configuration.

After you provision a database in cloud VM cluster, you can rotate the Vault encryption key or change the encryption management configuration for that database.

Note

- To ensure that your Exadata database uses the most current version of the Vault encryption key, rotate the key from the database details page on the Oracle Cloud Infrastructure Console. Do not use the Vault service.
- You can rotate Vault encryption keys only on databases that are configured with customer-managed keys.
- You can change encryption key management from Oracle-managed keys to customer-managed keys but you cannot change from customer-managed keys to Oracle-managed keys.
- Oracle supports administering encryption keys on databases after Oracle Database 11g release 2 (11.2.0.4).

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your compartment from the **Compartment** drop-down.
3. Navigate to the cloud VM cluster that contains the database for which you want to change encryption management or to rotate a key.

Cloud VM clusters: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, locate the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

4. In the **Databases** section, click the name of the database for which you want to change encryption management or to rotate a key to display its details page.
5. Click the **More Actions** drop-down.
6. Click **Manage encryption key**.
To rotate an encryption key on a database using customer-managed keys:

Note

Generate a new master encryption key version. Only the CDB root key version is changed or rotated to a new one. It doesn't generate a new key version for the dependent PDBs. Rotate customer-managed keys periodically to comply with security compliance and regulatory mandates.

- a. Click **Rotate Encryption Key** to display a confirmation dialog.
- b. Click **Update**.

To assign a new key version:

Assign a new key version (BYOK) to CDB while creating or after provisioning it.

- a. Click **Assign a new key version**.
- b. In the **Key version OCID** field, enter the OCID of the new key version you want to assign.
- c. Click **Update**.
To copy the Key version OCID:
 - i. Find the Vault and the Key details on the Key Details page (**Key Management & Secret Management >> Vault >> <Vault> >> Key Details**) by searching with the KMS key OCID provided in the CDB details page.
 - ii. Copy the OCID and paste it in the **Key version OCID** field.

To change key management type from Oracle-managed keys to customer-managed keys:

- a. Click **Change Key Management Type**.
- b. Select **Use customer-managed keys**.
You must have a valid encryption key in Oracle Cloud Infrastructure Vault service and provide the information in the subsequent steps. See [Key and Secret Management Concepts](#).
- c. Choose a vault from the **Vault in compartment** drop-down. You can change the compartment by clicking the **Change Compartment** link.
- d. Select an encryption key from the **Master encryption key in compartment** drop-down. You can change the compartment containing the encryption key you want to use by clicking the **Change Compartment** link.
- e. If you want to use an encryption key that you import into your vault, then select the **Choose the key version** check box and enter the OCID of the key you want to use in the **Key version OCID** field.

Note

If you do not choose a version, the latest version of the key is used.

7. Click **Update**.**Note**

Changing key management causes the database to become briefly unavailable.

Caution

After changing key management to customer-managed keys, do not delete the encryption key from the vault as this can cause the database to become unavailable.

On the database details page for this database, the **Encryption** section displays the encryption key name and the encryption key OCID.

Known Issues in Exadata Cloud Infrastructure

`rac stopdb failed`

rac stopdb failed to stop db

When GI version is 19.17 then creating a database against 11.2.0.4 Oracle home with patchsets July '22 RU or older will fail with error mentioned in bug#28326679

Example:

```
ERROR : rac stopdb, failed to stop db viacmd export ORACLE_HOME=/u02/app/oracle/product/11.2.0/dbhome_1 ;/u02/app/oracle/product/11.2.0/dbhome_1/bin/srvctl stop database -d db008077-o immediate, out : PRCD-1120 : The resource for database db008077 could not be found. PRCR-1001 : Resource ora.db008077.db does not exist, err :1 }
```

Solution:

Option 1: (Create new oracle home with Custom Image):

- Create custom image for 11.2.0.4 with patchsets July '22 RU or older along with bug#28326679 one off
- Create Oracle home using above customer image
- Create database against the home

Option 2 (Apply one-off to existing Oracle home) :

- Download the patch for bug#28326679
- Apply the patch using opatch

Applicability:

- For ExaCS and ExaCC-Gen2, Both options given above will work.
- For ExaCC – Gen1, Option 2 (Apply one-off to existing Oracle home) will work.

Using the API to manage Databases

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

Use these API operations to manage databases.

- [ListDatabases](#)
- [GetDatabase](#)
- [CreateDatabase](#)
- [UpdateDatabase](#) - Use this operation to move a database to another Database Home
- [DeleteDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Create and Manage Exadata Pluggable Databases

You can create and manage pluggable databases (PDBs) in Exadata Cloud Infrastructure using the Console and APIs.

In this documentation, "database" refers to a container database, also called a CDB. For more information on these resource types, see [Multitenant Architecture](#) in the Oracle Database documentation. See [Provisioning and Managing Exadata Databases](#) for information on container databases in Exadata Cloud Infrastructure.

Oracle 19c or later databases created in Exadata Cloud Infrastructure include an initial PDB that you can access from the Database Details page in the Console. You can create and manage additional PDBs in the database using the Console or APIs.

- **Backup**

You can take a backup of the PDB optionally during create, clone, or relocate operations when the CDB is configured with the auto-backup feature. The PDB backup destination will always be the same as CDB, and the backups cannot be accessed directly or created on demand. Oracle recommends immediately backing up the PDB after you create or clone it. This is because the PDB will not be recoverable until the next daily auto-backup completes successfully, leading to a possible data loss.
- **Restore**
 - **Base Database Service / Oracle Exadata Database Service on Dedicated Infrastructure:**
 - * **In place restore:** You can restore a PDB within the same CDB to last known good state or to a specified timestamp.
 - * **Out of place restore:** You can restore a PDB by creating a database (CDB) from the backup, then selecting a PDB or a subset of them you want to restore on the new database.
 - **Oracle Exadata Database Service on Cloud@Customer:**
 - * **In place restore:** You can restore a PDB within the same CDB to last known good state and specified timestamp.
 - * **Out of place restore:** It's not available.

You can perform an in-place restore when you want to move a PDB back to a specified state or time. Both the CDB and PDB must be up and running and only one PDB can be restored at a time.

- * If you have multiple PDBs in your CDB and want to restore multiple of them to the same CDB, then you could restore each individual PDB, one PDB at a time, from the CDB backup.
- * When the CDB is down, you could restore the complete CDB and all the PDBs in that CDB will also be restored.
- * You could either restore the database to the specified timestamp or to its last known good state.

- **Relocate**

You can relocate a PDB from one CDB to another CDB within the same availability domain (AD):

- Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before relocating.
- To the same or a higher database version.

During relocate, the PDB will be removed from the source CDB and moved to the destination CDB that is up and running. In a Data Guard association, a PDB relocated to the primary will be synchronized with the standby as well.

- **Clone**

Note

Cloning application root PDBs is not supported in the cloud.

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and refresh the cloned PDB.

The following types of clones are supported:

- **Local clone:** A copy of the PDB is created within the same CDB.
- **Remote clone:** A copy of the PDB is created in a different CDB.

You can perform a remote clone of a PDB from one CDB to another CDB within the same availability domain (AD):

- Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
- To the same or a higher database version.
- **Refreshable clone:** A copy of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.

You can perform a refreshable clone of a PDB from one CDB to another CDB within the same availability domain (AD):

- * Across compartments, VM clusters, DB system (for BaseDB only), or VCNs (not applicable to ExaDB-C@C). If two different VCNs are used, then both VCNs must be peered before cloning.
- * To the same or a higher database version.

- **Refreshable Clone**
A refreshable clone enables you to keep your remote clone updated with the source PDB. You can only refresh while the PDB is in mount mode. The only open mode you can have is read-only and refresh cannot be done while it is in read-only mode.
 - A database link user credential is required for creating a refreshable clone.
 - Clone, relocate, and in-place restore operations are not supported in the refreshable clone. Relocate and in-place restore operations are not supported in the source, and the source can only be deleted after disconnecting or deleting the refreshable clone.
 - In a Data Guard association, a refreshable clone cannot be created on standby, but it can be created on the primary. However, the primary will not be synced to the standby.

Note

A PDB in standby cannot be used as the source for a refreshable PDB.

- **Convert Refreshable PDB to Regular PDB**
You can convert a refreshable PDB to a regular PDB by disconnecting the refreshable clone (destination PDB) from the source PDB at any time. If the refresh PDB is in a Data Guard association, when it is converted to a regular PDB the PDB will be synced to the standby as part of the conversion process.
- **Open Modes**
On the Console, you can see the open modes of a PDB, such as read-write, read-only, and mounted. If the PDB status is the same across all nodes, the system displays the same status for all PDBs. If the PDB statuses are different across the nodes, the system displays a message indicating on which nodes the PDBs are opened in read-write mode. You cannot change the open mode of a PDB through the API or Console. However, you can start or stop a PDB. Starting the PDB will start it in read-write mode. Stopping the PDB will close it and it will remain in mount mode.
- [Limitations for Pluggable Database Management](#)
- [Creating an Exadata Pluggable Database](#)
- [Managing an Exadata Pluggable Database](#)
This topic includes the procedures to connect to, start, stop, and delete a pluggable database (PDB).
- [Cloning an Exadata Pluggable Database](#)
You can create local, remote, and refreshable clones.
- [Concurrently Create or Delete Pluggable Databases \(PDBs\)](#)
You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.

Limitations for Pluggable Database Management

- New PDBs created with SQL are not immediately discovered by OCI's control plane and displayed in the Console. However, OCI does perform a sync operation on a regular basis to discover manually-created PDBs, and they should be visible in the Console and with API-based tools within 45 minutes of creation. Oracle recommends using the Console or API-based tools (including the OCI CLI , SDKs, and Terraform) to create PDBs.

- Pluggable database operations are supported only for databases using Oracle Database 19c and later.
- PDBs are backed up at the CDB level when using the OCI Console or APIs, and each backup includes all the PDBs in the database. However, the dbaascli utility's [dbaascli database backup](#) command allows you to create backups of specified PDBs. See [Using the dbaascli Utility on Exadata Cloud Infrastructure](#) for more information.
- Restore operations are performed at the CDB level when using the OCI Console or APIs. However, the dbaascli utility's [dbaascli pdb recover](#) command allows you to restore backups of specified PDBs. See [Using the dbaascli Utility on Exadata Cloud Infrastructure](#) for more information.

Creating an Exadata Pluggable Database

You can create a pluggable database (PDB) in Exadata Cloud Service from the OCI Console, or with the APIs and API-based tools (the [OCI CLI](#), [SDKs](#), and Terraform). PDBs must be created one at a time. During the PDB create operation, the parent database (CDB) is in the "Updating" state. Creating a new PDB has no impact on existing PDBs in the database.

- [To create pluggable database](#)
- [To relocate a pluggable database](#)
- [Using the API to create pluggable database](#)

To create pluggable database

Note

If the databases are created directly on Guest VM, the attributed usage data would be delayed.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:
Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. On the Database Details page, click **Pluggable Databases** in the **Resources** section of the page.
5. Click **Create Pluggable Database**.
6. In the **Create Pluggable Database** dialog, enter the following:
 - **PDB Name:** Enter a name for the PDB. The name must begin with an alphabetic character and can contain a maximum of 30 alphanumeric characters.

- **Unlock my PDB Admin account:** *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
- **PDB Admin password:** If you clicked **Unlock my PDB Admin** account, create and enter a PDB admin password. The password must contain:
 - A minimum of 9 and a maximum of 30 characters
 - At least two uppercase characters
 - At least two lowercase characters
 - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of the same characters.
 - At least two numeric characters (0 - 9)
- **Confirm PDB Admin password:** Reenter the PDB admin password.
- **TDE wallet password:** *Applicable only to databases using Oracle-managed encryption keys.* Enter the TDE wallet password for the parent CDB.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the check box is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

7. Click Create.**Note**

- Create or delete a PDB while a Data Guard setup is running on another container database within the same Oracle home, and vice versa.
- Create or delete a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.
- Create or delete a PDB concurrently on different container databases within the same Oracle home.
- Create or delete a PDB while simultaneously updating VM Cluster tags.

WHAT NEXT?

After creating your PDB, you can get [connection strings](#) for the administrative service using the OCI Console.

To relocate a pluggable database

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.

3. Navigate to the database:

Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. On the Database Details page, click **Pluggable Databases** in the **Resources** section of the page.

5. Click the name of the PDB that you want to relocate.

From the Pluggable Database details page, click **More Actions**, and then select **Relocate**.

(or)

Click the Actions menu (three dots) and select **Relocate**.

6. In the resulting Relocate Pluggable Database window, enter the following:

- **VM Cluster:** Use the menu to select the destination VM cluster.
- **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
- **New PDB name for the clone:** The name must begin with an alphabetic character and can contain up to 30 characters. To keep the PDB name the same, just re-enter the source PDB name.
- **Database TDE wallet password:** Enter the TDE wallet password for the parent CDB of the source PDB.
- **Unlock my PDB Admin Account:**
 - To enter the administrator's password, check this check box.
 - * **PDB Admin Password:** Enter PDB admin password. The password must contain:
 - * a minimum of 9 and a maximum of 30 characters
 - * at least two uppercase characters
 - * at least two lowercase characters
 - * at least two special characters. The valid special characters are underscore (_), a pound or hash sign (#), and dash (-). You can use two of the same characters or any combination of two of the same characters.
 - * at least two numeric characters (0 - 9)
 - * **Confirm PDB Admin Password:** Enter the same PDB Admin password in the confirmation field.
 - To skip entering the administrator's password, uncheck this check box. If you uncheck this check box, then the PDB is created but you cannot use it. To use the PDB, you must reset the administrator password.

Note

When you create a new PDB, a local user in the PDB is created as the administrator and granted the PDB_DBA role locally to the administrator.

To reset the password:

- a. Connect to the container where your PDB exists using the SQL*Plus `CONNECT` statement.

```
SQL> show con_name;
CON_NAME
-----
CDB$ROOT
```

For more information, see *Administering a CDB* and *Administering PDBs* in the *Oracle® Multitenant Administrator's Guide*.

- b. Find the administrator name of your PDB:

```
SQL> select grantee from cdb_role_privs where con_id = (select
con_id from cdb_pdb_name where pdb_name = '<PDB_NAME>') and
granted_role = 'PDB_DBA';
```

- c. Switch into your PDB:

```
SQL> alter session set container=<PDB_NAME>;
Session altered.
SQL> show con_name;
CON_NAME
-----
<PDB_NAME>
```

- d. Reset the PDB administrator password:

```
SQL> alter user <PDB_Admin> identified by <PASSWORD>;
User altered.
```

- **Source database SYS password:** Enter the database admin password.
- **Database link:** Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Advanced Options:**
 - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

7. Click **Relocate**.

Note

Relocate will incur downtime during the process and that the time required is based on the size of the PDB.

Using the API to create pluggable database

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

Use the [CreatePluggableDatabase](#) API to create pluggable databases on Exadata Cloud Infrastructure.

For the complete list of APIs for the Database service, see [Database Service API](#).

Managing an Exadata Pluggable Database

This topic includes the procedures to connect to, start, stop, and delete a pluggable database (PDB).

It also includes instructions for getting PDB [connection strings](#) for the administrative service.

- [To start a pluggable database](#)
- [To stop a pluggable database](#)
- [To delete a pluggable database](#)
- [To get connection strings for a pluggable database](#)
- [Convert a Physical Standby Database to Snapshot Standby Database](#)
A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.
- [Convert a Snapshot Standby Database to Physical Standby Database](#)
Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.
- [Using the API to manage pluggable databases](#)
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or assign a new key version.

To start a pluggable database

Note

The PDB must be available and stopped to use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.

3. Navigate to the database:

Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to start. Click the PDB name to display details about it.
6. Click **Start**.
7. In the **Start PDB** dialog, click **Start** to confirm the start operation.

To stop a pluggable database

Note

The PDB must be available and running (started) to use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to stop. Click the PDB name to display details about it.
6. Click **Start**.
7. In the **Stop PDB** dialog, click **Stop** to confirm the stop operation.

To delete a pluggable database

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to delete. Click the PDB name to display details about it.
6. Click **More Actions**, then choose **Delete**.
7. In the **Delete PDB** dialog box, enter the name of the PDB that you want to delete to confirm the action, then click **Delete**.

Note

You can now delete a PDB when Data Guard setup is running on another container database in the same Oracle home, and vice versa.

To get connection strings for a pluggable database

Note

This topic explains how to get connection strings for the administrative service of a PDB. Oracle recommends that you connect applications to an application service, using strings created for the application service.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:
Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the PDB, and then click its name to display details about it.
6. Click **PDB Connection**.
7. In the **Pluggable Database Connection** dialog, use the **Show** and **Copy** links to display and copy connection strings, as needed.
8. Click **Close** to exit the dialog.

Convert a Physical Standby Database to Snapshot Standby Database

A snapshot standby database is a fully updateable standby database created by converting a physical standby database into a snapshot standby database.

Snapshot standby allows administrators to perform root cause analysis and test with production data in a very simple and risk-free manner, thereby drastically reducing the risk of errors and improving resiliency. Until now, customers were able to convert a standby database into snapshot mode outside of cloud automation. Customers will now be able to create a snapshot standby database via OCI Console/API/SDK, and Terraform.

The snapshot standby feature temporarily converts the standby database to read-write mode providing full access to update the database without any disruption to the business or risk of data loss. A snapshot standby database receives and archives redo logs from the primary database but does not apply it. The snapshot standby database can be converted back to a standby database at any point in time. When converting back, all the local updates made to the snapshot standby database will be discarded and the redo logs received from the primary database will be applied converting the standby database back into an exact physical replica of the primary database.

A snapshot standby database typically diverges from its primary database over time because redo data from the primary database is not applied as it is received. Local updates to the snapshot standby database cause additional divergence. The data in the primary database is fully protected however, because a snapshot standby can be converted back into a physical standby database at any time, and the redo data received from the primary is then applied.

A snapshot standby database provides disaster recovery and data protection benefits that are similar to those of a physical standby database. Snapshot standby databases are best used in scenarios where the benefit of having a temporary, updatable snapshot of the primary database justifies increased time to recover from primary database failures.

① Note

- You will have to ensure there is sufficient space in the Fast Recovery Area. Oracle recommends converting a snapshot standby database back to a physical standby database in the shortest amount of time.
- Automatic backups are suspended on a snapshot standby. Archive logs from the production database will continue to be backed up to preserve recoverability.
- If you have performed a database upgrade while in the snapshot standby mode, you must ensure you use the correct Oracle Home when converting back to a physical standby database.

Before converting a physical standby database to a snapshot standby database, ensure the following conditions are met:

- It is recommended that you have more than one standby database in a Data Guard Group before converting a standby database to a snapshot standby database.
 - Only one physical standby can be converted to a snapshot standby at a time. To convert multiple databases, perform the conversions sequentially.
1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
 2. Choose your **Compartment**.
 3. Click on the VM Cluster containing the databases you wish to view their roles in Data Guard associations.
 4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the Data Guard role column.

5. Choose a Standby database, select **Convert to snapshot standby** from the **More Actions** drop-down list.
6. Review the warning message on the Convert to snapshot standby window.
7. Click **Convert**.

After the conversion, the role of the standby database changes to **Snapshot Standby**, and the **Convert to Standby** option becomes available in the **More Actions** drop-down list.

Convert a Snapshot Standby Database to Physical Standby Database

Oracle recommends converting a snapshot standby database back to a physical standby database within 14 days.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Click on the VM Cluster containing the databases you wish to view their roles in Data Guard associations.
4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the Data Guard role column.
5. Choose a Snapshot standby database, select **Convert to Standby** from the **More Actions** drop-down list.

Note

If you convert your snapshot standby to physical standby database, all local updates to your snapshot standby database will be discarded and data from your primary database will be applied.

6. Click **Convert**.

Note

- You cannot perform a switchover operation on a snapshot standby database. You will need to convert the snapshot standby to a standby database before initiating a switchover.
- Oracle managed automatic backups will be disabled for a snapshot standby database. Automatic backups if configured will resume upon conversion from snapshot standby to standby database.
- While any database in the Data Guard configuration is in snapshot standby mode, there is no support to create additional PDBs on the primary or any of the standby databases.

Using the API to manage pluggable databases

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

Use these APIs to manage pluggable databases.

- [ListPluggableDatabases](#)
- [GetPluggableDatabase](#)
- [StartPluggableDatabase](#)
- [StopPluggableDatabase](#)
- [UpdatePluggableDatabase](#)
- [DeletePluggableDatabase](#)

Note

Use the [GetPluggableDatabase](#) API to get administration service connection strings and other details about a PDB.

For the complete list of APIs for the Database service, see [Database Service API](#).

To administer Vault encryption keys

Use this procedure to rotate the Vault encryption key or assign a new key version.

Note

Rotate Key is blocked on standby when KMS is configured in the current database. Also, you cannot change or update the encryption type once it is configured to KMS.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your compartment from the **Compartment** drop-down.
3. Navigate to the cloud VM cluster that contains the database for which you want to change encryption management or to rotate a key.
Cloud VM clusters: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, locate the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the **Databases** section, click the name of the database in which the pluggable database you want to change encryption management or to rotate a key exists.
5. Click the name of the database to view its details.
6. Under Resources, click **Pluggable Databases**.
7. From the list, click the name of a PDB to view its details.
8. Click **Manage encryption key**.
To rotate an encryption key on a database using customer-managed keys:
Note: Generate a new master encryption key version. Only the CDB root key version is changed or rotated to a new one. It doesn't generate a new key version for the dependent PDBs. Rotate customer-managed keys periodically to comply with security compliance and regulatory mandates. The rotation involves stopping and restarting the database.
 - a. Click **Rotate Encryption Key**.
 - b. Click **Update**.

To assign a new key version:

Assign a new key version (BYOK) to CDB while creating or after provisioning it.

- a. Click **Assign a new key version**.
- b. In the **Key version OCID** field, enter the OCID of the new key version you want to assign.
- c. Click **Update**.

To copy the Key version OCID:

- a. Find the Vault and the Key details on the Key Details page (**Key Management & Secret Management >> Vault >> <Vault> >> Key Details**) by searching with the KMS key OCID provided in the PDB details page.
- b. Copy the OCID and paste it in the **Key version OCID** field.

Cloning an Exadata Pluggable Database

You can create local, remote, and refreshable clones.

A clone is an independent and complete copy of the given database as it existed at the time of the cloning operation. You can create clones of your PDB within the same CDB or a different CDB and also refresh the cloned PDB.

Note

When cloning a PDB from 19c to 26ai, the cloned PDB is automatically upgraded to 26ai. For example, if you use refreshable clones to clone to 26ai and then convert it to regular PDB, all necessary upgrade steps are automatically handled, converting the refreshable clone into a fully upgraded 26ai PDB.

The following types of clones are supported:

- **Local clone:** A clone of the PDB is created within the same CDB.
- **Remote clone:** A clone of the PDB is created in a different CDB.
- **Refreshable clone:** A clone of the PDB is created in a different CDB, and you will be able to refresh the cloned PDB.
- [To create a local clone of a pluggable database \(PDB\)](#)
- [To create a remote clone of a pluggable database \(PDB\)](#)
- [To create a refreshable clone of a pluggable database \(PDB\)](#)
- [To refresh a cloned pluggable database \(PDB\)](#)
- [To convert a refreshable clone to a regular pluggable database \(PDB\)](#)
- [Using the API to clone a pluggable database](#)

To create a local clone of a pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.

3. Navigate to the database.
Cloud VM clusters ([new resource model](#)) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.
6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:
 - **Select clone type:** Select **Local clone** to create a copy of the source PDB to the same CDB.
 - **Exadata VM Cluster:** Use the menu to select the cloud VM cluster of the target database.

Note

The target VM Cluster may be on a different Exadata infrastructure.

- **Destination database:** This field is disabled.
- **PDB name:** Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.
- **Database TDE wallet password:** *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
- **Unlock my PDB Admin account:** *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
- **PDB Admin password:** Create and enter a new PDB Admin password. The password must contain:
 - 9–30 characters
 - At least two uppercase characters
 - At least two lowercase characters
 - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
 - At least two numeric characters (0-9)
- **Confirm PDB Admin password:** Enter the PDB Admin password again to confirm.
- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see [Thin Cloning a Pluggable Database](#) in the [Oracle® Exadata Exascale User's Guide](#). If you specifically want a thick clone (full copy), you need to deselect this option.
 - **Advanced Options:**
 - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
8. Click **Clone**.

To create a remote clone of a pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database.
Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.
6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:
 - **Select clone type:** Select **Remote clone** to create a copy of the source PDB to the same CDB.
 - **Exadata VM Cluster:** Use the menu to select the cloud VM cluster of the target database.

Note

The target VM Cluster may be on a different Exadata infrastructure.

- **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.

- **PDB name:** Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.
- **Database TDE wallet password:** *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
- **Unlock my PDB Admin account:** *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
- **PDB Admin password:** Create and enter a new PDB Admin password. The password must contain:
 - 9–30 characters
 - At least two uppercase characters
 - At least two lowercase characters
 - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
 - At least two numeric characters (0-9)
- **Confirm PDB Admin password:** Enter the PDB Admin password again to confirm.
- **Database link:** Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.

Note

If you do not provide database link information, cloud automation will create a database link using the common user. However, you can specify the database link information if you want cloud automation to use a specific database link.

- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see [Thin Cloning a Pluggable Database](#) in the [Oracle® Exadata Exascale User's Guide](#). If you specifically want a thick clone (full copy), you need to deselect this option.

Note

The thin clone option will be disabled (greyed out) if the source and target VM clusters do not share the same vault. In such cases, only thick clones are supported.

- **Advanced Options:**
 - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
8. Click **Clone**.

To create a refreshable clone of a pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database.
Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to clone, and then click its name to display details about it.
6. Click **Clone**.
7. In the **Clone PDB** dialog box, enter the following:
 - **Select clone type:** Select Refreshable clone to create a copy of the source PDB to the same CDB.
For more information about refreshable clones, see [About Refreshable Clone PDBs](#).
 - **Exadata VM Cluster:** Use the menu to select the cloud VM cluster of the target database.

Note

The target VM Cluster may be on a different Exadata infrastructure.

- **Destination database:** Use the menu to select an existing database where the PDB will be created. This database can be of the same version as the CDB the source PDB is in or of a higher version.
- **PDB name:** Provide a name for the new cloned PDB. The name must begin with an alphabetic character and can contain up to 30 characters.

- **Database TDE wallet password:** *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
- **Unlock my PDB Admin account:** *Optional.* Select this option to specify a PDB Admin password and configure the PDB to be unlocked at creation.
- **PDB Admin password:** Create and enter a new PDB Admin password. The password must contain:
 - 9–30 characters
 - At least two uppercase characters
 - At least two lowercase characters
 - At least two special characters. The valid special characters are: underscore (_), a hash sign (#), and a dash (-). You can use two of the same characters or any combination of two of these characters.
 - At least two numeric characters (0-9)
- **Confirm PDB Admin password:** Enter the PDB Admin password again to confirm.
- **Database link:** Enter the user name and password for the database link. Note that the user must be precreated in the source database. The DB link will be created in the destination using that username and password.

Note

If you do not provide database link information, cloud automation will create a database link using the common user. However, you can specify the database link information if you want cloud automation to use a specific database link.

- **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

- **Enable thin clone:** All clones are thin clones by default. For more information about thin clone, see [Thin Cloning a Pluggable Database](#) in the [Oracle® Exadata Exascale User's Guide](#). If you specifically want a thick clone (full copy), you need to deselect this option.

Note

The thin clone option will be disabled (greyed out) if the source and target VM clusters do not share the same vault. In such cases, only thick clones are supported.

- **Advanced Options:**

- **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see Resource Tags. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator.
8. Click **Clone**.

To refresh a cloned pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters ([new resource model](#)) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to refresh, and then click its name to display details about it.
6. Click **More Actions** and select **Refresh**.
7. In the resulting **Refresh** dialog box, click **Refresh** to confirm.

To convert a refreshable clone to a regular pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters ([new resource model](#)) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to convert to a regular PDB, and then click its name to display details about it.
6. In the resulting **Convert to regular PDB** dialog box, enter the following:
 - **Database TDE wallet password:** *Not applicable for databases using customer-managed keys from the Vault service.* Enter the TDE wallet password for the parent database (CDB) of the source PDB.
 - **Take a backup of the PDB immediately after creating it:** You must enable auto-backup on the CDB to back up a PDB immediately after creating it. This check box is checked by default if auto-backup was enabled on the CDB.

Note

If the checkbox is unchecked, the system displays a warning stating that PDB cannot be recovered until the next daily backup has been successfully completed.

7. Click **Convert**.

Using the API to clone a pluggable database

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

Use these APIs to clone pluggable databases:

- [LocalclonePluggableDatabase](#)
- [RemoteclonePluggabledatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Concurrently Create or Delete Pluggable Databases (PDBs)

You can now create or delete up to 10 PDBs concurrently, even when the container database (CDB) is in an updating state. However, you cannot create or delete PDBs while a CDB is in an updating state if other operations—excluding PDB creation or deletion—are modifying its metadata or structure.

The maximum number of CDBs and PDBs that can be created on a cluster is determined by the available memory on the VMs. By default, for Oracle Exadata Database Service on Dedicated Infrastructure, each CDB is allocated 12.6 GB of memory (7.6 GB for SGA and 5 GB for PGA) if the VM has more than 60 GB of total memory. For VMs with 60 GB or less, 6.3 GB is allocated (3.8 GB for SGA and 2.5 GB for PGA).

In addition, Grid Infrastructure or ASM typically consumes 2 to 4 GB of memory. You can adjust the VM's memory allocation if necessary. It's important to note that the memory allocated to a CDB is shared among its PDBs. If a CDB has insufficient memory, PDB creation will fail.

Additionally, PDBs can now be created or deleted concurrently in a Data Guard environment. There are no restrictions on concurrent PDB creation or deletion in non-Data Guard environments.

Note:

- Concurrent creation of PDBs in an Oracle Database configured with Data Guard and using a file-based wallet for TDE encryption keys (Oracle-managed keys) is also supported, however Oracle will serially create the PDBs on the standby database server in a serial manner. However, you can delete them concurrently.
- A PDB cannot be deleted if its creation is still in progress.

Operations you can perform in parallel:

- Create or delete CDBs in a VM cluster.
- Create a CDB in a VM cluster while the deletion of another CDB is in progress within the same cluster.
- Delete a CDB in a VM cluster while the creation of another CDB is in progress within the same cluster.

- In a CDB, delete a PDB (PDB1) while another PDB (PDB2) is being created.
- Create or delete additional PDBs if the CDB is in an updating state.
- Scale OCPU while a PDB is being created or deleted.

Related Topics

- [Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel](#)

Restoring an Exadata Pluggable Database

You can perform in-place and out of place restore of an Exadata pluggable database.

The following types of clones are supported:

- **In place restore:** You can restore a PDB within the same CDB to last known good state or to a specified timestamp.
- **Out of place restore:** You can restore a PDB by creating a database (CDB) from the backup, then selecting a PDB or a subset of them you want to restore on the new database.
- [To perform an in-place restore of a pluggable database \(PDB\)](#)
- [To perform an out-of-place restore of a pluggable database \(PDB\)](#)

To perform an in-place restore of a pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters (new resource model) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.

4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to restore, and then click its name to display details about it.
6. In the resulting Restore PDB dialog, enter the following:
 - **Restore to latest:** Select this option to restore and recover the database with zero, or least possible, data loss.
 - **Restore to a timestamp:** Select this option to restore and recover the database to the specified timestamp.
7. Click **Restore**.

To perform an out-of-place restore of a pluggable database (PDB)

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Choose your Compartment.
3. Navigate to the database:

Cloud VM clusters ([new resource model](#)) Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud VM cluster details page, in the **Databases** table, click the name of the database to display the Database Details page.
4. Click **Pluggable Databases** in the **Resources** section of the page.
5. In the list of pluggable databases, find the pluggable database (PDB) you want to restore, and then click its name to display details about it.
6. Under **Resources**, click **Backups**.
7. From the list of backups, choose a backup, click the Actions menu (three dots), and then select **Create Database**.
8. In the resulting Create database from backup dialog box, select either of these options, **Select all PDBs** or **Specify the PDBs to restore**.
 - [To create a database by selecting all pluggable databases](#)
Provide the requested information in the Create database from backup page:
 - [To create a database by specifying a subset of pluggable databases](#)
Provide the requested information in the Create database from backup page:

To create a database by selecting all pluggable databases

Provide the requested information in the Create database from backup page:

1. Click **Select all PDBs**.
2. Click **Next**.
3. Select the VM cluster where you want to create the database.
Click the **Change Compartment** hyperlink to choose your compartment.
4. **Configure Database Home**: Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
 - **Select an existing Database Home**: If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
 - **Create a new Database Home**: If no Database Homes exist for the database version you have selected, then this option is selected by default.
 - a. Enter **Database Home display name**.
 - b. Click **Change Database Image** to select your software version.
Select a **Database Software Image** window is displayed.
 - c. Select an **Image Type**, **Oracle Provided Database Software Images**, or **Custom Database Software Images**.
If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

Note

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

5. **Provide the database name:** Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters.

Review the following guidelines when selecting a database name.

- maximum of 8 characters
- contain only alphanumeric characters
- begin with an alphabetic character
- cannot be part of first 8 characters of a `db_unique_name` on the VM cluster
- unique within a VM cluster
- **DO NOT** use `grid` because `grid` is a reserved name
- **DO NOT** use `ASM` because `ASM` is a reserved name

6. **Provide a unique name for the database:** Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

The `db_unique_name` must contain only the permitted characters.

When choosing a database name, ensure the following:

- The name is a maximum of 30 characters.
- It contains only alphanumeric characters and underscores (`_`).
- It begins with an alphabetic character.
- The name is unique across your fleet/tenancy.

If a unique name is not provided, then the `db_unique_name` defaults to the following format `<db_name>_<3 char unique string>_<region-name>`.

If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

7. **Provide the administration password:** Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:

- The password for the Oracle Database `SYS` and `SYSTEM` users.
- The Transparent Data Encryption (TDE) Keystore password.

For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.

8. **Enter the source database's TDE wallet or RMAN password:** Password must match the TDE wallet or RMAN password of the source database contained in the backup.
9. Click **Create Backup**.

To create a database by specifying a subset of pluggable databases

Provide the requested information in the Create database from backup page:

1. Click **Specify the PDBs to restore**.
2. In the **Specify PDB to restore** field, provide a comma-delimited list of PDBs to restore.
3. Click **Next**.
4. Select the VM cluster where you want to create the database.

Click the **Change Compartment** hyperlink to choose your compartment.

5. **Configure Database Home:** Select an existing Database Home or create one as applicable. Note that this field is not available when you create a Database from the Database Home details page.
 - **Select an existing Database Home:** If one or more Database Homes already exist for the database version you have selected, then this option is selected by default. And, you will be presented with a list of Database Homes. Select a Database Home from the list.
 - **Create a new Database Home:** If no Database Homes exist for the database version you have selected, then this option is selected by default.
 - a. Enter **Database Home display name**.
 - b. Click **Change Database Image** to select your software version. Select a **Database Software Image** window is displayed.
 - c. Select an **Image Type, Oracle Provided Database Software Images, or Custom Database Software Images**.
If you choose **Oracle Provided Database Software Images**, then you can use the **Display all available version** switch to choose from all available PSUs and RUs. The most recent release for each major version is indicated with a **latest** label.

Note

For the Oracle Database major version releases available in Oracle Cloud Infrastructure, images are provided for the current version plus the three most recent older versions (N through N - 3). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, images available for provisioning are for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.

6. **Provide the database name:** Specify a user-friendly name that you can use to identify the database. The database name must contain only the permitted characters.

Review the following guidelines when selecting a database name.

- maximum of 8 characters
- contain only alphanumeric characters
- begin with an alphabetic character
- cannot be part of first 8 characters of a `db_unique_name` on the VM cluster
- unique within a VM cluster
- **DO NOT** use `grid` because `grid` is a reserved name
- **DO NOT** use `ASM` because `ASM` is a reserved name

- 7. Provide a unique name for the database:** Optionally, specify a unique name for the database. This attribute defines the value of the `db_unique_name` database parameter. The value is case insensitive.

The `db_unique_name` must contain only the permitted characters. Review the following guidelines when selecting a database name.

- maximum of 30 characters
- can contain alphanumeric and underscore (`_`) characters
- begin with an alphabetic character
- unique across the fleet/tenancy

If a unique name is not provided, then the `db_unique_name` defaults to the following format `<db_name>_<3 char unique string>_<region-name>`.

If you plan to configure the database for backup to a Recovery Appliance backup destination, then the unique database name must match the name that is configured in the Recovery Appliance.

- 8. Provide the administration password:** Provide and confirm the Oracle Database administration password. This password is used for administration accounts and functions in the database, including:
 - The password for the Oracle Database `SYS` and `SYSTEM` users.
 - The Transparent Data Encryption (TDE) Keystore password.For Oracle Database 12c Release 1 or later releases, the password for the PDB administration user in the first PDB (`PDBADMIN`) must be nine to 30 characters and contain at least two uppercase, two lowercase, two numeric, and two special characters. The special characters must be `_`, `#`, or `-`. In addition, the password must not contain the name of the tenancy or any reserved words, such as `Oracle` or `Table`, regardless of casing.
- 9. Enter the source database's TDE wallet or RMAN password:** Password must match the TDE wallet or RMAN password of the source database contained in the backup.
- 10. Click `Create Backup`.**

Cost and Usage Attribution for Pluggable Databases (PDBs)

Note

It is supported only on Oracle Databases 19c and higher running in a multitenant deployment.

With this enhancement to the Cost Analysis feature of the OCI Cost Management Service, you can view the attributed usage and cost for all the PDBs in a VM Cluster. This data will be available on the cost analysis dashboard and the reports.

Prerequisites:

- `dbaastools`: (minimum version) 24.2.1
 - To check the version of the `dbaastools` rpm on the guest VM, run: `rpm -qa | grep dbaastools`
 - To update the `dbaastools` rpm on the guest VM, run: `dbaascli admin updateStack`

Confirm you have the minimum version of `dbaastools` needed after you update the `dbaastools rpm` by running the `rpm -qa | grep dbaastools` command.

- `dbcsagent` needs to be running on the guest VM. Minimum version of `dbcsagent` needed is 23.3.2.
 - To check the version of the `dbcsagent` on the guest VM, run: `rpm -qa | grep dbcs-agent-update`
 - You will need to open a service request on [My Oracle Support](#) to update the `dbcsagent` on the guest VM.
 - To check the status of the `dbcsagent`, run: `systemctl status dbcsagent`
Run `systemctl start dbcsagent` if the `dbcsagent` is not in active (running) state.
Check the status of the agent again to confirm that it is running.
- [Generate Attributed Cost Analysis Report for Pluggable Databases](#)
Follow the steps below to view the attributed costs based on CPU utilization for all pluggable databases within a VM Cluster.

Generate Attributed Cost Analysis Report for Pluggable Databases

Follow the steps below to view the attributed costs based on CPU utilization for all pluggable databases within a VM Cluster.

1. Open the navigation menu and click **Billing & Cost Management**. Under **Cost Management**, click **Cost Analysis**.
2. From **Reports**, select one of the predefined reports, or use the default **Costs by Service report**.
3. Make your preferred query adjustments.
 - a. From **Start/End Date (UTC)**, select a time period.
 - b. From **Granularity**, select **Daily or Monthly**.
 - c. From **Show**, select **Attributed cost**.
 - d. From **Filters**, select **Tag**.
In the resulting **Tag** dialog, select **orcl-cloud** as the tag with the key **parent_resource_id_1** equal to the OCID of the VM Cluster.
 - e. From **Grouping dimensions**, select the preferred grouping dimension. For example, **Resource OCID**.
The VM Cluster OCID is the parent of the CDBs it contains, and the CDB OCID is the parent OCID of the PDBs it contains.
 - f. Click **Apply** to apply the changes and reload the chart and table with the selected filters.
The generated report will show the attributed costs for all the PDBs in the VM Cluster.
4. After you have made changes, the currently selected predefined report name from the **Reports** menu changes to (edited).
5. If you're done making changes and want to save a new report, click **Save** as new report.
6. In the **Save as new report** dialog, enter the report name in the **Name** field. Avoid entering confidential information..
7. Click **Save**.
A notification is displayed that your report has been saved, and the report is also selected in the **Reports** menu.

8. If you didn't already apply your custom report settings, click **Apply** to view your changes. The new saved report is now available for future selection from the **Reports** menu under **Saved Reports**.

For more information about generating a PDB attributed cost analysis report, see [Cost Analysis](#).

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Cloud Infrastructure instance or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB administrator credentials. Use the following procedures if you need to change passwords for an existing database.

Note

if you are enabling Data Guard for a database, then the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

Note

Using the `dbaascli` to change the SYS password will ensure the backup/restore automation can parallelize channels across all nodes in the cluster.

To Change the SYS Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli database changepassword --dbname database_name --user SYS
```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

```
dbaascli database changePassword --dbName <dbname> --user SYS --  
prepareStandbyBlob true --blobLocation <location to create the blob file>
```

2. Copy the blob file created to all the standby databases and update the file ownership to oracle user.
3. Run the following command on all the standby databases:

```
dbaascli database changePassword --dbName <dbname> --user SYS --  
standbyBlobFromPrimary <location of copies the blob file>
```

To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.

2. Run the following command:

```
sudo dbaascli tde changepassword --dbname database_name
```

Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure

Learn how to work with the backup and recovery facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure.

- [Oracle Recommended Options to Perform Backup and Recovery Operations](#)
Oracle offers the following options for Oracle Database Backup and Recovery operations. These options are mutually exclusive.
- [Managing Exadata Database Backups](#)
Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.
- [Managed Backup Types and Usage Information](#)
There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.
- [Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console](#)
- [Long-Term Retention Backup with Recovery Service](#)
Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.
- [Default Backup Channel Allocation](#)
The default settings for database backup channels when using "Oracle Managed Backup" or "User Configured Backup"
- [Prerequisites for Backups on Exadata Cloud Infrastructure](#)
- [Using the Console to Manage Backups](#)
- [To designate Autonomous Recovery Service as a Backup Destination for an Existing Database](#)
To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.
- [Recovering an Exadata Database from Backup Destination](#)
This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.
- [Managing Exadata Database Backups by Using dbaascli](#)
- [Using the API to Manage Backup and Recovery](#)
- [Alternative Backup Methods](#)
Learn about alternative backup methods that are available in addition to the OCI Console.
- [Recovering a Database Using Oracle Recovery Manager \(RMAN\)](#)

Oracle Recommended Options to Perform Backup and Recovery Operations

Oracle offers the following options for Oracle Database Backup and Recovery operations. These options are mutually exclusive.

Note

A hybrid configuration, that is, mixing the options is not supported. Mixing the options will break automation.

Option 1: Oracle Managed Backups

Oracle managed backups are entirely managed by Exadata Cloud Infrastructure (ExaDB-D) or Exadata Cloud@Customer (ExaDB-C@C) based on a one-time configuration. Besides being fully integrated into ExaDB-D or ExaDB-C@C cloud services Control Plane, these backups can also be accessed through OCI APIs. Oracle recommends this approach.

- The `dbaascli database backup` and `dbaascli database recover` commands can be used in conjunction with the automated backups for certain operations. For more information, see `dbaascli database backup` and `dbaascli database recover`.
- Customers are allowed to query RMAN views or issue RMAN restore and recovery commands, for example, `table`, `datafile`, or `tablespace recovery` commands.

Note

Do not use RMAN configuration to change any of the pre-tuned cloud RMAN settings.

Option 2: User Configured Backups

Customers can also configure backups from the host using the `dbaascli database backup` and `dbaascli database recover` commands. These backups, however, are not synchronized with the Control Plane nor are they integrated with the OCI APIs. Also, neither management nor lifecycle operations on these backups are supported from the service Control Plane console. Hence, this is not a recommended approach.

This approach is useful when direct access to Backup destinations is required to perform certain tasks. Accessing the OSS bucket, for example, to replicate backups across regions or monitor Backup Destinations.

If customers configure backups to Object Storage using RMAN without using the OCI Control Plane or OCI APIs, customers are responsible for manually configuring TDE Wallet backups. By default, Oracle cloud automation cleans up archive log files every 24 hours. When you use RMAN to perform manual backups, there is a risk of the archive logs being deleted. Refer to [dbaascli database backup](#) for information on how to configure the archive log cleanup. The recommendation is to use Oracle managed backups.

For more information, see *User Configured Backup*.

Option 3: Backups using RMAN

Backups can be directly taken using RMAN with customer-owned customized scripts. Oracle, however, does not recommend this approach.

It is not recommended to use RMAN backups in conjunction with Oracle Managed Backups or User Configured Backups.

Who can use this option:

- Customers who want to maintain their existing RMAN backup/restore scripts.
- Customers who want to configure backups from Standby database in Data Guard environments to offload the backup workload to Standby.

ExaDB-D:

If you plan to backup using RMAN, then you must unregister the database from backup automation. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

Related Topics

- [dbaascli database backup](#)
- [dbaascli database recover](#)
- [Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management](#)

Managing Exadata Database Backups

Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.

For unmanaged backups, see *Managing Exadata Database Backups by Using dbaascli*.

There are two destinations possible for automatic Exadata database backups: Autonomous Recovery Service, or Oracle Object Storage.

The Oracle-managed automatic backups feature is the preferred method for backing up Oracle Cloud databases because you can easily configure backup settings using the Console. The automatic backups feature supports Recovery Service and Object Storage as the backup destination to provide you with a fully automated cloud backup solution with the same cost. You do not need to perform any manual backups or backup storage administration tasks. You can also store backups in local storage. Each backup destination has its advantages and requirements that you should consider, as described below.

Recovery Service (Recommended)

A fully managed service based on the on-premises Oracle's [Zero Data Loss Recovery Appliance](#) technology which offers modern cybersecurity protection for Oracle Databases. Unique, automated capabilities protect Oracle Database changes in real time, validate backups without production database overhead, and enable fast, predictable recovery to any point in time.

If your backups are currently configured with Object Storage, you can seamlessly transition to Recovery Service to achieve advanced capabilities with the same cost.

For more information on Recovery Service, see [About Oracle Database Autonomous Recovery Service](#).

Object Storage

A secure, scalable, on-demand storage solution for databases.

Note

If you previously used `dbaascli` to configure backups and then you switch to using the Console or the API for backups:

- A new backup configuration is created and associated with your database. This means that you can no longer rely on your previously configured unmanaged backups to protect your database.

Related Topics

- [Managing Exadata Database Backups by Using dbaascli](#)

Managed Backup Types and Usage Information

There are two types of automatic Exadata database backups: Autonomous Recovery Service, and Oracle Object Storage.

The database and infrastructure (the VM cluster) must be in an “Available” state for a backup operation to run successfully. Oracle recommends that you avoid performing actions that could interfere with availability (such as patching operations) while a backup operation is in progress. If an automatic backup operation fails, the Database service retries the operation during the next day’s backup window. If an on-demand full backup fails, you can try the operation again when the Exadata Cloud Infrastructure instance and database availability are restored.

When you enable the Automatic Backup feature, either service creates daily incremental backups of the database to the selected Backup Destination.

If you choose to enable automatic backups, you can control the retention period. The system automatically deletes backups when the assigned retention period is expired.

Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30 days.

The automatic backup process starts at any time during your daily backup window. You can optionally specify a 2-hour scheduling window for your database during which the automatic backup process will begin. There are 12 scheduling windows to choose from, each starting on an even-numbered hour (for example, one window runs from 4:00-6:00 AM, and the next from 6:00-8:00 AM). Backups jobs do not necessarily complete within the scheduling window.

The default backup window of 00:00 to 06:00 in the time zone of the Exadata Cloud Infrastructure instance’s region is assigned to your database if you do not specify a window. Note that the default backup scheduling window is six hours long, while the windows you specify are two hours long.

Autonomous Recovery Service protection policy:

- **Bronze** :14 days
- **Silver**: 35 days
- **Gold**: 65 days
- **Platinum**: 95 days
- Custom defined by you
- **Default**: Silver - 35 days

The automatic backup process starts at any time or within the assigned window.

① Note

- **Data Guard:** You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association.
- **Backup Retention Changes:** If you shorten your database's backup retention period or your protection policy in the future, existing backups falling outside the updated retention period are deleted by the system.
- **Backup Storage Costs:** Automatic backups incur storage usage costs for either Autonomous Recovery Service or Object Storage depending on the backup destination selected.

You can create a full backup of your database at any time using either service.

When you terminate an Exadata Cloud Service instance database, all of its resources are deleted. Managed backups using the Object Storage destination will be deleted, and Managed backups using the Autonomous Recovery Service will be deleted according to the deletion option selected. Standalone backups created in Object Storage will remain after the database is terminated and must be manually deleted. You can use a standalone backup to create a new database.

To align with the Oracle recommended practice of using SYSBACKUP administrative privilege for Backup and Recovery operations, cloud automation creates a common administrative user C##DBLCMUSER with SYSBACKUP role at the CDB\$ROOT container level. Backup and Recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation. If the user is not found or is LOCKED and EXPIRED, then cloud automation will recreate or unlock this user during the backup or recovery operation. This change in the cloud automation is made starting with *dbaastools version 21.4.1.1.0*.

Related Topics

- [Release 21.4.1.1.0 \(220209\)](#)
- [To terminate a database](#)

Backup Destination Behavior When Enabling Automatic Backups and Standalone Backups Using the OCI Console

Effective August 06, 2025, when you enable automatic backups in the OCI Console, the Autonomous Recovery Service will be the only available backup destination under the following conditions:

- The tenancy was created on or after 06 August 2025.
- The database is deployed in OCI regions Frankfurt (FRA), Phoenix (PHX), and Tokyo (NRT).
- The Oracle Database version is later than 19.18 or 23.4.

If these conditions are not met, OCI Object Storage will be shown as a Backup destination.

Long-Term Retention Backup with Recovery Service

Long-term retention backup (LTR) allows you to store full backups for periods up to ten years for compliance, regulatory, or other business needs with complete LTR lifecycle management and immutability.

For LTR with Recovery Service, the retention period must be in Days (90 - 3,650) or years (1 - 10) from when the backup was created.

To create an LTR backup with the required retention period, Recovery Service does not require creating a new full production backup but does so by utilizing already existing operational backups in the system within the defined recovery window in the policy. For more information, see [To create an on-demand backup of a database](#).

You can change the retention period for a specific existing LTR backup within the retention period. For more information, see [To change the retention period of an LTR backup with Recovery Service](#).

You can restore an LTR backup to create a new database within the retention period. For more information, see [To create a database from a backup](#).

Upon terminating a database, the LTR backups will be deleted as per the 'Deletion options after database termination' value.

- **Delete backups in 72 hours:** All backups, including long-term backups, will be deleted.
- **Delete based on policy:** LTR backups will be retained according to the retention policy of each LTR backup.

Note: Oracle recommends choosing '**Delete based on policy**' option while terminating a database to ensure the long-term backups are retained.

Consider the following additional factors for long-term backups:

- LTR backups will continue to exist independently of any automatic backups configured on the database.
- **LTR backups will be automatically deleted after the specified retention period ends.**
- In-place restore is not supported for LTR.
- For databases in a Data Guard configuration, the long-term backup will be created only for the database where it is requested.
- The database must be in the AVAILABLE state to create a LTR.
- LTR is supported for databases with file-based TDE or KMS-based keystores.
- Encryption keys will be maintained for the entire retention period of the LTR.
- An LTR backup can be canceled while it is in the 'creating' state.
- An LTR backup can be deleted at any time after it is created.
- During restore:
 - If the backup is of a supported DBHome major version, it will be restored to the latest RU of that version.
 - If the backup is of an unsupported DBHome major version, it will be restored to a supported major version, after which the database must be upgraded to any of the supported major versions.

Default Backup Channel Allocation

The default settings for database backup channels when using "Oracle Managed Backup" or "User Configured Backup"

When a database is configured for backup using "Oracle Managed Backup" or "User Configured Backup", the tooling uses "default" for the backup channels. When default is used, dbaas will determine the number of channels to allocate at the time the backup or restore command is executed. The number of channels allocated is determined by the OCPU count of the node. The following table provides the values used and the OCPU range, both the OCPU and the channel values are per node. Restore operations are prioritized. The cluster-wide total channel count is the per node value multiplied by the number of nodes. The automation uses the SCAN to distribute RMAN channels across all nodes in the cluster.

OCPUs Per Node	Formula	Backup Channels Allocation Per Node	Restore Channels Allocation Per Node
Less than or equal to 12	$OCPU \leq 12$	2	4
Greater than 12 and less than or equal to 24	$OCPU > 12$ and $OCPU \leq 24$	4	8
Greater than 24	$OCPU > 24$	8	16

If needed, a static per node value can be set by using the DBAASCLI getConfig/configure to generate a bckup cfg and setting the parameter `bkup_channels_node` to the number of channels per node desired.

Valid values are 1 - 32: The total channel count will be the value times the number of nodes. This value cannot exceed the limit of 255 channels. A value of `default` for `bkup_channels_node` sets OCPU channel based allocation.

Prerequisites for Backups on Exadata Cloud Infrastructure

Note

Starting with dbaascli release 25.3.1.0.0, the Data Guard broker is mandatory when configuring Zero Data Loss Recovery Appliance or Autonomous Recovery Service as a backup destination in Data Guard-enabled databases.

Autonomous Recovery Service

Ensure that your tenancy is configured to use Autonomous Recovery Service.

Table 5-5 Review the prerequisite tasks before you use Recovery Service as the automatic backup destination

Task	More Information	Required or Optional
Create IAM policies	Policies to Enable Access to Recovery Service and Related Resources	Required

Table 5-5 (Cont.) Review the prerequisite tasks before you use Recovery Service as the automatic backup destination

Task	More Information	Required or Optional
Configure network resources and register a Recovery Service subnet	Creating a Recovery Service Subnet in the Database VCN	Required
Create protection policies	Review Protection Policies for Database Backup Retention	Optional

For more information about Recovery Service, see [Overview of Oracle Database Autonomous Recovery Service](#).

Object Storage

- The Exadata Cloud Service instance requires access to the Oracle Cloud Infrastructure Object Storage. Oracle recommends using a service gateway with the VCN to enable this access. For more information, see [Network Setup for Exadata Cloud Infrastructure Instances](#). In that topic, pay particular attention to:
 - [Service Gateway for the VCN](#)
 - [Node Access to Object Storage: Static Route](#)
 - Backup egress rule: Allows access to Object Storage*
 - [Subnet Size Requirements and Security Rules for Recovery Service Subnet](#)
- An existing Object Storage bucket to use as the backup destination. You can use the Console or the Object Storage API to create the bucket. For more information, see [Managing Buckets](#).
- An auth token generated by Oracle Cloud Infrastructure. You can use the Console or the IAM API to generate the password. For more information, see [Working with Auth Tokens](#).
- The user name specified in the backup configuration file must have tenancy-level access to Object Storage. An easy way to do this is to add the user name to the Administrators group. However, that allows access to all of the cloud services. Instead, an administrator should create a policy like the following that limits access to only the required resources in Object Storage for backing up and restoring the database:

```
Allow group <group_name> to manage objects in compartment
<compartment_name> where target.bucket.name = '<bucket_name>'
Allow group <group_name> to read buckets in compartment <compartment_name>
```

For more information about adding a user to a group, see [Managing Groups](#). For more information about policies, see [Getting Started with Policies](#).

Related Topics

- [Auth Token](#)

Using the Console to Manage Backups

You can use the Console to enable automatic incremental backups, create full backups on demand, and view the list of managed backups for a database. You can also use the Console to delete manual (on-demand) backups.

Note

- All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.
- Backups for a particular database are listed on the details page for that database. The Encryption Key column displays either Oracle-Managed Key or a key name if you are using your own encryption keys to protect the database. See [Backing Up Vaults and Keys](#) for more information.

Note

Do not delete any necessary encryption keys from the vault because this causes databases and backups protected by the key to become unavailable.

- [To configure automatic backups for a database](#)
- [To create an on-demand backup of a database](#)
- [To view backup status](#)
- [To cancel a backup](#)
- [To delete full backups from Object Storage](#)
- [To delete standalone backups from Object Storage](#)
- [To change the retention period of an LTR backup with Recovery Service](#)

To configure automatic backups for a database

When you create an Exadata Cloud Infrastructure instance, you can optionally enable automatic backups for the initial database. Use this procedure to enable or disable automatic backups after the database is created.

Note

Databases in a *security zone compartment* must have automatic backups enabled. See the *Security Zone Policies* topic for a full list of policies that affect Database service resources.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Exadata on Oracle Public Cloud**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to configure: **Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model)**: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database for which you want to enable or disable automatic backups, and click its name to display database details. The details indicate whether automatic backups are enabled.

5. Click **Enable automatic backups**.
6. In the Configure Automatic Backups dialog, enter the following details.

Note

Operational backups to two different backup destinations can create data loss scenarios. Therefore, before you enable automatic backups, you must disable manual backup scripts and processes to other storage destinations.

- **Backup Destination:** Your choices are **Autonomous Recovery Service** (default) or **Object Storage**.

Note

To use Autonomous Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

- **Scenario 1:** The customer enables automatic backups AND has available limits AND there is available capacity in the region for Autonomous Recovery Service.

Backup Destination: Your choices are Autonomous Recovery Service (default) or Object Storage. You can switch the backup destination from Autonomous Recovery Service to Object Storage.

- **Scenario 2:** Customer enables automatic backups AND has exhausted the default limits for the Recovery Service AND there is available capacity in the region for Autonomous Recovery Service.

Backup Destination: You can only use Object Storage. However, you can make an additional limits request and then use Autonomous Recovery Service.

The system displays the following message with a link to request an increase to the limits.

```
Tenancy has reached the limit for Autonomous Recovery Service. View your service limits and request an update.
```

- **Scenario 3:** Customer enables automatic backups AND there is no available capacity in the region for Autonomous Recovery Service.

Backup Destination: You can only use Object Storage. You can transition to Autonomous Recovery Service when there is sufficient capacity.

The system displays the following message

```
Autonomous Recovery Service has no available capacity in this region. Select Object Storage as your backup destination. You can transition from Object Storage to Autonomous Recovery Service when there is sufficient capacity.
```

Proactively check if Autonomous Recovery Service capacity is available. If the required capacity becomes available and if you had chosen Object Storage, then you can transition to Autonomous Recovery Service.

- **Backup Scheduling:**
 - **Object Storage (L0):**

- * **Full backup scheduling day:** Choose a day of the week for the initial and future L0 backups to start.
 - * **Full backup scheduling time (UTC):** Specify the time window when the full backups start when the automatic backup capability is selected.
 - * **Take the first backup immediately:** A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.

If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
- **Object Storage (L1):**
 - * **Incremental backup scheduling time (UTC):** Specify the time window when the incremental backups start when the automatic backup capability is selected.
 - **Autonomous Recovery Service (L0):**
 - * **Scheduled day for initial backup:** Choose a day of the week for the initial backup.
 - * **Scheduled time for initial backup (UTC):** Select the time window for the initial backup.
 - * **Take the first backup immediately:** A full backup is an operating system backup of all datafiles and the control file that constitute an Oracle Database. A full backup should also include the parameter file(s) associated with the database. You can take a full database backup when the database is shut down or while the database is open. You should not normally take a full backup after an instance failure or other unusual circumstances.
If you choose to defer the first full backup your database may not be recoverable in the event of a database failure.
 - **Autonomous Recovery Service (L1):**
 - * **Scheduled time for daily backup (UTC):** Specify the time window when the incremental backups start when the automatic backup capability is selected.
 - **Deletion options after database termination:** Options that you can use to retain protected database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
 - * **Retain backups for the period specified in your protection policy or backup retention period:** Select this option if you want to retain database backups for the entire period defined in the Object Storage Backup retention period or Autonomous Recovery Service protection policy after the database is terminated.
 - * **Retain backups for 72 hours, then delete:** Select this option to retain backups for a period of 72 hours after you terminate the database.
 - **Enable Real-Time Data Protection:** Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

7. Click **Save**.

The Database Details page displays the configuration details, **Health**, **Real-Time Data Protection**, and **Policy information** in the **Backup** section.

① Note

When Real-Time Data Protection is enabled with Autonomous Recovery Service as the backup destination, the tooling automatically provisions a VPC user on the protected database. This user has the same identifier as the VPC user on the recovery appliance and is granted `SYSOPER` privileges to enable redo transport and backup operations.

Related Topics

- [security zone compartment](#)
- [Security Zone Policies](#)
- [The New Exadata Cloud Infrastructure Resource Model](#)

To create an on-demand backup of a database

① Note

Object Storage creates a full backup of the database while Recovery Service creates an incremental backup.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to back up: *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database for which you want to create an on-demand full backup and click its name to display database details.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. Click **Create Backup**.
7. On the resulting Create backup window, do the following:
 - **Name**: Provide a descriptive name for the backup.
 - Select a **Backup retention** option:
 - **Retain backups per backup retention period**: Select this option to use the protection policy retention period for this backup.
 - **Specify long-term backup retention period**: Select this option to specify an LTR period with Autonomous Recovery Service. The retention period must be entered in Days (90 - 3,650) or Years (1 - 10) from when the backup was created.

8. Click **Create**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To view backup status

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup you want to view.
4. Click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
5. In the list of databases, find the database you are interested in and click its name to display database details.
6. Under **Resources**, click **Backups**.
A list of backups is displayed. The state column displays the status of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.

To cancel a backup

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup you want to view:
4. Click **Exadata VM Clusters**.
In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
5. In the list of databases, find the database you are interested in and click its name to display database details.
6. Under **Resources**, click **Backups**.
A list of backups is displayed. The state column displays the status of the backup: **Active**, **Creating**, **Canceled**, **Canceling**, or **Failed**.
7. A backup in the Creating state may be canceled by clicking the Actions icon (three dots) on the right of the backup row and clicking **Cancel Backup**.
A Cancel Backup confirmation dialog will appear.
8. Enter the name of the backup, and click **Cancel Backup**.
The state changes to **Canceling**.

The Cancel backup Work request can be viewed, by clicking **Work requests** under **Resources**.

If the Cancel backup fails:

- In the Work requests pane under Resources, you will see a line item called "**Cancel Database Backup**" with a state of "**Failed**". There will also be a work request for the backup "**Create Database Backup**" that will reflect the state of the Backup operation.

To delete full backups from Object Storage

Note

You cannot explicitly delete automatic backups. Unless you terminate the database, automatic backups remain in Recovery Service and Object Storage for the number of days specified by the user, after which time they are automatically deleted.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database backup you want to delete:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database you are interested in and click its name to display database details.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. Click the Actions icon (three dots) for the backup you are interested in, and then click **Delete**.
7. Confirm when prompted.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To delete standalone backups from Object Storage

1. Open the navigation menu. Click **Oracle AI Database**, then click **Standalone Backups** under **Resources**.
2. In the list of standalone backups, find the backup you want to use to delete.
3. Click the Actions menu for the backup you are interested in, and then click **Delete**.
4. In the **Delete** dialog, click **Delete** to confirm the backup deletion.

To change the retention period of an LTR backup with Recovery Service

1. Open the navigation menu. Select **Oracle AI Database**, then select **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to change the backup retention period:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the

VM cluster you want to access and click its highlighted name to view the details page for the cluster.

4. In the list of databases, click the name of the database for which you want to change the retention period.
5. Under **Resources**, click **Backups**.
A list of backups is displayed.
6. In the list of backups, click the **Actions** menu for the backup with type **Long-term backup** for which you want to change the retention period.
7. Click **Change retention period**.
8. In the resulting **Change retention period**, change the retention period.

Note

The retention period must be entered in Days (90 - 3,650) or Years (1 - 10) from when the backup was created.

9. Click **Save**.

To designate Autonomous Recovery Service as a Backup Destination for an Existing Database

To designate Autonomous Recovery Service as a Backup Destination for an existing database, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the database:
Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model): Under **Exadata on Oracle Public Cloud**, click **Exadata VM Clusters**.

In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

On the cloud **VM cluster** details page, in the Databases table, click the name of the database to display the **Database Details** page.

4. Click **Configure automatic backups**.
5. In the resulting window, provide the following details:
 - **Enable automatic backup:** Check the check box to enable automatic incremental backups for this database. If you are creating a database in a security zone compartment, you must enable automatic backups.
 - **Backup Destination:** Select **Autonomous Recovery Service**.
 - **Backup Scheduling:** If you enable automatic backups, you can choose a two-hour scheduling window to control when backup operations begin. If you do not specify a window, then a six-hour default window of 00:00 to 06:00 (in the time zone of the VM cluster's region) is used for your database.
 - **Protection Policy:** If you choose to enable automatic backups, you can choose a policy with one of the following preset retention periods, or a Custom policy.

Object Storage Backup retention period: 7, 15, 30, 45, 60. Default: 30. The system automatically deletes your incremental backups at the end of your chosen retention period.

Autonomous Recovery Service protection policy:

- **Bronze:** 14 days
 - **Silver:** 35 days
 - **Gold:** 65 days
 - **Platinum:** 95 days
 - Custom defined by you
 - **Default:** Silver - 35 days
- **Enable Real-Time Data Protection:** Real-time protection is the continuous transfer of redo changes from a protected database to **Autonomous Recovery Service**. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.

6. Click **Save**.

Recovering an Exadata Database from Backup Destination

This topic explains how to recover an Exadata database from a backup stored in either Object Storage or Autonomous Recovery Service by using the Console or the API.

- Object Storage service is a secure, scalable, on-demand storage solution in Exadata Cloud Infrastructure.
- OracleDatabase Autonomous Recovery Service is a centralized, fully managed, and standalone backup solution for Oracle Cloud Infrastructure (OCI) databases.

For more information about backing up your databases to Object Storage, see *Managing Exadata Database Backups*.

- [Using the Console to restore a database](#)
You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

Related Topics

- [Managing Exadata Database Backups](#)
Automatic Exadata database backups are managed by Oracle Cloud Infrastructure. You configure this by using the Console or the API.

Using the Console to restore a database

You can use the Console to restore the database from a backup in a backup destination that was created by using the Console.

Note

LTR backups represent a single point in time for the database, so the following options are not supported when restoring.

You can restore to:

- **Restore to latest:** Restores the database to the last known good state with the least possible data loss.
- **Restore to a timestamp:** Restores the database to the timestamp specified.
- **Restore to SCN:** Restores the database using the SCN specified. This SCN must be valid.

Note

You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

Note

The list of backups you see in the Console does not include any unmanaged backups (backups created directly by using `dbaascli`).

- [To restore a database](#)

To restore a database

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster containing the database you want to restore: **Cloud VM clusters (The New Exadata Cloud Infrastructure Resource Model)**: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the list of databases, find the database you want to restore, and click its name to display details about it.
5. Click **Restore**.
6. Select one of the following options, and click **Restore Database**:
 - **Restore to the latest:** Restores the database to the last known good state with the least possible data loss.
 - **Restore to the timestamp:** Restores the database to the timestamp specified.
 - **Restore to System Change Number (SCN):** Restores the database using the SCN specified. This SCN must be valid.

Note

You can determine the SCN number to use either by accessing and querying your database host, or by accessing any online or archived logs.

7. Confirm when prompted.
If the restore operation fails, the database will be in a **"Restore Failed"** state. You can try restoring again using a different restore option. However, Oracle recommends that you review the `RMAN` logs on the host and fix any issues before reattempting to restore the

database. These log files can be found in subdirectories of the `/var/opt/oracle/log` directory.

Managing Exadata Database Backups by Using dbaascli

You can use Exadata's backup utility, `dbaascli`, to back up databases on an Exadata Cloud Infrastructure instance to an existing bucket in the Oracle Object Storage service.

For backups managed by Oracle Cloud Infrastructure, see [Managing Exadata Database Backups](#).

This topic explains how to:

- Create a default backup configuration file and modify the parameters to match your requirements to backup the database to object storage service.
- Associate the backup configuration file with a database. Once the configuration is successful, the database will be backed up as scheduled, or you can create an on-demand backup with a tag.

Note

You must update the cloud-specific tooling on all the compute nodes in your Exadata Cloud Infrastructure instance before performing the following procedures. For more information, see [Patching and Updating an Exadata Cloud Infrastructure System Manually](#).

- [Default Backup Configuration](#)
Oracle best practice guidelines for default backup configuration.
- [To get the default backup configuration for a freshly provisioned database](#)
- [To create a backup configuration file](#)
- [To create an on-demand backup](#)
- [To remove the backup configuration](#)
- [To delete a backup in Object Storage](#)

Default Backup Configuration

Oracle best practice guidelines for default backup configuration.

The default backup configuration follows a set of Oracle best-practice guidelines:

- **Encryption:** All backups to Object storage are encrypted.
- **Compression for backups:** LOW
- **Default compression for archive logs:** false
- **RMAN Encryption Algorithm:** AES256
- **Optimization for backups:** ON

To get the default backup configuration for a freshly provisioned database

1. SSH to one of the database configured nodes in the VM cluster resource.

2. Log in as `opc` and then `sudo` to the `root` user.
3. Use the `dbaascli database backup --getConfig` command to generate a file containing the default backup settings for the freshly provisioned database deployment.

```
# dbaascli database backup --getConfig [--configFile <file_name>] --dbname
<database_name>
```

Where:

- **--getConfig** - returns database backup configuration.
- **--configFile** - specifies database backup configuration file.

Related Topics

- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To create a backup configuration file

Note

The following procedure must be performed on the first compute node in the Exadata Cloud Infrastructure VM cluster resource. To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

Note

In `dbaascli` Release 25.1.2.0.0, the backup configuration parameters have been renamed. However, you can still use the old parameter names, as they are retained for backward compatibility.

1. SSH to one of the database configured nodes in the VM cluster resource.

```
ssh -i <private_key_path> opc@<node_1_ip_address>
```

2. Log in as `opc` and then `sudo` to the `root` user.

```
login as: opc [opc@dbsys ~]
$ sudo su -
```

- Use the `dbaascli database backup --getConfig` command to generate a file containing the current backup settings for the database deployment:

```
# dbaascli database backup --getConfig [--configFile <file_name>] --dbname
<database_name>
```

- Modify the parameters in the file to meet your requirements.

Parameter	Description
<code>backupDestination=oss</code>	Whether to back up to Object Storage. If yes, you must also provide the parameters <code>bkup_oss_url</code> , <code>bkup_oss_user</code> , <code>bkup_oss_passwd</code> , and <code>bkup_oss_recovery_window</code> .
Old name: <code>bkup_oss_url=<swift_url></code> New name: <code>ossURL=<swift_url></code>	Required if <code>backupDestination=oss</code> . The Object Storage URL including the tenant and bucket you want to use. The URL is: <code>https:// swiftobjectstorage.<region_name>.ora clecloud.com/v1/<tenant>/<bucket></code> Where: <ul style="list-style-type: none"> <code><tenant></code> - lowercase tenant name (even if it contains uppercase characters) that you specify when signing in to the Console <code><bucket></code> - name of the existing bucket you want to use for backups.
Old name: <code>bkup_oss_user=<oci_user_name></code> New name: <code>ossUserName=<oci_user_name></code>	Required if <code>backupDestination=oss</code> . The user name for the Oracle Cloud Infrastructure user account. This is the user name you use to sign in to the Oracle Cloud Infrastructure Console. For example, <code>jsmith@example.com</code> for a local user or <code><identity_provider>/jsmith@example.com</code> for a federated user. To determine which type of user you have, see the following topics: <ul style="list-style-type: none"> Managing Users (for information on local users) Federating with Identity Providers (for information on federated users) Note that the user must be a member of the Administrators group, as described in Prerequisites for Backups on Exadata Cloud Infrastructure .
Old name: <code>bkup_oss_passwd=<auth_token></code> New name: <code>ossAuthToken=<auth_token></code>	Required if <code>backupDestination=oss</code> . The auth token generated by using the Console or IAM API, as described in Prerequisites . This is not the password for the Oracle Cloud Infrastructure user.

Parameter	Description
Old name: <code>bkup_oss_recovery_window=n</code> New name: <code>ossRecoveryWindow=n</code>	Required if <code>backupDestination=oss</code> . The number of days for which backups and archived redo logs are maintained in the Object Storage bucket. Specify 7 to 90 days.
Old name: <code>bkup_daily_time=hh:mm</code> New name: <code>autoBackupTime=hh:mm</code>	The time at which the daily backup is scheduled, specified in hours and minutes (hh:mm), in 24-hour format.

- Use the `dbaascli database backup --configure` to associate this backup configuration with a database name.

```
# dbaascli database backup --configure --configFile <file_name> --dbname <database_name>
```

- Use the `dbaascli database backup --status` to check the status of UUID generated for this command.

```
# dbaascli database backup --status --uuid <uuid> --dbname <database_name>
```

Note

A backup configuration file can contain the credentials to access the Object Storage bucket. For this reason, you might want to remove the file after successfully configuring the backup.

The following parameters can be modified to customize the backup configuration:

Note

`Compatible with Console Automatic Backups=Yes` indicates the parameter is safe to change, even when using console-based automatic backups. If using parameters with `Compatible with Console Automatic Backups=No`, then do not enable backups through the console.

Table 5-6 Backup Configuration Parameters - Schedule Parameters to dbaascli

Parameter	Description	Compatible with Console Automatic Backups*
Old name: <code>bkup_cron_entry</code> New name: <code>scheduleBackups</code>	Enables the automatic backup configuration. Valid values are <code>yes</code> and <code>no</code> .	No
Old name: <code>bkup_archlog_cron_entry</code> New name: <code>manageArchivelogs</code>	Enables automatic backups of archived database log files. Valid values are <code>yes</code> and <code>no</code> . Setting <code>manageArchivelogs</code> to <code>no</code> disables automatic archive log clean-up jobs. This setting is valid only when the associated database has no automatic database backups configured.	No

Table 5-6 (Cont.) Backup Configuration Parameters - Schedule Parameters to dbaascli

Parameter	Description	Compatible with Console Automatic Backups*
Old name: bkup_10_day New name: L0BackupDay	This parameter controls the Level 0 day of the week. Day of the week when a level 0 backup is taken. Valid values are mon, tue, wed, thu, fri, sat, and sun. Longer formats, for example, Monday, Tuesday are also supported. Default: sun.	No

Table 5-7 Backup Configuration Parameters - General RMAN Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
Old name: bkup_rman_compression New name: compressionLevel	Level of compression applied to automatic backups. Valid values are NONE, basic, low, medium, and high. Default value is low. A value of NONE disables RMAN compression. If RMAN compression is enabled, then any TDE encrypted datafile will be decrypted, compressed, and RMAN encrypted.	Yes
Old name: bkup_section_size New name: sectionSize	RMAN section size that is used for automatic backups. Default value is 64G.	Yes
Old name: bkup_channels_node New name: channelsPerNode	Number of RMAN channels per node used for automatic backups. Valid values are between 1 and 32. Default value is 2.	Yes
Old name: bkup_daily_time New name: autoBackupTime	Start time of the automatic daily backup expressed in 24-hour time as hh:mm.	Yes
Old name: bkup_archlog_frequency New Name: backupFrequencyAL	Interval in minutes between automatic backups of archived database log files. Valid values are 15, 20, 30, 60, 120 through 1440 in one-hour intervals expressed in minutes. Default value is 30 for ExaDB-D.	Yes

Table 5-7 (Cont.) Backup Configuration Parameters - General RMAN Configuration Parameters (valid for all backup destinations except Local Storage (FRA))

Parameter	Description	Compatible with Console Automatic Backups*
Old name: bkup_type New name: backupDestination	The type of the location where the backup resides. Specify OSS as the backup destination, which is the default and only option.	Yes
Old name: bkup_filesperset_regular New name: filesPerSet	Specifies the maximum number of data files that can be included in a backup set for Regular/ Archival backups.	Yes
Old name: bkup_filesperset_al New name: filesPerSetAL	Specifies the maximum number of archive log files that can be included in a backup set for Archivelog Backups.	Yes
Old name: bkup_encryption New name: encryption	Encryption specifies whether backups should be encrypted or not. By default, encryption is enabled for OSS and Recovery Service, and this setting cannot be changed.	Yes
Old name: rmanBackupOptimization New name: optimization	Optimization is a feature in that reduces the amount of data that needs to be backed up, transferred, and restored. Recommended value is ON.	Yes
Old name: rmanFraCleanupChannels New name: numberOfChannelsForFraCleanup	Specifies the number of channels used for FRA Cleanup job.	Yes
Old name: Compress_Archive_Logs New name: compressionAL	Specifies whether to compress the archive log backups are not. Not applicable to Recovery Service.	Yes
Old name: bkup_archlog_fra_retention New name: archivelogRetentionDays	Specifies the number of days archive log to be retained in FRA.	Yes

Table 5-8 Backup Configuration Parameters - Object Storage Service (OSS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
backupDestination=oss	Enables backups to cloud storage. Valid values are yes and no.	No

Table 5-8 (Cont.) Backup Configuration Parameters - Object Storage Service (OSS) Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: bkup_oss_recovery_window New name: ossRecoveryWindow	Retention period for backups to cloud storage, expressed as a number of days up to 90. Applicable only when bkup_oss is set to yes or backupdestination is set to OSS. Default value is 30.	No
Old name: bkup_oss_url New name: ossURL	Location of the storage container that is used for backup to cloud storage. Applicable only when bkup_oss is set to yes or backupdestination is set to OSS.	No
Old name: bkup_oss_user New name: ossUserName	User name of the Oracle Cloud user having write privileges on the cloud storage container specified in bkup_oss_url. Applicable only when bkup_oss is set to yes or backupdestination is set to OSS.	No
Old name: bkup_oss_passwd New name: ossAuthToken	Password of the Oracle Cloud user having write privileges on the cloud storage container specified in bkup_oss_url. Applicable only when bkup_oss is set to yes or backupdestination is set to OSS.	No

Table 5-9 Backup Configuration Parameters - RMAN Catalog Support Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: bkup_use_rcat New name: useCatalog	Enables the use of an existing RMAN recovery catalog. Valid values are yes and no.	Yes
Old name: bkup_rcat_user New name: catalogUserName	Recovery catalog user name. Applicable only when bkup_use_rcat is set to yes.	Yes

Table 5-9 (Cont.) Backup Configuration Parameters - RMAN Catalog Support Parameters

Parameter	Description	Compatible with Console Automatic Backups*
Old name: <code>bkup_rcat_passwd</code> New name: <code>catalogPassword</code>	Password for recovery catalog user specified in <code>bkup_rcat_user</code> . Applicable only when <code>bkup_use_rcat</code> is set to <code>yes</code> .	Yes
Old name: <code>bkup_rcat_conn</code> New name: <code>catalogConnectionString</code>	Connection string for the RMAN recovery catalog. Applicable only when <code>bkup_use_rcat</code> is set to <code>yes</code> .	Yes

Note

Only the above parameters noted with `Compatible with Console Automatic Backups = Yes` are safe to alter in conjunction with console-based automatic backups. If any other parameters are to be altered, then do not enable backups through the console.

Related Topics

- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To create an on-demand backup

You can use the `dbaascli` to create an on-demand backup of a database.

1. SSH to one of the database configured nodes in the VM cluster resource.

```
ssh -i <private_key_path> opc@<node_1_ip_address>
```

To determine the first compute node, connect to any compute node as the `grid` user and execute the following command:

```
$ $ORACLE_HOME/bin/olsnodes -n
```

The first node has the number 1 listed beside the node name.

2. Log in as `opc` and then `sudo` to the `root` user.

```
login as: opc [opc@dbsys ~]
$ sudo su -
```

3. You can let the backup follow the current retention policy, or you can create a long-term backup that persists until you delete it:

- To create a backup that follows the current retention policy, enter the following command:

```
# dbaascli database backup --start --dbname <database_name>
```

- To create a long-term backup, enter the following command:

```
# dbaascli database backup --start --archival --dbname --tag
<archival_tag>
```

4. Exit the root-user command shell and disconnect from the compute node:

```
# exit
$ exit
```

5. Use the `dbaascli database backup --status` to check the status of UUID generated for the backup command

```
# dbaascli database backup --status --uuid <uuid> --dbname <database_name>
```

Related Topics

- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To remove the backup configuration

1. SSH to one of the database configured nodes in the VM cluster resource.
2. Log in as `opc` and then `sudo` to the `root` user.
3. Create a `temp` file with following parameters:
 - `bkup_oss=no`
 - `bkup_cron_entry=no`
 - `bkup_archlog_cron_entry=no`
4. Use the above file with `dbaascli database backup --configure` to remove the backup configuration for a database.

```
# dbaascli database backup --configure --configFile <file_name> --dbname
<database_name>
```

5. Use the `dbaascli database backup --status` to check the status of UUID generated for this command.

```
# dbaascli database backup --status --uuid <uuid> --dbname <database_name>
```

This will disable all automatic backups.

Related Topics

- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

To delete a backup in Object Storage

You can delete an archival or long-term backup from the Object Storage.

```
# dbaascli database backup --delete --backupTag --dbname <database_name>
```

Where:

- `--dbname` - specifies Oracle Database name
- `--delete` - deletes Archival backup.
- `--backupTag` - specifies backup tag to delete.

Policy based backups are deleted with scheduled daily backups. Alternatively, you can use RMAN delete backup command to delete a backup from the Object store.

Related Topics

- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

Using the API to Manage Backup and Recovery

- [Using the API to manage backups](#)

Using the API to manage backups

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

Use these API operations to manage database backups:

- [ListBackups](#)
- [GetBackup](#)
- [CreateBackup](#)
- [DeleteBackup](#)
- [UpdateDatabase](#) - To enable and disable automatic backups.
- [RestoreDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Alternative Backup Methods

Learn about alternative backup methods that are available in addition to the OCI Console.

Backup for databases on Exadata Cloud Infrastructure can be accomplished through several methods in addition to the automatic backups configured in the console. Generally, the console (or the OCI API / CLI that correspond to it) is the preferred method as it provides the simplest and most automated method. In general, it is preferable to leverage the OCI Console, OCI API, or OCI command-line over alternative management methods. However, if required actions

cannot be completed through the preferred methods, two other options are available to manually configure backups: `dbaascli` and Oracle Recovery Manager (RMAN).

Note

Use the [dbaascli database backup](#), [dbaascli pdb backup](#), [dbaascli database recover](#), and [dbaascli pdb recover](#) commands to backup and recover container databases and pluggable databases. For more information, see *User Configured Backup* in [Oracle Recommended Options to Perform Backup and Recovery Operations](#).

RMAN is the backup tool included with the Oracle Database. For information about using RMAN, see the [Oracle Database Backup and Recovery User's Guide for Release 19](#). Using RMAN to back up databases on Exadata Cloud Infrastructure provides the most flexibility in terms of backup options, but also the most complexity.

Note

While using RMAN for restoring databases backed up through any method described herein is considered safe, RMAN should NEVER be used to set up backups in conjunction with either console (and OCI API / CLI), nor in conjunction with `dbaascli`. If you choose to orchestrate backups manually leveraging RMAN, you should not use either console automated backups, nor should you use `dbaascli`. You must first completely disable console based automated backups. For more information, see *Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management*.

The `dbaascli` method offers a middle ground between RMAN and console automated backups in terms of flexibility and simplicity. Use `dbaascli` if needed functionality is not supported with console automated backups, but when you wish to avoid complexity of using RMAN directly. In certain cases, `dbaascli` can be used to modify the console automated backup configuration, but this is not generally the case. Generally, `dbaascli` must be used instead of enabling backups in the console.

- [Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management](#)

Disabling Automatic Backups to Facilitate Manual Backup and Recovery Management

Backups, configured in the Exadata Cloud Infrastructure console, API or `dbaascli` work for a variety of backup and recovery use cases. If you require use cases not supported by the cloud-managed backups, then you can manage database backup and recovery manually, using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the *Oracle Database Backup and Recovery User's Guide for Release 19*.

Managing backup and recovery, using RMAN, on Exadata Cloud Infrastructure requires taking full ownership of both database and archive log backups, and the cloud-managed backups should no longer be used. Before manual backups are started, the cloud-managed backup functionality should be disabled. This is needed so the cloud backup jobs do not purge archive logs before they are manually backed up and do not conflict with the manual backups.

You can use the `dbaascli` utility to disable cloud-managed backups, including disabling the automatic archive log purge job.

Recovering a Database Using Oracle Recovery Manager (RMAN)

If you backed up your database using `dbaascli`, then you can manually restore that database backup by using the Oracle Recovery Manager (RMAN) utility. For information about using RMAN, see the [Oracle Database Backup and Recovery User's Guide for Release 19](#).

Note

While recovering using RMAN is safe, you must not use RMAN to initiate backups or edit backup setting in conjunction with either `dbaascli` usage or in conjunction with automated console backups. Doing so could result in conflicting conditions or overwrites of settings, and backups may not execute successfully.

Patch and Update an Exadata Cloud Infrastructure System

- [User-Managed Maintenance Updates](#)
Maintaining a secure Exadata Cloud Infrastructure instance in the best working order requires you to perform the following tasks regularly:
- [Patching and Updating an Exadata Cloud Infrastructure System](#)
Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.
- [Patching and Updating an Exadata Cloud Infrastructure System Manually](#)
This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.
- [Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade](#)
If you've installed non-Exadata software packages beyond those provided by Oracle, and the precheck fails during a Guest VM upgrade due to conflicts between and Oracle-installed RPMs, you can use the following procedure to resolve the conflicts and proceed with the upgrade.

User-Managed Maintenance Updates

Maintaining a secure Exadata Cloud Infrastructure instance in the best working order requires you to perform the following tasks regularly:

- Patching the Oracle Grid Infrastructure and Oracle Database software on the VM Cluster virtual machines. For information and instructions, see *Patching and Updating VM Cluster's GI and Database Homes*.
- Updating the operating system on the VM Cluster virtual machines. See *Updating an Exadata Cloud VM Cluster Operating System* for information and instructions.

Related Topics

- [Patching and Updating VM Cluster's GI and Database Homes](#)
This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.
- [Oracle Clusterware Configuration and Administration](#)
- [Updating an Exadata Cloud VM Cluster Operating System](#)
Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- [Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli](#)
Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

Patching and Updating an Exadata Cloud Infrastructure System

Learn how to perform patching operations on Exadata database virtual machines and Database Homes by using the Console, API, or the CLI.

For information and instructions on patching the system by using the `dbaascli` utility, see *Patching and Updating an Exadata Cloud Infrastructure System Manually*.

For more information and examples for applying database quarterly patches on Exadata Cloud Infrastructure refer to My Oracle Support note: *How to Apply Database Quarterly Patch on Exadata Cloud Service and Exadata Cloud at Customer Gen 2 (Doc ID 2701789.1)*.

For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- [Patching and Updating VM Cluster's GI and Database Homes](#)
This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.
- [Updating an Exadata Cloud VM Cluster Operating System](#)
Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.
- [Upgrading Exadata Grid Infrastructure](#)
This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.
- [Upgrading Exadata Databases](#)
This topic describes the procedures to upgrade an Exadata database instance to Oracle Database 19c and Oracle AI Database 26ai by using the Console and the API. The upgrade is accomplished by moving the Exadata database to a Database Home that uses the target software version.
- [Database Health Checks During Infrastructure and Guest VM OS Update Operations](#)
The Database Health Check feature is designed to identify potential issues that could cause database downtime or service degradation during update operations involving container databases (CDBs) and pluggable databases (PDBs).

Related Topics

- [Patching and Updating an Exadata Cloud Infrastructure System Manually](#)
This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.
- <https://support.oracle.com/epmos/faces/DocContentDisplay?id=2701789.1>
- [Application Checklist for Continuous Service for MAA Solutions](#)

Patching and Updating VM Cluster's GI and Database Homes

This topic explains how to perform patching operations on Exadata Cloud Infrastructure resources by using the Console, API, or the CLI.

Note

Oracle recommends patching databases by moving them to a Database Home that uses the target patching level. See [To patch a database by moving it to another Database Home](#) for instructions on this method of database patching.

For information and instructions on patching the system by using the `dbaascli` utility, see [Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli](#).

- [About Patching and Updating VM Cluster's GI and Database Homes](#)
This topic describes the types of patching performed on an Exadata Cloud Infrastructure instances and provides instructions for completing the patching operations.
- [Prerequisites for Patching and Updating an Exadata Cloud Infrastructure System](#)
The Exadata Cloud Infrastructure instance requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage
- [Using the Console to Patch and Update Exadata Cloud Infrastructure Instances](#)
You can use the Console to view the history of patch operations on Exadata Cloud Infrastructure instances, apply patches, and monitor the status of patch operations.
- [Using the API to Patch an Exadata Cloud Infrastructure Instance](#)
Use these API operations to manage patching the following Exadata resources: cloud VM clusters, databases, and Database Homes.

About Patching and Updating VM Cluster's GI and Database Homes

This topic describes the types of patching performed on an Exadata Cloud Infrastructure instances and provides instructions for completing the patching operations.

- [Oracle Grid Infrastructure \(GI\) Patching](#)
Patching an Exadata Cloud Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster patch updates the Oracle Grid Infrastructure (GI) on the resource.
- [Database Home Patching](#)
A Database Home patch updates the Oracle Database software shared by the databases in that home.
- [Best Practices for Patching Exadata Cloud Infrastructure Components](#)

Oracle Grid Infrastructure (GI) Patching

Patching an Exadata Cloud Infrastructure instance updates the components on all the compute nodes in the instance. A VM cluster patch updates the Oracle Grid Infrastructure (GI) on the resource.

Database Home Patching

A Database Home patch updates the Oracle Database software shared by the databases in that home.

Thus, you patch a database by either of the following methods:

- Move the database to a Database Home that has the correct patch version. This affects only the database being moved.
- Patching the Database Home the database is currently in. This affects all databases located in the Database Home being patched.

When patching a Database Home, you can use an Oracle-provided database software image to apply a generally-available Oracle Database software update, or you can use a custom database software image created by your organization to apply a specific set of patches required by your database. See [Oracle Database Software Images](#) for more information on creating and using custom images.

For instructions on performing patching operations, see [To patch the Oracle Database software in a Database Home \(cloud VM cluster\)](#).

Best Practices for Patching Exadata Cloud Infrastructure Components

Consider the following best practices:

- Back up your databases before you apply any patches. For information about backing up the databases, see [Managing Exadata Database Backups](#).
- Patch a VM cluster before you patch the Databases Homes and databases on that resource.
- Before you apply any patch, run the precheck operation to ensure your VM cluster or Database Home meets the requirements for that patch.
- To patch a database to a version other than the database version of the current home, move the database to a Database Home running the target version. This technique requires less downtime and allows you to easily roll back the database to the previous version by moving it back to the old Database Home. See [To move a database to another Database Home To patch a database by moving it to another Database Home](#).
- For the Oracle Database and Oracle Grid Infrastructure major version releases available in Oracle Cloud Infrastructure, patches are provided for the current version plus the three most recent older versions (N through $N - 3$). For example, if an instance is using Oracle Database 19c, and the latest version of 19c offered is 19.8.0.0.0, patches are available for versions 19.8.0.0.0, 19.7.0.0, 19.6.0.0 and 19.5.0.0.
- [dbaascli database runDatapatch](#)
To patch an Oracle Database, use the `dbaascli database runDatapatch` command.
- [Customer-Managed Keys in Exadata Cloud Infrastructure](#)
Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- [dbaascli database addInstance](#)
To add the database instance on the specified node, use the `dbaascli database addInstance` command.

- [dbaascli database convertToPDB](#)
To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.
- [dbaascli database getDetails](#)
This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.
- [dbaascli database modifyParameters](#)
To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.
- [dbaascli database upgrade](#)
To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

dbaascli database runDatapatch

To patch an Oracle Database, use the `dbaascli database runDatapatch` command.

Prerequisites

- Before performing a `runDatapatch` operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database runDatapatch --dbname
[--resume]
    [--sessionID]
[--skipPdfs | --pdfs]
[--executePrereqs]
[--patchList]
[--skipClosedPdfs]
[--rollback]
```

Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--skipPdfs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: `pdb1,pdb2...`
- `--pdfs` runs the datapatch only on a specified comma-delimited list of PDBs. For example: `pdb1,pdb2...`
- `--executePrereqs` runs prerequisite checks
- `--patchList` applies or rolls back the specified comma-delimited list of patches. For example: `patch1,patch2...`
- `--skipClosedPdfs` skips running the datapatch on closed PDBs
- `--rollback` rolls back the patches applied

```
dbaascli database runDatapatch --dbname db19
```

Customer-Managed Keys in Exadata Cloud Infrastructure

Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.

The OCI Vault service provides you with centralized key management capabilities that are highly available and durable. This key-management solution also offers secure key storage using isolated partitions (and a lower-cost shared partition option) in FIPS 140-2 Level 3-certified hardware security modules, and integration with select Oracle Cloud Infrastructure services. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Enable customer-managed keys when you create databases in Exadata Cloud Infrastructure
- Switch from Oracle-managed keys to customer-managed keys
- Rotate your keys to maintain security compliance

Requirements

To enable management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so, similar to the following: `allow dynamic-group dynamic_group_name to manage keys in tenancy`.

Another policy is needed if the Vault being used by the customer is replicated (<https://docs.oracle.com/en-us/iaas/Content/KeyManagement/Tasks/replicatingvaults.htm>). For vaults that are replicated, this policy is needed: `allow dynamic-group dynamic_group_name to read vaults in tenancy`

Limitations

To enable Data Guard on Exadata Cloud Infrastructure databases that use customer-managed keys, the primary and standby databases must be in the same [realm](#).

Task 1. Create a Vault and a Master Encryption Key

Create a vault in the Vault service by following the instructions in [To create a new vault](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, Oracle recommends that you create the vault in a compartment created specifically to contain the vaults containing customer-managed keys, as described in [Before You Begin: Compartment Hierarchy Best Practice](#).

After creating the vault, create at least one master encryption key in the vault by following the instructions in [To create a new master encryption key](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, make these choices:

- **Create in Compartment:** Oracle recommends that you create the master encryption key in the same compartment as its vault; that is, the compartment created specifically to contain the vaults containing customer-managed keys.
- **Protection Mode:** Choose an appropriate value from the drop-down list:
 - **HSM** to create a master encryption key that is stored and processed on a hardware security module (HSM).
 - **Software** to create a master encryption key that is stored in a software file system in the Vault service. Software-protected keys are protected at rest using an HSM-based

root key. You may export software keys to other key management devices or to a different OCI cloud region. Unlike HSM keys, software-protected keys are free of cost.

- **Key Shape Algorithm:** AES
- **Key Shape Length:** 256 bits

Oracle strongly recommends that you create a separate master encryption key for each of your container databases (CDBs). Doing so makes management of key rotation over time much simpler.

Task 2. Create a Service Gateway, a Route Rule, and an Egress Security Rule

Create a service gateway in the VCN (Virtual Cloud Network) where your Oracle Exadata Database Service on Dedicated Infrastructure resources reside by following the instructions in [Task 1: Create the service gateway](#) in Oracle Cloud Infrastructure Documentation.

After creating the service gateway, add a route rule and an egress security rule **to each subnet** (in the VCN) where Oracle Exadata Database Service on Dedicated Infrastructure resources reside so that these resources can use the gateway to access the Vault service:

1. Go to the **Subnet Details** page for the subnet.
2. In the **Subnet Information** tab, click the name of the subnet's **Route Table** to display its **Route Table Details** page.
3. In the table of existing **Route Rules**, check whether there is already a rule with the following characteristics:

- **Destination:** All IAD Services In Oracle Services Network
- **Target Type:** Service Gateway
- **Target:** The name of the service gateway you just created in the VCN

If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.

4. Return to the **Subnet Details** page for the subnet.
5. In the subnet's **Security Lists** table, click the name of the subnet's security list to display its **Security List Details** page.
6. In the side menu, under **Resources**, click **Egress Rules**.
7. In the table of existing **Egress Rules**, check whether there is already a rule with the following characteristics:

- **Stateless:** No
- **Destination:** All IAD Services In Oracle Services Network
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443

If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Task 3. Create a Dynamic Group and a Policy Statement

To grant your Oracle Exadata Database Service on Dedicated Infrastructure resources permission to access customer-managed keys, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the master encryption keys you created in the Vault service.

When defining the dynamic group, you identify your Oracle Exadata Database Service on Dedicated Infrastructure resources by specifying the OCID of the compartment containing your Exadata Infrastructure resource.

1. Copy the OCID of the compartment containing your Exadata Infrastructure resource. You can find this OCID on the **Compartment Details** page of the compartment.
2. Create a dynamic group by following the instructions in [To create a dynamic group](#) in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

```
ALL {resource.compartment.id = '<compartment-ocid>'}
```

where *<compartment-ocid>* is the OCID of the compartment containing your Exadata Infrastructure resource.

After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and keys. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group-name>
to manage keys
in compartment <vaults-and-keys-compartment>
where all {
target.key.id='<key_ocid>',
request.permission!='KEY_DELETE',
request.permission!='KEY_MOVE',
request.permission!='KEY_IMPORT',
request.permission!='KEY_BACKUP'
}
```

If you are using a replicated virtual private vault for the Oracle Data Guard deployment, add an additional policy statement in this format:

```
allow dynamic-group <dynamic-group>
to read vaults
in tenancy | compartment <vaults-and-keys-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-keys-compartment>* is the name of the compartment in which you created your vaults and master encryption keys.

Related Topics

- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.
- [Known Issues for Exadata Cloud Infrastructure and Data Guard](#)
Possible TDE key replication issue, and MRP and DG LCM operation failures.
- [To integrate customer-managed key management into Exadata Cloud Infrastructure](#)
If you choose to encrypt databases in an Exadata Cloud Infrastructure instance using encryption keys that you manage, then you may update the following two packages (using

Red Hat Package Manager) to enable DBAASSTOOLS to interact with the APIs that customer-managed key management uses.

dbaascli database addInstance

To add the database instance on the specified node, use the `dbaascli database addInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--newNodeSID <value>]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for the database instance
- `--newNodeSID` specifies SID for the instance to add in the new node

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--executePrereqs]
    {
        [--copyDatafiles [--keepSourceDB]]|[backupPrepared]
    }
    [--targetPDBName <value>] [--waitForCompletion <value>] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of Oracle Database
- `--cdbName` specifies the name of the target CDB in which the PDB will be created. If the CDB does not exist, then it will be created in the same Oracle home as the source non-CDB
- `--executePrereqs` specifies to run only the pre-conversion checks
- `--copyDatafiles` specifies to create a new copy of the data files instead of using the ones from the source database
- `--keepSourceDB` - to preserve the source database after completing the operation.
- `--backupPrepared` - flag to acknowledge that a proper database backup is in place for the non CDB prior to performing the conversion to PDB.
- `--backupPrepared` flag to acknowledge that a proper database backup is in place for the non-CDB prior to performing the conversion to PDB
- `--targetPDBName` specifies the name of the PDB that will be created as part of the operation

- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID

Example 5-3 dbaascli database convertToPDB

To run pre-conversion prechecks:

```
dbaascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

To run a full conversion with a copy of the data files from the non-CDB:

```
dbaascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

dbaascli database getDetails

This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

Prerequisites

Run the command as the `root` user or the `oracle` user

Syntax

```
dbaascli database getDetails --dbname <value>
```

Where :

- `--dbname` - Oracle database name.

dbaascli database modifyParameters

To modify or reset initialization parameters for an Oracle Database, use the `dbaascli database modifyParameters` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database modifyParameters --dbname <value> --setParameters <values> |
--resetParameters <values> | --responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
```

Where:

- `--dbname` specifies the name of the database.
- `--setParameters` specifies a comma-delimited list of parameters to modify with new values. For example: `parameter1=valueA,parameter2=valueB`, and so on. For blank values use `parameter1=valueA,parameter2=""`, etc.

- `--resetParameters` specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, `parameter1,parameter2`, and so on.
- `--responseFile` specifies the absolute location of the response JSON file to modify the database parameters
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--instance` specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.

Example 5-4 dbaascli database modifyParameters

```
dbaascli database modifyParameters --dbname dbname --setParameters
"log_archive_dest_state_17=ENABLE"
```

dbaascli database upgrade

To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database upgrade --dbname <value>
{--targetHome <value> | --targetHomeName <value>}
{ [--executePrereqs | --postUpgrade | --rollback]}
{[--standBy | --allStandbyPrepared]}
{[--upgradeOptions <value>] | [--standBy]}
[--removeGRP]
[--increaseCompatibleParameter]
[--resume [--sessionID <value>]]
[--waitForCompletion <value>]
```

Where:

- `--dbname` (mandatory) specifies the name of the database.
- `--targetHome` specifies the target Oracle home location
- `--targetHomeName` specifies the name of the target Oracle Database home
- `--standBy` use this option to upgrade standby databases in Data Guard configurations
- `--allStandbyPrepared` required for Data Guard configured primary databases. Flags to acknowledge that all the required operations are performed on the standby databases prior to upgrading primary database
- `--removeGRP` automatically removes the Guaranteed Restore Point (GRP) backup only if the database upgrade was successful
- `--increaseCompatibleParameter` automatically increases the compatible parameter as part of the database upgrade. The parameter will get increased only if the database upgrade was successful
- `--executePrereqs` runs only the preupgrade checks

- `--postUpgrade` use this option if postupgrade fails and needs to rerun the postupgrade steps
- `--rollback` reverts an Oracle Database to its original Oracle home
- `--upgradeOptions` use this option to pass DBUA-specific arguments to perform the Oracle Database upgrade. Refer to the corresponding Oracle documentation for the supported arguments and options.
- `--standby`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.
- `--waitForCompletion` specify false to run the operation in background. Valid values : true|false.

Example 5-5 dbaascli database upgrade pre-upgrade requisite checks

```
dbaascli database upgrade --dbname dbname --targetHome Target Oracle home
location --executePrereqs
```

Prerequisites for Patching and Updating an Exadata Cloud Infrastructure System

The Exadata Cloud Infrastructure instance requires access to the Oracle Cloud Infrastructure Object Storage service, including connectivity to the applicable Swift endpoint for Object Storage

Oracle recommends using a service gateway with the VCN to enable this access. For more information, see these topics:

- [Network Setup for Exadata Cloud Infrastructure Instances](#): For information about setting up your VCN for the Exadata Cloud Service instance, including the service gateway.
- [Object Storage FAQ](#)

Note

Ensure that the following conditions are met to avoid patching failures:

- The `/u01` directory on the database host file system has at least 15 GB of free space for the execution of patching processes.
- The Oracle Clusterware is up and running on the VM cluster.
- All nodes of the VM cluster are up and running.

Using the Console to Patch and Update Exadata Cloud Infrastructure Instances

You can use the Console to view the history of patch operations on Exadata Cloud Infrastructure instances, apply patches, and monitor the status of patch operations.

- [Patching Exadata Instances That use the New Resource Model](#)
The tasks in this section describe how to apply patches and monitor the status of patch operations on cloud VM clusters and their Database Homes.
- [Patching Individual Oracle Databases in an Exadata Cloud Infrastructure Instance](#)
This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

- [Viewing Patch History](#)
Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry.

Patching Exadata Instances That use the New Resource Model

The tasks in this section describe how to apply patches and monitor the status of patch operations on cloud VM clusters and their Database Homes.

- [To patch the Oracle Grid Infrastructure on an Exadata cloud VM cluster](#)
How to apply patches and monitor the status of patch operations on cloud VM clusters.
- [To patch the Oracle Database software in a Database Home](#)

To patch the Oracle Grid Infrastructure on an Exadata cloud VM cluster

How to apply patches and monitor the status of patch operations on cloud VM clusters.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Click **Updates (GI)**.
6. Review the list of available patches for the cloud VM cluster.
7. Click the Actions menu for the patch you are interested in, and then click one of the following actions:
 - **Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
 - **Apply Grid Infrastructure update:** Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
8. Confirm when prompted.

The patch list displays the status of the operation. While a patch is being applied, the patch's status displays as **Patching** and the cloud VM cluster's status displays as **Updating**. Lifecycle operations on the cluster and its resources might be temporarily unavailable. If patching completes successfully, the patch's status changes to **Applied** and the status of the cluster changes to **Available**. You can view more details about an individual patch operation by clicking **Update History**.

To patch the Oracle Database software in a Database Home

Note

This patching procedure updates the Oracle Database software for all databases located in the Database Home. To patch an individual database, you can [To move a database to another Database Home](#) that uses the desired Oracle Database software configuration.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**

2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Click **Database homes**.
6. Click the name of the Database Home you want to patch to display the Database Home details.
7. Click **Updates**.
8. Review the available patches for the Database Home. You can choose to patch using an Oracle-provided software image or a custom software image. Oracle-provide images are generally available release updates. Custom software images are created by your organization with a specified set of patches. See [Oracle Database Software Images](#) for information on creating custom software images. The image you use to patch must be based on either the latest version of the Oracle Database software release or one of the three prior versions of the release.
9. Click the Actions menu at the end of the table row that lists the patch you are interested in, and then click one of the following actions:
 - **Precheck:** Check for any prerequisites to make sure that the patch can be successfully applied.
 - **Apply Database Home update:** Applies the selected patch. Oracle highly recommends that you run the precheck operation for a patch before you apply it.
10. Confirm when prompted.

The patch list displays the status of the operation. While a patch is being applied, the status of the patch displays as **Patching** and the status of the Database Home and the databases in it display as **Updating**. During the operation, each database in the home is stopped and then restarted. If patching completes successfully, the patch's status changes to **Applied** and the Database Home's status changes to **Available**. You can view more details about an individual patch operation by clicking **Update History**.

Patching Individual Oracle Databases in an Exadata Cloud Infrastructure Instance

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

For information on patching Database Homes, see [To patch the Oracle Database software in a Database Home \(cloud VM cluster\)](#)

- [To move a database to another Database Home](#)
This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

To move a database to another Database Home

This task explains how to patch a single Oracle Database in your Exadata Cloud Infrastructure instance by moving it to another Database Home.

You can move a database to any Database Home that meets at either of the following criteria:

- The target Database Home uses the same Oracle Database software version (including patch updates) as the source Database Home
- The target Database Home is based on either the latest version of the Oracle Database software release used by the database, or one of the three prior versions of the release

Moving a database to a new Database Home brings the database up to the patch level of the target Database Home. For information on patching Database Homes, see [Database Home Patching](#).

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the database you want to move.
Cloud VM clusters ([The New Exadata Cloud Infrastructure Resource Model](#)): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you want to move.
4. Click **More Actions**, then click **Move to Another Home**.
5. Select the target Database Home.
6. Click **Move**.
7. Confirm the move operation.

The database is moved in a rolling fashion. The database instance will be stopped, node by node, in the current home and then restarted in the destination home. While the database is being moved, the Database Home status displays as **Moving Database**. When the operation completes, Database Home is updated with the current home. Datapatch is executed automatically, as part of the database move, to complete post-patch SQL actions for all patches, including one-offs, on the new Database Home. If the database move operation is unsuccessful, then the status of the database displays as `Failed`, and the Database Home field provides information about the reason for the failure.

Viewing Patch History

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry.

Patch history views in the Console do not show patches that were applied by using command line tools such as `dbaascli`.

If your service instance uses the [new resource model](#), the patch history available by navigating to the VM Cluster Details page.

- [To view the patch history of a cloud VM cluster](#)
Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed.
- [To view the patch history of a Database Home](#)
Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. If your service instance uses the new resource model, the patch history available by navigating to the VM Cluster Details page.

To view the patch history of a cloud VM cluster

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.

3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Click **Update History**.
The Update History page displays the history of patch operations for that cloud VM cluster and for the Database Homes on that cloud VM cluster.

To view the patch history of a Database Home

Each patch history entry represents an attempted patch operation and indicates whether the operation was successful or failed. You can retry a failed patch operation. Repeating an operation results in a new patch history entry. If your service instance uses the new resource model, the patch history available by navigating to the VM Cluster Details page.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Navigate to the cloud VM cluster that contains the Database Home.
 - *Cloud VM clusters (The Exadata Cloud Infrastructure Resource Model)* Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Click **Database homes**.
5. Click the name of the Database Home you want to view to display the Database Home details.
6. Click **Update History**.
The history page displays the history of patch operations for that Database Home and for the cloud VM cluster to which it belongs.

Using the API to Patch an Exadata Cloud Infrastructure Instance

Use these API operations to manage patching the following Exadata resources: cloud VM clusters, databases, and Database Homes.

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

Cloud VM clusters (*for systems using [the new resource model](#)*):

- [ListCloudVmClusterUpdates](#)
- [ListCloudVmClusterUpdateHistoryEntries](#)
- [GetCloudVmClusterUpdate](#)
- [GetCloudVmClusterUpdateHistoryEntry](#)
- [UpdateVmCluster](#)

Databases:

- [UpdateDatabase](#) - Use this operation to patch a database by moving it to another Database Home

Database Homes:

- [ListDbHomePatches](#)

- [ListDbHomePatchHistoryEntries](#)
- [GetDbHomePatch](#)
- [GetDbHomePatchHistoryEntry](#)
- [UpdateDbHome](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Updating an Exadata Cloud VM Cluster Operating System

Exadata VM cluster image updates allow you to update the OS image on your Exadata cloud VM cluster nodes in an automated manner from the OCI console and APIs.

This automated feature simplifies and speeds up VM cluster patching, makes patching less error-prone, and eliminates the need to use Patch Manager.

When you apply a patch, the system runs a precheck operation to ensure your cloud VM cluster or Database Home meets the requirements for that patch. If the precheck is not successful, the patch is not applied, and the system displays a message that the patch cannot be applied because the precheck failed. A separate precheck operation that you can run in advance of the planned update is also available.

- [Supported Software Versions and Update Restrictions](#)
Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):
- [Updating the Operating System using the Console](#)

Supported Software Versions and Update Restrictions

Minimum requirements for updating to Exadata image release 23.1.0.0.0 (Oracle Linux 8-based image):

Note

These are just the minimum requirements. If you want to update Grid Infrastructure and/or Oracle Database to meet the Exadata 23.1 requirements, then the recommendation is to update to the latest available versions of Grid Infrastructure and Oracle Database, and not to the minimum.

- **Exadata Image (Guest OS):** Exadata image release 22.1.0 (May 2022) or 21.2.10 (March 2022). Systems running versions older than 21.2.10 will first need to upgrade to at least 22.1.0 (May 2022) or 21.2.10 (March 2022) before updating to 23.1.0.0.0. This applies to both storage and database servers.
 - In addition to performing minor version updates to the Exadata VM Cluster images, you can update to a new major version if the currently installed version is 19.2 or higher. For example, if the VM cluster is on version 20, then you can update it to version 21.
 - The latest 4 (N to N-3) or more minor versions of each major version of the VM Cluster images are available through the console to apply.
- **Oracle Grid Infrastructure:** Exadata image release 23.1.0.0.0 supports the following minimum or newer Oracle Grid Infrastructure versions.
 - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)

- Release 21c: Version 21.6, April 2022 Release Update (RU) and newer
- **Oracle Database:** Exadata System Software 23.1 supports the following minimum versions or newer for new database installations.
 - Release 19c: Version 19.15, April 2022 Release Update (RU) and newer (Default)
 - Additional supported database releases under Market Driven Support or Quarterly Updates exception approval:
 - * Release 12.2.0.1, Release Update (RU) 12.2.0.1.220118 (Jan 2022)
 - * Release 12.1.0.2, Bundle Patch 12.1.0.2.220719 (Jul 2022) - requires patch 30159782
 - * Release 11.2.0.4, Bundle Patch 11.2.0.4.210119 (Jan 2021) - requires patch 30159782, patch 33991024
- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the Exadata Image update feature is not available.
- Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.

Note

Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2024 update cycle.

Updating the Operating System using the Console

Note

Once the VM cluster is upgraded to Exadata Database Service Guest VM OS 23.1, you will be able to add a new VM or a new database server to this VM cluster if Exadata Cloud Infrastructure is running an Exadata System Software version 22.1.16 and later.
Upgrade to Exadata System Software 23.1 for Exadata Cloud Infrastructure will be available with February 2024 update cycle.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. In the list of cloud VM clusters, click the name of the cluster that you want to patch to display the details page.
4. Click **Updates (OS)**.
5. Review the list of available software updates and locate the OS patch you are applying.
6. Click the Actions icon (three dots) at the end of the row listing the patch you are interested in, and then click one of the following actions:

- **Precheck:** Precheck checks the prerequisites to ensure that the patch can be successfully applied. Oracle highly recommends that you run the precheck operation before you apply a patch. The reason is that things can change in a database at any time, and the precheck you run just before running a patch may find errors that the previous precheck did not find.

Note

If the precheck fails, the system displays a message in the **Apply Exadata OS Image Update** dialog that the last precheck has failed. Oracle recommends that you run the precheck again. Click the Actions icon (three dots) at the end of the row listing the OS patch to view the dialog.

- **Apply Exadata OS Image Update:** This link displays the Apply Exadata Image Update dialog that you use to apply the patch. The dialog shows the name of the database system you are patching, the current version of the database, and the new version of the database after the patch is applied. To start the process, click **Apply Exadata OS Image Update**.
- **Copy OCID.** This copies the Oracle Cloud ID. This can be used when troubleshooting a patch or to give to Support when contacting them.

Note

While the patch is running:

- Run Precheck and Apply OS Image Update are not available. When the patch has completed, these actions are available again.
- If the Exadata infrastructure containing this VM cluster is scheduled for maintenance that conflicts with the patching operation, the patch fails and the system displays a message explaining why. After the infrastructure maintenance is complete, run the patch operation again.

7. Confirm when prompted.

The patch list displays the status of the operation in the Version section of the database details page. Click **View Updates** to view more details about an individual patch status and to display any updates that are available to run. If no new updates are available, the system displays a message that says **No Updates Available**.

Upgrading Exadata Grid Infrastructure

This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.

Upgrading allows you to provision Oracle Database Homes and databases that use the most current Oracle Database software. For more information on Exadata cloud VM clusters and the new Exadata resource model, see [Overview of X8M, X9M, and X11M Scalable Exadata Infrastructure](#).

- [Prerequisites for Upgrading Exadata Grid Infrastructure](#)
To upgrade your GI to Oracle Database 19c, you must be using the Oracle Linux 7 operating system for your VM cluster.

- [About Upgrading Oracle Grid Infrastructure](#)
Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.
- [Using the console to upgrade your Grid Infrastructure](#)
You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI), and to perform the GI upgrade operation.
- [Using the API to Upgrade the Grid Infrastructure in a VM Cluster](#)

Prerequisites for Upgrading Exadata Grid Infrastructure

To upgrade your GI to Oracle Database 19c, you must be using the Oracle Linux 7 operating system for your VM cluster.

For more information on upgrading the operating system, see the following document:

- [How to update the Exadata System Software \(DomU\) to 19 from 18 on the Exadata Cloud Service in OCI](#)(My Oracle Support Doc ID 2521053.1).

About Upgrading Oracle Grid Infrastructure

Upgrading the Oracle Grid Infrastructure (GI) on a VM cluster involves upgrading all the compute nodes in the instance. The upgrade is performed in a rolling fashion, with only one node being upgraded at a time.

- Oracle recommends running an upgrade precheck to identify and resolve any issues that would prevent a successful upgrade.
- You can monitor the progress of the upgrade operation by viewing the associated *work requests*.
- If you have an Exadata infrastructure maintenance operation scheduled to start within the next 24 hours, then the GI upgrade feature is not available.
- During the upgrade, you cannot perform other management operations such as starting, stopping, or rebooting nodes, scaling CPU, provisioning or managing Database Homes or databases, restoring a database, or editing IORM settings. The following Data Guard operations are not allowed on the VM cluster undergoing a GI upgrade:
 - Enable Data Guard
 - Switchover
 - Failover to the database using the VM cluster (a failover operation to standby on another VM cluster is possible)

Related Topics

- [Work Requests Integration](#)

Using the console to upgrade your Grid Infrastructure

You can use the Console to perform a precheck prior to upgrading your Oracle Grid Infrastructure (GI), and to perform the GI upgrade operation.

- [To precheck your cloud VM cluster prior to upgrading](#)
- [To upgrade the Oracle Grid Infrastructure of a cloud VM cluster](#)
Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

To precheck your cloud VM cluster prior to upgrading

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Click **Updates (GI)** to view the list of available patches and upgrades.
6. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Precheck**.
7. In the **Confirm** dialog, confirm you want to upgrade to begin the precheck operation.

To upgrade the Oracle Grid Infrastructure of a cloud VM cluster

Procedure for upgrading the Oracle Grid Infrastructure of a Cloud VM Cluster.

Note

- When planning to upgrade your Grid Infrastructure to 26ai, make sure that for each ASM diskgroup, `compatible.rdbms` has a value set to 19.0.0.0 and later.
- Minimum requirements for upgrading Grid Infrastructure from 19c to 26ai:
 - Exadata Guest VM running Exadata System Software 23.1.8
 - Exadata Infrastructure running Exadata System Software 23.1.x

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters**.
4. In the list of cloud VM clusters, click the name of the cluster you want to patch to display the cluster details.
5. Click **Updates (GI)** to view the list of available patches and upgrades.
6. Click the Actions icon (three dots) at the end of the row listing the Oracle Grid Infrastructure (GI) upgrade, then click **Apply Grid Infrastructure update**.
7. In the **Upgrade Grid Infrastructure** dialog, confirm you want to upgrade the GI by clicking **Upgrade Grid Infrastructure**. If you haven't run a precheck, you have the option of clicking **Run Precheck** in this dialog to precheck your cloud VM cluster prior to the upgrade.

Using the API to Upgrade the Grid Infrastructure in a VM Cluster

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

Use these API operations to upgrade the Oracle Grid Infrastructure in a cloud VM clusters and view the cluster's update history.

- [ListCloudVmClusterUpdates](#)
- [ListCloudVmClusterUpdateHistoryEntries](#)
- [GetCloudVmClusterUpdate](#)
- [GetCloudVmClusterUpdateHistoryEntry](#)
- [UpdateVmCluster](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Upgrading Exadata Databases

This topic describes the procedures to upgrade an Exadata database instance to Oracle Database 19c and Oracle AI Database 26ai by using the Console and the API. The upgrade is accomplished by moving the Exadata database to a Database Home that uses the target software version.

Note

This topic applies only to Exadata Cloud Infrastructure instances using the [new resource model](#).

For Oracle Database release and software support timelines, see [Release Schedule of Current Database Releases \(Doc ID 742060.1\)](#) in the My Oracle Support portal.

- [Prerequisites to Upgrade Oracle Databases](#)
Review the list of prerequisites to upgrade an Exadata Cloud Infrastructure Oracle Database instance.
- [About Upgrading a Database](#)
- [Using the Console to Upgrade a Database](#)
Procedures to precheck and upgrade a database, rollback a failed upgrade, and view the upgrade history.
- [Using the API to upgrade Databases](#)
Use the following APIs to manage database upgrades:

Prerequisites to Upgrade Oracle Databases

Review the list of prerequisites to upgrade an Exadata Cloud Infrastructure Oracle Database instance.

- To upgrade to 19c, Oracle Linux 7 is the minimum requirement, and to upgrade to 26ai, Oracle Linux 8 is the minimum requirement. For detailed instructions on manually updating the operating system, refer to *How to Update the Exadata System Software (DomU) to 19 from 18 on the Exadata Cloud Service in OCI (My Oracle Support Doc ID 2521053.1)*.
- The Oracle Grid Infrastructure can be version 19c or 26ai for Oracle Database 19c. However, the Oracle Grid Infrastructure must be version 26ai for Oracle AI Database 26ai. See [Upgrading Exadata Grid Infrastructure](#) for instructions on using the Oracle Cloud Infrastructure Console or API to upgrade Grid Infrastructure. If patches are available for your Grid Infrastructure, Oracle recommends applying them prior to performing a database upgrade.

- You must have an available Oracle Database Home that uses the four most recent versions of Oracle Database 19c or Oracle AI Database 26ai available in Oracle Cloud Infrastructure. See *To Create a new Oracle Database Home in an existing Exadata Cloud Infrastructure Instance* for information on creating a Database Home. You can use Oracle-published software images or a *custom database software image* based on your patching requirements to create Database Homes.
- You must ensure that all pluggable databases in the container database that is being upgraded can be opened. Pluggable databases that cannot be opened by the system during the upgrade can cause an upgrade failure.
- If you are upgrading databases in a manually-created Data Guard association (an association not created using the Console or APIs), the following apply:
 - The databases must be registered with the Cloud tooling. See [Updating Tooling on an Exadata Cloud Service Instance](#) for more information.
 - Redo apply needs to be disabled during the upgrade of both the primary and standby. For Oracle 11.2 and 12.1 databases, the Data Guard configuration also has to be disabled.
 - If you have configured an observer, the observer needs to be disabled prior to upgrade.

Your Oracle database must be configured with the following settings in order to upgrade:

- The database must be in archive log mode.
- The database must have flashback enabled.

See the *Oracle Database documentation* for your database's release version to learn more about these settings.

Related Topics

- [How to update the Exadata System Software \(DomU\) to 19 from 18 on the Exadata Cloud Service in OCI \(Doc ID 2521053.1\)](#)
- [Upgrading Exadata Grid Infrastructure](#)
This topic describes how to upgrade the Oracle Grid Infrastructure (GI) on an Exadata cloud VM cluster using the Oracle Cloud Infrastructure Console or API.
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)
To create an Oracle Database home in an existing VM cluster with the Console, be prepared to provide values for the fields required.
- [Oracle Database Software Images](#)
- [Oracle Database Documentation](#)

About Upgrading a Database

For database software version upgrades, note the following:

- Database upgrades involve database downtime. Keep this in mind when scheduling your upgrade.
- Oracle recommends that you back up your database and test the new software version on a test system or a cloned version of your database before you upgrade a production database. See *to create an on-demand full backup of a database* for information on creating an on-demand manual backup.
- Oracle recommends running an upgrade precheck operation for your database prior to attempting an upgrade so that you can discover any issues that need mitigation prior to the

time you plan to perform the upgrade. The precheck operation does not affect database availability and can be performed at any time that is convenient for you.

- If your databases uses Data Guard, you can upgrade either the primary or the standby first. To upgrade a primary, follow the steps in [To upgrade or precheck an Exadata database](#). To upgrade a standby, follow the steps in [To move a database to another Database Home](#)
- If your databases uses Data Guard, upgrading a primary or standby will disable redo apply during the upgrade operation. After you upgrade both the primary and standby, redo apply and open mode are re-enabled. Oracle recommends checking the redo apply and open mode configuration after upgrading.
- An upgrade operation cannot take place while an automatic backup operation is underway. Before upgrading, Oracle recommends disabling automatic backups and performing a manual backup. See *to configure automatic backups for a database* and *To create an on-demand full backup of a database* for more information.
- After upgrading, you cannot use automatic backups taken prior to the upgrade to restore the database to an earlier point in time.
- If you are upgrading an database that uses version 11.2 software, the resulting version 19c database will be a non-container database (non-CDB).
- [How the Upgrade Operation Is Performed by the Database Service](#)
During the upgrade process, the Database service does the following:
- [Rolling Back an Oracle Database Unsuccessful Upgrade](#)
If your upgrade does not complete successfully, then you have the option of performing a rollback.
- [After Upgrading an Oracle Database](#)
After a successful upgrade, note the following:

Related Topics

- [To create an on-demand backup of a database](#)
- [To configure automatic backups for a database](#)

How the Upgrade Operation Is Performed by the Database Service

During the upgrade process, the Database service does the following:

- Executes an automatic precheck. This allows the system to identify issues needing mitigation and to stop the upgrade operation.
- Sets a guaranteed restore point, enabling it to perform a flashback in the event of an upgrade failure.
- Moves the database to a user-specified Oracle Database Home that uses the desired target software version.
- Runs the Database Upgrade Assistant (DBUA) software to perform the upgrade.
- For databases in Data Guard associations, redo apply is disabled until both the primary and standby databases are successfully upgraded, at which point redo apply is re-enabled by the system. The system then enables Open Mode after redo apply is enabled.

Rolling Back an Oracle Database Unsuccessful Upgrade

If your upgrade does not complete successfully, then you have the option of performing a rollback.

Details about the failure are displayed on the **Database Details** page in the Console, allowing you to analyze and resolve the issues causing the failure.

A rollback resets your database to the state prior to the upgrade. All changes to the database made during and after the upgrade will be lost. The rollback option is provided in a banner message displayed on the database details page of a database following an unsuccessful upgrade operation. See *Using the Console to Roll Back a Failed Database Upgrade* for more information.

For standby databases in Oracle Data Guard associations, rollback is accomplished by moving the standby back to the original Database Home. See [To move a database to another Database Home](#) for instructions.

Related Topics

- [To roll back a failed database upgrade](#)

After Upgrading an Oracle Database

After a successful upgrade, note the following:

- Check that automatic backups are enabled for the database if you disabled them prior to upgrading. See *Customizing the Automatic Backup Configuration* for more information.
- Edit the Oracle Database `COMPATIBLE` parameter to reflect the new Oracle Database software version. See *What Is Oracle Database Compatibility?* for more information.
- If your database uses a `database_name.env` file, ensure that the variables in the file have been updated to point to the 19c Database Home. These variables should be automatically updated during the upgrade process.
- If you are upgrading a non-container database to Oracle Database version 19c, you can convert the database to a pluggable database after converting. See *How to Convert Non-CDB to PDB (Doc ID 2288024.1)* for instructions on converting your database to a pluggable database.
- If your old Database Home is empty and will not be reused, you can remove it. See *Using the Console to Delete an Oracle Database Home* for more information.
- For databases in Data Guard associations, check the open mode and redo apply status after the upgrade is complete.

Related Topics

- [Managing Exadata Database Backups by Using dbaascli](#)
- [What Is Oracle Database Compatibility?](#)
- [How to Convert Non-CDB to PDB - Step by Step Example \(Doc ID 2288024.1\)](#)
- [To delete a database home](#)

You cannot delete a Database Home that contains databases. You must first terminate the databases to empty the Database Home. See [To terminate a database](#) to learn how to terminate a database.

Using the Console to Upgrade a Database

Procedures to precheck and upgrade a database, rollback a failed upgrade, and view the upgrade history.

- [To upgrade or precheck an Exadata database](#)
Procedure to upgrade or precheck an Exadata database.

- [To roll back a failed database upgrade](#)
- [To view the upgrade history of a database](#)

To upgrade or precheck an Exadata database

Procedure to upgrade or precheck an Exadata database.

The following steps apply to databases for which either of the following apply:

- The database is the primary database in a Data Guard association
- The database is not part of a Data Guard association

To upgrade a standby database in a Data Guard configuration, move the standby to a Database Home using the Oracle Database version you are upgrading to. See [To move a database to another Database Home](#) for details.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you want to upgrade.

Note

If your database is in an Exadata Cloud Infrastructure instance that does not use the [new Exadata resource model](#), you will need to [switch the instance to the new model](#) before you can upgrade your database.

4. In the list of databases on the details page of the VM cluster, click the name of the database you want to upgrade to view the Database Details page.
5. Click **Actions**, and then select **Upgrade**.
6. In the **Upgrade Database** dialogue, select the following:
 - **Oracle Database version:** The drop-down selector lists only Oracle Database versions that are compatible with an upgrade from the current software version the database is using. The target software version must be higher than the database's current version.
 - **Target Database Home:** Select a Database Home for your database. The list of Database Homes is limited to those homes using the most recent versions of Oracle Database 19c software. Moving the database to the new Database Home results in the database being upgraded to the major release version and patching level of the new Database Home.
7. Click one of the following:
 - **Run Precheck:** This option starts an upgrade precheck to identify any issues with your database that need mitigation before you perform an upgrade.
 - **Upgrade Database:** This option starts upgrade operation. Oracle recommends performing an upgrade only after you have performed a successful precheck on the database.

To roll back a failed database upgrade

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the database with the failed upgrade.
4. Find the database that was unsuccessfully upgraded, and click its name to display details about it.
5. The database must display a banner at the top of the details page that includes a **Rollback** button and details about what issues caused the upgrade failure.
6. Click **Rollback**.
7. In the **Confirm rollback** dialog, confirm that you want to initiate a rollback to the previous Oracle Database version.

To view the upgrade history of a database

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
3. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, click the name of the VM cluster that contains the database you want to upgrade.

Note

If your database is in an Exadata Cloud Infrastructure instance that does not use the *new Exadata resource model*, you will need to *switch the instance to the new Exadata resource model* before you can upgrade your database.

4. In the list of databases on the details page of the VM cluster, click the name of the database for which you want to view the upgrade history.
5. Click **Upgrade History**.

Using the API to upgrade Databases

Use the following APIs to manage database upgrades:

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

Use these API operations to manage database upgrades:

- [ListDatabaseUpgradeHistoryEntries](#)
- [UpgradeDatabase](#)

For the complete list of APIs for the Database service, see [Database Service API](#).

Note

When using the `UpgradeDatabase` API to upgrade an Exadata Cloud Infrastructure database, you must specify `DB_HOME` as the upgrade source.

Database Health Checks During Infrastructure and Guest VM OS Update Operations

The Database Health Check feature is designed to identify potential issues that could cause database downtime or service degradation during update operations involving container databases (CDBs) and pluggable databases (PDBs).

This feature performs pre- and post-update validations to ensure the stability and health of database services throughout the update lifecycle. By proactively detecting issues before and after updates, it helps:

- Prevent downtime by identifying potential issues in advance.
- Ensure database availability by detecting degradations post-update.
- Improve reliability of database update workflows through automated health validations.

Note

- This is an opt-in feature, disabled by default and can be enabled by Oracle upon customer request.
- The feature performs validations during the following operations:
 - Infrastructure update (DB Server Update)
 - Guest VM OS update

Downtime Detection (Pre-Update)

Before initiating an update operation, the feature validates the running states of database, service, and listener CRS resources across the VM cluster to identify any risk of downtime.

- If at least one resource is running on remote nodes that are not being rebooted as part of the update, the operation proceeds.
- For PDBs, the feature verifies that all PDBs are in a consistent state, either restricted or non-restricted.
- If a mix of states is detected, the update operation stops immediately to prevent inconsistencies or service instability.

Degradation Detection (Post-Update)

After the update operation on the DB Server or Guest VM is complete, the feature compares the resource states before and after the update.

If any mismatch or degradation in database service state is detected, the update operation fails immediately and an appropriate error message is displayed.

Enabling the Feature

To enable this feature, customers can request Oracle to activate it for their environment. Once enabled, the Database Health Check runs automatically as part of the update process, validating database health before and after updates to ensure continuous availability.

Resolving Dependency Issues Associated with Additional Non-Exadata Software Packages for Guest VM Upgrade

If you've installed non-Exadata software packages beyond those provided by Oracle, and the precheck fails during a Guest VM upgrade due to conflicts between and Oracle-installed RPMs, you can use the following procedure to resolve the conflicts and proceed with the upgrade.

For updates that do not change the major Oracle Linux version, this integrated capability allows you to update additional non-Exadata software packages as part of the Exadata database server update. It simplifies the handling of package dependency issues that may arise when such non-Exadata software packages are present on the system.

You can iteratively run precheck to identify and resolve dependency issues associated with the additional non-Exadata software packages. Once the required updates are understood, you can confidently perform the Exadata database server update and update the additional packages in a single, coordinated operation.

Ensure that the configuration file exists on the target server to trigger the setup of a temporary YUM repository for non-Exadata software packages.

- **File Location:** `/etc/exadata/additional-packages.txt`
- **Ownership and Permissions:** This file must be owned and modifiable by the `root` user only.

If the file exists, it is used to collect information about the required non-Exadata software packages and to set up and enable a temporary YUM repository. If the file is not present, no repository is configured.

You may also create a symbolic link at `/etc/exadata/additional-packages.txt` that points to a configuration file located elsewhere—typically on a shared mount.

File Format

The file must contain a list of non-Exadata software packages, with each entry on a new line. Supported formats include:

- `http(s)://path/to/package.rpm`: Full URL to the RPM file
- `/full/path/to/package.rpm`: Absolute path to a local RPM file
- `repo:package.rpm`: Reference to a package in an existing YUM repository

Note

- If using the `repo:` format, ensure the referenced repository is defined in the target server's YUM configuration.
- Local files can reside in standard local directories, NFS mounts, or ACFS mounts.

Example: `additional-packages.txt`

```
/u01/elfutils-debuginfod-client-0.190-2.el8.x86_64.rpm
/u01/elfutils-libelf-devel-0.190-2.el8.x86_64.rpm
/u01/keyutils-libs-devel-1.5.10-9.0.1.el8.x86_64.rpm
https://example.com/packages/krb5-devel-1.18.2-28.0.1.el8_10.x86_64.rpm
https://example.com/packages/memstrack-0.2.5-2.el8.x86_64.rpm
/u01/pigz-2.4-4.el8.x86_64.rpm
/u01/sss-d-nfs-idmap-2.9.4-3.0.1.el8_10.x86_64.rpm
https://example.com/packages/timedatex-0.5-3.el8.x86_64.rpm
https://example.com/packages/zlib-devel-1.2.11-25.el8.x86_64.rpm
```

Patching and Updating an Exadata Cloud Infrastructure System Manually

This topic describes the procedures for patching and updating various components in Exadata Cloud Service outside of the cloud automation.

For information related to patching and updating with `dbaascli`, refer to "*Patching Oracle Grid Infrastructure and Oracle Databases Using `dbaascli`*".

Note

For more guidance on achieving continuous service during patching operations, see the *Application Checklist for Continuous Service for MAA Solutions* white paper.

- [Patching Oracle Database and Oracle Grid Infrastructure Software Manually](#)
For daylight savings time, and some routine or one-off patches, it can be necessary for you to patch software manually.
- [Updating the Exadata Cloud VM Cluster OS Manually](#)
You update the operating systems of Exadata compute nodes by using the `patchmgr` tool.
- [Updating Tooling on an Exadata Cloud Infrastructure Instance](#)
Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including `dbaascli` commands.

Related Topics

- [Patching Oracle Grid Infrastructure and Oracle Databases Using `dbaascli`](#)
Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.
- [Application Checklist for Continuous Service for MAA Solutions](#)

Patching Oracle Database and Oracle Grid Infrastructure Software Manually

For daylight savings time, and some routine or one-off patches, it can be necessary for you to patch software manually.

To perform routine patching of Oracle Database and Oracle Grid Infrastructure software, Oracle recommends that you use the facilities provided by Oracle Exadata Database Service on Dedicated Infrastructure. However, under some circumstances, it can be necessary for you to patch the Oracle Database or Oracle Grid Infrastructure software manually:

- **Daylight Savings Time (DST) Patching:** Because they cannot be applied in a rolling fashion, patches for the Oracle Database DST definitions are not included in the routine

patch sets for Exadata Cloud Infrastructure. If you need to apply patches to the Oracle Database DST definitions, you must do so manually. See My Oracle Support Doc ID 412160.1.

- **Non-routine or One-off Patching:** If you encounter a problem that requires a patch which is not included in any routine patch set, then work with Oracle Support Services to identify and apply the appropriate patch.

For general information about patching Oracle Database, refer to information about patch set updates and requirements in *Oracle Database Upgrade Guide* for your release.

Related Topics

- <https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=1929745.1>
- <https://support.oracle.com/epmos/faces/DocumentDisplay?cmd=show&type=NOT&id=412160.1>

Updating the Exadata Cloud VM Cluster OS Manually

You update the operating systems of Exadata compute nodes by using the `patchmgr` tool.

This utility manages the entire update of one or more compute nodes remotely, including running pre-reboot, reboot, and post-reboot steps. You can run the utility from either an Exadata compute node or a non-Exadata server running Oracle Linux. The server on which you run the utility is known as the "driving system." You cannot use the driving system to update itself. Therefore, if the driving system is one of the Exadata compute nodes on a system you are updating, you must run a separate operation on a different driving system to update that server.

The following two scenarios describe typical ways of performing the updates:

Scenario 1: Non-Exadata Driving System

The simplest way to run the update the Exadata system is to use a separate Oracle Linux server to update all Exadata compute nodes in the system.

Scenario 2: Exadata Node Driving System

You can use one Exadata compute node to drive the updates for the rest of the compute nodes in the system, and then use one of the updated nodes to drive the update on the original Exadata driver node.

For example: You are updating a half rack Exadata system, which has four compute nodes - node1, node2, node3, and node4. First, use node1 to drive the updates of node2, node3, and node4. Then, use node2 to drive the update of node1.

The driving system requires root user `SSH` access to each compute node the utility will update.

- [Preparing for the OS Updates](#)
Determine the latest software version available, and connectivity to the proper `yum` repository
- [To update the OS on all compute nodes of an Exadata Cloud Infrastructure instance](#)
Procedure to update all compute nodes using `patchmgr`.
- [Installing Additional Operating System Packages](#)
Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Dedicated Infrastructure.

Preparing for the OS Updates

Determine the latest software version available, and connectivity to the proper yum repository

Caution

Do not install NetworkManager on the Exadata Cloud Infrastructure instance. Installing this package and rebooting the system results in severe loss of access to the system.

- Before you begin your updates, review *Exadata Cloud Service Software Versions* ([Doc ID 2333222.1](#)) to determine the latest software version and target version to use.
- Some steps in the update process require you to specify a YUM repository. The YUM repository URL is:

```
http://yum-<region_identifier>.oracle.com/repo/EngineeredSystems/exadata/dbserver/<latest_version>/base/x86_64.
```

Region identifiers are text strings used to identify Oracle Cloud Infrastructure regions (for example, us-phoenix-1). You can find a complete list of region identifiers in *Regions*.

You can run the following `curl` command to determine the latest version of the YUM repository for your Exadata Cloud Service instance region:

```
curl -s -X GET http://yum-<region_identifier>.oracle.com/repo/EngineeredSystems/exadata/dbserver/ | egrep "18.1."
```

This example returns the most current version of the YUM repository for the US West (Phoenix) region:

```
curl -s -X GET http://yum-us-phoenix-1.oracle.com/repo/EngineeredSystems/exadata/dbserver/ | egrep "18.1."
<a href="18.1.4.0.0/">18.1.4.0.0/</a> 01-Mar-2018 03:36 -
```

- To apply OS updates, the system's **VCN** must be configured to allow access to the YUM repository. For more information, see *Option 2: Service Gateway to Both Object Storage and YUM repos.*

Related Topics

- [Regions](#)
- [Option 2: Service Gateway Access to OCI Service for Both the Client and Backup Subnets](#)
You configure both the client subnet and backup subnet to use the service gateway for access to the Oracle Services Network, which includes Object storage for backups, Oracle YUM repos for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS Integration).

To update the OS on all compute nodes of an Exadata Cloud Infrastructure instance

Procedure to update all compute nodes using `patchmgr`.

This example procedure assumes the following:

- The system has two compute nodes, `node1` and `node2`.
 - The target version is `18.1.4.0.0.180125.3`.
 - Each of the two nodes is used as the driving system for the update on the other one.
1. Gather the environment details.

- a. SSH to `node1` as `root` and run the following command to determine the version of Exadata:

```
[root@node1]# imageinfo -ver
12.2.1.1.4.171128
```

- b. Switch to the `grid` user, and identify all computes in the cluster.

```
[root@node1]# su - grid
[grid@node1]$ olsnodes
node1
node1
```

2. Configure the driving system.

- a. Switch back to the `root` user on `node1`, check whether a root ssh key pair (`id_rsa` and `id_rsa.pub`) already exists. If not, then generate it.

```
[root@node1 .ssh]# ls /root/.ssh/id_rsa*
ls: cannot access /root/.ssh/id_rsa*: No such file or directory
[root@node1 .ssh]# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/root/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /root/.ssh/id_rsa.
Your public key has been saved in /root/.ssh/id_rsa.pub.
The key fingerprint is:
93:47:b0:83:75:f2:3e:e6:23:b3:0a:06:ed:00:20:a5
root@node1.fraadlclient.exadataclientne.oraclevcn.com
The key's randomart image is:
+--[ RSA 2048 ]-----+
|O..    + .          |
|O.     o *         |
|E      . o o        |
| . .    =          |
| o .   S =         |
|  +    = .         |
|   +   o o         |
| . .   + .         |
|    . . .          |
+-----+

```

- b. Distribute the public key to the target nodes, and verify this step. In this example, the only target node is `node2`.

```
[root@node1 .ssh]# scp -i ~opc/.ssh/id_rsa ~root/.ssh/id_rsa.pub
opc@node2:/tmp/id_rsa.node1.pub
id_rsa.pub
```

```
[root@node2 ~]# ls -al /tmp/id_rsa.node1.pub
-rw-r--r-- 1 opc opc 442 Feb 28 03:33 /tmp/id_rsa.node1.pub
[root@node2 ~]# date
Wed Feb 28 03:33:45 UTC 2018
```

- c. On the target node (node2, in this example), add the root public key of node1 to the root `authorized_keys` file.

```
[root@node2 ~]# cat /tmp/id_rsa.node1.pub >> ~root/.ssh/authorized_keys
```

- d. Download `dbserver.patch.zip` as `p21634633_12*_Linux-x86-64.zip` onto the driving system (node1, in this example), and unzip it. See *dbnodeupdate.sh* and *dbserver.patch.zip: Updating Exadata Database Server Software using the DBNodeUpdate Utility and patchmgr (Doc ID 1553103.1)* for information about the files in this `.zip`.

```
[root@node1 patch]# mkdir /root/patch
[root@node1 patch]# cd /root/patch
[root@node1 patch]# unzip p21634633_181400_Linux-x86-64.zip
Archive:  p21634633_181400_Linux-x86-64.zip   creating:
dbserver_patch_5.180228.2/
  creating: dbserver_patch_5.180228.2/ibdiagtools/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/cable_check.pl
  inflating: dbserver_patch_5.180228.2/ibdiagtools/setup-ssh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/VERSION_FILE
  extracting: dbserver_patch_5.180228.2/ibdiagtools/xmonib.sh
  inflating: dbserver_patch_5.180228.2/ibdiagtools/monitord
  inflating: dbserver_patch_5.180228.2/ibdiagtools/checkbadlinks.pl
  creating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
VerifyTopologyUtility.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/
verifylib.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Node.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Rack.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Group.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topologies/Switch.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/topology-zfs
  inflating: dbserver_patch_5.180228.2/ibdiagtools/dcli
  creating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteScriptGenerator.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
CommonUtils.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
SolarisAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
LinuxAdapter.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteLauncher.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
remoteConfig.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/spawnProc.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/
runDiagnostics.pm
  inflating: dbserver_patch_5.180228.2/ibdiagtools/netcheck/OSAdapter.pm
```

```

inflating: dbserver_patch_5.180228.2/ibdiagtools/SampleOutputs.txt
inflating: dbserver_patch_5.180228.2/ibdiagtools/infinicheck
inflating: dbserver_patch_5.180228.2/ibdiagtools/ibping_test
inflating: dbserver_patch_5.180228.2/ibdiagtools/tar_ibdiagtools
inflating: dbserver_patch_5.180228.2/ibdiagtools/verify-topology
inflating: dbserver_patch_5.180228.2/installfw_exadata_ssh
  creating: dbserver_patch_5.180228.2/linux.db.rpms/
inflating: dbserver_patch_5.180228.2/md5sum_files.lst
inflating: dbserver_patch_5.180228.2/patchmgr
inflating: dbserver_patch_5.180228.2/xcp
inflating: dbserver_patch_5.180228.2/ExadataSendNotification.pm
inflating: dbserver_patch_5.180228.2/ExadataImageNotification.pl
inflating: dbserver_patch_5.180228.2/kernelupgrade_oldbios.sh
inflating: dbserver_patch_5.180228.2/cellboot_usb_pci_path
inflating: dbserver_patch_5.180228.2/exadata.img.env
inflating: dbserver_patch_5.180228.2/README.txt
inflating: dbserver_patch_5.180228.2/exadataLogger.pm
inflating: dbserver_patch_5.180228.2/patch_bug_26678971
inflating: dbserver_patch_5.180228.2/dcli
inflating: dbserver_patch_5.180228.2/patchReport.py
extracting: dbserver_patch_5.180228.2/dbnodeupdate.zip
  creating: dbserver_patch_5.180228.2/plugins/
inflating: dbserver_patch_5.180228.2/plugins/010-check_17854520.sh
inflating: dbserver_patch_5.180228.2/plugins/020-check_22468216.sh
inflating: dbserver_patch_5.180228.2/plugins/040-check_22896791.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_bash
inflating: dbserver_patch_5.180228.2/plugins/050-check_22651315.sh
inflating: dbserver_patch_5.180228.2/plugins/005-check_22909764.sh
inflating: dbserver_patch_5.180228.2/plugins/000-check_dummy_perl
inflating: dbserver_patch_5.180228.2/plugins/030-check_24625612.sh
inflating: dbserver_patch_5.180228.2/patchmgr_functions
inflating: dbserver_patch_5.180228.2/exadata.img.hw
inflating: dbserver_patch_5.180228.2/libxcp.so.1
inflating: dbserver_patch_5.180228.2/imageLogger
inflating: dbserver_patch_5.180228.2/ExaXMLNode.pm
inflating: dbserver_patch_5.180228.2/fwverify

```

- e. Create the `dbs_group` file that contains the list of compute nodes to update. Include the nodes listed after running the `olsnodes` command in step 1 except for the driving system node. In this example, `dbs_group` should include only `node2`.

```

[root@node1 patch]# cd /root/patch/dbserver_patch_5.180228
[root@node1 dbserver_patch_5.180228]# cat dbs_group
node2

```

3. Run a patching precheck operation.

Note

You must run the precheck operation with the `-nomodify_at_prereq` option to prevent any changes to the system that could impact the backup you take in the next step. Otherwise, the backup might not be able to roll back the system to its original state, should that be necessary.

```
patchmgr -dbnodes dbs_group -precheck -yum_repo <yum_repository> -
target_version <target_version> -nomodify_at_prereq
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -
precheck -yum_repo http://yum-phx.oracle.com/repo/EngineeredSystems/
exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3 -nomodify_at_prereq

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
*****
*****
2018-02-28 21:22:45 +0000      :Working: DO: Initiate precheck on 1
node(s)
2018-02-28 21:24:57 +0000      :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:26:15 +0000      :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:26:47 +0000      :Working: DO: dbnodeupdate.sh running a
precheck on node(s).
2018-02-28 21:28:23 +0000      :SUCCESS: DONE: Initiate precheck on
node(s).
```

4. Back up the current system.**Note**

This is the proper stage to take the backup, before any modifications are made to the system.

```
patchmgr -dbnodes dbs_group -backup -yum_repo <yum_repository> -
target_version <target_version> -allow_active_network_mounts
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# ./patchmgr -dbnodes dbs_group -
backup -yum_repo http://yum-phx.oracle.com/repo/EngineeredSystems/
exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3 -allow_active_network_mounts

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
*****
*****
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on 1
node(s).
2018-02-28 21:29:00 +0000      :Working: DO: Initiate backup on node(s)
2018-02-28 21:29:01 +0000      :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:30:18 +0000      :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:30:51 +0000      :Working: DO: dbnodeupdate.sh running a
backup on node(s).
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on
node(s).
2018-02-28 21:35:50 +0000      :SUCCESS: DONE: Initiate backup on 1
node(s).
```

5. Remove all custom RPMs from the target compute nodes that will be updated. Custom RPMs are reported in precheck results. They include RPMs that were manually installed after the system was provisioned.
 - If you are updating the system from version 12.1.2.3.4.170111, and the precheck results include `krb5-workstation-1.10.3-57.el6.x86_64`, remove it. (This item is considered a custom RPM for this version.)
 - Do **not** remove `exadata-sun-vm-computenode-exact` or `oracle-ofed-release-guest`. These two RPMs are handled automatically during the update process.
6. Run the `nohup` command to perform the update.

```
nohup patchmgr -dbnodes dbs_group -upgrade -nobackup -yum_repo
<yum_repository> -target_version <target_version> -
allow_active_network_mounts &
```

The output should look like the following example:

```
[root@node1 dbserver_patch_5.180228]# nohup ./patchmgr -dbnodes dbs_group -
upgrade -nobackup -yum_repo http://yum-phx.oracle.com/repo/
EngineeredSystems/exadata/dbserver/18.1.4.0.0/base/x86_64 -target_version
18.1.4.0.0.180125.3 -allow_active_network_mounts &
```

```

*****
*****
NOTE    patchmgr release: 5.180228 (always check MOS 1553103.1 for the
latest release of dbserver.patch.zip)
NOTE
NOTE    Database nodes will reboot during the update process.
NOTE
WARNING Do not interrupt the patchmgr session.
WARNING Do not resize the screen. It may disturb the screen layout.
WARNING Do not reboot database nodes during update or rollback.
WARNING Do not open logfiles in write mode and do not try to alter them.
*****
*****

```

```

2018-02-28 21:36:26 +0000      :Working: DO: Initiate prepare steps on
node(s).
2018-02-28 21:36:26 +0000      :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:37:44 +0000      :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:38:43 +0000      :SUCCESS: DONE: Initiate prepare steps on
node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on 1
node(s).
2018-02-28 21:38:43 +0000      :Working: DO: Initiate update on node(s)
2018-02-28 21:38:49 +0000      :Working: DO: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000      :SUCCESS: DONE: Get information about any
required OS upgrades from node(s).
2018-02-28 21:38:59 +0000      :Working: DO: dbnodeupdate.sh running an
update step on all nodes.
2018-02-28 21:48:41 +0000      :INFO    : node2 is ready to reboot.
2018-02-28 21:48:41 +0000      :SUCCESS: DONE: dbnodeupdate.sh running
an update step on all nodes.
2018-02-28 21:48:41 +0000      :Working: DO: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :SUCCESS: DONE: Initiate reboot on node(s)
2018-02-28 21:48:57 +0000      :Working: DO: Waiting to ensure node2 is
down before reboot.
2018-02-28 21:56:18 +0000      :Working: DO: Initiate prepare steps on
node(s).
2018-02-28 21:56:19 +0000      :Working: DO: Check free space and verify
SSH equivalence for the root user to node2
2018-02-28 21:57:37 +0000      :SUCCESS: DONE: Check free space and
verify SSH equivalence for the root user to node2
2018-02-28 21:57:42 +0000      :SEEMS ALREADY UP TO DATE: node2
2018-02-28 21:57:43 +0000      :SUCCESS: DONE: Initiate update on node(s)

```

7. After the update operation completes, verify the version of the kernel on the compute node that was updated.

```

[root@node2 ~]# imageinfo -ver
18.1.4.0.0.180125.3

```

8. If the driving system is a compute node that needs to be updated (as in this example), repeat steps 2 through 7 of this procedure using an updated compute node as the driving

system to update the remaining compute node. In this example update, you would use `node2` to update `node1`.

9. On each compute node, run the `uptrack-install` command as root to install the available `ksplice` updates.

```
uptrack-install --all -y
```

Installing Additional Operating System Packages

Review these guidelines before you install additional operating system packages for Oracle Exadata Database Service on Dedicated Infrastructure.

You are permitted to install and update operating system packages on Oracle Exadata Database Service on Dedicated Infrastructure as long as you do not modify the kernel or InfiniBand-specific packages. However, Oracle technical support, including installation, testing, certification and error resolution, does not apply to any non-Oracle software that you install.

Also be aware that if you add or update packages separate from an Oracle Exadata software update, then these package additions or updates can introduce problems when you apply an Oracle Exadata software update. Problems can occur because additional software packages add new dependencies that can interrupt an Oracle Exadata update. For this reason, Oracle recommends that you minimize customization.

If you install additional packages, then Oracle recommends that you have scripts to automate the removal and reinstallation of those packages. After an Oracle Exadata update, if you install additional packages, then verify that the additional packages are still compatible, and that you still need these packages.

For more information, refer to *Oracle Exadata Database Machine Maintenance Guide*.

Related Topics

- [Installing, Updating, and Managing Additional Software Packages](#)

Updating Tooling on an Exadata Cloud Infrastructure Instance

Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including `dbaascli` commands.

The cloud tooling is automatically updated by Oracle when new releases are made available. If needed, you can follow the steps in *Updating Cloud Tooling Using dbaascli* to ensure you have the latest version of the cloud tooling on all virtual machines in the VM cluster

Related Topics

- [Updating Cloud Tooling Using dbaascli](#)
To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

Interim Software Updates

For authorized environments, learn how to download interim software updates.

This feature enables cloud-only customers to download one-off patches from the OCI console and API. There is no option to apply the downloaded patch via console and API. To apply these patches, customers must log in to their VM and run the patch apply utility.

Note

To be able to download interim software update, you should at least have an ExaDB-D infrastructure provisioned.

Downloading one-off patches does not replace Database Software Image (DSI) creation. Customers must continue to use Database Software Images (DSI) to build and deploy their customized images.

- [Create an Interim Software Update](#)
- [Download an Interim Software Update](#)
- [Delete an Interim Software Update](#)
- [Move an Interim Software Update Resource to Another Compartment](#)
- [Using the API to Manage Interim Software Updates](#)

Create an Interim Software Update

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Resources**, click **Interim software updates**.
Interim software update page is displayed.
3. Click **Create interim software update**.
Create interim software update panel is displayed.
4. Enter the following details in the panel:
 - a. **Name**: Descriptive name for the patch download path.
 - b. **Compartment**: Select a compartment where you want to create the patch resource.
 - c. **Database version**: Choose the Database version for your image.
 - d. **Release Update**: Choose any supported Oracle Database release update (RU).
 - e. **Interim software update number**: Optionally, enter an interim patch number.
 - f. **Tag**: Apply a tag.
5. Click **Create**.

Download an Interim Software Update

The patch download path is valid for four days. Download the patch within the specified timeframe.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to download and then select **Download**.
The system starts downloading the patch.
2. You can also download a patch from the Interim Software Updates details page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Download**.

Note

You can only download the patches that are in **Available** state.

Interim Software Updates Lifecycle States:

- **Available:** Patch has been created successfully and the time-to-live (TTL) has not expired.
- **Creating:** The patch creation process is in progress.
- **Expired:** The lifetime of the patch download link has expired, which means you cannot download it.
- **Failed:** The patch create failed due to some error.
- **Terminating:** The patch deletion process is in progress.
- **Terminated:** The patch has been deleted.

Delete an Interim Software Update

Be discrete in deleting interim software updates. However, you can delete the interim software updates that have expired to free up space in the Object Store.

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to delete and then select **Delete**.
2. In the resulting dialog, enter the name of the patch to confirm and then click **Delete**.
3. You can also delete a patch from the Interim Software Updates details page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Delete**.

Move an Interim Software Update Resource to Another Compartment

1. On the Interim Software Updates page, click the Actions menu (three dots) of the Interim Software Update you want to move and then select **Move resource**.
2. In the resulting dialog, choose a new compartment, and click **Move Resource**.
3. You can also move a patch from the Interim Software Updates details page.
 - Click the Actions button (three dots) for the patch you're interested in, and select **Move Resource**.

Using the API to Manage Interim Software Updates

ExaDB-C@C and ExaDB-D use the same API to manage interim software updates.

For information about using the API and signing requests, see *REST APIs and Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Use these API operations to manage interim software updates:

- `CreateOneoffPatch`

- `DeleteOneoffPatch`
- `DownloadOneoffPatch`
- `UpdateOneoffPatch`
- `ListOneoffPatches`
- `GetOneoffPatch`
- `ChangeOneoffPatchCompartment`

Related Topics

- [REST APIs](#)
- [Security Credentials](#)
- [Software Development Kits and Command Line Interface](#)
- [OneoffPatch Reference](#)

Use Oracle Data Guard with Exadata Cloud Infrastructure

Learn to configure and manage Data Guard Groups in your VM cluster.

- [About Using Oracle Data Guard with Exadata Cloud Infrastructure](#)
Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions.
- [Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure](#)
An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.
- [Working with Data Guard](#)
Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.
- [Support for Concurrent Data Guard Operations in Multiple Standby Environments](#)
- [Enhanced Data Guard Health Status Reporting](#)
The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.
- [Using the Console to Manage an Oracle Data Guard Group](#)
Learn how to enable a Data Guard Group between databases, change the role of a database in a Data Guard Group using either a switchover or a failover operation, and reinstate a failed database.
- [Using the API to manage Data Guard associations](#)
Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:
- [Using the API to manage Data Guard Group](#)
Use these API operations to manage a Data Guard Group on an Exadata Cloud Infrastructure instance:

About Using Oracle Data Guard with Exadata Cloud Infrastructure

Oracle Data Guard provides a comprehensive set of services that create, maintain, manage, and monitor one or more standby databases to enable production Oracle databases to survive disasters and data corruptions.

Oracle Data Guard maintains these standby databases as copies of the production database. Then, if the production database becomes unavailable because of a planned or an unplanned outage, Oracle Data Guard can switch any standby database to the production role, minimizing the downtime associated with the outage. Oracle Data Guard can be used with traditional backup, restoration, and cluster techniques to provide a high level of data protection and data availability. Oracle Data Guard transport services are also used by other Oracle features such as Oracle Streams and Oracle GoldenGate for efficient and reliable transmission of redo from a source database to one or more remote destinations.

For complete information on Oracle Data Guard, see the [Oracle Data Guard Concepts and Administration](#) documentation and [Oracle Data Guard Broker Concepts](#) on the [Oracle Database Documentation](#) portal.

This topic explains how to use the Console or the API to configure and manage Data Guard resources in your VM cluster.

When you use the Console or the API to enable Data Guard for an Exadata database compute node database:

- The standby database that is created is a physical standby.
- The versions of peer databases (primary and standby) are identical.
- The standby database is deployed as an open, read-only database (Active Data Guard).
- A primary database can support up to a maximum of six standby databases.

Prerequisites for Using Oracle Data Guard with Exadata Cloud Infrastructure

An Exadata Cloud Infrastructure Oracle Data Guard implementation requires two existing Exadata VM Clusters: one containing an existing database that is to be duplicated by Data Guard, and one that will house the new standby database by Data Guard.

Note

Oracle strongly recommends the primary and standby databases for any production workloads be on different Exadata Cloud Infrastructures for better fault isolation and disaster protection. If you are adding a new standby in the same region with multiple availability domains, Oracle recommends choosing a separate availability domain for complete availability domain or data center fault isolation. If you are adding a new standby across regions, the standby will have fault isolation for a regional failure as well.

When enabling Data Guard, you must create a new Database Home on the standby instance to host the new standby database. Alternatively, you can provision the standby database within an existing Database Home on the standby instance. For information on creating the required resources for the standby system, see the following topics:

- [To create a Cloud Exadata infrastructure resource](#)
- [To create an ASM cloud VM cluster](#)
- [To create a new Database Home in an existing Exadata Cloud Infrastructure instance](#)

You can use a custom database software image to that contains the necessary patches for your databases when creating a Database Home on either the primary or the standby Exadata instance. See [Oracle Database Software Images](#) for information on working with custom Oracle Database software images.

If you choose to provision a standby database in an existing Database Home, ensure that the target Database Home on the standby instance has all required patches that are in use for the primary database before you provision the standby database. See the following topic for more information on patching an existing Database Home:

- [To patch the Oracle Database software in a Database Home \(cloud VM cluster\)](#)

If you are creating a Data Guard Group and you are using customer managed keys to encrypt the database, you must have configured the Vault Service and created a master key. See *To administer Vault encryption keys* and *Key and Secret Management Concepts*.

- [Network Requirements for Data Guard](#)
Before setting up Data Guard ensure that your Exadata Cloud Infrastructure environment meets the following network requirements:
- [Password Requirements](#)
To change the SYS password or rotate TDE keys, use OCI API.
- [Known Issues for Exadata Cloud Infrastructure and Data Guard](#)
Possible TDE key replication issue, and MRP and DG LCM operation failures.
- [Adding a Node to a VM Cluster](#)
- [Removing a Node from a VM Cluster](#)

Related Topics

- [Customer-Managed Keys in Exadata Cloud Infrastructure](#)
Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- [To administer Vault encryption keys](#)
Use this procedure to rotate the Vault encryption key or change the encryption management configuration.
- [Key and Secret Management Concepts](#)

Network Requirements for Data Guard

Before setting up Data Guard ensure that your Exadata Cloud Infrastructure environment meets the following network requirements:

- The primary and standby databases can be part of VM clusters in different compartments.
- If you want to configure Oracle Data Guard across regions, then you must configure remote virtual cloud network (VCN) peering between the primary and standby databases. Networking is configured on the cloud VM cluster resource for systems using the *The new Exadata Resource Model*. See *Remote VCN Peering using an RPC*.
For Exadata Data Guard configurations, OCI supports the use of hub-and-spoke network topology for the VCNs within each region. This means that the primary and standby databases can each utilize a "spoke" VCN that passes network traffic to the "hub" VCN

that has a remote peering connection. See *Transit Routing inside a hub VCN* for information on setting up this network topology.

- To set up Oracle Data Guard within a single region, both Exadata Cloud Infrastructure instances must use the same VCN. When setting up Data Guard within the same region, Oracle recommends that the instance containing the standby database be in a different **availability domain** from the instance containing the primary database to improve availability and disaster recovery.
- Configure the ingress and egress security rules for the subnets of both Exadata Cloud Infrastructure instances in the Oracle Data Guard association to enable TCP traffic to move between the applicable ports. Ensure that the rules you create are stateful (the default).

For example, if the subnet of the primary Exadata Cloud Infrastructure instance uses the source CIDR 10.0.0.0/24 and the subnet of the standby instance uses the source CIDR 10.0.1.0/24, then create rules as shown in the subsequent example.

Note

The egress rules in the example show how to enable TCP traffic only for port 1521, which is a minimum requirement for Oracle Data Guard to work. If TCP traffic is already enabled for all destinations (0.0.0.0/0) on all of your outgoing ports, then you need not explicitly add these specific egress rules.

Security Rules for Subnet of Primary Exadata Cloud Infrastructure instance

Ingress Rules:

```
Stateless: No
Source: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Egress Rules:

```
Stateless: No
Destination: 10.0.1.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Security Rules for Subnet of Standby Exadata Cloud Infrastructure instance

Ingress Rules:

```
Stateless: No
Source: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

Egress Rules:

```
Stateless: No
Destination: 10.0.0.0/24
IP Protocol: TCP
Source Port Range: All
Destination Port Range: 1521
Allows: TCP traffic for ports: 1521
```

For information about creating and editing rules, see *Security Lists*.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)
- [Remote VCN Peering using an RPC](#)
- [Transit Routing inside a hub VCN](#)
- [Security Lists](#)

Password Requirements

To change the SYS password or rotate TDE keys, use OCI API.

Related Topics

- [Changing the Database Passwords](#)
To change the SYS password, or to change the TDE wallet password, use this procedure.

Known Issues for Exadata Cloud Infrastructure and Data Guard

Possible TDE key replication issue, and MRP and DG LCM operation failures.

KMS RPM `libkmsdepkcs11_1.286-1.286-1-Linux.rpm` is the latest available which supports active replication of key between cross-region KMS vaults (source and target), and it is recommended to upgrade the RPM on clusters participating in Data Guard. OCI Vault cross-region Data Guard works with a lower version of RPM, but the older version does not guarantee active replication of keys. If the TDE keys have any replication issue between vaults, Data Guard replication might have an impact (MRP fails on standby cluster due to missing key on target vault) and MRP could resume only after the keys are replicated to the target vault. To avoid MRP and DG LCM operation failures, upgrade the `libkms` RPM on both the clusters, and restart the databases (only databases using customer-managed keys).

Adding a Node to a VM Cluster

When adding a node to a VM cluster, an instance of the Data Guard database is automatically created on the new node. However, metadata updation on the remote database, that is, the primary database if addition is done on the standby database and vice versa, must be done manually.

This can be done by copying over the `addinstance` JSON file, `/var/opt/oracle/dbaas_acfs/<dbname>/addInstance.json` created at the end of instance addition and running the `/var/opt/oracle/ocde/rops update_instance <dbname> <path to addInstance JSON>` command on any node of the remote cluster.

Related Topics

- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Removing a Node from a VM Cluster

When removing a node from a VM cluster, the instance and its metadata on the removing node is deleted automatically. However, deletion of the corresponding metadata on the remote database, that is, the primary database if removal is done on the standby database and vice versa, must be done manually.

This can be done by running the `/var/opt/oracle/ocde/rops remove_instance <dbname> <Instance Name>` command on any node of the remote cluster.

Related Topics

- [To add compute and storage resources to a flexible cloud Exadata infrastructure resource](#)
This task describes how to use the Oracle Cloud Infrastructure Console to scale a flexible cloud Exadata infrastructure resource.

Working with Data Guard

Oracle Data Guard ensures high availability, data protection, and disaster recovery for enterprise data.

The primary and standby databases constitute a Data Guard Group. Most of your applications access the primary database. A standby database is a transactionally consistent copy of the primary database.

Data Guard maintains the standby database by transmitting and applying redo data from the primary database. If the primary database becomes unavailable, you can use Data Guard to switchover or failover the standby database to the primary role. This is true even if you have more than one standby database.

- [Switchover](#)
A switchover reverses the primary and standby database roles.
- [Failover](#)
A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.
- [Reinstate](#)
Reinstates a database into the standby role in a Data Guard Group.

Switchover

A switchover reverses the primary and standby database roles.

Each database continues to be part of the Data Guard Group in its new role. A switchover ensures no data loss. You can use a switchover before you perform planned maintenance on the primary database. Performing planned maintenance on an Exadata database virtual machine with a Data Guard Group is typically done by switching the primary to the standby role, performing maintenance on the standby, and then switching it back to the primary role.

Failover

A failover transitions the standby database into the primary role after the existing primary database fails or becomes unreachable.

The failover may or may not result in data loss depending upon the protection mode and whether your primary and target standby databases were synchronized at the time of the primary database failure. For more information refer to [Manual Failover](#) in the Data Guard documentation.

Reinstate

Reinstates a database into the standby role in a Data Guard Group.

You can use the `reinstate` command to return a failed database into service after correcting the cause of failure.

Note

You cannot terminate a primary database that is part of a Data Guard Group that contains one or more standby databases. You will have to terminate the standby databases first. Alternatively, you can switch over the primary database to the standby role, and then terminate the former primary.

You cannot terminate a VM cluster that includes Data Guard enabled databases. You must first terminate the standby databases that are part of the Data Guard Group.

Support for Concurrent Data Guard Operations in Multiple Standby Environments

In addition to the [Enhancements to support concurrent Data Guard, Container Database \(CDB\), and Pluggable Database \(PDB\) Operations](#), these enhancements enable you to perform concurrent operations on CDBs and PDBs alongside Data Guard migration operations in environments with multiple standby databases.

You can now perform the following operations in parallel within the same Oracle Home:

- Perform Add Standby operations concurrently on different databases.
- Create or delete a CDB while an Add Standby operation is running on another database, and vice versa.
- Create or delete a PDB while an Add Standby operation is running on another database, and vice versa.
- Performing Data Guard actions (switchover, failover, reinstate, convert to snapshot and convert to physical) while an Add Standby operation is running on another database, and vice versa.

Similarly, you can perform the following Migrate Data Guard operations in parallel:

- Perform migrate Data Guard operations concurrently on different databases.
- Perform migrate Data Guard operations while an Add Standby operation is running on another database, and vice versa.

- Create or delete a CDB while a migrate Data Guard operations is running on another database, and vice versa.
- Create or delete a PDB while a migrate Data Guard operations is running on another database, and vice versa.
- Performing Data Guard actions (switchover, failover, reinstate, convert to snapshot and convert to physical) while a migrate Data Guard operations is running on another database, and vice versa.

Enhanced Data Guard Health Status Reporting

The enhanced Data Guard health status reporting provides comprehensive insights into protection mode, switchover and failover readiness, and data loss exposure across primary and standby databases.

With clear visual indicators (green, yellow, red, gray), you can quickly assess the readiness of your databases for role transitions and failover events. Additionally, detailed redo transport lag metrics help you evaluate potential data loss scenarios, enabling proactive disaster recovery planning and improved operational reliability.

Protection Mode

Data Guard associations provide two protection modes:

- Maximum Availability – Enables zero data loss failover and uses synchronous redo transport to at least one standby database.
- Maximum Performance – Enables near-zero data loss failover with no performance impact by using asynchronous transport to all standby targets.

Switchover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray.
 - Green (Healthy): All standby databases have passed all switchover readiness checks.
 - Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is switchover-ready. Address the failed checks.
 - Red (Critical): None of the standby databases are switchover-ready. Address failed checks to avoid extended downtime.
 - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database: status indicators – Green, Red, Gray.
 - Green (Healthy): The standby database has passed all switchover readiness checks.
 - Red (Critical): This standby database is not switchover-ready. Address failed checks to avoid extended downtime during switchover attempts.
 - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database and Snapshot Standby Database: status indicators – Gray
 - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

Failover Readiness

- Primary Database: status indicators – Green, Yellow, Red, Gray.
 - Green (Healthy): All standby databases have passed all failover readiness checks.

- Yellow (Warning): Applicable only if the primary database has multiple standby databases. A subset of standby databases is failover-ready. Address the failed checks.
- Red (Critical): None of the standby databases are failover-ready. Address failed checks to avoid extended downtime.
- Gray (Unknown): Shown when the health status cannot be determined at that point.
- Standby Database and Snapshot Standby: status indicators – Green, Red, Gray.
 - Green (Healthy): The standby database has passed all failover readiness checks.
 - Red (Critical): This standby database is not failover-ready. Address failed checks to avoid extended downtime during failover.
 - Gray (Unknown): Shown when the health status cannot be determined at that point.
- Disabled Standby Database: status indicators – Gray
 - Gray (Unknown): Disabled standby databases always appear as Unknown because their health status cannot be determined.

Note

When the health status cannot be determined for any database in a Data Guard Group at a given point in time, the status appears as null in the SDK/CLI, and Terraform, and as UNKNOWN (gray icon) in the Console.

Data Loss Exposure

- Primary Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.
- Standby Database: Data loss exposure is defined as the redo transport lag between the primary and standby databases. "Last Computed" if displayed represents the last time the metric could be retrieved and calculated.

Refresh Data Guard health status

Run an explicit refresh to get the latest status.

Using the Console to Manage an Oracle Data Guard Group

Learn how to enable a Data Guard Group between databases, change the role of a database in a Data Guard Group using either a switchover or a failover operation, and reinstate a failed database.

When you enable Data Guard, a separate Data Guard Group is created between the primary and the standby databases.

- [To Enable Data Guard on an Exadata Cloud Infrastructure System](#)
Learn to set up Data Guard Group between databases.
- [To view Data Guard Group details of databases in a Cloud VM Cluster](#)
To view the role of each database in a Data Guard Group in an Cloud VM Cluster, follow this procedure.
- [To enable automatic backups on a standby database](#)
Learn to enable automatic backups on a standby database.

- [To perform a database switchover](#)
You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.
- [To edit the Oracle Data Guard Group details](#)
- [To perform a database failover](#)
You can initiate a failover operation on a standby database that is a member of the Data Guard Group.
- [To reinstate a database](#)
After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary.
- [To terminate a Data Guard Group on an Exadata Cloud Infrastructure instance](#)
On an Exadata Cloud Infrastructure instance, you remove a Data Guard Group by terminating all the standby database.

To Enable Data Guard on an Exadata Cloud Infrastructure System

Learn to set up Data Guard Group between databases.

Note

At the time of this release, cross-service Oracle Data Guard between Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Exascale Infrastructure can be configured only with the Oracle AI Database 26ai release.

Note

- When you enable Data Guard, replication of data happens only over the client network.
- When you configure a Data Guard Group, the primary and standby databases must be on the same major release version while the standby database can be on a higher minor version.

Note

A parallel operation on the Standby, if it fails, should be retried after a 5-minute interval.

Note

You can create a Data Guard when the database is encrypted using OCI Virtual Vault.

As part of the latest release we are introducing an enhanced user experience and new APIs to improve performance and provide additional Data Guard capabilities including support for multiple standby databases via cloud automation.

- With the new API, your new Data Guard configuration will be created as a Data Guard Group resource.
 - If you have an existing Data Guard setup, you can continue to use current capabilities with no impact. However, if you wish to create multiple standby databases, you will need to migrate to the new API model, which can be done at any time.
 - If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you will need to update your applications to use the new API to take advantage of these new capabilities
Oracle currently supports both the existing Data Guard Association API and the new Data Guard Group API and the associated user interfaces.
1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
 2. Choose your **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard..
 3. Navigate to the cloud VM cluster that contains a database you want to assume the primary role:
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
 4. On the VM cluster details page, in the **Databases** section, click the name of the database you want to make primary.
 5. On the Database Details page, under **Resources**, click **Data Guard Associations**.
 6. Click **Add standby**.
 7. On the Add standby page, configure your Data Guard Group.
 - To explicitly run a Data Guard precheck, click **Run Precheck**. When you click Run Precheck, the following actions are performed:
 - The system closes the Add Standby page.
 - A message appears on the Primary Database Details page indicating that the Data Guard precheck is in progress, and the **Add Standby** button is disabled.
- Click the **View Details** link to navigate to the Work Requests page for the status:
- Review the precheck status: **Succeeded** or **Failed**.
 - If the precheck has failed, then resolve the reported errors and retry the Data Guard setup.

Note

While the precheck is running, the primary database and the VM Cluster remain in the **AVAILABLE** state.

- In the **Select peer VM Cluster** section, provide the following information for the standby database to obtain a list of available Exadata systems in which to locate the standby database:

- **Peer region:** Select a region where you want to locate the standby database. The region where the primary database is located is selected, by default. You can choose to locate the standby database in a different region. The hint text associated with this field tells you in which region the primary database is located.
 - **Availability domain:** Select an availability domain for the standby database. The hint text associated with this field tells you in which availability domain the primary database is located.
 - **Select a service:** Exadata Database Service on Dedicated Infrastructure or Exadata Database Service on Exascale Infrastructure.
 - **Select Exadata infrastructure:** Applicable only if you have selected Exadata Database Service on Dedicated Infrastructure from the Select a service drop-down list.
 - **Data Guard peer resource type:** Select **VM Cluster**.
 - Select a cloud VM cluster from the drop-down list.
 - **Choose the Data Guard experience:**
 - **Use the new Data Guard Group Resource** With this option, your new Data Guard configuration will be created as a Data Guard Group resource. This option with new APIs supports adding multiple standby databases and provides other enhancements. If you currently have automation that manages Data Guard operations using the existing Data Guard Association API, you can update your applications to use the new API to take advantage of these new capabilities.
 - **Use the existing Data Guard Association Resource** Choose this option if your automation for managing Data Guard operations relies on the existing Data Guard Association API. However, you will not be able to add multiple standby databases and will not get the enhancements provided by the new API.
 - **Data Guard Group details:**
 - **Data Guard Type:** Select Active Data Guard or Data Guard. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Far Sync, Global Data Services, and Application Continuity. Note that Active Data Guard requires an Oracle Active Data Guard license. For more information on Active Data Guard, see [Active Data Guard](#). For a complete overview of both Data Guard types, see [Introduction to Oracle Data Guard](#)
 - **Protection mode:** The protection mode can be **Maximum Performance** or **Maximum Availability**. See [Oracle Data Guard Protection Modes](#) for information on these options.
 - **Transport type:** The redo transport type used for this Data Guard Group. See [Redo Transport Services](#) for information on these options.
- Protection Mode and Transport Type: Rules for Standby Database Creation**
- * **Creating the first standby:** You cannot modify the Protection Mode or Transport Type for the first standby database during its creation. It is possible to modify it later.
 - * The default settings are:
 - * **Protection Mode:** Max Performance
 - * **Transport Type:** Async
 - * **Creating the second to Nth standby:** You cannot modify the Protection Mode or Transport Type for any subsequent standby databases.

- * The Protection Mode is inherited from the first standby.
- * The default Transport Type is set to Async.
- In the **Choose Database Home** section, choose one of the following:
 - **Select an existing Database Home:** If you use this option, select a home from the Database Home display name drop-down list.
 - **Create a new Database Home:** If you choose this option, enter a name for the new Database Home in the **Database Home display name** field. Click **Change Database Image** to select a database software image for the new Database Home. In the **Select a Database Software Image** panel, do the following:
 - a. Select the compartment containing the database software image you want to use to create the new Database Home.
 - b. Select the region containing the database software image you want to use to create the new Database Home. Region filter defaults to the currently connected region and lists all the software images created in that region. When you choose a different region, the software image list is refreshed to display the software images created in the selected region.
 - c. Select the Oracle Database software version that the new Database Home will use, then choose an image from the list of available images for your selected software version.
 - d. Click **Select**.

Note

- * Oracle recommends applying the same list of patches to the Database Homes of the primary and standby databases.
- * If you are using the new Data Guard Group resource, you must first create the database home before adding the standby database.

- In the **Configure standby database:** section, provide standby database details.

Note

You cannot modify the `db_unique_name` and `SID` prefix after creating the database.

- **Database unique name:** Optionally, specify a value for the `DB_UNIQUE_NAME` database parameter. This value must be unique across the primary and standby cloud VM clusters. The unique name must meet the requirements:
 - * Maximum of 30 characters
 - * Contain only alphanumeric or underscore (`_`) characters
 - * Begin with an alphabetic character
 - * Unique across the VM cluster. Recommended to be unique across the tenancy.

If not specified, the system automatically generates a unique name value, as follows:

```
<db_name>_<3_chars_unique_string>_<region-name>
```

- **Database password:** Enter the database administrator password of the primary database. Use this same database administrator password for the standby database.

Note

The administrator password and the TDE wallet password must be identical. If the passwords are not identical, then follow the instructions in [Changing the Database Passwords](#) to ensure that they are.

- **TDE wallet password:** Enter the TDE wallet password.

8. Click **Show advanced Options** to specify advanced options for the standby database:

- **Management:**

Oracle SID prefix: The Oracle Database instance number is automatically added to the SID prefix to create the `INSTANCE_NAME` database parameter. The `INSTANCE_NAME` parameter is also known as the SID. If not provided, then the SID prefix defaults to the first 12 characters of the `db_unique_name`.

Note

Entering an SID prefix is only available for Oracle 12.1 databases and above.

The SID prefix must meet the requirements:

- Maximum of 12 characters
- Contain only alphanumeric characters
- Begin with an alphabetic character
- Unique in the VM cluster and across primary and standby databases
- In the **Tags** tab, you can add tags to the database. 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 if you should apply tags, skip this option (you can apply tags later) or ask your administrator.

9. Click **Add standby**. When you create the association, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

A work request is issued to configure the Data Guard association. The progress of the request and the stages of provisioning can be viewed on the **Work Requests** page of the respective **Standby** database.

When the association is created, the details for a database and its peer display their respective roles as **Primary** or **Standby**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.
- [Changing the Database Passwords](#)
To change the SYS password, or to change the TDE wallet password, use this procedure.

To view Data Guard Group details of databases in a Cloud VM Cluster

To view the role of each database in a Data Guard Group in an Cloud VM Cluster, follow this procedure.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment.
3. Navigate to the cloud VM cluster that contains the databases you wish to view their roles in Data Guard associations.
4. In the **Databases** section under **Resources**, the role of each database in this VM Cluster is indicated in the **Data Guard role** column.
 - The role of each database in this VM Cluster is indicated in the **Data Guard role** column.
 - The service on which each database is running is indicated in the **Service name** column.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To enable automatic backups on a standby database

Learn to enable automatic backups on a standby database.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your Compartment that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable automatic database.
3. Navigate to the cloud VM cluster that contains the primary database.
 - *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. On the VM cluster details page, in the **Databases** section, click the name of the primary database.
5. On the Database Details page, under **Resources**, click **Data Guard Group**.
6. Click the name of the standby database for which you want to enable automatic backups. The system displays a banner if automatic backups are not enabled for this database.
7. Click **Enable automatic backups** on the banner.
8. On the resulting Configure Automatic Backups window, enter the following details:
 - **Enable automatic backup**: Check the check box to enable or disable automatic incremental backups for this database.

Note

- If your database is in a security zone compartment, you must enable automatic backups.
- If you are enabling automatic backups, you can select to configure Recovery Service or Object Storage as the Backup destination. However, if the backup was already configured on the primary database, then the standby must use the same backup destination.

- If **Recovery Service** is selected as the **Backup destination**, you can configure the following options:

Note

To use Autonomous Recovery Service as a backup destination, the Oracle Database version must be 19.18 or later.

- **Protection policy:** You can select from one of the preset protection policies or a custom policy. The system automatically deletes your backups at the end of your chosen protection policy recovery window.
- **Real-time data protection:** Real-time protection is the continuous transfer of redo changes from a protected database to Recovery Service. This reduces data loss and provides a recovery point objective (RPO) near 0. This is an extra cost option.
- **Deletion options after database termination:** You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
 - * **Retain backups according to the retention period:** When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
 - * **Retain backups for 72 hours, then delete:** When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
- **Scheduled day for initial backup:** Select a day of the week for the initial backup to begin.
- **Scheduled time for initial backup (UTC):** Select a time for the initial backup to begin. The initial backup could start at any time or within the chosen two-hour scheduling window.
- **Scheduled time for daily backup (UTC):** Select a time for the daily backup to begin. The daily backup could start at any time or within the chosen two-hour scheduling window.
- **Take the first backup immediately:** A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must not typically take a backup after an instance failure or

other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

- If **Object Storage** is selected as the **Backup destination**, you can configure the following options:
 - **Backup retention period:** If you select to enable automatic backups, you can select a policy with one of the preset retention periods. The system automatically deletes your incremental backups at the end of your chosen retention period. You can change the backup retention period after provisioning.
 - **Deletion options after database termination:** You can use the following options to retain managed database backups after the database is terminated. These options can also help restore the database from backups in case of accidental or malicious damage to the database.
 - * **Retain backups according to the retention period:** When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be removed at the end of the specified retention period.
 - * **Retain backups for 72 hours, then delete:** When a database is terminated, the automatic database backups associated with the terminated database and all of its resources will be retained for 72 hours and then deleted. The backups are retained for 72 hours to safeguard against accidental deletion by the user.
 - **Scheduled day for full backup:** Select a day of the week for the initial and future full backups to begin.
 - **Scheduled time for full backup (UTC):** Select a time for the full backup to begin. The full backup could start at any time or within the chosen two-hour scheduling window.
 - **Scheduled time for incremental backup (UTC):** Select a time for the incremental backup to begin. The incremental backup could start at any time or within the chosen two-hour scheduling window.
 - **Take the first backup immediately:** A full backup is an operating system backup of all data files and the control file that constitute an Oracle Database. A full backup must also include the parameter files associated with the database. You can take a database backup when the database is shut down or while the database is open. You must not typically take a backup after an instance failure or other unusual circumstances. If you select to defer the initial backup, your database may not be recoverable in the event of a database failure.

9. Click **Save Changes**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To perform a database switchover

You can initiate a switchover operation on a standby database that is a member of the Data Guard Group.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the Data Guard association:

Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.

4. Under **Resources**, click **Data Guard Group**.
5. Select the standby database in the Data Guard Group on which you want to perform a switchover. Click the Actions icon (three dots), and then click **Switchover**.
6. In the **Switchover database** dialog box, enter the database admin password, and then click **Switchover**.

This database should now assume the role of the standby, and the standby should assume the role of the primary in the Data Guard Group.

Note

You can now:

- Perform Data Guard actions (switchover, failover, and reinstate) while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Perform Data Guard setup concurrently on different databases within the same Oracle home.
- Perform Data Guard actions (switchover, failover, and reinstate) concurrently on different databases within the same Oracle home.
- Perform Data Guard setup while simultaneously updating VM Cluster tags.
- Create or delete a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To edit the Oracle Data Guard Group details

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Service instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the Data Guard association:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Data Guard Group**.
A list of databases that are members of the Data Guard Group is displayed with the **Data Guard type** you have chosen for each Data Guard Group member.
5. To edit Data Guard Group details, click the Actions icon (three dots), and then click **Edit**.
6. In the **Edit Data Guard Group** panel, configure the Data Guard Group:

- **Data Guard Type:** Select Active Data Guard or Data Guard. Active Data Guard provides additional features including: Real-Time Query and DML Offload, Automatic Block Repair, Standby Block Change Tracking, Global Data Services, and Application Continuity. Note that Active Data Guard requires an Oracle Active Data Guard license. For more information on Active Data Guard, see [Active Data Guard](#). For a complete overview of both Data Guard types, see [Introduction to Oracle Data Guard](#)
 - **Protection mode:** The protection mode can be **Maximum Performance** or **Maximum Availability**. See *Oracle Data Guard Protection Modes* for information on these options.
 - **Transport type:** The redo transport type used for this Oracle Data Guard Group.
 - **Database admin password:** Enter the ADMIN password for the database.
7. Click **Save**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)
- [Oracle Data Guard Protection Modes](#)

To perform a database failover

You can initiate a failover operation on a standby database that is a member of the Data Guard Group.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the Data Guard association:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Data Guard Group**.
5. Select the standby database in the Data Guard Group on which you want to perform a failover. Click the Actions icon (three dots), and then click **Failover**.
6. In the Failover database dialog box, enter the database admin password, and then click **Failover**.

Note

You can initiate a failover even if the primary database is in a healthy state; however, exercise caution when performing a failover.

This database should now assume the role of the primary, and the old primary's role should display as **Disabled Standby**.

Note

You can now:

- Perform Data Guard actions (switchover, failover, and reinstate) while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Perform Data Guard setup concurrently on different databases within the same Oracle home.
- Perform Data Guard actions (switchover, failover, and reinstate) concurrently on different databases within the same Oracle home.
- Perform Data Guard setup while simultaneously updating VM Cluster tags.
- Create or delete a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To reinstate a database

After you fail over a primary database to its standby, the standby assumes the primary role and the old primary is identified as a disabled standby. After you correct the cause of failure, you can reinstate the failed database as a functioning standby for the current primary.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the Data Guard association: *Cloud VM clusters (new resource model)*: Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. Under **Resources**, click **Data Guard Group**.
5. For the Data Guard Group on which you want to reinstate this database, click the Actions icon (three dots), and then click **Reinstate**.
6. In the Reinstate database dialog box, enter the database admin password, and then click **Reinstate**.

This database should now be reinstated as the standby in the Data Guard Group.

Note

You can now:

- Perform Data Guard actions (switchover, failover, and reinstate) while a Data Guard setup is running on another database within the same Oracle home, and vice versa.
- Perform Data Guard setup concurrently on different databases within the same Oracle home.
- Perform Data Guard actions (switchover, failover, and reinstate) concurrently on different databases within the same Oracle home.
- Perform Data Guard setup while simultaneously updating VM Cluster tags.
- Create or delete a PDB while concurrently performing Data Guard actions (switchover, failover, and reinstate) within the same Oracle home, and vice versa.

To terminate a Data Guard Group on an Exadata Cloud Infrastructure instance

On an Exadata Cloud Infrastructure instance, you remove a Data Guard Group by terminating all the standby database.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose the **Compartment** that contains the Exadata Cloud Infrastructure instance with the database for which you want to enable Oracle Data Guard.
3. Navigate to the cloud VM cluster that contains the standby database:
Cloud VM clusters (new resource model): Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**. In the list of VM clusters, find the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. For the standby database you want to terminate, click the Actions icon (three dots), and then click **Terminate**.
5. In the **Terminate database** dialog box, enter the name of the database, and then click **OK**.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

Using the API to manage Data Guard associations

Use these API operations to manage Data Guard associations on an Exadata Cloud Infrastructure instance:

Note

In February 2026, the Data Guard Association model and its associated APIs will be replaced by the new Data Guard Group model and APIs. Beginning February 2026, all new Data Guard configurations provisioned from the Oracle Cloud Infrastructure (OCI) Console will automatically use the Data Guard Group model.

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

- [CreateDataGuardAssociation](#)
- [ListDataGuardAssociations](#)
- [GetDataGuardAssociation](#)
- [UpdateDataGuardAssociation](#)
- [SwitchoverDataGuardAssociation](#)
- [FailoverDataGuardAssociation](#)
- [ReinstateDataGuardAssociation](#)
- [DeleteDatabase](#) - To terminate an Exadata Cloud Infrastructure instance Data Guard association, you delete the standby database.

For the complete list of APIs for the Database service, see [Database Service API](#).

Using the API to manage Data Guard Group

Use these API operations to manage a Data Guard Group on an Exadata Cloud Infrastructure instance:

Note

In February 2026, the Data Guard Association model and its associated APIs will be replaced by the new Data Guard Group model and APIs. Beginning February 2026, all new Data Guard configurations provisioned from the Oracle Cloud Infrastructure (OCI) Console will automatically use the Data Guard Group model.

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

Operation	REST API Endpoint	Comment
Create / Add Standby Database	CreateDatabase	<ul style="list-style-type: none"> • The same API is being used for creating a first standby and add more standby databases. • It uses the existing create database API with source as "DATAGUARD"
Update Data Guard Group Configuration	UpdateDataGuard	It takes either standby or primary database OCID to update the configuration.
Data Guard Action - Switchover	SwitchOverDataGuard	Switchover should be triggered on respective standby that to become primary.
Data Guard Action - Failover	FailoverDataGuard	Failover should be triggered on respective standby that to become primary.

Operation	REST API Endpoint	Comment
Data Guard Action - Reinstate	ReinstateDataGuard	Reinstate should be triggered on respective standby to be reinstated.
Delete Standby	DeleteDatabase	<ul style="list-style-type: none"> Delete standby remains same as existing.- DeleteDatabase Call goes on respective standby to be deleted.
Migrate Data Guard Association to multiple standby	MigrateDataGuardAssociationToMultiDataGuards	<ul style="list-style-type: none"> Migrate the existing data guard association to Data Guard Group model. New standby can be added only after migration is completed.

For the complete list of APIs for the Database service, see [Database Service API](#).

Configure Oracle Database Features for Exadata Cloud Infrastructure

This topic describes how to configure Oracle Multitenant, tablespace encryption, and Huge Pages for use with your Exadata Cloud Infrastructure instance.

- [Using Oracle Multitenant on an Exadata Cloud Infrastructure Instance](#)
- [Managing Tablespace Encryption](#)
- [Managing Huge Pages](#)

Using Oracle Multitenant on an Exadata Cloud Infrastructure Instance

When you create an Exadata Cloud Infrastructure Instance that uses Oracle Database 12c or later, an Oracle Multitenant environment is created.

The multitenant architecture enables an Oracle database to function as a multitenant container database (CDB) that includes zero, one, or many pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and non-schema objects that appears to an Oracle Net Services client as a non-CDB. All Oracle databases using versions earlier than Oracle Database 12c are non-CDBs.

To use Oracle Transparent Data Encryption (TDE) in a pluggable database (PDB), you must create and activate a master encryption key for the PDB.

In a multitenant environment, each PDB has its own master encryption key which is stored in a single keystore used by all containers.

You must export and import the master encryption key for any encrypted PDBs you plug into your Exadata Cloud Infrastructure Instance CDB.

If your source PDB is encrypted, you must export the master encryption key and then import it.

You can export and import all of the TDE master encryption keys that belong to the PDB by exporting and importing the TDE master encryption keys from within a PDB. Export and import of TDE master encryption keys support the PDB unplug and plug operations. During a PDB

unplug and plug, all of the TDE master encryption keys that belong to a PDB, as well as the metadata, are involved.

See "Exporting and Importing TDE Master Encryption Keys for a PDB" in *Oracle Database Advanced Security Guide* for Release [19](#), [18](#), [12.2](#) or [12.1](#).

See "ADMINISTER KEY MANAGEMENT" in *Oracle Database SQL Language Reference* for Release [19](#), [18](#), [12.2](#) or [12.1](#).

- [To determine if you need to create and activate an encryption key for the PDB](#)
- [To create and activate the master encryption key in a PDB](#)
- [To export and import a master encryption key](#)

To determine if you need to create and activate an encryption key for the PDB

1. Invoke SQL*Plus and log in to the database as the SYS user with SYSDBA privileges.
2. Set the container to the PDB:

```
SQL> ALTER SESSION SET CONTAINER = pdb;
```

3. Query V\$ENCRYPTION_WALLET as follows:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM v$encryption_wallet;
```

If the STATUS column contains a value of OPEN_NO_MASTER_KEY, you need to create and activate the master encryption key.

To create and activate the master encryption key in a PDB

1. Set the container to the PDB:

```
SQL> ALTER SESSION SET CONTAINER = pdb;
```

2. Create and activate a master encryption key in the PDB by executing the following command:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEY USING TAG 'tag' FORCE KEYSTORE  
IDENTIFIED BY keystore-password WITH BACKUP USING 'backup_identifier';
```

In the previous command:

- `keystore-password` is the keystore password. By default, the keystore password is set to the value of the administration password that is specified when the database is created.
- The optional `USING TAG 'tag'` clause can be used to associate a tag with the new master encryption key.
- The `WITH BACKUP` clause, and the optional `USING 'backup_identifier'` clause, can be used to create a backup of the keystore before the new master encryption key is created.

See also ADMINISTER KEY MANAGEMENT in *Oracle Database SQL Language Reference* for Release [19](#), [18](#) or [12.2](#).

Note

To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.

If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:

- a. Close the keystore.
- b. Open the password-based keystore.
- c. Create and activate a master encryption key in the PDB by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.
- d. Update the auto-login keystore by using `ADMINISTER KEY MANAGEMENT` with the `CREATE AUTO_LOGIN KEYSTORE FROM KEYSTORE` option.

3. Query `V$ENCRYPTION_WALLET` again to verify that the `STATUS` column is set to `OPEN`:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM v$encryption_wallet;
```

4. Query `V$INSTANCE` and take note of the value in the `HOST_NAME` column, which identifies the database server that contains the newly updated keystore files:

```
SQL> SELECT host_name FROM v$instance;
```

5. Copy the updated keystore files to all of the other database servers.

To distribute the updated keystore, you must perform the following actions on each database server that does not contain the updated keystore files:

- a. Connect to the root container and query `V$ENCRYPTION_WALLET`. Take note of the keystore location contained in the `WRL_PARAMETER` column:

```
SQL> SELECT wrl_parameter, status FROM v$encryption_wallet;
```

- b. Copy the updated keystore files.

You must copy all of the updated keystore files from a database server that is already updated. Use the keystore location observed in the `WRL_PARAMETER` column of `V$ENCRYPTION_WALLET`.

Open the updated keystore:

```
SQL> ADMINISTER KEY MANAGEMENT SET KEYSTORE open FORCE KEYSTORE IDENTIFIED BY keystore-password CONTAINER=all;
```

Note

To enable key management operations while the keystore is in use, Oracle Database 12c Release 2, and later, includes the `FORCE KEYSTORE` option to the `ADMINISTER KEY MANAGEMENT` command. This option is also available for Oracle Database 12c Release 1 with the October 2017, or later, bundle patch.

If your Oracle Database 12c Release 1 database does not have the October 2017, or later, bundle patch installed, you can perform the following alternative steps:

- a. Close the keystore before copying the updated keystore files.
- b. Copy the updated keystore files.
- c. Open the updated keystore by using `ADMINISTER KEY MANAGEMENT` without the `FORCE KEYSTORE` option.

6. Query `GV$ENCRYPTION_WALLET` to verify that the `STATUS` column is set to `OPEN` across all of the database instances:

```
SQL> SELECT wrl_parameter, status, wallet_type FROM gv$encryption_wallet;
```

To export and import a master encryption key

1. Export the master encryption key.
 - a. Invoke SQL*Plus and log in to the PDB.
 - b. Execute the following command:

```
SQL> ADMINISTER KEY MANAGEMENT EXPORT ENCRYPTION KEYS WITH SECRET  
"secret" TO 'filename' IDENTIFIED BY keystore-password;
```

2. Import the master encryption key.
 - a. Invoke SQL*Plus and log in to the PDB.
 - b. Execute the following command:

```
SQL> ADMINISTER KEY MANAGEMENT IMPORT ENCRYPTION KEYS WITH SECRET  
"secret" FROM 'filename' IDENTIFIED BY keystore-password;
```

Managing Tablespace Encryption

By default, all new tablespaces that you create in an Exadata database are encrypted.

However, the tablespaces that are initially created when the database is created may not be encrypted by default.

- For databases that use Oracle Database 12c Release 2 or later, only the `USERS` tablespaces initially created when the database was created are encrypted. No other tablespaces are encrypted including the non-`USERS` tablespaces in:
 - The root container (`CDB$ROOT`).
 - The seed pluggable database (`PDB$SEED`).
 - The first PDB, which is created when the database is created.

- For databases that use Oracle Database 12c Release 1 or Oracle Database 11g, none of the tablespaces initially created when the database was created are encrypted.

For further information about the implementation of tablespace encryption in Exadata, along with how it impacts various deployment scenarios, see [Oracle Database Tablespace Encryption Behavior in Oracle Cloud](#).

Creating Encrypted Tablespaces

User-created tablespaces are encrypted by default.

By default, any new tablespaces created by using the SQL `CREATE TABLESPACE` command are encrypted with the AES128 encryption algorithm. You do not need to include the `USING 'encrypt_algorithm'` clause to use the default encryption.

You can specify another supported algorithm by including the `USING 'encrypt_algorithm'` clause in the `CREATE TABLESPACE` command. Supported algorithms are AES256, AES192, AES128, and 3DES168.

Managing Tablespace Encryption

You can manage the software keystore (known as an Oracle wallet in Oracle Database 11g), the master encryption key, and control whether encryption is enabled by default.

Managing the Master Encryption Key

Tablespace encryption uses a two-tiered, key-based architecture to transparently encrypt (and decrypt) tablespaces. The master encryption key is stored in an external security module (software keystore). This master encryption key is used to encrypt the tablespace encryption key, which in turn is used to encrypt and decrypt data in the tablespace.

When a database is created on an Exadata Cloud Service instance, a local software keystore is created. The keystore is local to the compute nodes and is protected by the administration password specified during the database creation process. The auto-login software keystore is automatically opened when the database is started.

You can change (rotate) the master encryption key by using the `ADMINISTER KEY MANAGEMENT` SQL statement. For example:

```
SQL> ADMINISTER KEY MANAGEMENT SET ENCRYPTION KEY USING TAG 'tag'  
IDENTIFIED BY password WITH BACKUP USING 'backup';
```

keystore altered.

See "Managing the TDE Master Encryption Key" in *Oracle Database Advanced Security Guide* for Release [19](#), [18](#), [12.2](#) or [12.1](#) or "Setting and Resetting the Master Encryption Key" in *Oracle Database Advanced Security Administrator's Guide* for Release [11.2](#).

Controlling Default Tablespace Encryption

The `ENCRYPT_NEW_TABLESPACES` initialization parameter controls the default encryption of new tablespaces. In Exadata databases, this parameter is set to `CLOUD_ONLY` by default.

Values of this parameter are as follows.

Value	Description
ALWAYS	During creation, tablespaces are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the <code>ENCRYPTION</code> clause.
CLOUD_ONLY	Tablespaces created in an Exadata database are transparently encrypted with the AES128 algorithm unless a different algorithm is specified in the <code>ENCRYPTION</code> clause. For non-cloud databases, tablespaces are only encrypted if the <code>ENCRYPTION</code> clause is specified. <code>ENCRYPTION</code> is the default value.
DDL	During creation, tablespaces are not transparently encrypted by default, and are only encrypted if the <code>ENCRYPTION</code> clause is specified.

Note

With Oracle Database 12c Release 2 (12.2), or later, you can no longer create an unencrypted tablespace in an Exadata database. An error message is returned if you set `ENCRYPT_NEW_TABLESPACES` to `DDL` and issue a `CREATE TABLESPACE` command without specifying an `ENCRYPTION` clause.

Managing Huge Pages

Huge Pages provide considerable performance benefits for Oracle Database on systems with large amounts of memory. Oracle Database on an Exadata Cloud Infrastructure instance provides configuration settings that make use of Huge Pages by default; however, you can make manual adjustments to optimize the configuration of Huge Pages.

Huge Pages is a feature integrated into the Linux kernel 2.6. Enabling Huge Pages makes it possible for the operating system to support large memory pages. Using Huge Pages can improve system performance by reducing the amount of system CPU and memory resources required to manage Linux page tables, which store the mapping between virtual and physical memory addresses. For Oracle Databases, using Huge Pages can drastically reduce the number of page table entries associated with the System Global Area (SGA).

On Exadata Cloud Infrastructure instances, a standard page is 4 KB, while a Huge Page is 2 MB by default. Therefore, an Oracle Database on an Exadata VM cluster with a 50 GB SGA requires 13,107,200 standard pages to house the SGA, compared with only 25,600 Huge Pages. The result is much smaller page tables, which require less memory to store and fewer CPU resources to access and manage.

Adjusting the Configuration of Huge Pages

The configuration of Huge Pages for Oracle Database is a two-step process:

- At the operating system level, the overall amount of memory allocated to Huge Pages is controlled by the `vm.nr_hugepages` entry in the `/etc/sysctl.conf` file. This setting is made on each compute node in the environment and it is strongly recommended that the

setting is consistent across all of the compute nodes. To alter the Huge Page allocation, you can execute the following command on each compute node as the root user:

```
# sysctl -w vm.nr_hugepages=value
```

where `value` is the number of Huge Pages that you want to allocate.

On Exadata Cloud Infrastructure instances, each Huge Page is 2 MB by default. Therefore, to allocate 50 GB of memory to Huge Pages you can execute the following command:

```
# sysctl -w vm.nr_hugepages=25600
```

- At the Oracle Database level, the use of Huge Pages is controlled by the `USE_LARGE_PAGES` instance parameter setting. This setting applies to each database instance in a clustered database. Oracle strongly recommends a consistent setting across all of the database instances associated with a database. The following options are available:
 - `TRUE` — specifies that the database instance can use Huge Pages if they are available. For all versions of Oracle Database after 11.2.0.3, Oracle allocates as much of the SGA as it can, using Huge Pages. When the Huge Page allocation is exhausted, standard memory pages are used.
 - `FALSE` — specifies that the database instance does not use Huge Pages. This setting is generally not recommended if Huge Pages are available.
 - `ONLY` — specifies that the database instance must use Huge Pages. With this setting, the database instance fails to start if the entire SGA cannot be accommodated in Huge Pages.

If you make any adjustments at either the operating system or Oracle Database level, ensure that the overall configuration works.

For more information, see the *Oracle Database Administrator's Reference for Linux and UNIX-Based Operating Systems* for Release [19](#), [18](#), [12.1](#), or [11.2](#) for a general overview of Huge Pages and more information about configuring Huge Pages. Also, see `USE_LARGE_PAGES` in the *Oracle Database Reference* for Release [12.2](#), [12.1](#), or [11.2](#).

Managing Exadata Cloud Infrastructure I/O Resource Management (IORM)

This topic explains the I/O Resource Management (IORM) feature and how to enable it, modify the IORM settings, and disable it by using the Console or the API.

- [About IORM](#)
The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model.
- [Using the Console to Manage IORM](#)
- [Using the API to manage the I/O resources of an Exadata cloud VM cluster](#)

About IORM

The I/O Resource Management (IORM) feature allows you to manage how multiple databases share the I/O resources of an Oracle Exadata cloud VM cluster for systems using the new resource model.

On an Exadata VM cluster, all databases share dedicated storage servers which include flash storage. By default, the databases are given equal priority with respect to these resources. The Exadata storage management software uses a first come, first served approach for query processing. If a database executes a major query that overloads I/O resources, overall system performance can be slowed down.

IORM allows you to assign priorities to your databases to ensure critical queries are processed first when workloads exceed their resource allocations. You assign priorities by creating directives that specify the number of shares for each database. The number of shares corresponds to a percentage of resources given to that database when I/O resources are stressed.

Directives work together with an overall optimization objective you set for managing the resources. The following objectives are available:

- **Auto** - (Default objective). IORM determines the optimization objective and continuously and dynamically determines the optimal settings, based on the workloads observed, and resource plans enabled.
- **Balanced** - For critical OLTP and DSS workloads. This setting balances low disk latency and high throughput. This setting limits disk utilization of large I/Os to a lesser extent than low latency to achieve a balance between good latency and good throughput.
- **High throughput** - For critical DSS workloads that require high throughput.
- **Low latency** - For critical OLTP workloads. This setting provides the lowest possible latency by significantly limiting disk utilization.

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)
- [Managing I/O Resources](#)

Using the Console to Manage IORM

- [To enable IORM on your Exadata cloud VM cluster](#)
This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.
- [To modify the IORM configuration on your cloud VM cluster](#)
This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.

To enable IORM on your Exadata cloud VM cluster

This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.

Enabling IORM includes specifying an optimization objective and configuring your resource plan directives.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under the **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of VM clusters, find the VM cluster for which you want to enable IORM, and click its highlighted name. The cluster's details are displayed, showing the IORM status as "Disabled."

5. Click **Enable IORM**.

It might take a minute for the Enable I/O Resource Management dialog to retrieve the VM cluster information.

6. Select the objective to apply to the resource plan:

- **Auto** - (Default objective) Dynamically changes the objective based on the resource plan and observed workloads.
- **Balanced** - Weighs high throughput and low latency evenly.
- **High throughput** - Provides the best throughput for DSS workloads.
- **Low latency** - Provides the best latency for critical OLTP workloads.

7. Configure the resource plan default directive by setting the number of shares. This number of shares is assigned to each database not associated with a specific directive.

8. In the Resource Plan Directives section, add a directive for each database you want to assign a greater or lesser number of shares than the default directive.

To add a directive, click **+ Additional Directive**, then specify the database and the number of shares for that database.

9. When you are done adding directives, click **Enable**.

While the IORM configuration settings are being applied, the VM cluster details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your VM cluster. After a successful update, the IORM status shows as "Enabled."

Related Topics

- [The New Exadata Cloud Infrastructure Resource Model](#)

To modify the IORM configuration on your cloud VM cluster

This topic only applies to Exadata Cloud Infrastructure systems using the new resource model.

Use this procedure to change your IORM settings or to disable IORM.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Choose your **Compartment**.
3. Click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
4. In the list of VM clusters, find the VM cluster for which you want to update IORM, and click its highlighted name. The cluster's details are displayed, showing the IORM status as "Enabled."
5. Click **Update IORM**.
6. In the Update I/O Resource Management dialog, take one of the following actions:
 - Change your settings - Specify a new objective and adjust your directives, as applicable, and then click **Update**.
 - Disable IORM: Click **Disable IORM**.

Disabling IORM removes all your resource plan directives and restores the Auto (default) objective for I/O resource management.

While the new IORM configuration settings are being applied, the system details page shows the IORM status as "Updating." The update might take several minutes to complete but should have no impact on your ability to perform normal operations on your VM cluster. After a successful update, the IORM status shows as "Enabled" or "Disabled," depending on the action you took.

Using the API to manage the I/O resources of an Exadata cloud VM cluster

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

Use these API operations to manage the I/O resources of an Exadata cloud VM cluster. (see [The New Exadata Cloud Service Resource Model](#) for more information on this resource type).

- [ListCloudVmClusters](#)
- [GetCloudVmCluster](#)
- [GetCloudVmClusterIormConfig](#)
- [UpdateCloudVmClusterIormConfig](#)

Managing Encryption Keys on External Devices

Learn how to store and manage database encryption keys.

There are two options to store and manage database encryption keys for your databases on Oracle Exadata Database Service on Dedicated Infrastructure:

1. In an auto-login wallet file stored in an Oracle Advanced Cluster File System (Oracle ACFS) accessible by the customer VM operating system.
 2. Oracle Key Vault.
- [Customer-Managed Keys in Oracle Exadata Database Service on Dedicated Infrastructure](#)
Customer-managed keys for Oracle Exadata Database Service on Dedicated Infrastructure is a feature that enables you to migrate the Oracle Database TDE Master Encryption Key for an Oracle Database from the password-protected wallet file stored on the Oracle Exadata Database Service on Dedicated Infrastructure equipment to an OKV server that you control.
 - [About Oracle Key Vault](#)
Oracle Key Vault is a full-stack, security-hardened software appliance built to centralize the management of keys and security objects within the enterprise.
 - [Overview of Key Store](#)
Integrate your on-premises Oracle Key Vault (OKV) with customer-managed database cloud services to secure your critical data on ExaDB-D.
 - [Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Dedicated Infrastructure](#)
Review the identity access management (IAM) policy for managing OKV on Oracle Exadata Database Service on Dedicated Infrastructure systems.
 - [Tagging Resources](#)
You can apply tags to your resources to help you organize them according to your business needs.

- [Moving Resources to a Different Compartment](#)
You can move OKV Vault, Secret, and Keystore resources from one compartment to another.
- [Setting Up Your Oracle Exadata Database Service on Dedicated Infrastructure to Work With Oracle Key Vault](#)
Review the prerequisites to setup your Oracle Exadata Database Service on Dedicated Infrastructure to work with Oracle Key Vault.
- [Managing Your Key Store](#)
Learn how to manage your key store.
- [Administer Transparent Data Encryption \(TDE\) Keys](#)
Use this procedure to change the key management configuration.
- [How to Manually Clone a Pluggable Database \(PDB\) from a Remote Container Database \(CDB\) When Data is Encrypted with Master Encryption Key \(MEK\) in Oracle Key Vault \(OKV\)](#)
The dbaascli tool lets you clone PDBs when the source CDB and target CDB are the same (local clone) or if they are different (remote clone). However, you cannot clone a remote PDB if the data is encrypted with a MEK in OKV.
- [How to Upgrade Oracle Key Vault \(OKV\) Home in Oracle Exadata Database Service on Dedicated Infrastructure](#)

Customer-Managed Keys in Oracle Exadata Database Service on Dedicated Infrastructure

Customer-managed keys for Oracle Exadata Database Service on Dedicated Infrastructure is a feature that enables you to migrate the Oracle Database TDE Master Encryption Key for an Oracle Database from the password-protected wallet file stored on the Oracle Exadata Database Service on Dedicated Infrastructure equipment to an OKV server that you control.

The Oracle Key Vault (OKV) provides fault-tolerant, highly available and scalable key and secrets management for your encrypted ExaDB-D databases. Use customer-managed keys when you need security governance, regulatory compliance, and homogenous encryption of data, while centrally managing, storing, and monitoring the life cycle of the keys you use to protect your data.

You can:

- Switch from Oracle-managed keys to customer-managed keys for databases, whether they are enabled with Data Guard or not.
- Rotate your keys to maintain security compliance.
- Rotating the PDB key is also supported. Rotate CDB and PDB key operations are allowed only if the database is customer-managed.

Requirements

- To enable the management of customer-managed encryption keys, you must create a policy in the tenancy that allows a particular dynamic group to do so. For more information, see *Setting Up Your Oracle Exadata Database Service on Dedicated Infrastructure to Work With Oracle Key Vault*.
- Pluggable databases must be configured in United Mode. For more information about United Mode, see [Managing Keystores and TDE Master Encryption Keys in United Mode](#). Isolated Mode is not supported. For more information about Isolated Mode, see [Managing Keystores and TDE Master Encryption Keys in Isolated Mode](#)

- If an Exadata Database Service was configured for Oracle Key Vault using the procedures published at [Migration of File based TDE to OKV for Exadata Database Service on Cloud at Customer Gen2 \(Doc ID 2823650.1\)](#), then you should open a My Oracle Support (MOS) Service Request to have Oracle cloud operations update the control plane configuration to reflect the Oracle Key Vault information for the specific Exadata Database service

About Oracle Key Vault

Oracle Key Vault is a full-stack, security-hardened software appliance built to centralize the management of keys and security objects within the enterprise.

Note

The Oracle Key Vault is a customer-provisioned and managed system and it is not part of Oracle Cloud Infrastructure managed services.

Related Topics

- [Oracle Key Vault](#)

Overview of Key Store

Integrate your on-premises Oracle Key Vault (OKV) with customer-managed database cloud services to secure your critical data on ExaDB-D.

Oracle Key Vault integration enables you to take complete control of your encryption keys and store them securely on an external, centralized key management device.

OKV is optimized for Oracle wallets, Java keystores, and Oracle Advanced Security Transparent Data Encryption (TDE) master keys. Oracle Key Vault supports the OASIS KMIP standard. The full-stack, security-hardened software appliance uses Oracle Linux and Oracle Database technology for security, availability, and scalability, and can be deployed on your choice of compatible hardware.

OKV also provides a REST interface for clients to auto-enroll endpoints and setup wallets and keys. Oracle Exadata Database Service on Dedicated Infrastructure temporarily stores the OKV REST user administrator password required to connect to the OKV appliance in a password-protected wallet file so that the software running in the customer VM can connect to the OKV server. Following the migration of the TDE keys to OKV, the cloud automation software will remove the password from the wallet file. Ensure that you create a secret with Oracle's Vault Service, which will store the password required for autonomous databases to connect to OKV for key management.

For more information, see *Oracle Key Vault*.

Related Topics

- [Oracle Key Vault](#)

Required IAM Policy for Managing OKV on Oracle Exadata Database Service on Dedicated Infrastructure

Review the identity access management (IAM) policy for managing OKV on Oracle Exadata Database Service on Dedicated Infrastructure systems.

A **policy** is an IAM document that specifies who has what type of access to your resources. It is used in different ways: to mean an individual statement written in the policy language; to mean a collection of statements in a single, named "policy" document (which has an Oracle Cloud ID (OCID) assigned to it), and to mean the overall body of policies your organization uses to control access to resources.

A **compartment** is a collection of related resources that can be accessed only by certain groups that have been given permission by an administrator in your organization.

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

If you're new to policies, then see *Getting Started with Policies* and *Common Policies*. If you want to dig deeper into writing policies for databases, then see *Details for the Database Service*.

Related Topics

- [Getting Started with Policies](#)
- [Common Policies](#)
- [Details for the Database Service](#)

Tagging Resources

You can apply tags to your resources to help you organize them according to your business needs.

You can apply tags at the time you create a resource, or you can update the resource later with the desired tags. For general information about applying tags, see *Resource Tags*.

Related Topics

- [Resource Tags](#)

Moving Resources to a Different Compartment

You can move OKV Vault, Secret, and Keystore resources from one compartment to another.

After you move an OCI resource to a new compartment, inherent policies apply immediately and affect access to the resource. Moving an OKV Vault resource doesn't affect access to any OKV Vault Keys or OKV Vault Secrets that the OKV Vault contains. You can move an OCI Key or OCI Secret from one compartment to another independently of moving the OKV Vault it's associated with. For more information, see *Managing Compartments*.

Related Topics

- [Managing Compartments](#)

Setting Up Your Oracle Exadata Database Service on Dedicated Infrastructure to Work With Oracle Key Vault

Review the prerequisites to setup your Oracle Exadata Database Service on Dedicated Infrastructure to work with Oracle Key Vault.

Prerequisites

1. Ensure that OKV is set up and the network is accessible from the Exadata client network. Open ports 443, 5695, and 5696 for egress on the client network for the OKV client software and Oracle database instance to access the OKV server.
2. Ensure that the REST interface is enabled from the OKV user interface.
3. Create "OKV REST Administrator" user.
You can use any qualified username of your choice, for example, "okv_rest_user". For ADB-C@C, ExaDB-C@C, and ExaDB-D use the same or different REST users. Those databases can be key-managed in the same or different on-premises OKV clusters. ExaDB-C@C and ExaDB-D need REST user with create endpoint privilege. ADB-C@C needs REST user with create endpoint and create endpoint group privileges.
4. Gather OKV administrator credentials and IP address, which is required to connect to OKV.

For more information, see: "Network Port Requirements", "Managing Oracle Key Vault Users", and "Managing Administrative Roles and User Privileges".

- [Step 1: Create a Vault in OKV Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password](#)
Your Exadata Cloud Infrastructure communicates with OKV over REST each time an Oracle Database is provisioned to register the Oracle Database and request a wallet on OKV. Therefore, Exadata infrastructure needs access to the REST admin credentials to register with the OKV server.
- [Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault](#)
To grant your Key Store resources permission to access Secret in OCI Vault, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Secret you created in the OCI Vaults and Secrets.
- [Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store](#)
To grant your Exadata Cloud Infrastructure resources permission to access Key Store, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Key Store you created.
- [Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service](#)
To grant the Exadata Database service permission to use the secret in OCI Vault to log in to the OKV REST interface, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your OCI Vaults and Secrets.
- [Step 5: Create Key Store](#)
Follow these steps to create a Key Store to connect to an on-premises encryption key appliance such as Oracle Key Vault (OKV).

Related Topics

- [Network Port Requirements](#)
- [Managing Oracle Key Vault Users](#)
- [Managing Administrative Roles and User Privileges](#)

Step 1: Create a Vault in OKV Vault Service and Add a Secret to the Vault to Store OKV REST Administrator Password

Your Exadata Cloud Infrastructure communicates with OKV over REST each time an Oracle Database is provisioned to register the Oracle Database and request a wallet on OKV. Therefore, Exadata infrastructure needs access to the REST admin credentials to register with the OKV server.

These credentials are stored securely in the Oracle Vault Service in OCI as a Secret and accessed by your Exadata Cloud Infrastructure infrastructure only when needed. When needed, the credentials are stored in a password-protected wallet file.

To store the OKV administrator password in the OKV Vault service, create a vault by following the instructions outlined in *Managing Vaults* and create a Secret in that vault by following the instructions outlined in *Managing Secrets*.

Related Topics

- [Managing Vaults](#)
- [Managing Secrets](#)

Step 2: Create a Dynamic Group and a Policy Statement for Key Store to Access Secret in OCI Vault

To grant your Key Store resources permission to access Secret in OCI Vault, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Secret you created in the OCI Vaults and Secrets.

When defining the dynamic group, you identify your Key Store resources by specifying the OCID of the compartment containing your Key Store.

1. Copy the OCID of the compartment containing your Key Store resource. You can find this OCID on the Compartment Details page of the compartment.
2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

```
ALL {resource.compartment.id = '<compartment-ocid>'}
```

where *<compartment-ocid>* is the OCID of the compartment containing your Key Store resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your vaults and secrets. Then, add a policy statement of this format:

```
allow dynamic-group <dynamic-group> to use secret-family in compartment <vaults-and-secrets-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<vaults-and-secrets-compartment>* is the name of the compartment in which you created your vaults and secrets.

Related Topics

- [To create a dynamic group](#)

Step 3: Create a Dynamic Group and a Policy Statement for Exadata Infrastructure to Key Store

To grant your Exadata Cloud Infrastructure resources permission to access Key Store, you create an IAM dynamic group that identifies these resources and then create an IAM policy that grants this dynamic group access to the Key Store you created.

When defining the dynamic group, you identify your Exadata Cloud Infrastructure resources by specifying the OCID of the compartment containing your Exadata infrastructure.

1. Copy the OCID of the compartment containing your Exadata Cloud Infrastructure resource. You can find this OCID on the Compartment Details page of the compartment.
2. Create a dynamic group by following the instructions in "To create a dynamic group" in Oracle Cloud Infrastructure Documentation. When following these instructions, enter a matching rule of this format:

```
ALL {resource.compartment.id = '<compartment-ocid>'}
```

where *<compartment-ocid>* is the OCID of the compartment containing your Exadata infrastructure resource.

3. After creating the dynamic group, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your Key Store. Then, add a policy statement of this format:

```
Allow dynamic-group <dynamic-group> to use keystores in compartment <key-store-compartment>
```

where *<dynamic-group>* is the name of the dynamic group you created and *<key-store-compartment>* is the name of the compartment in which you created your Key Store.

Step 4: Create a Policy Statement for Database Service to Use Secret from OCI Vault Service

To grant the Exadata Database service permission to use the secret in OCI Vault to log in to the OKV REST interface, navigate to (or create) an IAM policy in a compartment higher up in your compartment hierarchy than the compartment containing your OCI Vaults and Secrets.

Then, add a policy statement of this format:

```
allow service database to read secret-family in compartment <vaults-and-secrets-compartment>
```

where *<vaults-and-secrets-compartment>* is the name of the compartment in which you created your OCI Vaults and Secrets.

Once the OKV Vault is set up and the IAM configuration is in place, you are now ready to deploy your Oracle Key Vault 'Key Store' in OCI and associate it with your Exadata VM Cluster.

Step 5: Create Key Store

Follow these steps to create a Key Store to connect to an on-premises encryption key appliance such as Oracle Key Vault (OKV).

1. Open the navigation menu. Under **Oracle Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Click **Key Stores**.

Key Stores page displays the list of name of key stores, the number of databases associated with each database, and the date on which each key store was created.

4. Click **Create Key Store**.
5. On the resulting Create Key Store dialog, enter the following general information:
 - **Name your key store:** A user-friendly description or other information that helps you easily identify the Key Store resource. Avoid entering confidential information.
 - **Oracle Key Vault connection settings**
 - **Connection IP addresses:** Enter at least one OKV cluster node IP address; multiple comma-separated IP addresses (of the same OKV cluster) are possible, for example, 193.10.20.1, 193.10.20.2.
 - **Administrator username:** Enter the user name of the `okv_rest_user`.
 - **Administrator Password Secret:** The administrator password is stored with the secret management service within OCI. Select the OCI Vault in your tenancy that contains `okv_rest_user` password stored as an OCI Secret.
 - **Tags:** Optionally, you can apply tags. If you have permission to create a resource, you also have permission to apply free-form tags to that resource. To apply a defined tag, you must have permission to use the tag namespace. For more information about tagging, see *Resource Tags*. If you are not sure if you should apply tags, skip this option (you can apply tags later) or ask your administrator. Avoid entering confidential information.
6. Click **Create Key Store**.
7. Ensure that you use the same `okv_rest_user` user credentials, while provisioning the database.

For more information, see: *Managing Vaults*, *Managing Keys*, and *Managing Secrets*.

Related Topics

- [Managing Vaults](#)
- [Managing Keys](#)
- [Managing Secrets](#)

Managing Your Key Store

Learn how to manage your key store.

- [View Key Store Details](#)
Follow these steps to view Key Store details that include Oracle Key Vault (OKV) connection details and the list of associated databases.

- [Edit Key Store Details](#)
You can edit a Key Store only if it is not associated with any CDBs.
- [Move a Key Store to Another Compartment](#)
Follow these steps to move a Key Store on an Oracle Exadata Database Service on Dedicated Infrastructure system from one compartment to another compartment.
- [Delete a Key Store](#)
You can delete a Key Store only if it is not associated with any CDBs.
- [View Key Store Associated Container Database Details](#)
Follow these steps to view details of the container database associated with a Key Store.
- [Using the API to Manage Key Store](#)
Learn how to use the API to manage key store.

View Key Store Details

Follow these steps to view Key Store details that include Oracle Key Vault (OKV) connection details and the list of associated databases.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Click **Key Stores**.
Key Stores page displays the list of names of key stores, the number of databases associated with each database, and the date on which each key store was created.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. Click the link in the **Administrator Password Secret** field to view secret details.
The **Associated Databases** section displays the list of CDBs associated with this Key Store.

Edit Key Store Details

You can edit a Key Store only if it is not associated with any CDBs.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Click **Key Stores**.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. On the Key Store Details page, click **Edit**.
6. On the Edit Key Store page, make changes as needed, and then click **Save Changes**.

Move a Key Store to Another Compartment

Follow these steps to move a Key Store on an Oracle Exadata Database Service on Dedicated Infrastructure system from one compartment to another compartment.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.

2. Select your Compartment.
3. Click **Key Stores**.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. On the Key Store Details page, click **Move Resource**.
6. On the Move Resource to a Different Compartment page, select the new compartment.
7. Click **Move Resource**.

Delete a Key Store

You can delete a Key Store only if it is not associated with any CDBs.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Click **Key Stores**.
4. Click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. On the Key Store Details page, click **Delete**.
6. On the Delete Key Store dialog, click **Delete**.

View Key Store Associated Container Database Details

Follow these steps to view details of the container database associated with a Key Store.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Click **Key Stores**.
4. On the resulting Key Stores page, click the name of the Key Store or click the Actions icon (three dots), and then click **View Details**.
5. Click the name of the associated database or click the Actions icon (three dots), and then click **View Details**.

Using the API to Manage Key Store

Learn how to use the API to manage key store.

For information about using the API and signing requests, see *REST APIs and Security Credentials*. For information about SDKs, see *Software Development Kits and Command Line Interface*.

Operation	REST API Endpoint
Create OKV Key Store	CreateKeyStore
View OKV Key Store	GetKeyStore
Update OKV Key Store	UpdateKeyStore
Delete OKV Key Store	DeleteKeyStore

Operation	REST API Endpoint
Change Key store compartment	ChangeKeyStoreCompartment
Choose between customer-managed and Oracle-managed encryption	CreateDatabase
Get the Key Store (OKV or Oracle-managed) and OKV wallet name	GetDatabase
Change Key store type	changeKeyStoreType
Rotate OKV and Oracle-managed key	RotateVaultKey

Related Topics

- [REST APIs](#)
- [Security Credentials](#)
- [Software Development Kits and Command Line Interface](#)

Administer Transparent Data Encryption (TDE) Keys

Use this procedure to change the key management configuration.

After provisioning a database on an ExaDB-D system, you can change the key management and perform operations such as rotating the TDE keys.

Note

- You can change the key management from Oracle Wallet to other available options.
- When you change the key management to OKV, the database will experience a shutdown abort operation followed by a restart. Plan to perform the migration in a planned maintenance window.
- You must rotate TDE keys only via OCI interfaces (Console, API).
- You cannot rotate an encryption key:
 - when a database restore is in progress in a given Oracle Home.
 - when a database patching or database home patching is in progress.

1. Open the navigation menu. Under **Oracle AI Database**, click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Select your Compartment.
3. Navigate to the VM Cluster that contains the database for which you want to change encryption management or rotate a key.
 - a. Exadata Database Service on Dedicated Infrastructure, click **Exadata VM Clusters**.
 - b. In the list of VM clusters, locate the VM cluster you want to access and click its highlighted name to view the details page for the cluster.
4. In the Databases section, click the name of the database for which you want to change encryption management or rotate a key to display its details page.
5. Go to the Encryption section on the database details page.

- **To change key management:**
If you have configured **Oracle Wallet** as the key management, then the system displays the **Change** option to change the key management to **OCI Vault** and **Oracle Key Vault**,
 - a. Click **Change**.
 - b. On the resulting Change key management page, select the **Key management**.
 - **OCI Vault:**
 - i. Select your Vault compartment that you are using, and then select your Vault that is available in the compartment.
 - ii. Select the Key compartment that you are using, and then select your Key from the dropdown list.
 - iii. Optionally, choose the key version and enter the OCID of the key version. By default, the latest key version is used.
 - iv. Click **Save changes**.
 - **Oracle Key Vault:**
You must have a valid encryption key in the Oracle Key Vault service and provide the information in the subsequent steps. For more information, see *Key and Secret Management Concepts*.
 - i. Choose a region.
 - ii. Choose a compartment.
You can change the compartment by clicking the **Change Compartment** link.
 - iii. Click **Save changes**.
- **To rotate an encryption key:**
If you have configured **OCI Vault** or **Oracle Key Vault** as the key management, then the system displays the **Rotate** option to rotate the encryption key.
 - a. Click **Rotate** to display a confirmation dialog.
 - b. Click **Rotate**.

Note

- Migration of TDE keys to Oracle Key Vault (OKV) requires 10 minutes of downtime. During the migration, the database state will be UPDATING and connections may fail due to multiple database restarts to enable OKV. Applications can resume operation after the migration completes and when the database returns to its original ACTIVE state.
- The OKV keystore password will be set to the TDE wallet password.

Caution

After changing key management, deleting the key from the OKV will cause the database to become unavailable.

On the database details page for this database, the Encryption section displays the encryption key name and the encryption key OCID.

Related Topics

- [Key and Secret Management](#)

How to Manually Clone a Pluggable Database (PDB) from a Remote Container Database (CDB) When Data is Encrypted with Master Encryption Key (MEK) in Oracle Key Vault (OKV)

The dbascli tool lets you clone PDBs when the source CDB and target CDB are the same (local clone) or if they are different (remote clone). However, you cannot clone a remote PDB if the data is encrypted with a MEK in OKV.

Note

To decrypt / encrypt the data during a remote clone, the container database must have access to MEK. The MEK must be made available to the target CDB when it is stored in the OKV server.

- [Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server](#)
- [Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server](#)

Related Topics

- [dbascli pdb localClone](#)
To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the dbascli pdb localClone command.
- [dbascli pdb remoteClone](#)
To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the dbascli pdb remoteClone command.

Source CDB and Target CDB are Encrypted with MEK in the Same OKV Server

1. Get the OKV object ID of the source PDB.
 - a. Get the latest encryption key of the source PDB using SQL*Plus.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Mon Jun 12 23:13:12 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> set heading off;
SQL> alter session set container=<SOURCE_PDB>;
```

```
Session altered.
```

```
SQL> select key_id,keystore_type,activation_time from v$encryption_keys
order by activation_time;
```

```
0648E5D8D5559B4F0EBFB8AA5EE730401A
SOFTWARE KEYSTORE
25-MAR-23 12.01.41.075932 AM +00:00
```

```
06AFF5B6E27A954F6EBFFC77296B27C9EC
SOFTWARE KEYSTORE
25-MAR-23 11.42.51.336955 AM +00:00
```

```
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver oracle]$
```

- b. Get the OKV object ID (uuid) of the newest MEK obtained from the step above. Enter the OKV Endpoint password when prompted and hit the **Enter** key on your keyboard.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil list | grep
06AFF5B6E27A954F6EBFFC77296B27C9EC
E5344379-8B16-4FE9-BF35-F8ECB057571A Symmetric Key TDE Master
Encryption Key: MKID 06AFF5B6E27A954F6EBFFC77296B27C9EC
[oracle@testserver oracle]$
```

2. Install OKV REST wallet in the source database.

- a. Create the `okv_rest_cli` directory if it does not exist.

```
[root@testserver newdbl]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
```

- b. Download and extract `okvrestclipackage.zip`.

Select ALL if prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<source_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current                                 Dload  Upload  Total  Spent  Left
Speed
100 3784k  100 3784k    0     0 19.0M    0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$
```

- c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<source_okv_server_ip1>
user=<source_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<source_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$
```

- d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

- e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the source OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet --wallet-user <source_okv_rest_user>
Password:
{
  "result" : "Success"
}
```

```
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
total 8
-rw----- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```

3. Create a new OKV wallet to store only the PDB MEK obtained in step #1.

a. Get the OKV wallet name from the source PDB in the format

EXA_DB_NAME_DBID_PDB_NAME_WL.

For example, the wallet name would be EXA_NEWDB1_37508325141_PDB_NAME_WL.

```
[root@testserver newdbl1]# su oracle
[oracle@testserver newdbl1]$ source ~/<source_db_name>.env
[oracle@testserver newdbl1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;

NAME          DB_UNIQUE_NAME          DBID
-----
NEWDB1        newdbl_uniq              3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
-----
newdbl1

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdbl1]$
```

b. Create a new wallet using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet create --
wallet <SOURCE_PDB_OKV_WALLET> --description "Wallet to clone
<source_pdb_name> pdb from <source_db_name>" --unique FALSE
{
```

```

    "result" : "Success",
    "value" : {
      "status" : "PENDING",
      "locatorID" : "BA5FBFE1-DB41-4425-8EE4-D58541A1E41A"
    }
  }
}
[root@testserver oracle]#

```

c. Check the status until it is ACTIVE.

```

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING"
  }
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "ACTIVE",
    "wallet" : "<SOURCE_PDB_OKV_WALLET>"
  }
}
[root@testserver oracle]#

```

4. Add **Read and Modify**, and **Manage Wallet** permissions from the source database OKV Endpoints to the OKV wallet created in step #3.

a. Get the Endpoint names from the source database. One per VM.

Usually, the structure is in the format, EXA_DB_UNIQUE_NAME_DBID_SID_EP.

For example, the Endpoint name of node 1 would be

EXA_NEWDB1_UNIQ_3750832514_NEWDB11_EP.

```

[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<source_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

```

```

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

```

Copyright (c) 1982, 2022, Oracle. All rights reserved.

```

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

```

```

SQL> select name,db_unique_name,dbid from v$database;

NAME          DB_UNIQUE_NAME          DBID
-----
NEWDB1        newdbl_uniq              3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
-----
newdbl1

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdbl1]$

```

- b. Add **Read and Modify**, and **Manage Wallet** permissions using the OKV REST command-line interface.

```

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP1> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP2> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]#

```

5. Store MEK from the source PDB obtained in step #1 into the OKV wallet created in step #3.
- Add MEK (uuid obtained in step #1.b) using the OKV REST command-line interface. Enter the source OKV Endpoint password when prompted.

```

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv managed-object wallet add-member
--uuid E5344379-8B16-4FE9-BF35-F8ECB057571A --wallet
<SOURCE_PDB_OKV_WALLET>
Password:

```

```
{
  "result" : "Success"
}
[root@testserver oracle]#
```

6. Install OKV REST wallet in the target database.

- a. Create the `okv_rest_cli` directory if it does not exist.**

```
[root@testserver newdbl]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli
```

- b. Download and extract `okvrestclipackage.zip`.**

Select ALL when prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<target_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current                                 Dload  Upload  Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$
```

- c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.**

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<target_okv_server_ip1>
user=<target_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<target_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
```

```

java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$

```

d. Create the `client_wallet` directory.

```

[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$

```

e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the target OKV REST password when prompted.

```

[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet --wallet-user <target_okv_rest_user>
Password:
{
  "result" : "Success"
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
total 8
-rw----- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$

```

7. Add **Read Only** and **Manage Wallet** permissions from the target database OKV Endpoints to the source PDB OKV wallet created in step #3.

a. Get the Endpoint names from the target database. One per VM.

Usually, the structure is in the format, `EXA_DB_UNIQUE_NAME_DBID_SID_EP`.

For example, the Endpoint name of node 1 would be

```
EXA_NEWDB1_UNIQ_3750832514_NEWDB11_EP.
```

```

[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<target_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

```

```

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

```

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:

```
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0
```

```
SQL> select name,db_unique_name,dbid from v$databases;
```

```
NAME          DB_UNIQUE_NAME          DBID
-----
NEWDB1        newdb1_uniq             3750832514
```

```
SQL> select value from v$parameter where name='instance_name';
```

```
VALUE
-----
newdb11
```

```
SQL> exit
```

```
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdb1]$
```

- b. Add **Read Only** and **Manage Wallet** permissions using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP1> --access
RO_MW
{
  "result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <TARGET_OKV_EP2> --access
RO_MW
{
  "result" : "Success"
}
[root@testserver oracle]#
```

8. Clone the PDB.

- Run dbaascli to clone the PDB.

Enter the source DB SYS user password when prompted.

```
[root@testserver oracle]# dbaascli pdb remoteClone --pdbName
<source_pdb_name> --dbName <target_db_name> --sourceDBConnectionString
<source_db_connection_string> --targetPDBName <target_pdb_name>
DBAAS CLI version 23.2.1.0.0
Executing command pdb remoteClone --pdbName <source_pdb_name> --dbName
<target_pdb_name> --sourceDBConnectionString scaqar06dvclu01-
```

```
scan1.us.oracle.com:1521/<source_db_unique_name>.us.oracle.com --
targetPDBName <target_pdb_name>
Job id: 197f30e9-209e-4ec5-9700-a13f7915f8b9
Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
dbaastools_2023-06-12_10-32-17-PM_188384.log
Enter REMOTE_DB_SYS_PASSWORD:

Enter REMOTE_DB_SYS_PASSWORD (reconfirmation):

Loading PILOT...
Session ID of the current execution is: 6848
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot_2023-06-12_10-32-35-PM_204184
-----
Running Plugin_initialization job
Enter REMOTE_DB_SYS_PASSWORD
*****
Completed Plugin_initialization job
-----
Running Validate_input_params job
Completed Validate_input_params job
-----
Running Perform_dbca_prechecks job
Completed Perform_dbca_prechecks job
-----
Running PDB_creation job
Completed PDB_creation job
-----
Running Load_pdb_details job
Completed Load_pdb_details job
-----
Running Configure_pdb_service job
Completed Configure_pdb_service job
-----
Running Configure_tnsnames_ora job
Completed Configure_tnsnames_ora job
-----
Running Set_pdb_admin_user_profile job
Completed Set_pdb_admin_user_profile job
-----
Running Lock_pdb_admin_user job
Completed Lock_pdb_admin_user job
-----
Running Register_ocids job
Skipping. Job is detected as not applicable.
-----
Running Prepare_blob_for_standby_in_primary job
Skipping. Job is detected as not applicable.
-----
Running Generate_dbssystem_details job
Completed Generate_dbssystem_details job
dbaascli execution completed
[root@testserver oracle]#
```

9. Delete the source PDB OKV wallet created in step #3 using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet delete --
wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success"
}
[root@testserver oracle]#
```

10. Delete the OKV REST wallet created in step #2.

- Delete the wallet files in the dbaas_acfs directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

11. Delete OKV REST wallet created in step #6.

- Delete the wallet files in the dbaas_acfs directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

Source CDB and Target CDB are Encrypted with MEK in a Different OKV Server

1. Get the OKV object ID of the source PDB.
 - a. Get the latest encryption key of the source PDB using SQL*Plus.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Mon Jun 12 23:13:12 2023
Version 19.19.0.0.0
```

```
Copyright (c) 1982, 2022, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0
```

```
SQL> set heading off;
SQL> alter session set container=<SOURCE_PDB>;
```

```
Session altered.
```

```
SQL> select key_id,keystore_type,activation_time from v$encryption_keys
order by activation_time;
```

```

0648E5D8D5559B4F0EBFB8AA5EE730401A
SOFTWARE KEYSTORE
25-MAR-23 12.01.41.075932 AM +00:00

06AFF5B6E27A954F6EBFFC77296B27C9EC
SOFTWARE KEYSTORE
25-MAR-23 11.42.51.336955 AM +00:00

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver oracle]$

```

- b. Get the OKV object ID (uuid) of the newest MEK obtained from the step above. Enter the OKV Endpoint password when prompted and hit the **Enter** key on your keyboard.

```

[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<source_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil list | grep
06AFF5B6E27A954F6EBFFC77296B27C9EC
E5344379-8B16-4FE9-BF35-F8ECB057571A Symmetric Key TDE Master
Encryption Key: MKID 06AFF5B6E27A954F6EBFFC77296B27C9EC
[oracle@testserver oracle]$

```

2. Install OKV REST wallet in the source database.

- a. Create the `okv_rest_cli` directory if it does not exist.

```

[root@testserver newdbl]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli

```

- b. Download and extract `okvrestclipackage.zip`.

Select **ALL** if prompted for replacement.

```

[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli
[oracle@scaqar06dv0101 okv_rest_cli]$ curl -O -k https://
<source_okv_server_ip1>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current                                 Dload  Upload  Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M    0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$

```

- c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<source_okv_server_ip1>
user=<source_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<source_db_name>/okv_home/conf/
okvclient.ora
```

```
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$
```

- d. Create the `client_wallet` directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$
```

- e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the source OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<source_db_name>/okv_rest_cli/
client_wallet --wallet-user <source_okv_rest_user>
Password:
{
  "result" : "Success"
```

```

}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet
total 8
-rw----- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$

```

3. Create a new OKV wallet to store only the PDB MEK obtained in step #1.

a. Get the OKV wallet name from the source PDB in the format

EXA_DB_NAME_DBID_PDB_NAME_WL.

```

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet create --
wallet <SOURCE_PDB_OKV_WALLET> --description "Wallet to clone
<source_pdb_name> pdb from <source_db_name>" --unique FALSE
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING",
    "locatorID" : "BA5FBFE1-DB41-4425-8EE4-D58541A1E41A"
  }
}
[root@testserver oracle]#

```

b. Check the status until it is ACTIVE.

```

[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "PENDING"
  }
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet check-status
--wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success",
  "value" : {
    "status" : "ACTIVE",
    "wallet" : "<SOURCE_PDB_OKV_WALLET>"
  }
}
[root@testserver oracle]#

```

4. Add **Read and Modify**, and **Manage Wallet** permissions from the source database OKV Endpoints to the OKV wallet created in step #3.

- a. Get the Endpoint names from the source database. One per VM.

Usually, the structure is in the format, EXA_DB_UNIQUE_NAME_DBID_SID_EP.

For example, the Endpoint name of node 1 would be

```
EXA_NEWDB1_UNIQ_3750832514_NEWDB11_EP.
```

```
[root@testserver newdbl]# su oracle
[oracle@testserver newdbl]$ source ~/<source_db_name>.env
[oracle@testserver newdbl]$ sqlplus / as sysdba
```

```
SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0
```

```
Copyright (c) 1982, 2022, Oracle. All rights reserved.
```

```
Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0
```

```
SQL> select name,db_unique_name,dbid from v$database;
```

NAME	DB_UNIQUE_NAME	DBID
NEWDB1	newdbl_uniq	3750832514

```
SQL> select value from v$parameter where name='instance_name';
```

```
VALUE
```

```
-----
newdbl1
```

```
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdbl]$
```

- b. Add **Read and Modify**, and **Manage Wallet** permissions using the OKV REST command-line interface.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP1> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
```

```
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet add-access --
wallet <SOURCE_PDB_OKV_WALLET> --endpoint <SOURCE_OKV_EP2> --access
RM_MW
{
  "result" : "Success"
}
[root@testserver oracle]#
```

5. Store MEK from the source PDB obtained in step #1 into the OKV wallet created in step #3.
 - Add MEK (uuid obtained in step #1.b) using the OKV REST command-line interface. Enter the source OKV Endpoint password when prompted.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv managed-object wallet add-member
--uuid E5344379-8B16-4FE9-BF35-F8ECB057571A --wallet
<SOURCE_PDB_OKV_WALLET>
Password:
{
  "result" : "Success"
}
[root@testserver oracle]#
```

6. Download OKV wallet created in step #3 from the OKV server to the local filesystem.
 - a. Create a new directory with permissions for oracle user.

This directory will store the wallet that will contain only the MEK of the source PDB.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ mkdir /home/oracle/<source_pdb_wallet_dir>
[oracle@testserver oracle]$
```

- b. Download the OKV wallet created in step #3 to the directory created in step #6.a using okvutil.

It will prompt twice for a password to encrypt the local wallet. Use the same password as the source Endpoint password. Also, enter the source Endpoint password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/nfsa.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil download -l /home/
oracle/<source_pdb_wallet_dir> -t wallet -g <SOURCE_PDB_OKV_WALLET>
Enter new wallet password (<enter> for auto-login):
Confirm new wallet password:
Enter Oracle Key Vault endpoint password:
Download succeeded
[oracle@testserver oracle]$
```

- c. Zip the wallet directory.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /home/oracle
```

```
[oracle@testserver oracle]$ zip -r <source_pdb_wallet_dir>.zip
<source_pdb_wallet_dir>
  adding: <source_pdb_wallet_dir>/ (stored 0%)
  adding: <source_pdb_wallet_dir>/ewallet.pl2 (stored 0%)
[oracle@testserver oracle]$
```

7. Delete the source PDB OKV wallet created in step #3.

```
[root@testserver oracle]# export JAVA_HOME=/usr/java/latest; export
OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<source_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/bin/okv manage-access wallet delete --
wallet <SOURCE_PDB_OKV_WALLET>
{
  "result" : "Success"
}
[root@testserver oracle]#
```

8. Delete the OKV REST wallet created in step #1.

- Delete the wallet files in the `dbaas_acfs` directory.

```
[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<source_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#
```

9. Copy the source PDB wallet downloaded to the focal filesystem in step #6 to the target Cluster VM.

10. Delete the source PDB wallet from the source local filesystem created in step #6.

- a. Delete the wallet directory.**

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ rm -rf /home/oracle/<source_pdb_wallet_dir>
[oracle@testserver oracle]$
```

- b. Delete the wallet zip file.**

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ rm -f /home/oracle/
<source_pdb_wallet_dir>.zip
[oracle@testserver oracle]$
```

11. Install OKV REST wallet in the target database.

- a. Create the `okv_rest_cli` directory if it does not exist.**

```
[root@testserver newdbl]# su oracle
[oracle@testserver oracle]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli
```

- b. Download and extract `okvrestclipackage.zip`.**

Select ALL when prompted for replacement.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /var/opt/oracle/dbaas_acfs/
```

```

<target_db_name>/okv_rest_cli
[oracle@testserver okv_rest_cli]$ curl -O -k https://
<target_okv_server_ip>:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current                                  Dload  Upload   Total   Spent    Left
Speed
100 3784k  100 3784k    0     0  19.0M      0 --:--:-- --:--:--
--:--:-- 19.1M
[oracle@testserver okv_rest_cli]$ unzip -q okvrestclipackage.zip
[oracle@testserver okv_rest_cli]$

```

- c. Modify the `okvrestcli.ini` and `okvrestcli_logging.properties` files as follows.

```

[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli.ini
[Default]
server=<target_okv_server_ip>
user=<target_okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet
log_property=/var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
conf/okvrestcli_logging.properties
okv_client_config=/u02/app/oracle/admin/<target_db_name>/okv_home/conf/
okvclient.ora

[oracle@testserver okv_rest_cli]$ vi /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
[oracle@testserver okv_rest_cli]$ cat /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/conf/okvrestcli_logging.properties
handlers=java.util.logging.FileHandler
java.util.logging.FileHandler.pattern=/var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/logs/okvrest.log
java.util.logging.FileHandler.limit=200000
java.util.logging.FileHandler.count=1
java.util.logging.FileHandler.formatter=com.oracle.okv.rest.log.OkvForma
tter
java.util.logging.ConsoleHandler.level=FINER
java.util.logging.ConsoleHandler.formatter=com.oracle.okv.rest.log.OkvFo
rmatter
[oracle@testserver okv_rest_cli]$

```

- d. Create the `client_wallet` directory.

```

[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ mkdir /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet
[oracle@testserver okv_rest_cli]$

```

- e. Create OKV REST wallet using the OKV REST command-line interface.

Enter the target OKV REST password when prompted.

```
[root@testserver oracle]# su oracle
[oracle@testserver okv_rest_cli]$ export JAVA_HOME=/usr/java/latest;
export OKV_RESTCLI_CONFIG=/var/opt/oracle/dbaas_acfs/<target_db_name>/
okv_rest_cli/conf/okvrestcli.ini; /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/bin/okv admin client-wallet add --client-
wallet /var/opt/oracle/dbaas_acfs/<target_db_name>/okv_rest_cli/
client_wallet --wallet-user <target_okv_rest_user>
/var/opt/oracle/dbaas_acfs/newdb1/okv_rest_cli/logs/okvrest.log.lck
Password:
{
  "result" : "Success"
}
[oracle@testserver okv_rest_cli]$ ls -ltr /var/opt/oracle/dbaas_acfs/
newdb1/okv_rest_cli/client_wallet
total 8
-rw----- 1 oracle oinstall    0 Jun 16 01:29 ewallet.p12.lck
-rw----- 1 oracle oinstall    0 Jun 16 01:29 cwallet.sso.lck
-rw----- 1 oracle oinstall  976 Jun 16 01:29 ewallet.p12
-rw----- 1 oracle oinstall 1021 Jun 16 01:29 cwallet.sso
[oracle@testserver okv_rest_cli]$
```

12. Upload the source PDB wallet created in step #6 and copied to the target Cluster VM in step #9.

- a. Unzip the source PDB wallet.

```
[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ cd /home/oracle/
[oracle@testserver ~]$ unzip <source_pdb_wallet_dir>.zip
Archive:  nfsa_1672104454_NFSPDB_wallet.zip
  creating: <source_pdb_wallet_dir>/
  extracting: <source_pdb_wallet_dir>/ewallet.p12
[oracle@testserver ~]$
```

- b. Get OKV wallet name from the target database in the format EXA_DB_NAME_DBID_WL.

For example, the wallet name would be EXA_NEWDB1_37508325141_WL.

```
[root@testserver newdb1]# su oracle
[oracle@testserver newdb1]$ source ~/<target_db_name>.env
[oracle@testserver newdb1]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Tue Jun 20 21:26:54 2023
Version 19.19.0.0.0

Copyright (c) 1982, 2022, Oracle. All rights reserved.

Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.19.0.0.0

SQL> select name,db_unique_name,dbid from v$database;

NAME          DB_UNIQUE_NAME          DBID
-----
-----
```

```

NEWDB1      newdbl_uniq      3750832514

SQL> select value from v$parameter where name='instance_name';

VALUE
-----
newdbl1

SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release
19.0.0.0.0 - Production
Version 19.19.0.0.0
[oracle@testserver newdbl]$

```

- c. Upload the source PDB wallet to the target OKV wallet using `okvutil`.

Enter the source PDB wallet password when prompted. Use the same password as the source Endpoint password.

Also, enter the target Endpoint password when prompted.

```

[root@testserver oracle]# su oracle
[oracle@testserver oracle]$ source ~/<target_db_name>.env
[oracle@testserver oracle]$ $OKV_HOME/bin/okvutil upload -t WALLETT -l /
home/oracle/<source_pdb_wallet_dir> -g <TARGET_OKV_WALLETT>
Enter source wallet password:
Enter Oracle Key Vault endpoint password:
WARNING: Object ORACLE.SECURITY.ID.ENCRIPTION. already exists; use -o
to overwrite
Upload succeeded
[oracle@testserver oracle]$

```

13. Clone the PDB.

- Run `dbaascli` to clone the PDB.

```

[root@testserver oracle]# dbaascli pdb remoteClone --pdbName
<source_pdb_name> --dbName <target_db_name> --sourceDBConnectionString
<source_db_connection_string> --targetPDBName <target_pdb_name>
DBAAS CLI version 23.2.1.0.0
Executing command pdb remoteClone --pdbName <source_pdb_name> --dbName
<target_db_name> --sourceDBConnectionString scaqar06dvclu01-
scan1.us.oracle.com:1521/<source_db_unique_name>.us.oracle.com --
targetPDBName <target_pdb_name>
Job id: 7d4f638a-1f3a-4219-a05a-0215588dcae8
Session log: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
dbaastools_2023-06-13_01-29-09-AM_179996.log
Enter REMOTE_DB_SYS_PASSWORD:

Enter REMOTE_DB_SYS_PASSWORD (reconfirmation):

Loading PILOT...
Session ID of the current execution is: 6857
Log file location: /var/opt/oracle/log/alyokv1/pdb/remoteClone/
pilot_2023-06-13_01-29-21-AM_196991
-----

```

```

Running Plugin_initialization job
Enter REMOTE_DB_SYS_PASSWORD
*****
Completed Plugin_initialization job
-----
Running Validate_input_params job
Completed Validate_input_params job
-----
Running Perform_dbca_prechecks job
Completed Perform_dbca_prechecks job
-----
Running PDB_creation job
Completed PDB_creation job
-----
Running Load_pdb_details job
Completed Load_pdb_details job
-----
Running Configure_pdb_service job
Completed Configure_pdb_service job
-----
Running Configure_tnsnames_ora job
Completed Configure_tnsnames_ora job
-----
Running Set_pdb_admin_user_profile job
Completed Set_pdb_admin_user_profile job
-----
Running Lock_pdb_admin_user job
Completed Lock_pdb_admin_user job
-----
Running Register_ocids job
Skipping. Job is detected as not applicable.
-----
Running Prepare_blob_for_standby_in_primary job
Skipping. Job is detected as not applicable.
-----
Running Generate_dbssystem_details job
Completed Generate_dbssystem_details job
dbaascli execution completed
[root@testserver oracle]#

```

14. Delete the OKV REST wallet created in step #1.

- Delete the wallet files in the `dbaas_acfs` directory.

```

[root@testserver oracle]# rm -f /var/opt/oracle/dbaas_acfs/
<target_db_name>/okv_rest_cli/client_wallet/*
[root@testserver oracle]#

```

How to Upgrade Oracle Key Vault (OKV) Home in Oracle Exadata Database Service on Dedicated Infrastructure

After the encryption type is migrated from Oracle Managed Keys to Customer Managed Keys (Oracle Key Vault), the OKV home in the DomUs remains with the same version used for the migration.

In case the OKV Server is upgraded the functionality would keep working because of backward compatibility. However, the customer might want to get the new features for the client tools. In that case, upgrade the OKV home and the PKCS#11 Library.

1. Validate current OKV Home version is minor to OKV Server version.
 - a. Get the OKV Home version by running `okvutil`. In this case the value is 21.6.0.0.0.

```
# su oracle
$ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil
okvutil version 21.6.0.0.0
Usage: okvutil <command> [-v <verbosity>] [<command args>]
  <command> := list | upload | download | sign | sign-verify |
  changepwd | diagnostics
Options:
  -v, --verbose <verbosity>
    Print extra information to standard out.
    Possible verbosity values are 0, 1 and 2 (more detailed information
    with higher verbosity level).
  For help on a particular command, use [okvutil <command> -h].
  You have new mail in /var/spool/mail/root
```

- b. Get the OKV Server version by logging in to the OKV Server console through the browser. In this case the Version is 21.7.0.0.0.

2. Install OKV REST wallet in the source database. This step must be done in only one node.
 - a. If not existing, create the `okv_rest_cli` directory.

```
# su oracle
$ mkdir /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli
```

- b. Download and extract `okvrestclipackage.zip`. If prompted for replacement, select ALL.

```
$ cd /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli
$ curl -O -k https://100.75.59.249:5695/okvrestclipackage.zip
  % Total    % Received % Xferd  Average Speed   Time    Time     Time
Current                                  Dload  Upload  Total   Spent    Left
Speed
100 3865k  100 3865k    0     0  5102k      0 --:--:-- --:--:--
--:--:-- 5106k
$ unzip -q okvrestclipackage.zip
```

- c. Modify `okvrestcli.ini` with the next info.

```
$ vi /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/
okvrestcli.ini

$ cat !$
cat /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini
#Provide absolute path for log_property, okv_client_config properties
[Default]
#log_property=./conf/okvrestcli_logging.properties
#server=[OKV IP Address]
#okv_client_config=./conf/okvclient.ora
#user=[OKV username]
```

```
#password=[user password]

#[Profile1]
#server=
#okv_client_config=
#user=

#[Profile2]
#server=
#okv_client_config=
#user=

server=<okv_server_ip>
user=<okv_rest_user>
client_wallet=/var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/
client_wallet
```

- d. Create the `client_wallet` directory.

```
$ mkdir /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/client_wallet
```

- e. Create OKV REST Wallet using OKV REST CLI. It will prompt for the source OKV REST password.

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin client-wallet add
--client-wallet /var/opt/oracle/dbaas_acfs/<dbname>/okv_rest_cli/
client_wallet --wallet-user <okv_rest_user>
Password:
{
  "result" : "Success"
}
```

3. Prepare the OKV Home directories in DomU 1.

- a. Rename OKV Home directory as the current OKV Home version.

```
$ mv /u02/app/oracle/admin/<dbname>/okv_home/okv /u02/app/oracle/admin/
<dbname>/okv_home/okv<current_okv_home_version>
```

- b. Create a new OKV Home directory as the OKV Server version.

```
$ mkdir /u02/app/oracle/admin/<dbname>/okv_home/okv<okv_server_version>
```

- c. Create a symlink of regular OKV Home name to the directory created in step 3.b.

```
$ ln -s /u02/app/oracle/admin/<dbname>/okv_home/
okv<okv_server_version> /u02/app/oracle/admin/<dbname>/okv_home/okv
```

4. Upgrade OKV Home in DomU 1.

- a. Get OKV Endpoint name via `okvutil`. It will prompt for OKV Endpoint Password (TDE Password) The entry is the one called "Template".

Note that the hostname should be the same as the current DomU Hostname. (Usually, this is named as EXA_<DBNAME>_<resourceID>_<CURRENT_DOMU_HOST_NAME>_EP. <resourceID> can be get by listing the DB via dbaascli system getDatabases).

```
$ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil list
Enter Oracle Key Vault endpoint password:
Unique ID                               Type           Identifier
DC690343-5694-4FC8-BFE4-6C7F1A550F67   Opaque Object TDE Wallet
Metadata
9E317DDB-0542-553B-A47D-FCC31AB6DD7C    Symmetric Key TDE Master
Encryption Key: MKID AaTAGyAWyk/fv7pnl8qx4s0AAAAAAAAAAAA
D9D840AF-A60E-5850-AA86-8C9F216F5501    Symmetric Key TDE Master
Encryption Key: MKID AUP0Tq+un08Mv1+onNhT4RUAAAAAAAAAAAA
364EFC2F-1909-4F34-BF1B-90D3D03DA7EB    Private Key Private Key
A9D0134F-C895-4F33-BF85-351B754E9FF9    Opaque Object TDE Wallet
Metadata
E1AC8D2F-90E9-4F88-BFEE-2883FCBB7271    Opaque Object TDE Wallet
Metadata
25B7DE14-3849-4F67-BFBE-1934BFE3559B    Opaque Object TDE Wallet
Metadata
4ED713ED-FE2B-4F35-BF7D-BCBEA8327A0B    Symmetric Key TDE Master
Encryption Key: MKID 06EA813441C26B4F53BFD58E55C4BE90F4
6162E200-EF0A-4F89-BF25-A8596B3AD7B0    Opaque Object Certificate
Request
85A55486-28E5-4FFB-BF1C-B93C4C0BAD74    Secret Data Oracle Secret Data:
ID HSM_PASSWORD
67E74D97-56F6-407A-A035-009D953F907A    Template Default template
for EXA_DB1902_7274B2A2-6F71-4516-B2BB-6D67CC3824FC_SCAQAE08DV0308_EP
E621EA72-5DD1-4F4F-BFD4-451E5B7DB8A9    Symmetric Key TDE Master
Encryption Key: MKID 0625BA455B03CD4F57BFA5D2290FD379A1
```

b. Re-enroll the Endpoint in DomU 1.

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin endpoint re-
enroll --endpoint <endpoint_name>
{
  "result" : "Success"
}
```

c. Provision Endpoint in DomU 1. It will prompt for OKV Endpoint password (TDE password).

```
$ export JAVA_HOME=/usr/java/latest; export OKV_RESTCLI_CONFIG=/var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/conf/okvrestcli.ini; /var/opt/
oracle/dbaas_acfs/<dbname>/okv_rest_cli/bin/okv admin endpoint
provision --endpoint <endpoint_name> --location /u02/app/oracle/admin/
<dbname>/okv_home/okv --auto-login FALSE
Enter Oracle Key Vault endpoint password:
{
  "result" : "Success"
}
```

5. Validate the OKV Home Upgrade was successfully done.

- a. Validate the OKV Endpoint lists the entries in OKV Wallet. It will prompt for OKV Endpoint password (TDE password).

```
$ /u02/app/oracle/admin/db1902/okv_home/okv/bin/okvutil list
Enter Oracle Key Vault endpoint password:
Unique ID                               Type                               Identifier
DC690343-5694-4FC8-BFE4-6C7F1A550F67   Opaque Object                     TDE Wallet
Metadata
9E317DDB-0542-553B-A47D-FCC31AB6DD7C     Symmetric Key                     TDE Master
Encryption Key: MKID AaTAGyAWyk/fv7pnl8qx4s0AAAAAAAAAAAA
D9D840AF-A60E-5850-AA86-8C9F216F5501     Symmetric Key                     TDE Master
Encryption Key: MKID AUP0Tq+un08Mv1+onNhT4RUAAAAAAAAAAAA
364EFC2F-1909-4F34-BF1B-90D3D03DA7EB     Private Key                       Private Key
A9D0134F-C895-4F33-BF85-351B754E9FF9     Opaque Object                     TDE Wallet
Metadata
E1AC8D2F-90E9-4F88-BFEE-2883FCBB7271     Opaque Object                     TDE Wallet
Metadata
25B7DE14-3849-4F67-BFBE-1934BFE3559B     Opaque Object                     TDE Wallet
Metadata
4ED713ED-FE2B-4F35-BF7D-BCBEA8327A0B     Symmetric Key                     TDE Master
Encryption Key: MKID 06EA813441C26B4F53BFD58E55C4BE90F4
6162E200-EF0A-4F89-BF25-A8596B3AD7B0     Opaque Object                     Certificate
Request
85A55486-28E5-4FFB-BF1C-B93C4C0BAD74     Secret Data                       Oracle Secret Data:
ID HSM_PASSWORD
67E74D97-56F6-407A-A035-009D953F907A     Template                           Default template
for EXA_DB1902_7274B2A2-6F71-4516-B2BB-6D67CC3824FC_SCAQAE08DV0308_EP
E621EA72-5DD1-4F4F-BFD4-451E5B7DB8A9     Symmetric Key                     TDE Master
Encryption Key: MKID 0625BA455B03CD4F57BFA5D2290FD379A1
You have new mail in /var/spool/mail/root
```

- b. Get the OKV Home version by running `okvutil`. The version should be the same as OKV Server version. In this case the value should be 21.7.0.0.0.

```
# su oracle

$ /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/okvutil
okvutil version 21.7.0.0.0
Usage: okvutil <command> [-v <verbosity>] [<command args>]
  <command> := list | upload | download | sign | sign-verify |
changepwd | diagnostics
Options:
  -v, --verbose <verbosity>
    Print extra information to standard out.
    Possible verbosity values are 0, 1 and 2 (more detailed information
with higher verbosity level).
For help on a particular command, use [okvutil <command> -h].
You have new mail in /var/spool/mail/root
```

6. Repeat the Steps 3 - 5 in the rest of the DomUs.
7. Repeat the Steps 1 - 6 for any other DB that needs to upgrade its OKV Home.
8. Stop the DomU 1 instances of all Databases with OKV-based TDE. This can be done via the Console, `srvctl` command or SQL* Plus.

9. Run `root.sh` from the selected OKV Home. It usually should be the one with the newer OKV version. It will prompt for replace PKCS11 Library, YES should be selected.

```
# /u02/app/oracle/admin/<dbname>/okv_home/okv/bin/root.sh
```

10. Start the DomU 1 instances of all databases with OKV-based TDE. This can be done via the Console, `srvctl` command or SQL* Plus.
11. Repeat the Steps 8 - 10 in the rest of the DomUs.

Migrate to Exadata Cloud Infrastructure

For general guidance on methods and tools to migrate databases to Oracle Cloud Infrastructure database services, including Exadata Cloud Infrastructure see "Migrating Databases to the Cloud".

A recommended approach for migrating to Exadata Cloud Infrastructure is using Zero Downtime Migration

- [Moving to Oracle Cloud Using Zero Downtime Migration](#)

Related Topics

- [Migrating Databases to the Cloud](#)

Moving to Oracle Cloud Using Zero Downtime Migration

Oracle now offers the Zero Downtime Migration service, a quick and easy way to move on-premises databases to Oracle Cloud Infrastructure.

Zero Downtime Migration leverages Oracle Active Data Guard to create a standby instance of your database in an Oracle Cloud Infrastructure system. You switch over only when you are ready, and your source database remains available as a standby. Use the Zero Downtime Migration service to migrate databases individually or at the fleet level. See *Move to Oracle Cloud Using Zero Downtime Migration* for more information.

Related Topics

- [Move to Oracle Cloud Using Zero Downtime Migration](#)

Connect Identity and Access Management (IAM) Users to Oracle Exadata Database Service on Dedicated Infrastructure

You can configure Oracle Exadata Database Service on Dedicated Infrastructure to use Oracle Cloud Infrastructure Identity and Access Management (IAM) authentication and authorization to allow IAM users to access an Oracle Database with IAM credentials.

- [Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication with Oracle Database](#)

Learn to enable an Oracle Database instance on Oracle Exadata Database Service on Dedicated Infrastructure to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.

- [Prerequisites for Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication on Oracle Database](#)
Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.
- [Enabling the Database and Clients for IAM Integration](#)
Follow the appropriate link below to configure IAM users to access your database.

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

Learn to enable an Oracle Database instance on Oracle Exadata Database Service on Dedicated Infrastructure to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier), or SSO tokens.

- [About Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Authentication with Oracle Database](#)
IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.
- [Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) Database Password Verifier Authentication](#)
You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).
- [Oracle Cloud Infrastructure \(OCI\) Identity and Access Management \(IAM\) SSO Token Based Authentication](#)
For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

About Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication with Oracle Database

IAM users can connect to the database instance by using either an IAM database password verifier or an IAM token.

Using the IAM database password verifier is similar to the database password authentication process. However, instead of the password verifier (encrypted hash of the password) being stored in the database, the verifier is instead stored as part of the OCI IAM user profile.

The second connection method, the use of an IAM token for the database, is more modern. The use of token-based access is a better fit for Cloud resources such as Oracle Databases in the Exadata Cloud Infrastructure. The token is based on the strength that the IAM endpoint can enforce. This can be multi-factor authentication, which is stronger than the use of passwords alone. Another benefit of using tokens is that the password verifier (which is considered sensitive) is never stored or available in memory.

Note

Oracle Database supports the Oracle DBaaS integration for Oracle Cloud Infrastructure (OCI) IAM with identity domains as well as the legacy IAM, which does not include identity domains. Both default and non-default domain users and groups are supported when using IAM with Identity Domains.

Support for non-default custom domains are only available with Oracle Database Release 19c, Version 19.21 and higher (but not Oracle Database Release 21c).

Oracle Cloud Infrastructure IAM integration with Oracle Exadata Database Service on Dedicated Infrastructure supports the following:

- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication*
- *Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication*

For complete details about the architecture for using IAM users on Oracle Exadata Database Service on Dedicated Infrastructure, see *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* in the [Oracle Database 19c Security Guide](#) and [Oracle AI Database 26ai Security Guide](#).

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Database Password Verifier Authentication

You can enable an Oracle Database instance to allow user access with an Oracle Cloud Infrastructure IAM database password (using a password verifier).

Note

Any supported 12c and above database client can be used for IAM database password access to Oracle Database.

An Oracle Cloud Infrastructure IAM database password allows an IAM user to log in to an Oracle Database instance as Oracle Database users typically log in with a username and password. The user enters their IAM username and IAM database password. An IAM database password is a different password than the Oracle Cloud Infrastructure Console password. Using an IAM user with a password verifier, you can log in to Oracle Database with any supported database client.

For password verifier database access, you create the mappings for IAM users and OCI applications to the Oracle Database instance. The IAM user accounts themselves are managed in IAM. The user accounts and user groups can be in either the default domain or in a custom, non-default domain.

For more information about managing IAM database password, see *Managing User Credentials*.

Related Topics

- [Managing User Credentials](#)

Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) SSO Token Based Authentication

For IAM token access to the database, the client application or tool requests a database token from IAM for the IAM user.

The client application will pass the database token directly to the database client through the database client API.

If the application or tool has not been updated to request an IAM token, then the IAM user can use OCI CLI to request and store the database token. You can request a database access token (`db-token`) using the following credentials:

- Security tokens (with IAM authentication), delegation tokens (in the OCI cloud shell) and API-keys, which are credentials that represent the IAM user to enable the authentication
- Instance principal tokens, which enable instances to be authorized actors (or principals) to perform actions on OCI resources after authentication
- Resource principal token, which is a credential that enables the application to authenticate itself to other OCI services
- Using an IAM user name and IAM database password (can only be requested by database client)

When the IAM user logs into the client with a slash / login and the `OCI_IAM` parameter is configured (`sqlnet.ora`, `tnsnames.ora`, or as part of a connect string), then the database client retrieves the database token from a file. If the IAM user submits a user name and password, the connection will use the IAM database verifier access described for client connections that use IAM database password verifiers. If the parameter `PASSWORD_AUTH=OCI_TOKEN`, then the database driver will instead use the username and password to connect directly to IAM and request a database token. The instructions in this guide show how to use the OCI CLI as a helper for the database token. If the application or tool has been updated to work with IAM, then follow the instructions for the application or tool. Some common use cases include the following: SQL*Plus on-premises, SQLcl on-premises, SQL*Plus in Cloud Shell, or applications that use SEP wallets.

There are several ways a database client can obtain an IAM database token:

- A client application or tool can request the database token from IAM for the user and can pass the database token through the client API. Using the API to send the token overrides other settings in the database client. Using IAM tokens requires the latest Oracle Database client 19c (at least 19.16). Some earlier clients (19c and 21c) provide a limited set of capabilities for token access. Oracle Database client 21c does not fully support the IAM token access feature:
 - JDBC-thin on all platforms
 - * See *Support for IAM Token-Based Authentication and JDBC and UCP Downloads* for more information.
 - SQL*Plus and Oracle Instant Client OCI-C on Linux:
See *Identity and Access Management (IAM) Token -Based Authentication* for more information
 - Oracle Data Provider for .NET (ODP.NET) Core: .NET clients (latest version of Linux or Windows). .NET software components are available as a free download from the following sites:
 - * *Oracle Data access Components - .NET Downloads*

- * *NuGet Gallery*
- * *Visual Studio Code Market Place*
- If the application or tool does not support requesting an IAM database token through the client API, the IAM user can first use the Oracle Cloud Infrastructure command line interface (CLI) to retrieve the IAM database token and save it in a file location. For example, to use SQL*Plus and other applications and tools using this connection method, you first obtain the database token using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI). For more information, see [db-token get](#). If the database client is configured for IAM database tokens, when a user logs in with the slash login form, the database driver uses the IAM database token that has been saved in default or specified file location.
- Some Oracle AI Database 26ai clients can also get a token directly from OCI IAM instead of using the OCI command line interface. Please review the client documentation to see which clients support this native IAM integration..
- A client application or tool can use an Oracle Cloud Infrastructure IAM instance principal or resource principal to get an IAM database token and use the IAM database token to authenticate itself to an Oracle Database instance. For more information, see *Mapping Instance and Resource Principals*.
- IAM users and OCI applications can request a database token from IAM with several methods, including using an API key. See *Configuring a Client Connection for SQL*Plus That Uses an IAM Token* for an example. See *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* for a description of other methods such as using a delegation token within an OCI cloud shell.

Note

If your database is in Restricted Mode, only DBAs with the `RESTRICTED SESSION` privilege can connect to the database.

If a user enters a username/password to log in, then the database driver uses the password verifier method to access the database. If the parameter `PASSWORD_AUTH=OCI_TOKEN`, then the database driver will instead user the username and password to connect directly to IAM and request a database token.

Related Topics

- [Support for IAM Token-Based Authentication](#)
- [JDBC and UCP Downloads](#)
- [Identity and Access Management \(IAM\) Token-Based Authentication](#)
- [db-token get](#)
- [Oracle Data access Components - .NET Downloads](#)
- [NuGet Gallery](#)
- [Visual Studio Code Marketplace](#)
- [Mapping Instance and Resource Principals](#)
- [Configuring a Client Connection for SQL*Plus That Uses an IAM Token](#)
- [Authenticating and Authorizing IAM Users for Oracle DBaaS Databases](#)

Prerequisites for Oracle Cloud Infrastructure (OCI) Identity and Access Management (IAM) Authentication on Oracle Database

Review the prerequisites for Identity and Access Management (IAM) authentication on an Oracle Database.

- [Prerequisites for IAM Authentication on Oracle Database](#)
Before using IAM authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.
- [Disable External Authentication Scheme](#)
Review the prerequisites for enabling IAM user access to Oracle Database.
- [Configure TLS to Use IAM Tokens](#)
When sending IAM tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB_GUID>/tls`.

Prerequisites for IAM Authentication on Oracle Database

Before using IAM authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.

1. Create a service gateway in the VCN where your database resources reside by following the instructions in *Task 1: Create the service gateway* in OCI documentation.
2. After creating the service gateway, add a route rule and an egress security rule to each subnet (in the VCN) where the database resources reside so that these resources can use the gateway to use IAM authentication:
 - a. Go to the **Subnet Details** page for the subnet.
 - b. In the **Subnet Information** tab, click the name of the subnet's Route Table to display its **Route Table Details** page.
 - c. In the table of existing Route Rules, check whether there is already a rule with the following characteristics:
 - **Destination:** All IAD Services In Oracle Services Network
 - **Target Type:** Service Gateway
 - **Target:** The name of the service gateway you just created in the VCNIf such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
 - d. Return to the Subnet Details page for the subnet.
 - e. In the subnet's Security Lists table, click the name of the subnet's security list to display its Security List Details page.
 - f. In the side menu, under **Resources**, click **Egress Rules**.
 - g. In the table of existing Egress Rules, check whether there is already a rule with the following characteristics:
 - **Stateless:** No

- **Destination:** All IAD Services In Oracle Services Network
- **IP Protocol:** TCP
- **Source Port Range:** All
- **Destination Port Range:** 443

- h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Related Topics

- [Task 1: Create the service gateway](#)

Disable External Authentication Scheme

Review the prerequisites for enabling IAM user access to Oracle Database.

If the database is enabled for another external authentication scheme, verify that you want to use IAM on the Oracle Database instance. There can only be one external authentication scheme enabled at any given time.

If you want to use IAM and another external authentication scheme is enabled, you must first disable the other external authentication scheme.

Configure TLS to Use IAM Tokens

When sending IAM tokens from the database client to the database server, a TLS connection must be established. The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS)
- Client with or without a wallet

Self-Signed Certificate

Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see *Configuring Transport Layer Security Authentication* in the *Oracle Database Security Guide*.

If you choose to use self-signed certificates and for additional wallet related tasks, see *Managing Public Key Infrastructure (PKI) Elements* in the *Oracle Database Security Guide*.

Related Topics

- [Configuring Transport Layer Security Authentication](#)
- [Managing Public Key Infrastructure \(PKI\) Elements](#)

Enabling the Database and Clients for IAM Integration

Follow the appropriate link below to configure IAM users to access your database.

For complete details about the architecture for using IAM users on Oracle Exadata Database Service on Dedicated Infrastructure, see *Authenticating and Authorizing IAM Users for Oracle DBaaS Databases* in the [Oracle Database 19c Security Guide](#) and [Oracle AI Database 26ai Security Guide](#).

Authenticating and Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure

An Oracle Database can be configured for Microsoft Azure users of Microsoft Entra ID to connect using single sign-on authentication.

- [About Authorizing Microsoft Entra ID \(MS-EI\) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure](#)
Users for Oracle Exadata Database Service on Dedicated Infrastructure can be centrally managed in an MS-EI service.
- [Configuring the Oracle Database for Microsoft Entra ID \(MS-EI\) Integration](#)
The MS-EI integration with the Oracle Database instance requires the database to be registered with MS-EI so that the database can request the MS-EI public key.

About Authorizing Microsoft Entra ID (MS-EI) Users for Oracle Databases on Oracle Exadata Database Service on Dedicated Infrastructure

Users for Oracle Exadata Database Service on Dedicated Infrastructure can be centrally managed in an MS-EI service.

The Oracle Database integration with MS-EI is supported for on-premises databases and most Oracle OCI DBaaS platforms.

The instructions for configuring MS-EI use the term "Oracle Database" to encompass these environments.

This type of integration enables the MS-EI user to access an Oracle Exadata Database Service on Dedicated Infrastructure instance. MS-EI users and applications can log in with MS-EI Single Sign On (SSO) credentials to get an MS-EI OAuth2 access token to send to the database.

The administrator creates and configures the application registration (app registration) of the Oracle Exadata Database Service on Dedicated Infrastructure instance with MS-EI. The administrator also creates application (app) roles for the database app registration in MS-EI, and assigns these roles to MS-EI users, groups, and applications. These app roles will be mapped to the database global schemas and global roles. An MS-EI principal that is assigned to an app role will be mapped to either a database global schema or database global role. An Oracle global schema can also be mapped exclusively to an MS-EI user. When the principal is a guest user or service principal, they can only be mapped to the database schema through an MS-EI app role. An Oracle global role can only be mapped to an MS-EI app role.

Tools and applications that are updated to support MS-EI tokens can authenticate users directly with MS-EI and pass the database access token to the Oracle Exadata Database Service on Dedicated Infrastructure instance. You can configure existing database tools such as SQL*Plus to use an MS-EI token from a file location or get the token directly from MS-EI. When using a utility to get the token to pass to the database client driver through a file location, MS-EI tokens can be retrieved using tools like Microsoft PowerShell or Azure CLI and put into a file location. An MS-EI OAuth2 database access token is a bearer token with an expiration time. The Oracle Database client driver will ensure that the token is in a valid format and that it has not expired before passing it to the database. The token is scoped for the database. Assigned app roles for the Azure AD principal are included as part of the access token. The directory location for the MS-EI token should only have enough permission for the user to write the token file to the location and the database client to retrieve these files (for example, just read and write by the process user). Because the token allows access to the database, it should be protected within the file system.

MS-EI users can request a token as a client registered with MS-EI app registration by using methods such as the following:

- Entering the MS-EI credentials into an MS-EI authentication screen with or without multi-factor authentication

Oracle Exadata Database Service on Dedicated Infrastructure supports the following MS-EI authentication flows:

- Interactive flow (authorization code), which is used when a browser can be used to enter credentials for the user
- Client credentials, which are for applications that connect as themselves (and not the end-user)
- On-Behalf-Of (OBO), where an application requests an access token on behalf of a logged-in user to send to the database
- ROPC is also supported for test and development environments

Oracle Exadata Database Service on Dedicated Infrastructure accepts tokens representing the following MS-EI principals:

- MS-EI user, who is registered user in the MS-EI tenancy
- Guest user, who is registered as a guest user in the MS-EI tenancy
- Service, which is the registered application connecting to the database as itself with the client credential flow (connection pool use case)

Configuring the Oracle Database for Microsoft Entra ID (MS-EI) Integration

The MS-EI integration with the Oracle Database instance requires the database to be registered with MS-EI so that the database can request the MS-EI public key.

For information on configuring MS-EI, configuring the database, and configuring the database client, see:

- [Authenticating and Authorizing Microsoft Azure Active Directory Users for Oracle Databases](#) in the Oracle Database 19c Security Guide.
- [Authenticating and Authorizing Microsoft Azure Users for Oracle Databases](#) in the Oracle AI Database 26ai Security Guide.
- [Prerequisites for Microsoft Entra ID \(MS-EI\) Authentication](#)
Review the prerequisites for integrating Oracle Database with MS-EI.
- [Networking Prerequisites for Microsoft Entra ID \(MS-EI\) Authentication](#)
Before using Azure AD authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.
- [Configure TLS to Use Microsoft Entra ID \(MS-EI\) Tokens](#)
When sending MS-EI tokens from the database client to the database server, a TLS connection must be established.

Prerequisites for Microsoft Entra ID (MS-EI) Authentication

Review the prerequisites for integrating Oracle Database with MS-EI.

The MS-EI integration with the Oracle Database on Oracle Exadata Database Service on Dedicated Infrastructure requires:

1. The Oracle Database to be version 19.18 or higher.
2. Connectivity to the database on TLS port 2484. Non TLS connections are not supported.

3. Outbound network connectivity to MS-EI so that the database can request the MS-EI public key.
4. The Oracle Database to be registered with MS-EI.
5. Users and applications that need to request an MS-EI token must also be able to have network connectivity to MS-EI. You may need to configure a proxy setting for the connection.
6. For Oracle Exadata Database Service on Dedicated Infrastructure deployments, the HTTP Proxy settings in your environment must allow the database to use MS-EI. These settings are defined by your fleet administrator while creating the Oracle Exadata Database Service on Dedicated Infrastructure infrastructure as described in [To create a Cloud Exadata infrastructure resource](#).

Note

The network configuration including the HTTP Proxy can only be edited until the Exadata Infrastructure is in Requires Activation state. Once it is activated, you cannot edit those settings.

Setting up an HTTP Proxy for an already provisioned Exadata Infrastructure needs a Service Request (SR) in My Oracle Support. See [Create a Service Request in My Oracle Support](#) for details.

Related Topics

- [Task 1: Create the service gateway](#)

Networking Prerequisites for Microsoft Entra ID (MS-EI) Authentication

Before using Azure AD authentication on databases in the Exadata Cloud Infrastructure, you must use the Networking service to add a service gateway, a route rule, and an egress security rule to the Virtual Cloud Network (VCN) and subnets where your database resources reside.

1. Create a service gateway in the VCN where your database resources reside by following the instructions in *Task 1: Create the service gateway* in OCI documentation.
2. After creating the service gateway, add a route rule and an egress security rule to each subnet (in the VCN) where the database resources reside so that these resources can use the gateway to use Azure AD authentication:
 - a. Go to the **Subnet Details** page for the subnet.
 - b. In the **Subnet Information** tab, click the name of the subnet's Route Table to display its **Route Table Details** page.
 - c. In the table of existing Route Rules, check whether there is already a rule with the following characteristics:
 - **Destination:** 0.0.0.0/0
 - **Target Type:** NAT Gateway
 - **Target:** The name of the NAT gateway you just created in the VCN
 If such a rule does not exist, click **Add Route Rules** and add a route rule with these characteristics.
 - d. Return to the **Subnet Details** page for the subnet.

- e. In the subnet's Security Lists table, click the name of the subnet's security list to display its **Security List Details** page.
- f. In the side menu, under **Resources**, click **Egress Rules**.
- g. In the table of existing Egress Rules, check whether there is already a rule with the following characteristics:
 - **Destination Type:** CIDR
 - **Destination:** 0.0.0.0/0
 - **IP Protocol:** TCP
 - **Source Port Range:** 443
 - **Destination Port Range:** All
- h. If such a rule does not exist, click **Add Egress Rules** and add an egress rule with these characteristics.

Related Topics

- [Task 1: Create the service gateway](#)

Configure TLS to Use Microsoft Entra ID (MS-EI) Tokens

When sending MS-EI tokens from the database client to the database server, a TLS connection must be established.

The TLS wallet with the database certificate for the ExaDB-D service instance must be stored under the `WALLET_ROOT` location. Create a `tls` directory so it looks like: `WALLET_ROOT/<PDB GUID>/tls`.

When configuring TLS between the database client and server there are several options to consider.

- Using a self-signed database server certificate vs a database server certificate signed by a commonly known certificate authority
- One-way TLS (TLS) vs Mutual or two-way TLS (mTLS)
- Client with or without a wallet

Self-Signed Certificate

Using a self-signed certificate is a common practice for internally facing IT resources since you can create these yourself and it's free. The resource (in our case, the database server) will have a self-signed certificate to authenticate itself to the database client. The self-signed certificate and root certificate will be stored in the database server wallet. For the database client to be able to recognize the database server certificate, a copy of the root certificate will also be needed on the client. This self-created root certificate can be stored in a client-side wallet or installed in the client system default certificate store (Windows and Linux only). When the session is established, the database client will check to see that the certificate sent over by the database server has been signed by the same root certificate.

A Well-Known Certificate Authority

Using a commonly known root certificate authority has some advantages in that the root certificate is most likely already stored in the client system default certificate store. There is no extra step for the client to store the root certificate if it is a common root certificate. The disadvantage is that this normally has a cost associated with it.

One-Way TLS

In the standard TLS session, only the server provides a certificate to the client to authenticate itself. The client doesn't need to have a separate client certificate to authenticate itself to the server (similar to how HTTPS sessions are established). While the database requires a wallet to store the server certificate, the only thing the client needs to have is the root certificate used to sign the server certificate.

Two-Way TLS (also called Mutual TLS, mTLS)

In mTLS, both the client and server have identity certificates that are presented to each other. In most cases, the same root certificate will have signed both of these certificates so the same root certificate can be used with the database server and client to authenticate the other certificate. mTLS is sometimes used to authenticate the user since the user identity is authenticated by the database server through the certificate. This is not necessary for passing IAM tokens but can be used when passing IAM tokens.

Client with a Wallet

A client wallet is mandatory when using mTLS to store the client certificate. However, the root certificate can be stored either in the same wallet or in the system default certificate store.

A Client without a Wallet

Clients can be configured without a wallet when using TLS under these conditions: 1) One-way TLS is being configured where the client does not have its own certificate and 2) the root certificate that signed the database server certificate is stored in the system default certificate store. The root certificate would most likely already be there if the server certificate is signed by a common certificate authority. If it's a self-signed certificate, then the root certificate would need to be installed in the system default certificate store to avoid using a client wallet.

For details on how to configure TLS between the database client and database server including the options described above, see:

- [Configuring Transport Layer Security Authentication](#) in the Oracle Database 19c Security Guide.
- [Configuring Transport Layer Security Encryption](#) in the Oracle AI Database 26ai Security Guide.

If you choose to use self-signed certificates and for additional wallet related tasks, see:

- [Managing Public Key Infrastructure \(PKI\) Elements](#) in the Oracle Database 19c Security Guide.
- [Configuring TLS Connection With a Client Wallet](#) in the Oracle AI Database 26ai Security Guide.

Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure

Exadata Database Service on Oracle Database@Azure enables you to store your database's transparent data encryption (TDE) keys, also known as master encryption keys (MEKs) in either a file-based Oracle wallet or in the OCI Vault.

This feature enables Exadata Database Service on Oracle Database@Azure users to utilize Azure Key Vault (AKV) Managed HSM, AKV Premium and AKV Standard for managing TDE MEKs. This integration allows applications, Azure services, and databases to use a centralized key management solution for enhanced security and simplified key lifecycle management.

- [Prerequisites](#)
The following steps must be completed before you can configure Azure Key Vault as the key management for your databases.
- [Network Requirements for Creating an Identity Connector and KMS Resources](#)
Azure Key Management Service (KMS) resources support both public and private connectivity. Azure Key Vault Managed HSM requires private connectivity, whereas Azure Key Vault Premium and Standard tiers support both public and private connectivity options.
- [Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure](#)
Learn how to manage Azure Key Vault integration for Exadata Database Service on Oracle Database@Azure.
- [Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure](#)
- [Updating the Multicloud PKCS#11 Driver](#)
- [Release Notes](#)
Review the changes made across the different releases of the Azure driver.

Prerequisites

The following steps must be completed before you can configure Azure Key Vault as the key management for your databases.

The following steps must be completed before you can configure Azure Key Vault as Key Management Service at the Exadata VM Cluster level.

1. You must first complete the registration required for delegated subnets to use advanced network features mentioned in [Network planning for Oracle Database@Azure](#), and then [create an Azure Virtual Network](#) with at least one delegated subnet in it to be used by Exadata VM cluster.
2. Provision an Exadata VM Cluster via the Azure interface. See [Provisioning an Exadata VM Cluster for Azure](#) for step-by-step instructions.
3. Review the networking requirements to determine whether the VM Cluster will connect to Azure KMS via a public network or through private connectivity. For more information, see [Connected Machine agent network requirements](#) or [Network Requirements for Creating an Identity Connector and KMS Resources](#) for specific steps to follow.
4. Ensure that the following policy is created before creating the database.

```
allow any-user to manage oracle-db-azure-vaults IN tenancy where ALL
{ request.principal.type in ('cloudvmcluster') }
```

Network Requirements for Creating an Identity Connector and KMS Resources

Azure Key Management Service (KMS) resources support both public and private connectivity. Azure Key Vault Managed HSM requires private connectivity, whereas Azure Key Vault Premium and Standard tiers support both public and private connectivity options.

The following sections outline the network requirements for public network access.

Configuration Using Private Network

- **Arc agent network configuration**

To create an Identity Connector over a private network, an Azure Arc Private Link Scope and a Private Endpoint must be configured through the Azure portal. Refer to the [Azure documentation](#) for detailed steps on setting up private connectivity for Azure Arc-enabled servers.

! Important

The Private Endpoint must be created in a non-delegated subnet within the Azure Virtual Network (VNet) that hosts the Oracle Exadata VM Cluster. Private Endpoints are not supported in delegated subnets. By default, Exadata VM Clusters are provisioned in delegated subnets.

Managed HSM requires private connectivity and is supported only in Azure regions that offer Advanced Networking features. For a list of supported regions, see [Network planning for Oracle Database@Azure](#).

To allow communication with private agent resources over the private network, a private DNS zone and corresponding A records must be created within the VCN's DNS configuration in your Oracle Cloud Infrastructure (OCI) tenancy.

The DNS configuration for the Private Endpoint associated with the Private Link Scope must include the necessary private agent resource addresses. For more information, see [URLs](#) section in [Connected Machine agent network requirements](#).

First, retrieve the list of required addresses from the Azure portal. Then, update the DNS zone entry in OCI to complete the configuration.

Steps to retrieve the list of required IP addresses:

1. Sign in to the Azure portal.
2. Search for "Azure Arc Private link scopes".
3. Choose any private link scope from the list.
4. Under the Configure menu, click Private endpoint connections.
5. Click the Private endpoint link.
6. Under Settings, select DNS configuration to view the required addresses.

Example: Add a Private Agent Resource (for example, `gbl.his.arc.azure.com`)

The IP address associated with `gbl.his.arc.azure.com`, along with any other required agent resources, must be defined in the private DNS zone.

Steps:

1. Create a Private Zone
For more information, see [Creating a Private DNS Zone](#).
 - Zone type: Primary
 - Zone name: <Descriptive name>
 - Compartment: <Compartment name or OCID>
2. Add DNS Records
 - Navigate to the Records tab on the zone details page.

- Click Manage Records, then Add Record:
 - * Name: gbl.his.arc.azure.com
 - * Type: A (IPv4 address)
 - * TTL (seconds): 3600
 - * RDATA mode: Basic
 - * Address: <Private IP address>
- 3. Publish the Zone
- 4. Ensure the record appears on the zone's page after publishing.
- 5. Verify that connectivity to Azure services from the VM cluster is routed through the private network.

Even with Private Connectivity, the following endpoints must be routed through the Azure NAT gateway.

Agent resources:

- `packages.microsoft.com`
- `login.microsoftonline.com`
- `pas.windows.net`
- `management.azure.com`
- **Azure Key Vault private endpoints configuration**

To access endpoints of Azure Key Vaults over private connectivity, you must create a DNS zone. Additionally, an A record mapping the fully qualified domain name (FQDN) of the resource to the IP address of the corresponding private endpoint must be added in the OCI tenancy.

To access the Managed HSM service over a Private Endpoint in your virtual network hosting an Exadata VM Cluster, you can establish a private link connection to Managed HSM and associate it with either the default subnet or a non-delegated subnet. Follow the steps outlined in the "Configuration Using Private Network" section of the "Network Requirements for Creating an Identity Connector and KMS Resources" topic.. For more information, see [Integrate Managed HSM with Azure Private Link](#).

Configuration Using Public Network

Create a NAT gateway in Azure portal and associate it with the delegated subnet of the Exadata VM Cluster. For more information, see [Create a NAT gateway and associate it with an existing subnet](#).

Using the Console to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure

Learn how to manage Azure Key Vault integration for Exadata Database Service on Oracle Database@Azure.

- [Create an Identity Connector from the OCI Console](#)

Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.
- [View the details of an Identity Connector](#)

To view the details of an identity connector, use this procedure.

- [Enable or Disable the Azure Key Management](#)
This step installs the required library on the VM cluster to support Azure Key Vault integration. Ensure that an identity connector is created before enabling Azure Key Management on the Exadata VM Cluster.
- [Create Azure Key Vault \(Managed HSM, Premium, and Standard\) and Assign Required Permissions](#)
Create Azure Key Vault Managed HSM, Azure Key Vault Premium, or Azure Key Vault Standard, then assign the permission.
- [Register Azure Key Vaults in the OCI console](#)
This is an alternative way to register your Azure key vaults from the OCI console. If you have already registered your vault during the creation of a database in your existing Exadata VM Cluster, you can skip this step.
- [Create a Database and Use Azure Key Vault as the Key Management Solution](#)
This topic describes only the steps for creating a database and using Azure Key Vault as the key management solution.
- [Change the Key Management from Oracle Wallet to Azure Key Vault](#)
Learn to change encryption keys between different encryption methods.
- [Rotate the Keys Managed by Azure Key Vault for a Container Database](#)
To rotate the Azure key vault encryption key of a container database (CDB), use this procedure.
- [Rotate the Keys Managed by Azure Key Vault for a Pluggable Database](#)
To rotate the Azure key vault encryption key of a pluggable database (PDB), use this procedure.

Create an Identity Connector from the OCI Console

Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.

This enables secure communication with the Azure Key Management Service (KMS) using the Azure identity generated by the Arc agent. The Azure Arc agent can communicate with Azure services over either a public network or a private connectivity setup. Learn more about [Azure Arc](#).

Each Exadata VM cluster must have an identity connector enabled to access Azure resources. The identity connector establishes either a public or private connection between the Exadata VM cluster and Azure Key Management resources, depending on the roles assigned.

To generate an **access token** for your current Azure account, see [az account get-access-token](#).

You can create an identity connector in one of two ways—using the Oracle Exadata Database Service on Dedicated Infrastructure interface or the Database Multicloud Integrations interface.

Oracle Exadata Database Service on Dedicated Infrastructure

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. From the left menu, click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. From the list of **Exadata VM Clusters**, select the cluster you are using.
4. Select **VM Cluster information**, and then navigate to **Identity connector** located under **Multicloud information**. Click the **Create** link.

Note

If an identity connector has not been created previously, it is displayed as **None**.

5. The **Identity connector name**, **Exadata VM cluster**, **Azure subscription id**, and **Azure resource group name** are read-only fields and will be populated with values.
6. Enter your **Azure tenant id**, and **Access token**.
7. Expand the **Show advanced options** section.
The **Private connectivity information** and **Tags** sections populate.
To enable a private endpoint connection, enter the **Azure arc private link scope** name.
8. To add tags for your resources, click **Add tag**, and then enter required values.
9. Review your selections, and then click **Create** to create the identity connector.

Database Multicloud Integrations

1. Open the navigation menu. Click **Oracle Database**, then click **Database Multicloud Integrations**.
2. Select **Identity Connectors** from the left navigation menu.
3. From the **Compartment** drop-down list, select your compartment that you are using.
4. Once you select your compartment, the **Identity connector name** automatically populates a name.
By default, the identity connector type is selected as **Azure**.
5. Select **ARC agent** as an identity mechanism.
6. Select your compartment from the **Choose an Exadata VM cluster compartment** list, and then select your Exadata VM Cluster from the **Choose an Exadata VM cluster** list.
7. Enter your **Azure tenant id**. The **Azure subscription id** and **Azure resource group name** fields populate values based on your Exadata VM Cluster selection.
8. Enter an **Access token**.
9. Expand the **Show advanced options** section. The **Private connectivity information** and **Tags** sections populate. These fields are optional.
10. To add tags for your resources, click **Add tag**, and then enter required values.
11. Review your selections, and then click **Create**.

View the details of an Identity Connector

To view the details of an identity connector, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster of your choice.
4. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, confirm that the Identity connector field displays the identity connector created previously.
5. Click the name of the **Identity Connector** to view its details.

Enable or Disable the Azure Key Management

This step installs the required library on the VM cluster to support Azure Key Vault integration. Ensure that an identity connector is created before enabling Azure Key Management on the Exadata VM Cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster of your choice.
4. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, click the **Enable** link next to **Azure key store**.
5. On the resulting **Enable Azure key management** dialog, click **Enable** to confirm the operation.
Confirming the action will install a library on your Exadata VM Cluster.
The status of Azure key store changes from **Disabled** to **Enabled**.
6. To disable Azure key store, click the **Disable** link.
7. On the resulting **Disable Azure key management** dialog, click **Disable** to confirm the operation.
Disabling the Azure key management removes the library installed during enablement, which will impact the availability of databases configured to use it.

Note

Azure key management is configured at the VM cluster level, requiring all databases in the cluster to use the same key management solution. However, databases that use Oracle Wallet can coexist alongside those that use Azure Key Vault within the same cluster.

Create Azure Key Vault (Managed HSM, Premium, and Standard) and Assign Required Permissions

Create Azure Key Vault Managed HSM, Azure Key Vault Premium, or Azure Key Vault Standard, then assign the permission.

For more information, see [Create a key vault using the Azure portal](#).

Note

There are specific roles that must be assigned to the group to grant the necessary permissions for accessing and managing **Azure Key Vault Managed HSM**, **Azure Key Vault Premium**, and **Azure Key Vault Standard** resources.

1. Create a group and add members.
Azure groups allow you to manage users by assigning them the same access and permissions to resources.

- To manage group in Azure, you are required to have **User Administrator** or **Groups Administrator** role. For more information, see [Manage Microsoft Entra groups and group membership](#).
 - You must create a security group and add members from the Microsoft Azure portal. You must select the **Security** option as your **Group type**. For more information, see [Create a basic group and add members](#)
2. Assign the following roles based on the type of Azure Key Vault:
 - **For Managed HSM:**
 - IAM: Reader
 - Local RBAC: Managed HSM crypto officer + Managed HSM crypto user
 - **For Key Vault Premium and Standard**
 - IAM : Reader + Key Vault Crypto Officer

For detailed steps, refer to [Assign Azure roles using the Azure portal](#).

Register Azure Key Vaults in the OCI console

This is an alternative way to register your Azure key vaults from the OCI console. If you have already registered your vault during the creation of a database in your existing Exadata VM Cluster, you can skip this step.

1. From the OCI console, navigate to **Database Multicloud Integrations**, and then select **Microsoft Azure Integration**. From the **Microsoft Azure Integration** section, select Azure Key Vaults.

Note

At least one key must be created in the vault on the Azure portal for the registration to succeed.

2. Select the **Register Azure key vaults** button.
3. From the drop-down list, select your **Compartment**.
4. To discover Azure Key Vault or Managed HSM within the same subscription in an Azure tenancy (Identity Connector and Azure Key Vault or Managed HSM in the same subscription) and register keys:
 - a. From the **Azure key vaults** section, select an **identity connector** from the **Discover Azure key vaults using connector** list.
 - b. Click **Discover**.
The list of **Vault name(s)** is displayed.
 - c. Select the check box located next to **Vault name**.
5. To discover a single Azure Key Vault or Managed HSM within the same subscription in an Azure tenancy, provide the full resource ID in one of the following formats:
For Azure Key Vault:

```
/subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
providers/Microsoft.KeyVault/vaults/<key-vault-name>
```

For Azure Managed HSM:

```
/subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
providers/Microsoft.KeyVault/managedHSMs/<hsm-name>
```

6. To discover Azure Key Vault or Managed HSM across subscriptions in an Azure tenancy (Identity Connector and Azure Key Vault or Managed HSM in different subscriptions) and register keys:
 - a. From the **Azure key vaults** section, select an **identity connector** from the **Discover Azure key vaults using connector** list.
 - b. Provide the full resource ID of the Azure Key Vault in the following format.


```
/subscriptions/<subscription-id>/resourceGroups/<resource-groups>/
providers/Microsoft.KeyVault/vaults/<key-vault-name>
```
 - c. Click **Discover**.
7. If you want to add tags for your resources, expand the **Advanced options** section, and then click **Add tag**.
8. Click **Register** to register the vault(s) locally in the OCI.
9. Once you register the vault, you can view the **Display name**, **State**, **Type**, **Azure resource group**, and **Created** information of the vaults in the list.
10. Select the vault you are using, and then click the **Identity connector associations** tab, which lists identity connector associations in the current compartment.

Note

A default association is automatically created between the vault and the Identity Connector used during the vault registration process. This allows the vault to be used on the Exadata VM cluster associated with that specific Identity Connector.

If you want to use the same vault in other clusters that are registered with different Identity Connectors (i.e., not the one used during vault discovery), you must explicitly create an association between the vault and those additional Identity Connectors.

11. Click **Create association**.
12. From the drop-down list, select your **Compartment**, **Azure key vault association name**, and **Identity connector**.
13. If you expand the **Advanced options** section, you can add **Tags** for organizing your resources.
14. Review your selections, and then click **Create**.

Create a Database and Use Azure Key Vault as the Key Management Solution

This topic describes only the steps for creating a database and using Azure Key Vault as the key management solution.

For the generic database creation procedure, see [To create a database in an existing VM Cluster](#).

Prerequisites

Before creating your first database and selecting Azure Key Vault for key management, ensure the following prerequisites are met:

- All network prerequisites outlined in the [Network Requirements for Creating an Identity Connector and KMS Resources](#) section are fulfilled
- The identity connector is created and available for use
- Azure key management is enabled at the VM cluster level
- The VM cluster has the necessary permissions to access the vaults
- The vaults are registered as OCI resources

Limitations

- **Virtual Machines Restriction:** Scaling out a VM cluster does not automatically extend databases that use Azure Key Vault to the newly added virtual machine. To complete the extension, you must update the existing Identity Connector for the Exadata VM Cluster by supplying the Azure access token. After updating the Identity Connector, run the [dbaascli database addInstance](#) command to add the database instance to the new VM.
- **Data Guard Restrictions:**
 - When creating a standby database for a primary that uses Azure Key Vault, ensure that the target VM cluster has an active Identity Connector, Azure key management is enabled, and the required association between the Identity Connector and the Key Vault is properly configured.
 - Cross-region Data Guard and database restore operations are not supported for databases that use Azure Key Vault for key management.
- **PDB Operations Restriction:** Remote PDB operations—such as clone, refresh, and relocate—are supported only if both the source and destination databases use the same Transparent Data Encryption (TDE) key.

Steps

Note

If Azure key management is enabled at the VM cluster level, you will have two key management options: **Oracle Wallet** and **Azure Key Vault**.

1. In the **Encryption** section, choose **Azure Key Vault**.
2. Select a registered **Vault** available in your compartment.

Note

- The Vault list populates only registered vaults. Click the **Register new vaults** link to register your vault. From the Register Azure key vaults page, select your vault, and then click **Register**.
- At least one key must be registered in your vaults.

3. Select the Key available in your compartment.

Change the Key Management from Oracle Wallet to Azure Key Vault

Learn to change encryption keys between different encryption methods.

1. Navigate to your existing Exadata VM Cluster in the OCI console. Select the **Databases** tab. Then, select the database resource that you are using.
2. Select the **Database information tab**, and then scroll down to **Key management** section.
3. In the **Encryption** section, verify that **Key management** is set to **Oracle Wallet**, and then select the **Change** link.
4. Enter the following information on the **Change key management** page.
 - a. Select your **Key management** as **Azure key vault** from the drop-down list.
 - b. Select your **Vault** compartment that you are using, and then select your **Vault** that is available in the compartment.
 - c. Select the **Key** compartment that you are using, and then select your **Key** from the drop-down list.
 - d. Click **Save changes**.

Note

Changing key management from Azure Key Vault to Oracle Wallet cannot be performed using the API or OCI Console—it is only supported through the [dbaascli tde fileToHsm](#) command. Additionally, switching between Azure Key Vault and OCI Vault or Oracle Key Vault (OKV) are not supported.

Rotate the Keys Managed by Azure Key Vault for a Container Database

To rotate the Azure key vault encryption key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM Clusters, click the name of the VM cluster that contains the database that you want to rotate encryption keys.
4. Click **Databases**.
5. Click the name of the database that you want to rotate encryption keys.
The Database Details page displays information about the selected database.
6. In the **Encryption** section, verify that the **Key Management** is set to **Azure Key Vault**, and then click the **Rotate** link.
7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Note

Key rotation must be performed through the OCI interface. Rotating the key directly from the Azure interface has no effect on the database.

Rotate the Keys Managed by Azure Key Vault for a Pluggable Database

To rotate the Azure key vault encryption key of a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the PDB you want to start, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click Pluggable Databases in the Resources section of the page.
A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as Azure Key Vault.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Using the API to Manage Azure Key Vault Integration for Exadata Database Service on Oracle Database@Azure

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

Use these API operations to manage Azure Key Vault integration for Exadata Database Service on Oracle Database@Azure.

Table 5-10 API operation to manage Azure Key Vault integration for Exadata Database Service on Oracle Database@Azure

API	Description
<code>createOracleDbAzureConnector</code>	Captures Azure-specific details from the customer and automates the installation of the ARC Agent on the ExaDB-D VM Cluster.
<code>deleteOracleDbAzureConnector</code>	Deletes the Azure Connector resource and uninstalls the Arc Agent from the ExaDB-D VM Cluster.

Table 5-10 (Cont.) API operation to manage Azure Key Vault integration for Exadata Database Service on Oracle Database@Azure

API	Description
getOracleDbAzureConnector	Fetches the details of a specific Azure Connector resource.
listOracleDbAzureConnectors	Lists Azure Connector resources based on the specified filters.
CreateMultiCloudResourceDiscovery	Creates a new multi-cloud resource discovery resource.
GetMultiCloudResourceDiscovery	Retrieves details of a specific multi-cloud resource discovery resource.
ListMultiCloudResourceDiscoveries	Retrieves a list of all multi-cloud resource discovery resources.
CreateOracleDbAzureVaultAssociation	Creates a new association between an Oracle DB and an Azure vault.
GetOracleDbAzureVaultAssociation	Retrieves details of a specific Oracle DB Azure vault association.
ListOracleDbAzureVaultAssociations	Retrieves a list of all Oracle DB Azure vault associations.
CreateCloudVMCluster	Creates a cloud VM cluster.
GetCloudVmCluster	Gets information about the specified cloud VM cluster. Applies to Exadata Cloud Service instances and Autonomous Database on dedicated Exadata infrastructure only.
ListCloudVmClusters	Gets a list of the cloud VM clusters in the specified compartment. Applies to Exadata Cloud Service instances and Autonomous Database on dedicated Exadata infrastructure only.
DeleteCloudVMCluster	Deletes the specified cloud VM cluster. Applies to Exadata Cloud Service instances and Autonomous Database on dedicated Exadata infrastructure only.
CreateDatabase	Creates a new database in the specified Database Home. If the database version is provided, it must match the version of the Database Home. Applies to Exadata and Exadata Cloud@Customer systems.
CreateDatabaseFromBackup	Details for creating a database by restoring from a database backup. Warning: Oracle recommends that you avoid using any confidential information when you supply string values using the API.
MigrateVaultKey	Changes encryption key management from customer-managed, using the Vault service , to Oracle-managed.
RotateVaultKey	Creates a new version of an existing Vault service key.
RestoreDatabase	Restores a Database based on the request parameters you provide.

Updating the Multicloud PKCS#11 Driver

Note

The PKCS#11 driver for Azure is installed automatically when Azure key management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:

```
rpm -qa | grep pkcs
```

This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.

Example output:

```
pkcs-multicloud-driver-maz-0.2-250814.1708.x86_64
```

This confirms that the `pkcs-multicloud-driver-maz` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require different actions depending on the version, particularly when updating the PKCS#11 driver.

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one node at a time):

1. Shutdown the database instances that are using the Azure KMS. All databases that use the Azure KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

```
srvctl stop instance -d <db_unique_name> -i <instance_name> [-o  
<stop_option>]
```

- `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
- `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

For details on additional options, refer to the SRVCTL command-line help or reference manual.

2. Verify that the databases have been shut down.

Ensure that all database instances that are using the Azure KMS have been stopped before proceeding.

3. Update the RPM.
As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider  
AZURE --nodeList <node> --databasesStopped
```

This command downloads and installs the latest RPM on the specified node.

4. Verify driver installation.
Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Start the database instances.
After the driver update, restart the database instances using:

```
<oracle_home>/bin/srvctl start instance -node "<node>"
```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

1. Stop databases using the Azure KMS.
Stop all databases that use the Azure KMS for master encryption keys (MEK) before updating the driver:

```
dbaascli database stop [--dbname <value>]
```

2. Verify that the databases have been shut down.
Ensure that all databases using the Azure KMS have been stopped before proceeding.
3. Upgrade the RPM.
As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider  
AZURE --databasesStopped
```

This command downloads and installs the latest RPM.

4. Verify driver installation.
Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Restart the databases.
Once the driver has been updated, restart the databases:

```
dbaascli database start [--dbname <value>]
```

Release Notes

Review the changes made across the different releases of the Azure driver.

- [Release 0.2 \(250910.0501\)](#)

Release 0.2 (250910.0501)

- Various bug fixes and stability improvements.

Google Cloud Key Management Integration for Exadata Database Service on Oracle Database@Google Cloud

Exadata Database Service on Oracle Database@Google Cloud now supports integration with Google Cloud Platform's Key Management Service (KMS).

This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using GCP Customer-Managed Encryption Keys (CMEKs).

Previously, Transparent Data Encryption (TDE) master encryption keys (MEKs) could only be stored in a file-based Oracle Wallet, Oracle Cloud Infrastructure (OCI) Vault, or Oracle Key Vault (OKV).. With this update, users can now store and manage MEKs directly in GCP KMS, providing improved key lifecycle control and alignment with organization-specific security policies.

This integration enables applications, Google Cloud services, and databases to benefit from a centralized key management solution that offers enhanced security and simplified key lifecycle management.

- [Prerequisites](#)
Before configuring GCP Customer Managed Encryption Keys (CMEK) as the key management service for your databases, ensure the following prerequisites are met.
- [Using the Console to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud](#)
Learn how to manage GCP KMS integration for Exadata Database Service on Oracle Database@Google Cloud.
- [Using the API to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud](#)
- [Updating the Multicloud PKCS#11 Driver](#)
- [Release Notes](#)
Review the changes made across the different releases of the GCP driver.

Prerequisites

Before configuring GCP Customer Managed Encryption Keys (CMEK) as the key management service for your databases, ensure the following prerequisites are met.

1. Provision an Exadata VM Cluster via the Google Cloud console. See [Provisioning an Exadata VM Cluster for Google Cloud](#) for step-by-step instructions.
2. Review the Identity Connector connection to ensure it is correctly configured and active. For more information, see *Verify the Default Identity Connector Attached to the VM Cluster*.

3. Prerequisites for Configuring GCP Customer Managed Encryption Keys (CMEK) at the Exadata VM Cluster Level.
To enable Google Cloud Platform (GCP) Customer Managed Encryption Keys (CMEK) for databases deployed with Exadata Database Service on Oracle Database@Google Cloud, you must configure CMEK as the key management option at the VM cluster level. Once CMEK is enabled, all database encryption and decryption operations will use the specified GCP-managed key.

Before enabling CMEK, ensure that:

- The required GCP key rings and encryption keys are already created in GCP.
- These keys are mirrored as anchor resources in Oracle Cloud Infrastructure (OCI), ensuring synchronization between GCP and OCI.
- The anchor resources are in place for database provisioning and for managing the encryption key lifecycle, including key rotation, revocation, and auditing.

4. IAM Policy Requirements for Accessing GCP Key Resources.
The database uses the cluster resource principal to securely retrieve GCP key resources. To enable this functionality, you must define the appropriate IAM policies in your OCI tenancy.

Read-Only Access to Oracle GCP Keys:

```
Allow any-user to read oracle-db-gcp-keys in compartment id <your-  
compartment-OCID>  
where all { request.principal.type = 'cloudvmcluster', }
```

This policy grants read-only access to GCP key resources for the VM cluster resource principal.

5. Required Google Cloud Services Connectivity GCP CMEK Integration.
Google Cloud VPCs typically include default routes to the services listed below. Ensure that no firewall egress rules block access to these endpoints.

- <https://iamcredentials.googleapis.com/>
- <https://sts.googleapis.com/>
- <https://cloudkms.googleapis.com/>

Using the Console to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud

Learn how to manage GCP KMS integration for Exadata Database Service on Oracle Database@Google Cloud.

- [To create an ASM cloud VM cluster](#)
To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.
- [Verify the Default Identity Connector Attached to the VM Cluster](#)
To view the details of an identity connector attached to a VM cluster, use this procedure.
- [Create a Key Ring in Google Cloud Console](#)
To create a key ring, use this procedure.
- [Create a Key in Google Cloud Console](#)
To create a raw symmetric encryption key in the specified key ring and location, use this procedure.

- [Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure \(OCI\)](#)
To allow a key to be discoverable in Oracle Cloud Infrastructure (OCI), use this procedure.
- [Register GCP Key Ring in Oracle Cloud Infrastructure \(OCI\)](#)
To enable Google Cloud Customer Managed Encryption Keys (CMEK) for your VM cluster, you must first register the GCP Key Ring in OCI.
- [Enable or Disable Google Cloud Key Management](#)
To enable GCP CMEK for your Exadata VM Cluster, use this procedure.
- [Create a Database and Use GCP Customer-Managed Encryption Key \(CMEK\) as the Key Management Solution](#)
This topic describes only the steps for creating a database and using GCP Customer-managed encryption key (CMEK) as the key management solution.
- [Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key \(CMEK\)](#)
To change encryption keys between different encryption methods, use this procedure.
- [Rotate the GCP Customer Managed Encryption Key of a Container Database \(CDB\)](#)
To rotate the GCP Customer Managed Encryption Key of a container database (CDB), use this procedure.
- [Rotate the GCP Customer Managed Encryption Key of a Pluggable Database \(PDB\)](#)
To rotate the GCP Customer Managed Encryption Key of a pluggable database (PDB), use this procedure.

To create an ASM cloud VM cluster

To create your ASM VM cluster, be prepared to provide values for the fields required for configuring the infrastructure.

Note

To create a cloud VM cluster in an Exadata Cloud Infrastructure instance, you must have first [created a Cloud Exadata infrastructure resource](#).

Note

Multi-VM enabled Infrastructure will support creating multiple VM Clusters. Infrastructures created before the feature [Create and Manage Multiple Virtual Machines per Exadata System \(MultiVM\) and VM Cluster Node Subsetting](#) was released only support creating a single cloud VM cluster.

Note

When you provision an Exadata VM cluster in Exadata Database Service on Oracle Database@Google Cloud, an Identity Connector is automatically created and associated with the VM cluster.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

3. Click **Create Exadata VM Cluster**.
The **Create Exadata VM Cluster** page is displayed. Provide the required information to configure the VM cluster.
4. **Compartment:** Select a compartment for the VM cluster resource.
5. **Display name:** Enter a user-friendly display name for the VM cluster. The name doesn't need to be unique. An Oracle Cloud Identifier (OCID) will uniquely identify the VM cluster. Avoid entering confidential information.
6. **Select Exadata infrastructure:** Select the infrastructure resource that will contain the VM cluster. You must choose an infrastructure resource that has enough resources to create a new VM cluster. Click **Change Compartment** and pick a different compartment from the one you are working in to view infrastructure resources in other compartments.

Note

Multiple VM clusters may be created only in a Multi-VM enabled Infrastructure.

7. **VM Cluster Type:**

Note

You cannot change the VM cluster type after deploying the VM cluster. If you wish to change the VM cluster type, you must create a new VM cluster and migrate the database to the new cluster.

- **Exadata Database:** Standard Database VM with no restrictions, suitable for all workloads.
 - **Exadata Database-Developer:** Developer Database VM with restrictions, suitable for application development only.
8. **Configure the VM cluster:** Specify the DB servers to used for new VM cluster (by default all DB Servers are selected). Click **Select DB Servers** to select from the available DB servers, and then click **Save**.
VM Cluster Type - Exadata Database: Select a minimum of one database server for VM placement. If you require a high availability database service that remains available during maintenance and unplanned outages, select at least two database servers. Maximum resources available for allocation per VM are based on the number of database servers selected.
VM Cluster Type - Exadata Database-Developer: Select one database server for VM placement. Only one database server may be selected.

In the **Resource allocation per VM** pane:

- Specify the number of OCPU/ECPU you want to allocate to each of the VM cluster's virtual machine compute nodes. For VM clusters created on X11M Exadata infrastructure specify ECPUs. For VM Clusters created on X10M and earlier Exadata infrastructure, specify OCPUs. The minimum is 2 OCPU per VM for X10M and earlier infrastructure or 8 ECPUs per VM for VM clusters created on X11M Exadata infrastructure. The read-only **Requested OCPU count for the Exadata VM cluster** field displays the total number of OCPU or ECPU cores you are allocating.
- Specify the **Memory per VM** to allocate to each VM. The minimum per VM is 30 GB.
- Specify the **Local Storage per VM** to allocate local storage to each VM. The minimum per VM is 60 GB.
Each time when you create a new VM cluster, the space remaining out of the total available space is utilized for the new VM cluster.

In addition to `/u02`, you can specify the size of additional local file systems.

For more information and instructions to specify the size for each individual VM, see [Introduction to Scale Up or Scale Down Operations](#).

- Click **Show additional local file systems configuration options**.
- Specify the size of `/`, `/u01`, `/tmp`, `/var`, `/var/log`, `/var/log/audit`, and `/home` file systems as needed.

Note

- * You can only expand these file systems and cannot reduce the size once expanded.
- * Due to backup partitions and mirroring, the `/` and `/var` file systems will consume twice the space they were allocated, which is indicated in the read-only **Total allocated storage for / (GB) due to mirroring** and **Total allocated storage for /tmp (GB) due to mirroring** fields.

- After creating the VM Cluster, check the **Exadata Resources** section on the **Exadata Infrastructure Details** page to check the file size allocated to the local storage (`/u02`) and local storage (additional file systems).

9. Exadata storage:

- **Specify the usable Exadata storage TB.** Specify the storage in multiples of 1 TB. Minimum: 2 TB
- **Allocate storage for Exadata sparse snapshots:** Select this configuration option if you intend to use snapshot functionality within your VM cluster. If you select this option, the SPARSE disk group is created, which enables you to use VM cluster snapshot functionality for PDB sparse cloning. If you do not select this option, the SPARSE disk group is not created and snapshot functionality will not be available on any database deployments that are created in the environment.

Note

The storage configuration option for sparse snapshots cannot be changed after VM cluster creation.

- **Allocate storage for local backups:** Select this option if you intend to perform database backups to the local Exadata storage within your Exadata Cloud Infrastructure instance. If you select this option, more space is allocated to the RECO disk group, which is used to store backups on Exadata storage. If you do not select this option, more space is allocated to the DATA disk group, which enables you to store more information in your databases.

Note

The storage configuration option for local backups cannot be changed after VM cluster creation.

10. Version:

- **Oracle Grid Infrastructure version:** From the list, choose the Oracle Grid Infrastructure release (19c and 26ai) that you want to install on the VM cluster. The Oracle Grid Infrastructure release determines the Oracle Database releases that can be supported on the VM cluster. You cannot run an Oracle Database release that is later than the Oracle Grid Infrastructure software release.

Note

Minimum requirements for provisioning a VM Cluster with Grid Infrastructure 26ai:

- Exadata Guest VM running Exadata System Software 23.1.8
- Exadata Infrastructure running Exadata System Software 23.1.x

- **Exadata guest version:**
 - **Exadata infrastructure with Oracle Linux 7 and Exadata image version 22.1.10.0.0.230422:**
 - * The **Change image** button is not enabled.
 - * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0.
 - * The Exadata guest version will be the same as that of the host OS.
 - **Exadata infrastructure with Oracle Linux 8 and Exadata image version 23.1.3.0.0.230613:**
 - * The Exadata guest version defaults to the latest (23.1.3.0).
 - * The Oracle Grid Infrastructure version defaults to 19.0.0.0.0
 - * The **Change image** button is enabled.
 - * Click **Change image**.
The resulting Change image panel displays the list of available major versions of Exadata image (23.1.3.0 and 22.1.3.0).
The most recent release for each major version is indicated by "(latest)".
 - * Slide **Display all available versions**.
Six past versions including the latest versions of Exadata images 23.1.3.0 and 22.1.3.0 are displayed.
 - * Choose a version.

* Click **Save Changes**.

11. **SSH Keys:** Add the public key portion of each key pair you want to use for SSH access to the VM cluster:
 - **Generate SSH key pair** (Default option) Select this radio button to generate an SSH keypair. Then in the dialog below click **Save private key** to download the key, and optionally click **Save public key** to download the key.
 - **Upload SSH key files:** Select this radio button to browse or drag and drop .pub files.
 - **Paste SSH keys:** Select this radio button to paste in individual public keys. To paste multiple keys, click **+ Another SSH Key**, and supply a single key for each entry.
12. **Network settings:** Specify the following:

Note

IP addresses (100.64.0.0/10) are used for Exadata Cloud Infrastructure X8M interconnect

You do not have the option to choose between IPv4 (single stack) and IPv4/IPv6 (dual stack) if both configurations exist. For more information, see [VCN and Subnet Management](#).

- **Virtual cloud network:** The VCN in which you want to create the VM cluster. Click **Change Compartment** to select a VCN in a different compartment.
- **Client subnet:** The subnet to which the VM cluster should attach. Click **Change Compartment** to select a subnet in a different compartment. Do not use a subnet that overlaps with 192.168.16.16/28, which is used by the Oracle Clusterware private interconnect on the database instance. Specifying an overlapping subnet causes the private interconnect to malfunction.
- **Backup subnet:** The subnet to use for the backup network, which is typically used to transport backup information to and from the **Backup Destination**, and for Data Guard replication. Click **Change Compartment** to select a subnet in a different compartment, if applicable. Do not use a subnet that overlaps with 192.168.128.0/20. This restriction applies to both the client subnet and backup subnet.

If you plan to back up databases to Object Storage or Autonomous Recovery service, see the network prerequisites in [Managing Exadata Database Backups](#).

Note

In case Autonomous Recovery Service is used, a new dedicated subnet is highly recommended. Review the network requirements and configurations required to backup your Oracle Cloud databases to Recovery Service. See, [Configuring Network Resources for Recovery Service](#).

- **Network Security Groups:** Optionally, you can specify one or more network security groups (NSGs) for both the client and backup networks. NSGs function as virtual firewalls, allowing you to apply a set of ingress and egress **security rules** to your Exadata Cloud Infrastructure VM cluster. A maximum of five NSGs can be specified.

For more information, see **Network Security Groups** and *Network Setup for Exadata Cloud Infrastructure Instances*.

Note that if you choose a subnet with a **security list**, the security rules for the VM cluster will be a union of the rules in the security list and the NSGs.

To use network security groups:

- Check the **Use network security groups to control traffic** check box. This box appears under both the selector for the client subnet and the backup subnet. You can apply NSGs to either the client or the backup network, or to both networks. Note that you must have a virtual cloud network selected to be able to assign NSGs to a network.
- Specify the NSG to use with the network. You might need to use more than one NSG. If you're not sure, contact your network administrator.
- To use additional NSGs with the network, click **+;Another Network Security Group**.

Note

To provide your cloud VM Cluster resources with additional security, you can use Oracle Cloud Infrastructure Zero Trust Packet Routing to ensure that only resources identified with security attributes have network permissions to access your resources. Oracle provides Database policy templates that you can use to assist you with creating policies for common database security use cases. To configure it now, you must already have created security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. Click **Show Advanced Options** at the end of this procedure.

Be aware that when you provide security attributes for a cluster, as soon as it is applied, all resources require a Zero Trust Packet policy to access the cluster. If there is a security attribute on an endpoint, then it must satisfy both network security group (NSG) and Oracle Cloud Infrastructure Zero Trust Packet Routing policy (OCI ZPR) rules.

- **To use private DNS Service**

Note

A Private DNS must be configured before it can be selected. See *Configure Private DNS*

- Check the **Use private DNS Service** check box.
- Select a private view. Click **Change Compartment** to select a private view in a different compartment.
- Select a private zone. Click **Change Compartment** to select a private zone in a different compartment.
- **Hostname prefix:** Your choice of hostname for the Exadata VM cluster. The host name must begin with an alphabetic character and can contain only alphanumeric characters and hyphens (-). The maximum number of characters allowed for an Exadata VM cluster is 12.

 **Caution**

The hostname must be unique within the subnet. If it is not unique, the VM cluster will fail to provision.

- **Host domain name:** The domain name for the VM cluster. If the selected subnet uses the Oracle-provided Internet and VCN Resolver for DNS name resolution, this field displays the domain name for the subnet and it can't be changed. Otherwise, you can provide your choice of the domain name. Hyphens (-) are not permitted. If you plan to store database backups in Object Storage or Autonomous Recovery service, Oracle recommends that you use a VCN Resolver for DNS name resolution for the client subnet because it automatically resolves the Swift endpoints used for backups.
 - **Host and domain URL:** This read-only field combines the host and domain names to display the fully qualified domain name (FQDN) for the database. The maximum length is 63 characters.
13. **Choose a license type:** The type of license you want to use for the VM cluster. Your choice affects metering for billing.
- **License Included** means the cost of the cloud service includes a license for the Database service.
 - **Bring Your Own License (BYOL)** means you are an Oracle Database customer with an Unlimited License Agreement or Non-Unlimited License Agreement and want to use your license with Oracle Cloud Infrastructure. This removes the need for separate on-premises licenses and cloud licenses.
14. **Diagnostics Collection:** By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

 **Note**

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt out of this feature at any time

- **Enable Diagnostic Events:** Allow Oracle to collect and publish critical, warning, error, and information events to me.
- **Enable Health Monitoring:** Allow Oracle to collect health metrics/events such as Oracle Database up/down, disk space usage, and so on, and share them with Oracle Cloud operations. You will also receive notification of some events.
- **Enable Incident Logs and Trace Collection:** Allow Oracle to collect incident logs and traces to enable fault diagnosis and issue resolution.

 **Note**

You are opting in with the understanding that the above list of events (or metrics, log files) can change in the future. You can opt-out of this feature at any time.

All three checkboxes are selected by default. You can leave the default settings as is or clear the check boxes as needed. You can view the Diagnostic Collection settings on the **VM Cluster Details** page under **General Information >> Diagnostics Collection**.

- **Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Disabled:** When you choose not to collect diagnostics, health metrics, incident logs, and trace files (all three options).
- **Partially Enabled:** When you choose to collect diagnostics, health metrics, incident logs, and trace files (one or two options).

15. Click **Show Advanced Options** to specify advanced options for the VM cluster:

- **Time zone:** This option is located in the **Management** tab. The default time zone for the VM cluster is UTC, but you can specify a different time zone. The time zone options are those supported in both the `Java.util.TimeZone` class and the Oracle Linux operating system.

Note

If you want to set a time zone other than UTC or the browser-detected time zone, and if you do not see the time zone you want, try selecting the **Select another time zone**, option, then selecting "Miscellaneous" in the **Region or country** list and searching the additional **Time zone** selections.

- **SCAN Listener Port:** This option is located in the **Network** tab. You can assign a SCAN listener port (TCP/IP) in the range between 1024 and 8999. The default is 1521.

Note

Manually changing the SCAN listener port of a VM cluster after provisioning using the backend software is not supported. This change can cause Data Guard provisioning to fail.

- **Zero Trust Packet Routing (ZPR):** This option is located in the **Security attributes** tab. Select a namespace, and provide the key and value for the security attribute. To complete this step during configuration, you must already have set up security attributes with Oracle Cloud Infrastructure Zero Trust Packet Routing. You can also add security attributes after configuration, and add them later. For more information about adding Oracle Exadata Database Service on Dedicated Infrastructure specific policies, see [Policy Template Builder](#).
- **Cloud Automation Update:** Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. You can configure your preferred time window for these updates to be applied to your VM Cluster. Set the start time for cloud automation updates.

Note

Oracle will check for latest VM Cloud Automation updates every day between the configured time window and apply updates when applicable. If automation is unable to start applying updates within the configured time window due to some underlying long running process, Oracle will automatically check the following day during the configured time window to start applying cloud automation updates to the VM Cluster.

Enable early access for cloud tools update: VM clusters designated for early access receive updates 1-2 weeks before they are available to other systems. Check this check box if you want early adoption for this VM cluster.

Cloud Automation Update Freeze Period: Oracle periodically applies updates to the database tools and agent software necessary for cloud tooling and automation. Enable a freeze period to define a time window during which Oracle automation will not apply cloud updates.

Move the slider to set the freeze period.

Note

- The freeze period can extend for a maximum of 45 days from the start date.
- Oracle automation will automatically apply updates with critical security fixes (CVSS \geq 9) even during a configured freeze period.

- **Tags:** 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.

16. Click Create.**WHAT NEXT?**

After your VM cluster is successfully created and in the **Available** state.

- You can view the VM Cluster Details page by clicking the name of the VM cluster in the list of clusters. From the VM Cluster Details page, you can **create your first database** in the cluster by clicking **Create Database**
- The **SCAN IP address (IPv4)** and **SCAN IP address (IPv6)** fields in the **Network** section on the VM Cluster Details page displays the dual stack IP address details.
- The **Cloud Automation Update** field in the **Version** section on the VM Cluster Details page displays the freeze period you have set.

Related Topics

- [Network Security Groups](#)
- [Network Setup for Exadata Cloud Infrastructure Instances](#)
This topic describes the recommended configuration for the VCN and several related requirements for the Exadata Cloud Infrastructure instance.

- [Security Lists](#)
- [Configure Private DNS](#)
Prerequisites needed to use Private DNS.
- [Resource Tags](#)
- [To create a database in an existing VM Cluster](#)
This topic covers creating your first or subsequent databases.
- [Oracle Cloud Infrastructure Zero Trust Packet Routing](#)
- [Getting Started with Events](#)
- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.
- [Overview of Automatic Diagnostic Collection](#)
By enabling diagnostics collection and notifications, Oracle Cloud Operations and you will be able to identify, investigate, track, and resolve guest VM issues quickly and effectively. Subscribe to Events to get notified about resource state changes.

Verify the Default Identity Connector Attached to the VM Cluster

To view the details of an identity connector attached to a VM cluster, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Click the name of the VM cluster of your choice.
4. On the resulting **VM Cluster Details** page, in the **Multicloud Information** section, confirm that the Identity connector field displays the identity connector attached to this VM cluster.
5. Click the name of the **Identity Connector** to view its details.
You will be redirected to the Database Multicloud Integrations portal.

Create a Key Ring in Google Cloud Console

To create a key ring, use this procedure.

1. Open the **Google Cloud Console**, navigate to the **Key Management** page.
2. Click **Create key ring**.
3. Provide the following details:
 - **Name:** Enter a descriptive name for the key ring.
 - **Location:** Select a location for your key ring.**Important:**
 - Key rings with the same name can exist in different locations, so you must always specify the location.
 - Choose a location close to the resources you want to protect.
 - For Customer Managed Encryption Keys, ensure the key ring is in the same location as the resources that will use it.

Choosing a location for your Key Ring:

When creating a key ring in Google Cloud Key Management Service (KMS), selecting the right location is crucial. Your choice affects where your cryptographic keys are stored and how they're replicated. For more information, see [Cloud KMS locations](#).

– **Region:**

- * Data is stored in a specific geographic region.
- * Keys remain within the boundaries of this single region.
- * Ideal for:
 - * Low-latency applications
 - * Compliance with data residency requirements
 - * Region-specific workloads

– **Multi-region:**

- * Data is replicated across multiple regions within a larger geographical area.
- * Google manages distribution and replication automatically.
- * You cannot select individual data centers or regions.
- * Ideal for:
 - * High availability
 - * Resilient, fault-tolerant applications
 - * Services serving a wide regional area

– **Global:**

- * A special type of multi-region.
- * Keys are distributed across Google data centers worldwide.
- * Location selection and control are not available.
- * Ideal for:
 - * Applications with global users
 - * Use cases needing maximum redundancy and reach

4. Click **Create**.

Once the key ring is created, you can begin creating and managing encryption keys within it.

Create a Key in Google Cloud Console

To create a raw symmetric encryption key in the specified key ring and location, use this procedure.

1. Open the **Google Cloud Console**, navigate to the **Key Management** page.
2. Click the name of the key ring where you want to create the key.
3. Click **Create key**.
4. Provide the following details:
 - **Key name:** Enter a descriptive name for your key.
 - **Protection level:** Choose **Software** or **HSM** (Hardware Security Module). The protection level of a key can't be changed after the key is created. For more information, see [Protection levels](#).

- **Key material:** Select **Generate key** or **Import key**.
Generate key material in Cloud KMS or import key material that is maintained outside of Google Cloud. For more information, see [Customer-managed encryption keys \(CMEK\)](#).
 - **Purpose and Algorithm:**
For more information, see [Key purposes and algorithms](#).
 - Set **Purpose** to **Raw encryption/decryption**.
 - For **Algorithm**, select **AES-256-CBC**.
5. Click **Create**.

After creation, you can use this key for cryptographic operations that require AES-CBC encryption and decryption.

Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure (OCI)

To allow a key to be discoverable in Oracle Cloud Infrastructure (OCI), use this procedure.

1. In **Google Cloud KMS**, select the key you want to make discoverable.
2. Navigate to the **Permissions** tab and click **Add principal**.
3. In the **New principals** field, enter the service account associated with your **Workload Resource Service Agent**.

Note

You can find this service account on the **Identity Connector details** page, under the **GCP Information** section. Look for the **Workload resource service agent** and note its ID — this is the required service account.

4. Under **Assign roles**, add a role of your choice.

Note

Create a custom role with the following minimum permissions and assign it to the key ring of your choice.

These permissions together allow OCI to:

- Discover KMS resources like key rings and keys.
- Access metadata about keys and their versions.
- Use the keys for cryptographic operations (encryption/decryption).
- Create key versions.

Minimum Required Permissions:

- `cloudkms.cryptoKeyVersions.get`
Allows retrieval of metadata for a specific key version.
- `cloudkms.cryptoKeyVersions.manageRawAesCbcKeys`
Enables management of raw AES-CBC key material (import, rotation, etc.).

- `cloudkms.cryptoKeyVersions.create`
Allows creation of new key versions within a key.
 - `cloudkms.cryptoKeyVersions.list`
Lists all versions of a given key.
 - `cloudkms.cryptoKeyVersions.useToDecrypt`
Grants permission to use a key version for decrypting data.
 - `cloudkms.cryptoKeyVersions.useToEncrypt`
Grants permission to use a key version for encrypting data.
 - `cloudkms.cryptoKeys.get`
Allows retrieval of metadata for a key.
 - `cloudkms.cryptoKeys.list`
Lists all keys within a key ring.
 - `cloudkms.keyRings.get`
Allows retrieval of metadata for a key ring.
 - `cloudkms.locations.get`
Retrieves information about supported key locations.
5. Click **Save** to apply the changes.
 6. Click **Refresh** to confirm that the updated permissions have taken effect.

Register GCP Key Ring in Oracle Cloud Infrastructure (OCI)

To enable Google Cloud Customer Managed Encryption Keys (CMEK) for your VM cluster, you must first register the GCP Key Ring in OCI.

Note

Before proceeding, ensure that the permissions outlined in [Grant Permissions in Google Cloud KMS for Key Discovery by Oracle Cloud Infrastructure \(OCI\)](#) have been granted.

1. In the **Database Multicloud Integrations** portal, navigate to: **Google Cloud Integration > GCP Key Rings**.
2. Click **GCP Key Ring**,
3. Click **Register GCP key rings**
4. On the resulting **Register GCP key rings** page, provide the following details:
 - **Compartment:** Select the compartment where the VM cluster resides.
 - **Identity Connector:** Choose the Identity Connector attached to the VM cluster.
 - **Key Ring:** Enter the name of the GCP key ring to register.
To discover all available key rings through a single identity connector, you must grant the following permissions to that identity connector. These permissions should be assigned at the appropriate project or folder level to ensure the connector can access all key rings across the intended scope.
 - `cloudkms.keyRings.list`
Allows listing all key rings within a project.
 - `cloudkms.locations.get`

Allows retrieving metadata for a specific key ring.

5. Click **Discover** to verify if the key ring exists in GCP. If successful, the key ring's details will be displayed.

Note

Only key rings can be registered — not individual keys. All supported keys associated with a registered key ring will be available, provided the required permissions are in place.

6. Click **Register**.

Enable or Disable Google Cloud Key Management

To enable GCP CMEK for your Exadata VM Cluster, use this procedure.

Note

When you provision an Exadata VM Cluster, GCP CMEK is disabled by default.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Under **Oracle Exadata Database Service on Dedicated Infrastructure**, click **Exadata VM Clusters**.
3. Select the name of the **VM cluster** you want to configure.
4. On the **VM Cluster Details** page, scroll to the **Multicloud Information** section and click **Enable** next to GCP CMEK.
5. To disable GCP CMEK, click **Disable**.

Create a Database and Use GCP Customer-Managed Encryption Key (CMEK) as the Key Management Solution

This topic describes only the steps for creating a database and using GCP Customer-managed encryption key (CMEK) as the key management solution.

For the generic database creation procedure, see [To create a database in an existing VM Cluster](#).

Prerequisites

- Enable Google Cloud Key Management at the VM cluster level.
- Register the GCP key rings in OCI.

Steps

If Google Cloud Key Management is enabled at the VM cluster, you will have two key management options: **Oracle Wallet** and **GCP Customer Managed Encryption Key**.

1. In the Encryption section, choose **GCP Customer Managed Encryption Key**.

2. Select a registered Key ring available in your compartment. **Note**
 - Only registered key rings are listed.
 - If your desired key ring is not visible, it may not have been registered yet. Click **Register Key Rings** to discover and register it. For detailed instructions, refer to [Register GCP Key Ring in Oracle Cloud Infrastructure \(OCI\)](#).
3. Select the key within the selected key ring in your compartment.

Change the Key Management from Oracle Wallet to GCP Customer Managed Encryption Key (CMEK)

To change encryption keys between different encryption methods, use this procedure.

Note

- You cannot migrate from GCP Customer Managed Encryption Key to Oracle Wallet.
- Your database will experience a brief downtime while the key management configuration is being updated.

1. Navigate to your database details page in the OCI console.
2. In the **Encryption** section, verify that **Key management** is set to **Oracle Wallet**, and then click the **Change** link.
3. Enter the following information on the **Change key management** page.
 - a. Select your **Key management** as **GCP Customer Managed Encryption Key** from the drop-down list.
 - b. Select the compartment you are using, and then choose the Key Ring available in that compartment.
 - c. Next, select the Key compartment you are using, and then choose the desired Key from the drop-down list.
 - d. Click **Save changes**.

Rotate the GCP Customer Managed Encryption Key of a Container Database (CDB)

To rotate the GCP Customer Managed Encryption Key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM Clusters, click the name of the VM cluster that contains the database that you want to rotate encryption keys.
4. Click **Databases**.
5. Click the name of the database that you want to rotate encryption keys.
The Database Details page displays information about the selected database.

6. In the **Encryption** section, verify that the **Key Management** is set to **GCP Customer Managed Encryption Key**, and then click the **Rotate** link.
7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Rotate the GCP Customer Managed Encryption Key of a Pluggable Database (PDB)

To rotate the GCP Customer Managed Encryption Key of a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the PDB you want to start, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click **Pluggable Databases** in the **Resources** section of the page.
A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as GCP Customer Managed Encryption Key.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Using the API to Manage GCP KMS Integration for Exadata Database Service on Oracle Database@Google Cloud

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

The following resources will be made available to customers through OCI SDK, CLI, and Terraform. These APIs will be used by customers who wish to integrate Oracle Database on Exadata with Google Cloud Services.

Table 5-11 OracleDbGcpIdentityConnectors

API	Description
ListOracleDbGcpIdentityConnectors	Lists all GCP Identity Connector resources based on the specified filters.
GetOracleDbGcpIdentityConnector	Retrieves detailed information about a specific GCP Identity Connector resource.
CreateOracleDbGcpIdentityConnector	Creates a new GCP Identity Connector resource for the specified ExaDB-D VM Cluster.

Table 5-11 (Cont.) OracleDbGcpIdentityConnectors

API	Description
UpdateOracleDbGcpIdentityConnector	Updates the configuration details of an existing GCP Identity Connector resource.
ChangeOracleDbGcpIdentityConnectorCompartment	Moves the GCP Identity Connector resource to a different compartment.
DeleteOracleDbGcpIdentityConnector	Deletes the specified GCP Identity Connector resource.

Table 5-12 OracleDbGcpKeyRings

API	Description
ListOracleDbGcpKeyRings	Lists all GCP Key Ring resources based on the specified filters.
CreateOracleDbGcpKeyRing	Creates a new GCP Key Ring resource.
ChangeOracleDbGcpKeyRingCompartment	Moves the GCP Key Ring resource to a different compartment.
RefreshOracleDbGcpKeyRing	Refreshes the details of a GCP Key Ring resource.
GetOracleDbGcpKeyRing	Retrieves detailed information about a specific GCP Key Ring resource.
UpdateOracleDbGcpKeyRing	Updates the configuration details of an existing GCP Key Ring resource.
DeleteOracleDbGcpKeyRing	Deletes the specified GCP Key Ring resource.

Table 5-13 OracleDbGcpKeyKeys

API	Description
ListOracleDbGcpKeys	Lists all GCP Key Ring resources based on the specified filters.
GetOracleDbGcpKey	Retrieves detailed information about a specific GCP Key resource.

Updating the Multicloud PKCS#11 Driver

Note

The PKCS#11 driver for Google Cloud is installed automatically when Google Cloud Key Management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:

```
rpm -qa | grep pkcs
```

This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.

Example output:

```
pkcs-multicloud-driver-gcp-0.1-250723.0511.x86_64
```

This confirms that the `pkcs-multicloud-driver-gcp` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one VM at a time):

1. Shutdown the database instances that are using the Google Cloud KMS. All databases that use the Google Cloud KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

```
srvctl stop instance -d <db_unique_name> -i <instance_name> [-o <stop_option>]
```

- `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
- `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

For details on additional options, refer to the SRVCTL command-line help or reference manual.

2. Verify that the databases have been shut down. Ensure that all database instances that are using the Google Cloud KMS have been stopped before proceeding.

3. Upgrade the RPM.

As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider  
GOOGLE --nodeList <node> --databasesStopped
```

This command downloads and installs the latest RPM on the specified node.

4. Verify driver installation.

Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Start the database instances.

After the driver update, restart the database instances using:

```
<oracle_home>/bin/srvctl start instance -node "<node>"
```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

1. Stop databases using the Google Cloud KMS.

Stop all databases that use the Google Cloud KMS for master encryption keys (MEK) before updating the driver:

```
dbaascli database stop [--dbname <value>]
```

2. Verify that the databases have been shut down.

Ensure that all databases using the Google Cloud KMS have been stopped before proceeding.

3. Update the RPM.

As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider  
GOOGLE --databasesStopped
```

This command downloads and installs the latest RPM.

4. Verify driver installation.

Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Restart the databases

Once the driver has been updated, restart the databases:

```
dbaascli database start [--dbname <value>]
```

Release Notes

Review the changes made across the different releases of the GCP driver.

- [Release 2.185 \(251201.0551\)](#)

Release 2.185 (251201.0551)

- Various bug fixes and stability improvements.

AWS Key Management Service Integration for Exadata Database Service on Oracle Database@AWS

Exadata Database Service on Oracle Database@AWS supports integration with AWS Key Management Service (KMS). This enhancement allows users to manage Transparent Data Encryption (TDE) master encryption keys (MEKs) using AWS customer managed keys.

For Exadata Database Service on Oracle Database@AWS TDE MEKs can be stored in a file-based Oracle Wallet, Oracle Cloud Infrastructure (OCI) Vault, Oracle Key Vault (OKV), or AWS KMS, providing options to align with organization-specific security policies. Integration with AWS KMS enables applications, AWS services, and databases on Exadata VM Clusters to leverage a single centralized key management solution.

- [Prerequisites](#)
Before using AWS KMS as the key management solution for your databases, you must complete the following steps.
- [Using the Console to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS](#)
Learn how to manage AWS KMS integration for Exadata Database Service on Oracle Database@AWS.
- [Using the API to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS](#)
- [Updating the Multicloud PKCS#11 Driver](#)
- [Release Notes](#)
Review the changes made across the different releases of the AWS driver.

Prerequisites

Before using AWS KMS as the key management solution for your databases, you must complete the following steps.

- [Configure OCI Identity Domain](#)
- [Create or Use an Existing ODB Network](#)
- [Create an Exadata VM Cluster](#)
- [Configure Identity Provider](#)
- [Associate an IAM Role to an Exadata VM Cluster](#)
- [Create a Key](#)
- [Register AWS KMS Key](#)
- [OCI IAM Policy Requirements](#)

Configure OCI Identity Domain

OCI Identity domain is automatically configured during the Oracle Database@AWS onboarding process, and no action is required. Complete the following steps only for accounts that were linked before the general availability of AWS KMS integration (November 18, 2025).

Configure an OCI identity domain to enable AWS integration for your Exadata VM Clusters. It allows you to associate an AWS Identity and Access Management (IAM) service role with AWS integrations.

1. From the AWS console, select **Oracle Database@AWS**, and then select **Settings**.
2. Select the **Configure** button to configure your **OCI identity domain**.
3. Once it is complete, you can review the **Status**, **OCI identity domain ID** and **OCI identity domain URL** information from the **Settings** page.

Create or Use an Existing ODB Network

1. If you do not have an existing ODB Network, you can provision one by following the [ODB Network](#) step-by-step instructions. If you have an existing ODB Network, you can modify it by selecting the **Modify** button, follow [Modify an ODB Network](#) for step-by-step instructions.
2. From the **Configure service integrations** section,
 - a. Select the **Security Token Service (STS)** option to setup the networking for AWS KMS and AWS STS access from the database.
 - b. Select the **AWS KMS** checkbox to enable AWS KMS to use KMS keys in your authentication polices.

Create an Exadata VM Cluster

Exadata VM Cluster creation is only available through the AWS Console and AWS CLI. For more information, see [Exadata VM Cluster](#).

Configure Identity Provider

Create a stack for each cluster. Use the following steps to create a CloudFormation stack. This needs to be performed for each Exadata VM cluster.

When your CloudFormation stack creates an OIDC (OpenID Connect) Provider, it establishes a trust relationship so that an external identity source (like OCI) can authenticate to AWS. The associated IAM Role defines what that trusted identity is allowed to do.

1. Select your existing Exadata VM Cluster from the list to open its details page in the AWS Management Console.

Note

If you do not have an existing Exadata VM Cluster, you can provision one by following the [Exadata VM Cluster](#) step-by-step instructions.

2. Select the **IAM service roles** tab and then click the **CloudFormation** link.
3. From the Quick create stack page, review the pre-filled information.

- a. Under the **Parameters** section, parameters are defined in your template and allow you to enter custom values when you create or update a stack.
 - i. Review your pre-filled **OCIIdentityDomainUrl** information.
 - ii. The **OIDCProviderArn** field is optional. When you create the CloudFormation stack for the first time, it creates an OIDC Provider and an associated IAM Role. If you are executing the stack for the first time, you do not need to provide a value for the **OIDCProviderArn** field.
 - iii. For any additional execution time, you must provide the ARN of the existing OIDC Provider and specify it in the **OIDCProviderArn** field. To obtain your **OIDCProviderArn** information, complete the following steps:
 - i. From the AWS console, navigate to **IAM**, and then select **Identity providers**.
 - ii. From the **Identity providers** list, you can filter your provider selecting **Type** as **OpenID Connect**.
 - iii. Select the **Provider** link that was created when the CloudFormation stack was initially created.
 - iv. From the **Summary** page, copy the **ARN** information.
 - v. Paste the **ARN** information into the **OIDCProviderArn** field.
 - b. From the **Permissions** section, select IAM role name from the drop-down list, and then select the IAM role for CloudFormation to use for all operations performed on the stack.
 - c. Select the **Create stack** button to apply the changes.
4. After the CloudFormation stack deployment is complete, navigate to the **Resources** section of the stack and verify the created IdP and IAM Role resources. Copy the IAM Role ARN for future use.

Associate an IAM Role to an Exadata VM Cluster

You must attach an IAM role that you created previously to the Exadata VM Cluster to assign an identity connector that enables access to AWS resources.

1. From the AWS console, select Oracle Database@AWS.
2. From the left menu, select **Exadata VM clusters**, and then select your Exadata VM Cluster from the list.
3. Select the **IAM roles** tab, and then select the **Associate** button.
 - a. The **AWS integration** field is read-only.
 - b. Enter the Amazon Resource Name (ARN) of the IAM role you want to associate with the VM cluster in the **Role ARN** field. You can obtain the **ARN** information from the **Summary** section of the role that you previously created.
 - c. Select the **Associate** button to attach the role.

Note

Once you associate an IAM role to your VM cluster, an identity connector gets attached to your Exadata VM Cluster.

Once you complete the prerequisites section, continue with the following actions:

1. From the **AWS Console**, create a customer-managed key in AWS KMS.
2. From the **OCI Console**, enable Exadata VM Cluster(s) and databases to utilize the AWS KMS key created in the step 1.

Create a Key

1. From the AWS console, select **Key Management Service (KMS)**.
2. From the left menu, select **Customer managed keys**, and then select the **Create key** button.
3. In the **Configure key** section, enter the following information.
 - a. Choose the **Symmetric** option as the **Key type**.
 - b. Choose the **Encrypt and decrypt** option as the **Key usage**.
 - c. Expand the **Advanced options** section. Both **AWS KMS** and **AWS CloudHSM** are supported.
 - i. If you want to use KMS standard key store, choose the **KMS - recommended** option as **Key material origin**, and then choose either the **Single-Region key** option or the **Multi-Region key** option from the **Regionality** section.

Note

Cross-region Data Guard and restoring databases to a different region are currently not supported for databases that use AWS customer-managed keys for key management.

- ii. If you want to use **AWS CloudHSM**, choose the **AWS CloudHSM key store** option as **Key material origin**.
- d. Select the **Next** button to continue the creation process.
4. In the **Add labels** section, enter the following information.
 - a. Enter a descriptive display name in the **Alias** field. Maximum 256 characters. Use alphanumeric and '_'-' characters.

Note

- The alias name cannot begin with `aws/`. The `aws/` prefix is reserved by AWS to represent AWS managed keys in your account.
- An alias is a display name that you can use to identify the KMS key. We recommend that you choose an alias that indicates the type of data you plan to protect or the application you plan to use with the KMS key
- Aliases are required when you create a KMS key in the AWS console.

- b. The **Description** field is optional.
- c. The **Tags** section is optional. You can use tags to categorize and identify your KMS keys and help you track your AWS costs.
- d. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.

5. In the **Define key administrative permissions** section, complete the following substeps.
 - a. Search the role that you previously created, and then select the checkbox. Select the IAM users and roles that can administer the KMS key.

Note

This key policy gives the AWS account full control of this KMS key. It allows account administrators to use IAM policies to give other principals permission to manage the KMS key.

- b. From the **Key deletion** section, the **Allow key administrators to delete this key** checkbox is selected by default. To prevent the selected IAM users and roles from deleting the KMS key, you can deselect the checkbox.
 - c. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.
6. In the **Define key usage permissions** section, complete the following substeps.
 - a. Search the role that you previously created, and then select the check box.
 - b. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.
7. In the **Edit key policy** section, complete the following substeps.
 - a. From the **Preview** section, you can review the key policy. If you want to make a change, select the **Edit** tab.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "KMSKeyMetadata",
      "Effect": "Allow",
      "Principal": {
        "AWS": "<arn>"
      },
      "Action": [
        "kms:DescribeKey"
      ],
      "Resource": "*"
    },
    {
      "Sid": "KeyUsage",
      "Effect": "Allow",
      "Principal": {
        "AWS": "<arn>"
      },
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt"
      ],
      "Resource": "*"
    }
  ]
}
```

```
    ]
  }
```

- b. Select the **Next** button to continue the creation process or the **Previous** button to return the previous page.
8. In the **Review** section, review your information and then select the **Finish** button.

For more information, see [Create a symmetric encryption KMS key](#).

Register AWS KMS Key

To enable AWS KMS for your Exadata VM Cluster(s), you must first register the AWS KMS key in OCI.

1. From the OCI console, select **Oracle AI Database**, and then select **Database Multicloud Integrations**.
2. From the left menu, select **AWS Integration**, and then select **AWS Keys**.
3. Select the **Register AWS keys** button, and then complete the following substeps.
 - a. From the dropdown list, select the **Compartment** where your Exadata VM Cluster resides.
 - b. Under the **AWS keys** section, select your identity connector from the dropdown list.

Note

Ensure that the role associated with the connector has the **DescribeKey** permission on the key. This permission is required to successfully perform discovery.

- c. The **Key ARN** field is optional.
- d. Click the **Discover** button.
4. Once the key is discovered, select the **Register** button to register the key in OCI.

OCI IAM Policy Requirements

```
Allow any-user to read oracle-db-aws-keys in compartment id <your-compartment-OCID>
where all { request.principal.type = 'cloudvmcluster' }
```

This policy allows Oracle-managed Cloud VM Clusters to read the `oracle-db-aws-keys` resource in your compartment. It grants the necessary permissions for the VM cluster (the principal of type `cloudvmcluster`) to access the AWS key metadata required for integrating with external key management or performing cross-cloud operations. Without this policy, the VM cluster would not be able to retrieve the keys needed to complete the configuration.

Using the Console to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS

Learn how to manage AWS KMS integration for Exadata Database Service on Oracle Database@AWS.

- [Enable or Disable AWS Key Management](#)
- [Create a Database and Use AWS KMS as the Key Management Solution](#)
- [Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS](#)
- [Rotate the AWS KMS Key of a Container Database \(CDB\)](#)
- [Rotate the AWS KMS Key of a Pluggable Database \(PDB\)](#)
- [\(Optional\) Disable a CDB or PDB Key](#)

Enable or Disable AWS Key Management

When you enable AWS key management for your database, only AWS keys that are authorized for use with Exadata VM Cluster and registered with OCI can be used.

1. From the OCI console, select **Oracle AI Database** and then select **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. From the left menu, select **Exadata VM Clusters**, and then select your Exadata VM Cluster.
3. Select the **VM Cluster information** tab, and select the **Enable** button next to **AWS Customer Managed Encryption Key**.

Note

Once you enable AWS Key Management for your Exadata VM Cluster, you can disable it using the **Disable** button. Disabling this feature will affect the availability of databases that use AWS Key Management Service for encryption and decryption operations. Ensure that no database is currently using AWS key management before disabling it at Exadata VM Cluster level.

Create a Database and Use AWS KMS as the Key Management Solution

This topic describes only the steps for creating a database and using AWS KMS as the key management solution.

For the generic database creation procedure, see [To create a database in an existing VM Cluster](#).

Prerequisites

Refer to [Prerequisites](#).

Steps

If AWS KMS key management is enabled at the VM cluster level, you will have two key management options: **Oracle Wallet** and **AWS key management**.

1. In the **Encryption** section, choose **AWS key management**.
2. Select the encryption Key available in your compartment.

Note

- Only registered keys are listed.
- If your desired key is not visible, it may not have been registered yet. Click **Register AWS Keys** to discover and register it. For detailed instructions, refer to [Register AWS KMS Key](#).
- You can select a key alias from the drop-down list. If no key alias is available, the drop-down list will display the key ID.

Steps to Change the Key Management of Existing Databases from Oracle Wallet to AWS KMS

If you want to change the key management of your existing databases in the Exadata VM Cluster from Oracle Wallet to AWS KMS, complete the following steps.

1. From your **Exadata VM Clusters**, navigate to **Databases** tab, and then select the database that you are using.
2. From the **Encryption** section, confirm that **Key management** is set to **Oracle Wallet**, and then select the **Change** link.
3. From the **Change key management** page, enter the following information.
 - a. Select your **Key management** as **AWS key management** from the dropdown list.
 - b. Select the key compartment you are using, and then select the desired key from the dropdown list.
 - c. Select the **Save changes** button.

Rotate the AWS KMS Key of a Container Database (CDB)

To rotate the AWS KMS Key of a container database (CDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM Clusters, click the name of the VM cluster that contains the database that you want to rotate encryption keys.
4. Click **Databases**.
5. Click the name of the database that you want to rotate encryption keys.
The Database Details page displays information about the selected database.
6. In the **Encryption** section, verify that the **Key Management** is set to **AWS key management**, and then click the **Rotate** link.

7. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Note

Rotating the AWS KMS key generates a new encryption context for the same key.

Rotate the AWS KMS Key of a Pluggable Database (PDB)

To rotate the AWS KMS Key of a pluggable database (PDB), use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of VM Clusters is displayed for the chosen Compartment.
3. In the list of VM clusters, click the name of the VM cluster that contains the PDB you want to start, and then click its name to display the details page.
4. Under **Databases**, find the database containing the PDB you want to rotate encryption keys.
5. Click the name of the database to view the Database Details page.
6. Click **Pluggable Databases** in the **Resources** section of the page.
A list of existing PDBs in this database is displayed.
7. Click the name of the PDB that you want to rotate encryption keys.
The pluggable details page is displayed.
8. In the **Encryption** section displays that the Key management is set as AWS key management.
9. Click the **Rotate** link.
10. On the resulting **Rotate Key** dialog, click **Rotate** to confirm the action.

Note

Rotating the AWS KMS key generates a new encryption context for the same key.

(Optional) Disable a CDB or PDB Key

Optionally, complete the following steps to disable a specific CDB or PDB key.

1. Navigate to the AWS console, select **Key Management Service (KMS)**.
2. From the left menu, select **Customer managed keys**, and then select the **Key ID** of the key that you want to edit.
3. Select the **Edit policy** button.
4. Run the following query to obtain the **EncryptionContext MKID** of the key.

Use the following command to view the current master key IDs (MKIDs) for enabling or disabling operations:

```
dbaascli tde getHSMKeys --dbname <DB-Name>
```

You can also run the following SQL query to list all key versions:

```
SELECT key_id, con_id, creation_time, key_use FROM v$encryption_keys;
```

5. Add a deny policy using the retrieved **EncryptionContext** MKID as shown below:

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt"
  ],
  "Resource": "arn:aws:kms:us-east-1:867344470629:key/7139075d-
a006-4302-92d5-48ecca31d48e",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:MKID":
"ORACLE.TDE.HSM.MK.06AE5EFB528D9D4F21BFEDA63C6C8738D9"
    }
  }
}
```

6. Choose **Save changes** to apply the policy update.

Using the API to Manage AWS KMS Integration for Exadata Database Service on Oracle Database@AWS

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

The following resources will be made available to customers through OCI SDK, CLI, and Terraform. These APIs will be used by customers who wish to integrate Oracle Database on Exadata with AWS KMS.

- [OracleDbAwsIdentityConnector](#)
- [OracleDbAwsKey](#)

OracleDbAwsIdentityConnector

Table 5-14 OracleDbAwsIdentityConnector

API	Description
ListOracleDbAwsIdentityConnectors	Lists all AWS Identity Connector resources based on the specified filters.

Table 5-14 (Cont.) OracleDbAwsIdentityConnector

API	Description
GetOracleDbAwsIdentityConnector	Retrieves detailed information about a specific AWS Identity Connector resource.
CreateOracleDbAwsIdentityConnector	Creates a new AWS Identity Connector resource for the specified ExaDB-D VM Cluster.
UpdateOracleDbAwsIdentityConnector	Updates the configuration details of an existing AWS Identity Connector resource.
ChangeOracleDbAwsIdentityConnectorCompartment	Moves the AWS Identity Connector resource to a different compartment.
DeleteOracleDbAwsIdentityConnector	Deletes the specified AWS Identity Connector resource.
RefreshOracleDbAwsIdentityConnector	Refreshes the configuration details of the specified AWS Identity Connector resource.

OracleDbAwsKey

Table 5-15 OracleDbAwsKey

API	Description
ListOracleDbAwsKeys	Lists all AWS Key resources based on the specified filters.
CreateOracleDbAwsKey	Creates a new AWS Key resource.
ChangeOracleDbAwsKeyCompartment	Moves the AWS Key resource to a different compartment.
GetOracleDbAwsKey	Retrieves detailed information about a specific AWS Key resource.
UpdateOracleDbAwsKey	Updates the configuration details of an existing AWS Key resource.
DeleteOracleDbAwsKey	Deletes the specified AWS Key resource.
RefreshOracleDbAwsKey	Refreshes the configuration details of a AWS Key resource.

Updating the Multicloud PKCS#11 Driver

Note

The PKCS#11 driver for AWS is installed automatically when AWS key management is enabled on a VM cluster. To verify that the driver is installed and check its version, run the following command:

```
rpm -qa | grep pkcs
```

This command lists the PKCS package installed on the VM cluster. Only one PKCS#11 driver can be active at a time, corresponding to either Azure, Google Cloud, or AWS.

Example output:

```
pkcs-multicloud-driver-maz-0.2-250814.1708.x86_64
```

This confirms that the `pkcs-multicloud-driver-maz` package is installed, along with its version and architecture.

The PKCS#11 driver can be updated using either a rolling or non-rolling mechanism, depending on customer preference.

In rolling mode, the update is applied to one VM at a time, and the databases on that VM are restarted before proceeding to the next. In non-rolling mode, all databases in the cluster are shut down, the update is applied across all nodes, and the databases are then brought back online. Users must be aware of the `dbaastools` version in use, as subsequent steps, specifically step 3 in both scenarios, require different actions depending on the version, particularly when updating the PKCS#11 driver.

The version of `dbaastools` can be determined by running the following command:

```
/var/opt/oracle/dbaascli/dbaascli admin showLatestStackVersion
```

Rolling update

On each Guest VM, apply the PKCS#11 driver update as follows (one node at a time):

1. Shutdown the database instances that are using the AWS KMS. All databases that use the AWS KMS for master encryption keys (MEK) must be stopped before updating the driver. Use:

```
srvctl stop instance -d <db_unique_name> -i <instance_name> [-o <stop_option>]
```

- `-d <db_unique_name>`: The unique name of the database (same as `DB_UNIQUE_NAME` in the database).
- `-i <instance_name>`: The name of the instance to stop, for example, `orcl1`.

For details on additional options, refer to the SRVCTL command-line help or reference manual.

2. Verify that the databases have been shut down.

Ensure that all database instances that are using the AWS KMS have been stopped before proceeding.

3. Update the RPM.

As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider AWS
--nodeList <node> --databasesStopped
```

This command downloads and installs the latest RPM on the specified node.

4. Verify driver installation.

Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Start the database instances.

After the driver update, restart the database instances using:

```
<oracle_home>/bin/srvctl start instance -node "<node>"
```

Non-Rolling update

In non-rolling mode, all databases on the cluster are stopped before applying the PKCS#11 driver update.

1. Stop databases using the AWS KMS.

Stop all databases that use the AWS KMS for master encryption keys (MEK) before updating the driver:

```
dbaascli database stop [--dbname <value>]
```

2. Verify that the databases have been shut down.

Ensure that all databases using the AWS KMS have been stopped before proceeding.

3. Upgrade the RPM.

As the `root` user, run the following command (for `dbaastools 25.4.1` and later):

```
/var/opt/oracle/dbaascli/dbaascli admin updateMCKMS --keystoreProvider AWS
--databasesStopped
```

This command downloads and installs the latest RPM.

4. Verify driver installation.

Run the following command as the `root` user:

```
/bin/rpm -qa | grep -i pkcs-multicloud-driver
```

5. Restart the databases.

Once the driver has been updated, restart the databases:

```
dbaascli database start [--dbname <value>]
```

Release Notes

Review the changes made across the different releases of the AWS driver.

- [Release 2.185 \(251201.0551\)](#)

Release 2.185 (251201.0551)

- Various bug fixes and stability improvements.

Database Multicloud Integration for Oracle Database Cloud Services

- [Using the Console to Manage Database Multicloud Integration for Oracle Database Cloud Services](#)
Learn how to manage database multicloud integration for Oracle Database Cloud Services.
- [Using the API to Manage Multi-Cloud Data Integration for Oracle Database Cloud Services](#)

Using the Console to Manage Database Multicloud Integration for Oracle Database Cloud Services

Learn how to manage database multicloud integration for Oracle Database Cloud Services.

- [Integrate Oracle Exadata with Azure Cloud services](#)
Oracle Database@Azure Identity Connectors, Storage Containers, and Storage Mounts are key components in the integration of Oracle Exadata with Azure Cloud services.
- [Create an Identity Connector from the OCI Console](#)
Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.
- [View the list of Identity Connector Resources](#)
- [View the details of an Identity Connector Resource](#)
- [Create a Storage Container Resource](#)
- [Discover Azure Storage Containers](#)
- [View the list of Storage Container Resources](#)
- [View the details of a Storage Container Resource](#)
- [Create a Storage Mount Resource](#)
- [View the list of Storage Mount Resources](#)
- [View the details of a Storage Mount Resource](#)

Integrate Oracle Exadata with Azure Cloud services

Oracle Database@Azure Identity Connectors, Storage Containers, and Storage Mounts are key components in the integration of Oracle Exadata with Azure Cloud services.

Oracle Database@Azure Identity Connectors: These connectors enable seamless interaction between Oracle Exadata Database and Azure Cloud services by linking the Oracle database with Azure identity and access management. It allows for Azure-specific authentication and authorization to be utilized directly within the Exadata VM Cluster.

Azure Storage Containers: Azure Blob Storage is used to store large amounts of unstructured data such as text, images, videos, and backups. By connecting Oracle Exadata to Azure Blob Storage, users can leverage cloud storage for scalable and secure data storage.

Azure Storage Mounts: Azure Blob Storage Mount refers to the process of attaching an Azure Storage Container as a mounted file system on an Oracle Exadata Database Service VM. This mount allows database utilities like Data Pump (impdp/expdp) and RMAN to interact with Azure storage as if it were a local directory.

Create an Identity Connector from the OCI Console

Creating an Identity Connector installs the Azure Arc agent on the Exadata VM Cluster VMs, registering them as Azure Arc-enabled virtual machines.

This enables secure communication with the Azure Key Management Service (KMS) using the Azure identity generated by the Arc agent. The Azure Arc agent can communicate with Azure services over either a public network or a private connectivity setup. Learn more about [Azure Arc](#).

Each Exadata VM cluster must have an identity connector enabled to access Azure resources. The identity connector establishes either a public or private connection between the Exadata VM cluster and Azure Key Management resources, depending on the roles assigned.

To generate an **access token** for your current Azure account, see [az account get-access-token](#).

You can create an identity connector in one of two ways—using the Oracle Exadata Database Service on Dedicated Infrastructure interface or the Database Multicloud Integrations interface.

Oracle Exadata Database Service on Dedicated Infrastructure

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. From the left menu, click **Exadata VM Clusters** under **Oracle Exadata Database Service on Dedicated Infrastructure**.
3. From the list of **Exadata VM Clusters**, select the cluster you are using.
4. Select **VM Cluster information**, and then navigate to **Identity connector** located under **Multicloud information**. Click the **Create** link.

Note

If an identity connector has not been created previously, it is displayed as **None**.

5. The **Identity connector name**, **Exadata VM cluster**, **Azure subscription id**, and **Azure resource group name** are read-only fields and will be populated with values.
6. Enter your **Azure tenant id**, and **Access token**.
7. Expand the **Show advanced options** section. The **Private connectivity information** and **Tags** sections populate.

To enable a private endpoint connection, enter the **Azure arc private link scope** name.

8. To add tags for your resources, click **Add tag**, and then enter required values.
9. Review your selections, and then click **Create** to create the identity connector.

Database Multicloud Integrations

1. Open the navigation menu. Click **Oracle Database**, then click **Database Multicloud Integrations**.
2. Select **Identity Connectors** from the left navigation menu.
3. From the **Compartment** drop-down list, select your compartment that you are using.
4. Once you select your compartment, the **Identity connector name** automatically populates a name.
By default, the identity connector type is selected as **Azure**.
5. Select **ARC agent** as an identity mechanism.
6. Select your compartment from the **Choose an Exadata VM cluster compartment** list, and then select your Exadata VM Cluster from the **Choose an Exadata VM cluster** list.
7. Enter your **Azure tenant id**. The **Azure subscription id** and **Azure resource group name** fields populate values based on your Exadata VM Cluster selection.
8. Enter an **Access token**.
9. Expand the **Show advanced options** section. The **Private connectivity information** and **Tags** sections populate. These fields are optional.
10. To add tags for your resources, click **Add tag**, and then enter required values.
11. Review your selections, and then click **Create**.

View the list of Identity Connector Resources

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Identity Connectors** to view the list of Identity Connectors.

View the details of an Identity Connector Resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Identity Connectors** to view the list of Identity Connectors.
3. Click the name of the identity connector to view its details.
The resulting details page displays information, including the Arc Agent status, indicating whether it is connected or disconnected.

Create a Storage Container Resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Containers**.
3. Click **Create Storage container**.
4. Enter the following on the resulting **Create storage container** page.

- **Compartment:** Select a compartment where you want to create this storage container.
- **Name:** Enter a descriptive name for the storage container.
- **Account name:** Enter the Azure storage container user name. This typically refers to the identity or credentials used to authenticate and access an Azure Storage Container. Depending on the authentication method, it could be an Azure Active Directory (Azure AD) identity or an access key-based authentication.
- **Storage container name:** Enter the name of the Azure storage container. This is the unique name assigned to a specific storage container within an Azure Storage Account. The container name helps organize and store blobs (files and data) within the account.
- **IP address:** Enter the IP address of the Azure Storage Container.
- **DNS Alias:** Enter the fully qualified domain name (FQDN) that points to the Azure Storage Container.

Discover Azure Storage Containers

To discover Azure storage container using an identity connector, use this procedure.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Containers**.
3. Click **Discover Azure storage containers**.
4. Enter the following on the resulting Discover Azure storage containers page.
 - **Compartment:** Select a compartment where you want to create this storage container.
 - **Azure storage containers:** Choose an Identity Connector.
5. Click **Discover**.
The list of Storage Containers name(s) is displayed.
6. Select the check box located next to Storage Container name.
7. Click **Create**.

View the list of Storage Container Resources

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Containers** to view the list of storage containers.

View the details of a Storage Container Resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Containers** to view the list of storage containers.
3. Click the name of the storage container to view it's details.

Create a Storage Mount Resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Mounts**.
3. Enter the following on the resulting **Create storage mount** page.
 - **Compartment**: Select a compartment where you want to create this storage mount.
 - **Storage mount name**: Enter a descriptive name for the storage mount.
 - **Identity connector**: Select an identity connector.
 - **Blob container**: Select a storage container.

View the list of Storage Mount Resources

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Mounts** to view the list of storage mounts.

View the details of a Storage Mount Resource

1. Open the navigation menu. Click **Oracle AI Database**, then click **Database Multicloud Integrations**.
2. Click **Storage Mounts** to view the list of storage mounts.
3. Click the name of the storage mount to view its details.

Using the API to Manage Multi-Cloud Data Integration for Oracle Database Cloud Services

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

The following resources will be made available to customers through OCI SDK, CLI, and Terraform. These APIs will be used by customers who wish to integrate Oracle Database on Exadata with Azure Cloud Services.

- [oracle-db-azure-connectors](#)
- [oracle-db-azure-blob-containers](#)
- [oracle-db-azure-blob-mounts](#)
- [oracle-db-mci-work-requests](#)
- [multi-cloud-resource-discoveries](#)
- [oracle-db-azure-vaults](#)
- [oracle-db-azure-keys](#)
- [oracle-db-azure-vault-associations](#)

oracle-db-azure-connectors

Table 5-16 oracle-db-azure-connectors

API	Description
<code>createOracleDbAzureConnector</code>	Captures Azure-specific details from the customer and automates the installation of the ARC Agent on the ExaDB-D VM Cluster.
<code>updateOracleDbAzureConnector</code>	Updates Azure-specific details from the customer and modifies the Arc Agent configuration if required.
<code>changeOracleDbAzureConnectorCompartment</code>	Changes the compartment of the Azure Connector resource but does not affect the Arc Agent installation.
<code>deleteOracleDbAzureConnector</code>	Deletes the Azure Connector resource and uninstalls the Arc Agent from the ExaDB-D VM Cluster.
<code>getOracleDbAzureConnector</code>	Fetches the details of a specific Azure Connector resource.
<code>listOracleDbAzureConnectors</code>	Lists Azure Connector resources based on the specified filters.

oracle-db-azure-blob-containers

Table 5-17 oracle-db-azure-blob-containers

API	Description
<code>createOracleDbAzureBlobContainer</code>	Capture the Azure Storage Account and Container details, which will be used when mounting the Azure Container in the ExaDB-C@C VM Cluster. Additionally, this operation will update the Azure Storage Account and Container details as needed.
<code>updateOracleDbAzureBlobContainer</code>	Updates the Azure Storage Account and Container details.
<code>changeOracleDbAzureBlobContainerCompartment</code>	Changes the compartment of the Azure Blob Container resource.
<code>deleteOracleDbAzureBlobContainer</code>	Deletes the Azure Blob Container Resource.
<code>getOracleDbAzureBlobContainer</code>	Retrieves a specific Azure Blob Container resource.
<code>listOracleDbAzureBlobContainers</code>	Lists Azure Blob Container resources based on the specified filters.

oracle-db-azure-blob-mounts

Table 5-18 oracle-db-azure-blob-mounts

API	Description
<code>createOracleDbAzureBlobMount</code>	Captures the Azure Storage Account and Container details and mounts the Azure Container on the ExaDB-D VM Cluster.

Table 5-18 (Cont.) oracle-db-azure-blob-mounts

API	Description
updateOracleDbAzureBlobMount	Updates the Azure Storage Account and Container details and, if required, updates the mount details of the Azure Container on the ExaDB-D VM Cluster.
changeOracleDbAzureBlobMountCompartment	Changes the compartment of the Azure Blob Mount Resource without making any changes to the ExaDB-C@C VM Cluster.
deleteOracleDbAzureBlobMount	Deletes the Azure Blob Mount Resource and unmounts the Azure Container from the ExaDB-D VM Cluster.
getOracleDbAzureBlobMount	Retrieves a specific Azure Blob Mount Resource.
listOracleDbAzureBlobMounts	Retrieves a list of Azure Blob Mount Resources.

oracle-db-mci-work-requests

Table 5-19 oracle-db-mci-work-requests

API	Description
ListWorkRequests	Retrieves a list of all work requests.
GetWorkRequest	Retrieves details of a specific work request.
CancelWorkRequest	Cancels the specified work request.
ListWorkRequestErrors	Retrieves a list of errors associated with a work request.
ListWorkRequestLogs	Retrieves a list of logs related to a work request.

multi-cloud-resource-discoveries

Table 5-20 multi-cloud-resource-discoveries

API	Description
ListMultiCloudResourceDiscoveries	Retrieves a list of all multi-cloud resource discovery resources.
CreateMultiCloudResourceDiscovery	Creates a new multi-cloud resource discovery resource.
ChangeMultiCloudResourceDiscoveryCompartment	Moves a multi-cloud resource discovery resource to a different compartment.
GetMultiCloudResourceDiscovery	Retrieves details of a specific multi-cloud resource discovery resource.
UpdateMultiCloudResourceDiscovery	Updates the details of a specific multi-cloud resource discovery resource.
DeleteMultiCloudResourceDiscovery	Deletes a specific multi-cloud resource discovery resource.

oracle-db-azure-vaults

Table 5-21 oracle-db-azure-vaults

API	Description
ListOracleDbAzureVaults	Retrieves a list of all Oracle DB Azure vault resources.
CreateOracleDbAzureVault	Creates a new Oracle DB Azure vault resource.
ChangeOracleDbAzureVaultCompartment	Moves an Oracle DB Azure vault resource to a different compartment.
RefreshOracleDbAzureVault	Refreshes the details of an Oracle DB Azure vault resource.
GetOracleDbAzureVault	Retrieves details of a specific Oracle DB Azure vault resource.
UpdateOracleDbAzureVault	Updates the details of a specific Oracle DB Azure vault resource.
DeleteOracleDbAzureVault	Deletes a specific Oracle DB Azure vault resource.

oracle-db-azure-keys

Table 5-22 oracle-db-azure-keys

API	Description
ListOracleDbAzureKeys	Retrieves a list of all Oracle DB Azure keys.
GetOracleDbAzureKey	Retrieves details of a specific Oracle DB Azure key.

oracle-db-azure-vault-associations

Table 5-23 oracle-db-azure-vault-associations

API	Description
ListOracleDbAzureVaultAssociations	Retrieves a list of all Oracle DB Azure vault associations.
CreateOracleDbAzureVaultAssociation	Creates a new association between an Oracle DB and an Azure vault.
ChangeOracleDbAzureVaultAssociationCompartment	Moves an Oracle DB Azure vault association to a different compartment.
GetOracleDbAzureVaultAssociation	Retrieves details of a specific Oracle DB Azure vault association.
UpdateOracleDbAzureVaultAssociation	Updates the details of a specific Oracle DB Azure vault association.
DeleteOracleDbAzureVaultAssociation	Deletes a specific Oracle DB Azure vault association.
CascadingDeleteOracleDbAzureVaultAssociation	Deletes an Oracle DB Azure vault association and any dependent resources.

6

Reference Guides for Exadata Cloud Infrastructure

- [Using the dbaascli Utility on Exadata Cloud Infrastructure](#)
Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.
- [Monitoring and Managing Exadata Storage Servers with ExaCLI](#)
The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.
- [Monitor Metrics for VM Cluster Resources](#)
- [Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service](#)
This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.
- [Oracle Exadata Database Service on Dedicated Infrastructure Events](#)
Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.
- [Policy Details for Exadata Cloud Infrastructure](#)
This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.
- [Managing Exadata Resources with Oracle Enterprise Manager Cloud Control](#)
To manage and monitor Exadata Cloud Infrastructure and Exadata Database Service on Cloud@Customer resources, use Oracle Enterprise Manager Cloud Control.
- [Observability and Management for Exadata Database Service on Dedicated Infrastructure](#)
- [Security Guide for Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [EU Data Act Compliance](#)
Oracle supports the goal of fair and transparent data sharing and access within the European Union.
- [Troubleshooting Exadata Cloud Infrastructure Systems](#)
These topics cover some common issues you might run into and how to address them.

Using the dbaascli Utility on Exadata Cloud Infrastructure

Learn to use the `dbaascli` utility on Exadata Cloud Infrastructure.

- [About Using the dbaascli Utility on Exadata Cloud Infrastructure](#)
You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.
- [Creating Databases Using dbaascli](#)
Using `dbaascli`, you can create an Oracle Database by first creating an Oracle Database home of desired version, followed by creating a database in that Oracle Database home

- [Changing the Database Passwords](#)
To change the SYS password, or to change the TDE wallet password, use this procedure.
- [Managing Exadata Cloud Infrastructure Software Images Using the Dbaascli Utility](#)
You can list and download the Oracle database software images on an Exadata Cloud Infrastructure instance, which can then be used for provisioning a database home.
- [Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli](#)
Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.
- [Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli](#)
Using the `dbaascli diag` command allows you to collect Guest VM `dbaas` tooling logs for Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Cloud@Customer systems. You can use these logs to troubleshoot issues related to `dbaas` tooling.
- [Updating Cloud Tooling Using dbaascli](#)
To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.
- [Creating a Duplicate Database](#)
- [Release Notes](#)
Review the changes made in various releases of `dbaascli`.
- [dbaascli Command Reference](#)
You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

About Using the dbaascli Utility on Exadata Cloud Infrastructure

You can use the `dbaascli` utility to perform various database lifecycle and administration operations on Exadata Cloud Infrastructure such as changing the password of a database user, starting a database, managing pluggable databases (PDBs), and more.

You must use the Oracle Cloud Infrastructure console or command-line interface to scale resources. The capabilities of the `dbaascli` utility are in addition to, and separate from, the Console, API, or command-line interface (CLI). Unless specified differently, you need `root` access to `dbaascli` to run all administration commands.

To use the utility, you must be connected to an Exadata Cloud Infrastructure virtual machine. For detailed instructions, see *Connecting to an Exadata Cloud Infrastructure Instance*.

To get possible commands available with `dbaascli`, run `dbaascli --help`.

To get command-specific help, run `dbaascli command --help`. For example, `dbaascli database create --help`.

See *dbaascli Command Reference* in the document for commands and command specific information.

Related Topics

- [Connecting to an Exadata Cloud Infrastructure Instance](#)
This topic explains how to connect to an Exadata Cloud Infrastructure instance using SSH or SQL Developer.
- [dbaascli Command Reference](#)
You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

Creating Databases Using dbaascli

Using `dbaascli`, you can create an Oracle Database by first creating an Oracle Database home of desired version, followed by creating a database in that Oracle Database home

- [Listing Available Software Images and Versions for Database and Grid Infrastructure](#)
To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- [Creating Oracle Database Home](#)
To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.
- [Creating Oracle Database In the Specified Oracle Database Home](#)
To create an Oracle Database in the specified Oracle Database home of desired version, use the `dbaascli database create` command.

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli cswlib showImages --product database
```

The command output lists the available database software images.

```
dbaascli cswlib showImages --product grid
```

The command output lists the available grid software images.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-1 dbaascli cswlib showImages

```
[root@dg1llrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
      AM_46941.log

#####
```

```
List of Available Database Images
```

```
#####
```

```
17.IMAGE_TAG=18.17.0.0.0  
   VERSION=18.17.0.0.0  
   DESCRIPTION=18c JAN 2022 DB Image
```

```
18.IMAGE_TAG=19.10.0.0.0  
   VERSION=19.10.0.0.0  
   DESCRIPTION=19c JAN 2021 DB Image
```

```
19.IMAGE_TAG=19.11.0.0.0  
   VERSION=19.11.0.0.0  
   DESCRIPTION=19c APR 2021 DB Image
```

```
20.IMAGE_TAG=19.12.0.0.0  
   VERSION=19.12.0.0.0  
   DESCRIPTION=19c JUL 2021 DB Image
```

```
21.IMAGE_TAG=19.13.0.0.0  
   VERSION=19.13.0.0.0  
   DESCRIPTION=19c OCT 2021 DB Image
```

Images can be downloaded using their image tags. For details, see help using 'dbaascli cswlib download --help'.
dbaascli execution completed

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli cswlib showImages](#)
To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Creating Oracle Database Home

To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Note

You can create an Oracle Database home with a specified Oracle home name. If you do not specify, then this is computed automatically (recommended).

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli dbhome create --version Oracle Home Version --imageTag image Tag Value
```

Where:

- `--version` specifies the Oracle Database version
- `--imageTag` specifies the Image Tag of the image to be used

For example:

```
dbaascli dbhome create --version 19.9.0.0.0
```

Note

Specifying `imageTag` is optional. To view the Image Tags, refer to command `dbaascli cswlib showImages`. Image Tags are typically same as the version of the database. However, it is kept as a provision for cases where multiple images may need to be released for the same version - each catering to a specific customer requirement.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli dbhome create`.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli dbhome create](#)
To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Creating Oracle Database In the Specified Oracle Database Home

To create an Oracle Database in the specified Oracle Database home of desired version, use the `dbaascli database create` command.

You can use the `dbaascli database create` command to:

- Create a Container Database (CDB) or non-Container Database
- Create a CDB with pluggable databases (PDBs)
- Create an Oracle Database with the specified Character Set
- Create Oracle Databases on a subset of cluster nodes

Note

Databases created on a subset of nodes will not be displayed in the OCI console.

- Create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.
- 1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- 2. Start a `root` user command shell:

```
sudo -s
```

- 3. Run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path
```

Where:

- `--dbName` specifies the name of the database
- `--oracleHome` specifies Oracle home location

To create a CDB, run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path
```

To create a non-CDB, run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path --createAsCDB false
```

When prompted, enter the `sys` and `tde` passwords.

- 4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli database create`.

- [Running Prerequisite Checks Prior to Creating Oracle Database](#)
To run prerequisite checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database creation.
- [Resuming or Reverting Oracle Database Creation Operation](#)
To resume or revert a failed database creation operation, use the `--resume` or `--revert` command option.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

- [dbaascli database create](#)
To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Running Prerequisite Checks Prior to Creating Oracle Database

To run prerequisites checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database creation.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path --executePrereqs
```

Where:

- `--dbName` specifies the name of the database
 - `--oracleHome` specifies the Oracle home location
4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli database create`.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli database create](#)
To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Resuming or Reverting Oracle Database Creation Operation

To resume or revert a failed database creation operation, use the `--resume` or `--revert` command option.

For example:

```
dbaascli database create --dbName database name --oracleHome Oracle Home Path --resume
```

Note

- While using the `--resume` or `--revert` command options, ensure that you use the same command from the same node that was used for actual create operation flow.
- You can resume database creation only if there is a failure in the post database creation step.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli database create](#)
To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Changing the Database Passwords

To change the SYS password, or to change the TDE wallet password, use this procedure.

The password that you specify in the **Database Admin Password** field when you create a new Exadata Cloud Infrastructure instance or database is set as the password for the SYS, SYSTEM, TDE wallet, and PDB administrator credentials. Use the following procedures if you need to change passwords for an existing database.

Note

if you are enabling Data Guard for a database, then the SYS password and the TDE wallet password of the primary and standby databases must all be the same.

Note

Using the `dbaascli` to change the SYS password will ensure the backup/restore automation can parallelize channels across all nodes in the cluster.

To Change the SYS Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli database changepassword --dbname database_name --user SYS
```

To Change Database Passwords in a Data Guard Environment

1. Run the following command on the primary database:

```
dbaascli database changePassword -dbName <dbname> --user SYS --  
prepareStandbyBlob true --blobLocation <location to create the blob file>
```

2. Copy the blob file created to all the standby databases and update the file ownership to oracle user.
3. Run the following command on all the standby databases:

```
dbaascli database changePassword -dbName <dbname> --user SYS --  
standbyBlobFromPrimary <location of copies the blob file>
```

To Change the TDE Wallet Password for an Exadata Cloud Infrastructure Database

1. Log onto the Exadata Cloud Infrastructure virtual machine as `opc`.
2. Run the following command:

```
sudo dbaascli tde changepassword --dbname database_name
```

Managing Exadata Cloud Infrastructure Software Images Using the Dbaascli Utility

You can list and download the Oracle database software images on an Exadata Cloud Infrastructure instance, which can then be used for provisioning a database home.

Note

You can create custom database software images for your Exadata Cloud Infrastructure instances using the Console or API. These images are stored in Object Storage, and can be used to provision a Database Home in your Exadata instance. See [Oracle Database Software Images](#) more information.

You can control the version of Oracle binaries that is installed when you provision a new database on an Exadata Cloud Infrastructure instance by maintaining the software images on the system. Oracle provides a library of cloud software images that you can view and download onto your instance by using the `dbaascli` utility.

- [Listing Available Software Images and Versions for Database and Grid Infrastructure](#)
To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- [To download a software image](#)
You can download available software images onto your Exadata Cloud Infrastructure instance by using the `cswlib download` subcommand of the `dbaascli` utility.

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli cswlib showImages --product database
```

The command output lists the available database software images.

```
dbaascli cswlib showImages --product grid
```

The command output lists the available grid software images.

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-2 dbaascli cswlib showImages

```
[root@dg1llrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
  showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
  /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
  AM_46941.log

#####
List of Available Database Images
#####

17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
```

```
VERSION=19.12.0.0.0
DESCRIPTION=19c JUL 2021 DB Image
```

```
21.IMAGE_TAG=19.13.0.0.0
VERSION=19.13.0.0.0
DESCRIPTION=19c OCT 2021 DB Image
```

Images can be downloaded using their image tags. For details, see help using 'dbaascli cswlib download --help'.
dbaascli execution completed

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli cswlib showImages](#)
To view the list of available Database and Grid Infrastructure images, use the dbaascli cswlib showImages command.

To download a software image

You can download available software images onto your Exadata Cloud Infrastructure instance by using the cswlib download subcommand of the dbaascli utility.

1. Connect to a compute node as the `opc` user. For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a root-user command shell:

```
$ sudo -s
#
```

3. Execute the dbaascli command with the cswlib download subcommand:

```
# dbaascli cswlib download [--version <software_version>] [--imageTag
<image tag
value>]
```

The command displays the location of software images that are downloaded to your Exadata Cloud Infrastructure environment.

The optional parameters are:

- **version:** specifies an Oracle Database software version. For example, 19.14.0.0.0.
 - **imageTag:** specifies the image tag of the image.
4. Exit the root-user command shell:

```
# exit
$
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli

Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

- [Patching Databases using dbaascli](#)
Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.
- [Patching Oracle Grid Infrastructure](#)
To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.
- [Listing Available Software Images and Versions for Database and Grid Infrastructure](#)
To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.
- [Performing a Precheck Before Patching Databases and Grid Infrastructure](#)
You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.
- [Resuming or Rolling Back a Patching Operation](#)
You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

Patching Databases using dbaascli

Using `dbaascli`, you can choose to patch a database by patching Oracle home, or by moving the database to an Oracle home with the desired patch level.

- Patching an Oracle home (in-place patching). This updates all databases located in the Oracle home.
- Moving a database to a different Oracle home that has the desired Oracle Database software version (out-of-place patching).
- [Patching a Database Home \(In-Place Database Patching\)](#)
To patch an Oracle home, use the `dbaascli dbHome patch` command.
- [Moving a Database to a Different Oracle Home \(Out-of-Place Patching\)](#)
To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

Patching a Database Home (In-Place Database Patching)

To patch an Oracle home, use the `dbaascli dbHome patch` command.

This will patch all databases running in the specified home, and the databases will remain in the home after the patching is complete. The following apply to using the `dbHome patch` command for in-place patching operations:

- You can patch all of your database nodes or a subset of nodes.
- Multi-node patching takes place in a rolling fashion.
- Optionally, you can perform a software-only patch operation. Then, when you are ready, you can run `datapatch` to perform post-patch SQL actions.
- You can patch an Oracle home containing one or more databases.

To patch an Oracle Home (dbhome):

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli dbhome patch --oracleHome dbhome_path --targetVersion  
Oracle_Database_version
```

Where:

- `--oracleHome` identifies the path of the Oracle home to be patched.
- `--targetVersion` specifies the target Oracle Database version to use for patching, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).

For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/  
dbhome_2 --targetVersion 19.9.0.0.0
```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli dbHome patch`.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Moving a Database to a Different Oracle Home (Out-of-Place Patching)

To patch an Oracle Database by moving it to an Oracle home that is already at the desired patch level, use the `dbaascli database move` command.

After the database move operation is complete, the database runs using the Oracle Database software version of the target Oracle Home.

To patch a database by moving it to a different Oracle Home:

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli database move --oracleHome path_to_target_oracle_home --dbname database_name
```

Where:

- `--oracleHome` identifies the path of the target Oracle home that uses the desired Oracle Database software version. Note that the target Oracle home must exist in your system prior to using the `database move` command.
- `--dbname` specifies the name of the database that is being moved.

For example:

```
dbaascli database move --oracleHome /u02/app/oracle/product/19.0.0.0/dbhome_2 --dbname xyz
```

4. Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli database move`.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli database move](#)
To move the database from one home to another, use the `dbaascli database move` command.

Patching Oracle Grid Infrastructure

To apply a patch to Oracle Grid Infrastructure, use the `grid patch` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli grid patch --targetVersion target_software_version_number
```

Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to.

For example:

```
dbaascli grid patch --targetVersion 19.11.0.0.0
```

- Exit the `root` user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli grid patch`.

- [Patching Oracle Grid Infrastructure \(GI\) Using GI Software Image](#)
To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Patching Oracle Grid Infrastructure (GI) Using GI Software Image

To patch Oracle Grid Infrastructure (GI) using GI software image, use this procedure.

Oracle Grid Infrastructure can also be patched by first creating a patched software image, and then using that image to perform the patching operation. This provides the advantage that an image can be created ahead of time outside of the patching window. It also helps in conflict resolution as any conflicts among the patches are highlighted during the image creation process without impacting the patching window.

- Create a patched software image.

```
dbaascli grid patch --targetVersion <target_software_version_number> --  
createImage
```

Once the patched software image creation is completed, the image can then be used for performing the patching operation.

- Perform the patching operation.

```
dbaascli grid patch --targetVersion <target_software_version_number> --  
imageLocation <location_of_patched_software_image>
```

Listing Available Software Images and Versions for Database and Grid Infrastructure

To produce a list of available supported versions for patching, use the `dbaascli cswlib showImages` command.

- Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
- Start a `root` user command shell:

```
sudo -s
```

- Run the following command:

```
dbaascli cswlib showImages --product database
```

The command output lists the available database software images.

```
dbaascli cswlib showImages --product grid
```

The command output lists the available grid software images.

4. Exit the root user command shell:

```
exit
```

For more details on advanced supported options, see `dbaascli cswlib showImages`.

Example 6-3 dbaascli cswlib showImages

```
[root@dg11lrg1 dbhome_1]# dbaascli cswlib showImages
DBAAS CLI version <version>
Executing command cswlib
      showImagesJob id: 00e89b1a-1607-422c-a920-22f44bec1953Log file location:
      /var/opt/oracle/log/cswLib/showImages/dbaastools_2022-05-11_08-49-12-
AM_46941.log
```

```
#####
List of Available Database Images
#####
```

```
17.IMAGE_TAG=18.17.0.0.0
   VERSION=18.17.0.0.0
   DESCRIPTION=18c JAN 2022 DB Image

18.IMAGE_TAG=19.10.0.0.0
   VERSION=19.10.0.0.0
   DESCRIPTION=19c JAN 2021 DB Image

19.IMAGE_TAG=19.11.0.0.0
   VERSION=19.11.0.0.0
   DESCRIPTION=19c APR 2021 DB Image

20.IMAGE_TAG=19.12.0.0.0
   VERSION=19.12.0.0.0
   DESCRIPTION=19c JUL 2021 DB Image

21.IMAGE_TAG=19.13.0.0.0
   VERSION=19.13.0.0.0
   DESCRIPTION=19c OCT 2021 DB Image
```

```
Images can be downloaded using their image tags. For details, see help using
'dbaascli cswlib download --help'.
dbaascli execution completed
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

- [dbaascli cswlib showImages](#)
To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Performing a Precheck Before Patching Databases and Grid Infrastructure

You can perform a prerequisites-checking operation (also called a "precheck") for the commands in this topic using the applicable precheck flag.

Running prechecks allows you to run only the precheck portion of the patching operation without performing actual patching. Oracle recommends running prechecks to discover software issues that could prevent successful patching.

To perform patching prechecks, first, connect to a virtual machine in your Exadata Cloud Infrastructure instance as the `root` user.

- [Precheck for Oracle Home Patching \(In-Place Patching\)](#)
Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.
- [Precheck for Database Move Patching \(Out-of-Place Patching\)](#)
Use the `--executePrereqs` flag with the `dbaascli database move` command.
- [Precheck for Oracle Grid Infrastructure Patching](#)
Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

Precheck for Oracle Home Patching (In-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli dbhome patch` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli dbhome patch --oracleHome dbhome_path --targetVersion  
Oracle_Database_version --executePrereqs
```

Where:

- `--oracleHome` identifies the path of the Oracle home to be prechecked.
 - `--targetVersion` specifies the target Oracle Database version to be patched to, specified as five numeric segments separated by periods (e.g. 19.12.0.0.0).
4. Exit the `root` user command shell:

```
exit
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Precheck for Database Move Patching (Out-of-Place Patching)

Use the `--executePrereqs` flag with the `dbaascli database move` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli database move --oracleHome path_to_target_oracle_home --dbname  
database_name --executePrereqs
```

Where:

- `--oracleHome` identifies the path of the target Oracle Home that uses the desired Oracle Database software version. Note that the target Oracle Home must exist in your system prior to using the `database move` command.
 - `--dbname` specifies the name of the database that is being moved
4. Exit the `root` user command shell:

```
exit
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli database move](#)
To move the database from one home to another, use the `dbaascli database move` command.

Precheck for Oracle Grid Infrastructure Patching

Use the `--executePrereqs` flag with the `dbaascli grid patch` command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the following command:

```
dbaascli grid patch --targetVersion target_software_version_number --  
executePrereqs
```

Where `--targetVersion` identifies target software version that the Oracle Grid Infrastructure will be patched to, specified as five numeric segments separated by periods, for example, 19.12.0.0.0

4. Exit the `root` user command shell:

```
exit
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Resuming or Rolling Back a Patching Operation

You can resume or revert a failed patching operation. Reverting a patch is known as a rollback.

- [Resuming a Patch Operation](#)
To resume a patching operation, use the `--resume` flag with the original patching command.
- [Rolling Back a Patch Operation](#)
Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

Resuming a Patch Operation

To resume a patching operation, use the `--resume` flag with the original patching command.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the original patching command to resume a patching operation:
For example:

```
dbaascli dbhome patch --oracleHome /u02/app/oracle/product/19.0.0.0/  
dbhome_2 --targetVersion 19.9.0.0.0 --resume
```

4. Exit the `root` user command shell:

```
exit
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Rolling Back a Patch Operation

Use the `--rollback` flag with the original patching command to roll back (revert) a patching operation.

1. Connect to the virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. Run the original patching command to roll back (revert) a patching operation:
For example:

```
dbaascli grid patch --targetVersion 19.11.0.0.0 --rollback
```

Note

- Resume and Rollback operations are supported for Oracle Home patching, Oracle Grid Infrastructure patching, and database move operations.
- When resuming or rolling back a patching operation, you must run the resume or rollback command from the same node that was used to run the original patching command, and you must run the original command with the addition of the `--resume` or `--rollback` flag.

4. Exit the `root` user command shell:

```
exit
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check Using dbaascli

Using the `dbaascli diag` command allows you to collect Guest VM dbaas tooling logs for Exadata Database Service on Dedicated Infrastructure and Exadata Database Service on Cloud@Customer systems. You can use these logs to troubleshoot issues related to dbaas tooling.

You can use the `diag` command to collect dbaastools logs and perform a health check on all nodes in an Exadata cluster. Note that the `--waitForCompletion` options is supported starting in version 22.4.1

Note

- `dbaascli diag` commands must be run as the `root` user
- Running the `dbaascli diag collect` command on a single node will collect log data for all nodes
- We recommend running the commands documented in this topic using the `--waitForCompletion` option for long-running commands. Refer to the examples for sample usage.

For information on updating Exadata Cloud Tooling, see *dbaascli admin updateStack*.

- [Collecting Tooling Log Data Examples](#)
The `dbaascli dbaascli diag collect` command uses the syntax shown below to collect tooling log data:
- [Performing a Health Check Examples](#)
Use `dbaascli dbaascli diag healthcheck` command to perform a health check on all system nodes.

Related Topics

- [dbaascli diag collect](#)
To collect diagnostics, use the `dbaascli diag collect` command.
- [dbaascli admin updateStack](#)
To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Collecting Tooling Log Data Examples

The `dbaascli dbaascli diag collect` command uses the syntax shown below to collect tooling log data:

See `dbaascli diag collect` in the *dbaascli Command Reference* for syntax details

NOT_SUPPORTED

```
# dbaascli diag collect
DBAAS CLI version 24.1.1.0.0
Executing command diag collect
Job id: 92f33125-aa70-4ce2-94fb-64d8f1cbdc93
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_07-20-44-
```

```

PM_83383.log
Loading PILOT...
Session ID of the current execution is: 10
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_07-20-48-
PM_83856
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-1920/
diag_cloudlogs_20231214-1920_node1.zip=a0d049b87ab9e9cec2ab7d95ded4903bac818c8
1c8b6a46d295e1e75f4630e19}"
}
dbaascli execution completed

```

NOT_SUPPORTED

```

# dbaascli diag collect --waitForCompletion false
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --waitForCompletion false
Job id: 5b556976-dba1-4be9-a4fe-4b58e69c1d96
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_07-23-26-
PM_98107.log
Job accepted. Use "dbaascli job getStatus --jobID 5b556976-dba1-4be9-
a4fe-4b58e69c1d96" to check the job status.

```

Note

Use the job status command to monitor progress.

NOT_SUPPORTED

```

# dbaascli diag collect --dbnames myOracleDatabase19cName
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --dbnames myOracleDatabase19cName
Job id: 8e1d2667-4649-4384-8610-b6348d6548ac
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_08-41-41-
PM_88831.log
Loading PILOT...
Session ID of the current execution is: 12
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_08-41-45-
PM_89361
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2041/
diag_cloudlogs_20231214-2041_node1.zip=9e50500089a74ca7cd8ae08550c06868e26e1cd
9c52e808194256594f63397e4}"
}
dbaascli execution completed

```

NOT_SUPPORTED

```
# dbaascli diag collect --destLocation /tmp/test/
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --destLocation /tmp/test/
Job id: f992afdf-415e-4b58-ab5b-9e38f8c2079d
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-42-54-
PM_16270.log
Loading PILOT...
Session ID of the current execution is: 14
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-42-58-
PM_16777
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/tmp/test/diag_collect/
artifacts_diag_cloudlogs_20231214-2143/
diag_cloudlogs_20231214-2143_node1.zip=8a26cfffcd72c261660d4f736c615981856e35
7749d90751b94f3eda19a9a70}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --startTime 2023-12-05T10:00:00 --endTime
2023-12-05T11:00:00
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --startTime 2023-12-05T10:00:00 --endTime
2023-12-05T11:00:00
Job id: 70b03e50-98cc-4c2b-9684-1f82070bac88
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-45-17-
PM_42856.log
Loading PILOT...
Session ID of the current execution is: 15
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-45-21-
PM_43526
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2145/
diag_cloudlogs_20231214-2145_node1.zip=b44cf3bfca1ab7a1629dd83098a7772790ab949
e50dbb3950f0017e427d7bd05}"
}
dbaascli execution completed
```

NOT_SUPPORTED

```
# dbaascli diag collect --nodes node1,node2
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --nodes node1,node2
Job id: fa70da09-3de6-4cc8-854c-a739b4fc2ceb
```

```

Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-46-58-
PM_55884.log
Loading PILOT...
Session ID of the current execution is: 16
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-47-02-
PM_56418
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2147/
diag_cloudlogs_20231214-2147_node1.zip=de2805c9c6c2af2d602395a84d37747935327b7
3a6c73052282665a8410eb41f}"
}

```

NOT_SUPPORTED

```

# dbaascli diag collect --components dbaastools
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --components dbaastools
Job id: da941d3c-5191-4ced-b1bb-9b083fa75865
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-47-23-
PM_68256.log
Loading PILOT...
Session ID of the current execution is: 17
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-47-27-
PM_68729
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2147/
diag_cloudlogs_20231214-2147_node1.zip=d1f290fb42c981935e1142ec059c2dbba8be2e0
a9ffebc9eea83a6336abe2eed}"
}
dbaascli execution completed

```

NOT_SUPPORTED

```

# dbaascli diag collect --objectStoreBucketUri https://objectstorage.us-
phoenix-1.oraclecloud.com/p/aL-IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--
BeVuXaipSEEMISrSCfFrVEolG/n/intexadatateam/b/diag_collect_test/o/
DBAAS CLI version 24.1.1.0.0
Executing command diag collect --objectStoreBucketUri https://
objectstorage.us-phoenix-1.oraclecloud.com/p/aL-
IbIKQ1j6lWNftJc2rLoLh6o9bJgbZm8z0S--BeVuXaipSEEMISrSCfFrVEolG/n/
intexadatateam/b/diag_collect_test/o/
Job id: 028151b7-cbc4-409a-9ec6-69affe10f3bb
Session log: /var/opt/oracle/log/diag/collect/dbaastools_2023-12-14_09-51-36-
PM_2963.log
Loading PILOT...
Session ID of the current execution is: 20
Log file location: /var/opt/oracle/log/diag/collect/pilot_2023-12-14_09-51-40-

```

```

PM_3555
-----
..
----- DIAG COLLECT PLUGIN RESULT -----
{
  "collectedArchive with SHA256 CheckSum" : "{/var/opt/oracle/dbaas_acfs/
diag_collect/artifacts_diag_cloudlogs_20231214-2151/
diag_cloudlogs_20231214-2151_node1.zip=71633e13ccd06de15cb26850bb0266cf0d869e2
59550515c5b1fb734c487b470}"
}
dbaascli execution completed

```

Related Topics

- [dbaascli diag collect](#)
To collect diagnostics, use the `dbaascli diag collect` command.

Performing a Health Check Examples

Use `dbaascli dbaascli diag healthcheck` command to perform a health check on all system nodes.

See *dbaascli diag healthcheck* for the syntax details in the *dbaascli Command Reference*.

NOT_SUPPORTED

```

# dbaascli diag healthcheck
DBAAS CLI version MAIN
Executing command diag healthcheck
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /var/opt/oracle/dbaas_acfs/
diag_cloudlogs_20210322-2246.tar.gz

```

NOT_SUPPORTED

```

# dbaascli diag healthcheck --destLocation /tmp/test
DBAAS CLI version MAIN
Executing command diag healthcheck --destLocation /tmp/test
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /tmp/test/diag_cloudlogs_20210322-2250.tar.gz

```

NOT_SUPPORTED

```

# dbaascli diag healthcheck --nodes rbcl1,rbcl2
DBAAS CLI version MAIN
Executing command diag healthcheck --nodes rbcl1,rbcl2
INFO: Starting diag healthcheck
INFO: Collected diag logs at: /var/opt/oracle/dbaas_acfs/
diag_cloudlogs_20210421-1915.tar.gz

```

NOT_SUPPORTED

```

# dbaascli diag healthcheck --objectStoreBucketUri https://objectstorage.us-
phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-

```

```
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/  
DBAAS CLI version MAIN  
Executing command diag healthcheck --objectStoreBucketUri https://  
objectstorage.us-phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-  
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/  
INFO: Collected diag logs at: https://objectstorage.us-  
phoenix-1.oraclecloud.com/p/t0Z-kRV5pSmFzqnf-  
y5XhaAbM4LS82epeBnulKnCr31IeHVjxI9tOkntLF2kq7fP/n/MyNamespace/b/MyParBucket/o/  
diag_cloudlogs_20210421-1839.tar.gz
```

Related Topics

- [dbaascli diag collect](#)
To collect diagnostics, use the `dbaascli diag collect` command.
- [dbaascli diag healthCheck](#)
To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Updating Cloud Tooling Using dbaascli

To update the cloud tooling release for Oracle Exadata Database Service on Dedicated Infrastructure, complete this procedure.

Cloud-specific tooling is used on the Exadata Cloud Infrastructure Guest VMs for local operations, including `dbaascli` commands.

The cloud tooling is automatically updated by Oracle when new releases are made available. If needed, you can follow the steps below to ensure you have the latest version of the cloud-specific tooling on all of the virtual machines in the VM cluster.

Note

You can update the cloud-specific tooling by downloading and applying a software package containing the updated tools.

1. Connect to a virtual machine as the `opc` user.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.
2. Start a `root` user command shell:

```
sudo -s
```

3. To update to the latest available cloud tooling release, run the following command:

```
dbaascli admin updateStack
```

The command takes care of updating the cloud tooling release on all the nodes of the cluster.

For more details and other available options, refer to `dbaascli admin updateStack --help`.

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.
- [dbaascli admin updateStack](#)
To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Creating a Duplicate Database

- [Using dbaascli to Duplicate a Cloud Database](#)
- [Considerations When Using OCI Vault for the Key Management](#)
- [Duplicate an On-Premises Database](#)

Using dbaascli to Duplicate a Cloud Database

You can create a duplicate database using `dbaascli`. This new database can be in the same cloud region as the source region or across the regions. The following steps describe how to create a duplicate database on cloud.

Note

If a database is configured with OCI Vault for TDE encryption and you want to duplicate a database, then refer to the following sections.

Prepare for duplication

Ensure that the following prerequisites are met:

- Make sure that there is a network path setup to access the source database through the `EZConnect` string.
- Copy the TDE wallet file (`ewallet.p12`) to the target database node. The node where you decide to run the `dbaascli` command.
- Create an Oracle home on the target node if required. Oracle home version must be the same version as the source or of higher RU version.

Run prerequisite checks

To run prerequisite checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database duplication.

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home Path> --sourceDBConnectionString <source database EZConnect string> --sourceDBTDEWalletLocation <location of copied wallet> --sourceDBTdeConfigMethod FILE --tdeConfigMethod FILE --executePrereqs
```

Duplicate the database

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home Path> --sourceDBConnectionString <source database EZConnect string> --
```

```
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod FILE --tdeConfigMethod FILE
```

Considerations When Using OCI Vault for the Key Management

This section is applicable only in the case of database is configured with OCI Vault for TDE encryption and you want to duplicate a database.

Duplicating a database within the same region

- Additional prerequisite steps
Make sure to setup OCI Vault access policies for target database nodes. Target database nodes should be able to access both source database's OCI key vault along with its new key vault (if it is decided to use separate key vault).
- Run prerequisite checks

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID> --
executePrereqs
```

- Duplicate the database

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID>
```

Upon successful completion of this command, the database is duplicated.

Duplicating a database across regions

- Additional prerequisite steps
 - Setup a new OCI Vault for target database on the corresponding region by following the steps outlined in [Prepare to Use Customer-Managed Keys in the Vault Service](#). Complete Tasks 1 through 3.
 - Setup OCI Vault replication from source region to target region. For more information, see [Replicating Vaults and Keys](#).
 - Update Dynamic group policy, which is created in step 2 to allow access to replicated OCI Vault key.
- Run prerequisite checks

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID> --
executePrereqs
```

- Duplicate the database

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle
Home Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --
sourceDBTdeConfigMethod KMS --sourceDBKMSKeyOCID <Source Database OCI
Vault key OCID> --tdeConfigMethod KMS --kmsKeyOCID <OCI Vault key OCID>
```

Upon successful completion of this command, the database is duplicated.

Duplicate an On-Premises Database

Using `dbaascli`, you can duplicate an on-prem database onto the cloud. This can be done with the `dbaascli database duplicate` command. This command creates a new database on the cloud, which is a duplicate of an on-prem database along with its data. While this process is going on, the on-prem database remains still operational. You can migrate your applications to the duplicated database on the cloud after due verification.

Prepare for duplication

The migration process includes the following prerequisites to be met.

- Make sure that there is a network path setup to access an on-prem database from the OCI node through the `EZConnect` string.
- If an on-prem database is configured with TDE, copy the TDE wallet file (`ewallet.p12`) to the OCI node, where you decide to run the `dbaascli` command.
- Create an Oracle home on the OCI node if required. The Oracle home version must be the same as the source or of a higher RU version.

Verify the necessary RPMs

This process requires a minimum `dbaastools` RPM version of 23.3.2.0.0 but updating to the latest `dbaastools` rpm is always recommended.

- To check the currently installed version, run:

```
dbaascli --version
DBAAS CLI version 23.3.2.0.0
```

- To apply the latest tools RPM, as the `root` user, run:

```
# dbaascli admin updateStack
```

Run the prerequisite checks

To run the prerequisite checks, use the `--executePrereqs` command option. This will perform only the prerequisite checks without performing the actual Oracle Database duplication.

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home
Path> --sourceDBConnectionString <source database EZConnect string> --
sourceDBTDEWalletLocation <location of copied wallet> --executePrereqs
```

Duplicate the database

Duplicate the database using the following command:

```
dbaascli database duplicate --dbName <database name> --oracleHome <Oracle Home Path> --sourceDBConnectionString <source database EZConnect string> --sourceDBTDEWalletLocation <location of copied wallet>
```

For example:

```
dbaascli database duplicate --sourceDBConnectionString xyzhost.oracle.com:1521/dbuniquename.oracle.com --dbName orcl --oracleHome /u02/app/oracle/product/19.0.0.0/dbhome_1 --sourceDBTDEWalletLocation /tmp/wallet_copy/tde --waitForCompletion false
```

Upon successful completion of this command, the database is duplicated to Cloud and ready for sanity checks for application usage. Once verification is done, application connections can be migrated to the Cloud database.

Refer to `dbaascli database duplicate -help` for additional configuration options.

Few considerations for migration

- If you prefer to allocate multiple channels for RMAN duplicate, you could do so by specifying the `--rmanParallelism` argument.
- Exadata Cloud Service configures database memory as Automatic Shared Memory Management (ASMM). If your on-prem database is configured with different memory management, make sure to adjust memory parameter values accordingly on the OCI side by providing values for `--sgaSizeInMB` and `--pgaSizeInMB`.
- Verify that the on-prem database does not contain any deprecated or invalid initialization parameters.
- Database initialization parameters related to database storage (datafile location, redo location, recovery area destination, control file multiplexing) may be changed using the `--initParams` argument.
For example, to override `db_create_online_log_dest` value for the duplicate database: `--initParams db_create_online_log_dest_1=+DATA1,db_create_online_log_dest_2=+RECOC1`

Troubleshooting the database duplication

- `dbaascli` operation log file can be found under `/var/opt/oracle/log/<dbname>/database/duplicate`
- One of the jobs of the duplicate is to run `dbca`. Its log file can be found under `/u02/app/oracle/cfgtoollogs/dbca` and `/u02/app/oracle/cfgtoollogs/dbca/<dbuniquename>`.

If the operation fails, you will have an option to resume the operation by providing the `--resume` argument to the same command. Alternatively, clean up the database using `dbaascli database delete -dbname <dbname> -force`, and then rerun the database duplicate command.

Release Notes

Review the changes made in various releases of `dbaascli`.

- [Release 25.4.1.0.0 \(251107\)](#)
- [Release 25.3.1.0.0 \(250826\)](#)
- [Release 25.2.1.0.0 \(250522\)](#)

- [Release 25.1.2.0.0 \(250325\)](#)
- [Release 25.1.1.0.0 \(250107\)](#)
- [Release 24.4.1.0.0 \(241104\)](#)
- [Release 24.3.2.0.0 \(240828\)](#)
- [Release 24.3.1.0.0 \(240711\)](#)
- [Release 24.2.1.0.0 \(240530\)](#)
- [Release 24.1.2.0.0 \(240306\)](#)
- [Release 24.1.1.0.0 \(231219\)](#)
- [Release 23.4.1.0.0 \(231102\)](#)
- [Release 23.3.2.0.0 \(230921\)](#)
- [Release 23.3.1.0.0 \(230712\)](#)
- [Release 23.2.1.0.0 \(230503\)](#)
- [Release 23.1.2.0.0 \(230305\)](#)
- [Release 23.1.1.0.1 \(230113\)](#)
- [Release 22.4.1.0.1 \(221122\)](#)
- [Release 22.3.1.1.0 \(221003\)](#)
- [Release 22.3.1.0.1 \(220721\)](#)
- [Release 22.2.1.1.0 \(220623\)](#)
- [Release 22.2.1.1.0 \(220609\)](#)
- [Release 22.2.1.0.1 \(220423\)](#)
- [Release 22.1.1.2.0 \(220405\)](#)
- [Release 22.1.1.1.0 \(220317\)](#)
- [Release 22.1.1.0.1 \(220223\)](#)
- [Release 21.4.1.1.0 \(220209\)](#)
- [Release 21.4.1.1.0](#)
- [Release 21.3.1.2.0](#)
- [Release 21.3.1.1.0](#)
- [Release 21.3.1.0.1](#)
- [Release 21.2.1.x.x](#)

Release 25.4.1.0.0 (251107)

- Multicloud: AWS — Support for cross-region database creation from irestore backups in Amazon S3 (Simple Storage Service).
- Various bug fixes and stability improvements.

Release 25.3.1.0.0 (250826)

- Includes AHF 25.6.2
- Includes syslens 25.2.3.0

- Includes `exacs-configs-release-master_25-3-1`
- The `dbaascli dbhome patch` command now includes the `skipDBForDatapatch` option. **Purpose:** Allows you to specify a list of databases on which the datapatch step will be skipped during in-place DB home patching.
- Various bug fixes and stability improvements

Release 25.2.1.0.0 (250522)

- Various bug fixes and stability improvements

Release 25.1.2.0.0 (250325)

- Various bug fixes and stability improvements

Release 25.1.1.0.0 (250107)

- Various bug fixes and stability improvements

Release 24.4.1.0.0 (241104)

- Enhanced `dbaascli database getDetails` to include tde encryption details. The command `dbaascli tde status` is now deprecated.
- Various bug fixes and stability improvements

Release 24.3.2.0.0 (240828)

- Various bug fixes and stability improvements

Release 24.3.1.0.0 (240711)

- TDE migration from `sqlnet.ora` to `wallet_root` on database upgrade to 19c release
- Grid patch in-place to use image-based patching as default mode
- Various bug fixes and stability improvements

Release 24.2.1.0.0 (240530)

- Added support for Oracle AI Database 26ai.
- Improvements in backup and recovery area with Zero Data Loss Autonomous Recovery Service (ZRCV) as backup destination.
- Various bug fixes and stability improvements.

Release 24.1.2.0.0 (240306)

- Introduced a new optimized workflow for Data Guard operations
- Various bug fixes and stability improvements

Release 24.1.1.0.0 (231219)

- Various bug fixes and stability improvements

Release 23.4.1.0.0 (231102)

- **Backup and Recovery:** Minimum Backup recovery window has been changed to 7 days. While obsoleting backup pieces automation considers recovery window as 7 days if it discovers any value less than 7 from the system.
- Various bug fixes and stability improvements

Release 23.3.2.0.0 (230921)

- Pluggable Database Operations
 - Added support to set custom key version OCID (Bring Your Own Key - BYOK) of OCI Vault during create and clone operations. For details, see respective PDB commands help.
- Grid Infrastructure (GI) Patching
 - Enhanced the patching workflow to improve patching time, especially in environments having high number of databases.
 - A new option `--patchInParallel` is introduced that can be used to perform patching remote nodes in parallel.
- Database Patching
 - Provided option to run `datapatch` on a specific node of cluster.
- Various bug fixes and stability improvements

Release 23.3.1.0.0 (230712)

- New dbaascli commands
 - `dbaascli gridHome create` - This command can be used to create a Grid Infrastructure home of a supported version. For details, see `dbaascli gridHome create --help`.
 - `dbaascli system getGridHomes` - This command gives details on the available Grid Infrastructure homes on the system. For details, see `dbaascli system getGridHomes --help`.
 - `dbaascli admin updateAHF` - This command can be used to update the AHF to a specified cloud certified version of AHF release. It is however recommended that AHF updates be implicitly handled by cloud automation. For details, see `dbaascli admin updateAHF --help`.
- Pluggable Database Operations
 - Improvements in the area of refreshable pluggable database lifecycle.
- Database Backup and Recovery
 - Added support to configure backups on standby sites in case of dataguard configurations. The backups configuration are Data Guard site-specific, that is, the change of roles (for example, with Data Guard switchover operation) will not impact the backup operations of the database on primary or standby sites. Backups, if configured on primary site or stand-by site, will continue regardless of the role-change.
- Various bug fixes and stability improvements

Release 23.2.1.0.0 (230503)

- Database Lifecycle related improvements
 - Introduced `dbaascli grid removeTCPSCert` to remove expired TCPS certificates. For details, see `dbaascli grid removeTCPSCert --help`.
 - Added option to exclude specific PDBs during database duplicate. For details, see `skipPDBs` argument in `dbaascli database duplicate --help`.
- Database Backup and Recovery
 - Changed the default for `FILES_PER_SET` to 64 for OSS backups. This can be changed with `dbaascli database backup --configure`. For details, see `dbaascli database backup --help`.
 - Archive log backups continue from the standby site after the role switchover in data guard environments.
 - For backups that are not managed by Oracle, the schedules for L0 and L1 backups are not created by default. They must be created explicitly by using `dbaascli database backup --configure` command.
- sysLens
A framework that collects, analyzes, and reports system resource data for ExaDB-D fleets is included in 23.2.1.0.0 (235503). For more information, see [Manage sysLens](#).
- Various bug fixes and stability improvements

Release 23.1.2.0.0 (230305)

- Database Lifecycle related improvements
 - Added option to create database templates (DBCA templates) to object store. DBCA templates can subsequently be used to create databases. For details, see `dbaascli database createTemplate --help`.
- Pluggable Database Operations
 - Introduced `dbaascli pdb refresh` to refresh a pluggable database that was created using manual refresh option. For details, see `dbaascli pdb refresh --help`.
 - Added option to convert refreshable pluggable database to a regular pluggable database. For details, see `dbaascli pdb open --help`.
 - Creation of a refreshable pluggable database now requires existing source database user for creation of database link to the source pluggable database. For details, see `dblinkUserName` argument in `dbaascli pdb remoteClone --help`.
- Various bug fixes and stability improvements

Release 23.1.1.0.1 (230113)

- Database Lifecycle related improvements
 - Added support to create a duplicate database from a source database which uses OCI Vault Services for encryption key management.
- Various bug fixes and stability improvements

Release 22.4.1.0.1 (221122)

- Pluggable Database Operations
 - Added option to not open the PDB at the end of relocate. For details, see `skipOpenPDB` argument in `dbaascli pdb relocate --help`. After using this option, the `pdb relocate` can be completed by running the command using `completePDBRelocate` argument.
 - Added option to clean up the relocated PDB metadata/services at the source location. For details, see `cleanupRelocatedPDB` argument in `dbaascli pdb delete --help`
- New dbaascli commands
 - `dbaascli database createTemplate` - This command can be used to create database templates (DBCA templates) that can subsequently be used to create databases. DBCA templates are widely used for creating a clone database with DBCA - a tool that is shipped with Oracle Database server software. For details, see `dbaascli database createTemplate --help`
 - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
- Database Lifecycle related improvements
 - Added support to use dbca templates in database creation workflows. For details, see `dbcaTemplateFilePath` argument in `dbaascli database create --help`
 - Improved performance for duplicate database creation. For details on how to create duplicate database, see `dbaascli database duplicate --help`
 - Added support to create a duplicate database from a source database which is not TDE-encrypted.
- TDE management
 - Introduced `dbaascli tde rotateMasterKey` to rotate the master key for database encryption. For details, see `dbaascli tde rotateMasterKey --help`. The command `dbaascli tde rotate masterkey` is now deprecated.
 - Revamped workflow for all TDE operations. For details, see `dbaascli tde --help`
- Grid Infrastructure (GI) Patching
 - Added support to allow parallel execution of patching operation on nodes. This option needs to be carefully exercised as it results into reduced database availability.
- Database Backup and Recovery
 - Revamped workflow for creating database from standalone backups
- Includes AHF version 22.2.4
- Various bug fixes and stability improvements

Release 22.3.1.1.0 (221003)

- New dbaascli commands
 - `dbaascli database getDetails` - This command shows the detailed information of a given database, for example, `dbname`, node information, pluggable databases information, and so on. For details, see `dbaascli database getDetails --help`.
- Pluggable Database Operations

- Added support for creating pluggable databases as refreshable clone using `refreshablePDB` argument. For details, see `dbaascli pdb remoteClone --help`
- Various bug fixes and stability improvements

Release 22.3.1.0.1 (220721)

- New database lifecycle commands
 - `dbaascli database addInstance` - This command can be used to add a database instance to one of the nodes of the cluster where database is not already configured. For details, see `dbaascli database addInstance --help`.
 - `dbaascli database deleteInstance` - This command can be used to delete a database instance from one of the nodes of the cluster where database is configured. For details, see `dbaascli database deleteInstance --help`.
 - `dbaascli database duplicate` - This command can be used to create a new database from an already existing database within a cluster, or across clusters, provided network connection exists between the clusters. For details, see `dbaascli database duplicate --help`.
- Cloud Software Library
 - Introduced `dbaascli cswlib listLocal` command to list images that are downloaded from software library locally on the system. For details, see `dbaascli cswlib listLocal --help` The command `dbaascli dbimage list` is now deprecated.
 - Introduced `dbaascli cswlib deleteLocal` command to delete images that are downloaded from cloud software library. For details, see `dbaascli cswlib deleteLocal --help` The command `dbaascli dbImage purge` is now deprecated.
- The log location for the command `dbaascli admin updateStack` has been changed to follow the convention of other `dbaascli` commands. The logs can be conveniently found under `/var/opt/oracle/log/admin/updateStack` directory. The earlier location was `/var/opt/oracle/log/tooling/Update`.
- `dbaascli help` is now cloud platform aware in that it will list help output for commands applicable for the cloud environment it is operating on.
- Added support for changing TDE password in dataguard environments. For details, see `dbaascli tde changePassword --help`. This support is currently not available for 11.2.0.4 release.
- Included AHF version 22.1.5.
- Revamped workflow for database upgrade operation.
- Revamped workflow for database home create operation.
- Various bug fixes and stability improvements

Release 22.2.1.1.0 (220623)

- Included AHF version 22.1.1
- Fixed an issue where update of `dbaastools rpm` on the system may have resulted into database downtime with ORA-600 error
- Various bug fixes and stability improvements

Release 22.2.1.1.0 (220609)

- New dbaascli commands:
 - `dbaascli dbHome getDatabases` - This command lists all the databases running from a given database Oracle home. The output is returned in JSON format to facilitate automation. For details, see `dbaascli dbHome getDatabases --help`.
 - `dbaascli database getPDBs` - This command lists all the pluggable databases of a given container database. The output is returned in JSON format to facilitate automation. For details, see `dbaascli database getPDBs --help`.
 - `dbaascli dbHome delete` - This command deletes a given database Oracle home. For details, see `dbaascli dbHome delete --help`.
 - `dbaascli dataguard prepareStandbyBlob` - This command generates a blob file containing various files that are required on the standby site for a Data Guard environment. For details, see `dbaascli dataguard prepareStandbyBlob --help`.
- Grid Infrastructure (GI) Patching:
 - New optimized workflow
 - Introduced a way to create the Grid Infrastructure (GI) software image prior to patching. This GI image can be subsequently used for performing the GI patching operation. The advantage of this approach is that it results in reduced patching window as the image is already prepared. The GI stack on the node is not brought down to create the image. For details, see `createImage` option in `dbaascli grid patch --help`
 - Introduced a way to perform the Grid Infrastructure patching through the use of user specified GI software image, created using `createImage` option of the `dbaascli grid patch` command. For details, see `imageLocation` option in `dbaascli grid patch --help`.
- Change Password support in Data Guard environment:
 - Added support to change password in Data Guard environments. For details, see `dbaascli database changePassword --help` and `dbaascli dataguard prepareStandbyBlob --help`
- Data Guard configuration:
 - Added support to update Data Guard Automation Attributes (in the `/var/opt/oracle/dg/dg.conf` file). For details, see `dbaascli dataguard --help`.
- Various bug fixes and stability improvements

Release 22.2.1.0.1 (220423)

- New dbaascli commands
 - Introduced `dbaascli admin showLatestStackVersion` to show the latest dbaastools version available for customers to download and install. The installation of dbaastools rpm can be performed by using the command `dbaascli admin updateStack`. For details see “dbaascli Command Reference” section.
- Cloud Software Library
 - Deprecated the support for BP activation (`dbaascli cswlib activateBP`) as BPs (Bundle Patches) are now replaced with RUs (“Release Updates”). Cloud deployment consumes RUs in the form of software images, identified with “Image Tags”. It is

therefore recommended to use image tags while interfacing with Cloud Software Library (cswlib) commands. For details, see *dbaascli cswlib download -help*.

- Eliminated the need to download Non-CDB images to create nonCDB databases. Now users can create the nonCDB database using regular images. For details, see *createAsCDB* option in *dbaascli database create -help*.
- Non-CDB Database Creation
 - Enhanced database creation workflow to create a nonCDB database using standard database software image. For details, see *createAsCDB* option in *dbaascli database create -help*.
- Database Home Patching
 - New optimized workflow
- Grid Infrastructure Upgrade
 - New optimized workflow
- Pluggable Database (PDB) Operations
 - Deletion of PDB in DataGuard environments requires explicit acknowledgement to indicate that operations necessary on standby site are completed, by passing of additional argument *-allStandByPrepared*. For details, see *dbaascli pdb delete --help*
- Provided rolling capability for database bounce operation. For details, see *dbaascli database bounce -help*.
- Various bug fixes and stability improvements

Release 22.1.1.2.0 (220405)

- Added support for ExaDB-D X9M
- Various bug fixes and stability improvements

Release 22.1.1.1.0 (220317)

- New *dbaascli* commands:
 - Introduced *dbaascli system getDBHomes* to get all the database Oracle homes on the cluster. The output is returned in JSON format to facilitate automation.
 - Introduced *dbaascli dbhome getDetails* to get detailed information on a specific Oracle home. The output is returned in JSON format to facilitate automation.
- Cloud Software Library (cswlib):
 - Deprecated the support for *dbaascli cswlib list* command for cloud software library listing operations. The new command is *dbaascli cswlib showImages* that lists the images along with its of *ImageTag*. It is recommended to use *Image tags* to download the images from the cloud software library. For details on downloads using image tags, see *dbaascli cswlib download -help*.
 - Various bug fixes and stability improvements

Release 22.1.1.0.1 (220223)

- Grid Infrastructure Upgrade
 - New optimized workflow

- Database Backup And Recovery
 - Internal update to metadata repository for backup metadata
 - Introduced deprecation messages for `bkup_api` commands as they are now replaced with `dbaascli` commands. For details, see '`dbaascli database backup --help`' and '`dbaascli database recover --help`'
- Pluggable Database (PDB) Operations
 - Relocate operation of PDB is now supported. For details, see '`dbaascli pdb relocate --help`'.
 - Revamped workflow for nonCDB to PDB conversion. For details, see '`dbaascli database convertToPDB --help`'.
- Encryption Key Management
 - Transparent Data Encryption (TDE) heartbeat specific initialization parameters are set to the cloud recommended values for databases with 'Customer Managed Keys'.
- Cloud Software Library Management
 - Revamped software library download of artifacts through `imageTags`. It is recommended to use `imageTags` to download the database and grid software images. For details, see '`dbaascli cswlib showimages`' and '`dbaascli cswlib download --help`'
- Included AHF version 21.4.2
- Various bug fixes and stability improvements

Release 21.4.1.1.0 (220209)

- Included AHF version 21.4.1
- Bug fixes and stability improvements

Release 21.4.1.1.0

- Enabled encryption of the system level tablespaces (`SYSTEM`, `SYSAUX`, `UNDO`, and `TEMP`) for databases that will get created with this version of `dbaastools` onwards. This feature is enabled for Oracle Database version 19.6.0.0.0 and above.
- Grid Patching:
 - Prerequisite condition added to check for following file ownership to be owned by `grid` user.
 - * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jdev-rt.jar`
 - * `<gi_home>/suptools/tfa/release/tfa_home/jlib/jewt4.jar`
- Database Patching:
 - Simultaneous `database move` operation is disallowed by default. A new option `-allowParallelDBMove` is introduced that can be used to override the default behavior for Oracle Database releases 12.2 and above.
 - Fixed issues related to move of standby databases being in `MOUNT` mode.
- Database Backup and Recovery:

- Added new command-line options for database backup. For more details, refer to *dbaascli database backup* command reference.
- Added new command-line options for database recovery. For more details, refer to *dbaascli database recover* command reference.
- `bkup_api` usage for backup and recovery operations will be deprecated in future.
- To align with the Oracle recommended practice of using `SYSBACKUP` administrative privilege for Backup and Recovery operations, cloud automation creates a common administrative user `C##DBLCMUSER` with `SYSBACKUP` role at the `CDB$ROOT` container level. Backup and Recovery operations are therefore performed with the user having the least required privileges. Credentials for this user are randomly generated and securely managed by cloud automation. If the user is not found or is `LOCKED` and `EXPIRED`, then cloud automation will recreate or unlock this user during the backup or recovery operation. This change in the cloud automation is made starting with `dbaastools` version 21.4.1.1.0.
- Enhanced `dbaascli resume` functionality to resume any previous session by specifying the `-sessionID <value>` argument to the resume command. The session ID is shared in the `dbaascli` output as well as in the logs.
- Enhanced `dbaascli help` output to show the command usage.
- Deprecated the usage of `dbaascli shell` (interactive session). This will be completely unsupported after March 2022. It is recommended to execute complete `dbaascli` commands on command prompt as suggested in all document examples.
- Included Autonomous Health Framework (AHF) version 21.2.8.
- Various bug fixes and stability improvements.

Release 21.3.1.2.0

- Improved the timing of `dbaascli` operations with enhanced Control Plane metadata synchronization logic.
- Enhanced `dbaascli` logs to have millisecond-level information along with the associated thread.
- Introduced more prerequisite checks in database home patching and database move operations to catch potential failures scenarios with suggestions to corrective action.
- Database patching operations now retain the state of the databases to be same as it was prior to patching. For pluggable databases, `pdb` saved state is honored.
- Various bug fixes and stability improvements.

Release 21.3.1.1.0

- Added support to unlock PDB Admin user account as part of PDB creation, `localClone`, or `remoteClone` operation. For details, see option `--lockPDBAdminAccount` in `dbaascli pdb create --help`.
- Fixed an issue that updates the database resource registered with Oracle Grid Infrastructure in existing environments with the correct value of database name.
- Enhanced PDB lifecycle operations.
- Various bug fixes and stability improvements.

Release 21.3.1.0.1

- Support for the following dbaascli commands to be run as oracle user.
 - dbaascli pdb bounce
 - dbaascli pdb close
 - dbaascli pdb connectString
 - dbaascli pdb create
 - dbaascli pdb delete
 - dbaascli pdb getDetails
 - dbaascli pdb list
 - dbaascli pdb localClone
 - dbaascli pdb open
 - dbaascli pdb remoteClone
- Revamped out-of-place patching of database. For details, see dbaascli database move -help.
- Timing related enhancements in Oracle Grid Infrastructure patching workflow. For details, see dbaascli grid patch -help.
- Deprecated the support for exadbcpatchmulti / dbaascli patch for patching operations. The dbaascli dbhome patch and dbaascli grid patch commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure* and *Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for dbaascli tools patch command to bring consistency in the dbaascli command conventions. The new command is dbaascli admin updateStack. For details, see section *Updating Cloud Tooling using dbaascli*.
- Ability to run dbaascli in disconnected mode for long running operations. Executing dbaascli command with --waitForCompletion false gets you a job ID that can be queried later to get the status of the operation, using dbaascli job getStatus -jobid job_id. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for dbaascli database create command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of --waitForCompletion option.
- Deprecated the support for dbaascli shell. It is recommended that users run the complete dbaascli commands on the command prompt as suggested in all the document examples. Execution of just dbaascli will show the output of its usage help instead of entering into a dbaascli shell.
- Various bug fixes and stability improvements.

Release 21.2.1.x.x

- Redesigned Oracle Grid Infrastructure patching operation and added ability to resume from failed point, patch on subset of nodes, instance draining, and other enhancements. For

details, see `dbaascli grid patch --help`. Also refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section.

- Deprecated the support for `exadbcpatchmulti / dbaascli patch` for patching operations. `dbaascli dbhome patch` and `dbaascli grid patch` commands are provided for patching operation for database homes and Oracle Grid Infrastructure. Refer to the *Patching Oracle Grid Infrastructure and Oracle Database Using dbaascli* section for details. Also see, *dbaascli Command Reference* section.
- Deprecated the support for `dbaascli tools patch` command to bring consistency in the command conventions. The new command is `dbaascli admin updateStack`.
- Redesigned PDB management APIs for create, local clone, and remote clone operations. For details, see `dbaascli pdb --help`.
- Redesigned database delete API. For details, see `dbaascli database delete --help`.
- Revamped dbhome creation (support for custom software image, scale-out operation). For details, see `dbaascli dbhome create --help`.
- Support for database creation on subset of cluster nodes. For details, see `dbaascli database create --help`.
- Ability to run `dbaascli` in disconnected mode for long running operations. Executing `dbaascli` command with `--waitForCompletion false` gets you a job ID that can be queried later to get the status of the operation, using `dbaascli job getStatus -jobid job_id`. This is useful for long running operations where users may want to get the control back immediately after command execution. In this release, this option is available only for `dbaascli database create` command. More commands will be added in subsequent releases to have this support. The help output for those commands will reflect the support of `--waitForCompletion` option.
- Enhanced dbhome patching experience with introduction of multiple options like `skipPDBs`, `continueWithDowntime`, and so on. For details, see `dbaascli dbhome patch --help`.
- Support for better diagnostic collection. For details, see `dbaascli diag collect --help`.
- Minor improvements in the area of database upgrade automation.
- Various bug fixes and stability improvements.

dbaascli Command Reference

You must use `dbaascli` to create databases and integrate them with the cloud automation framework.

`dbaascli` is a cloud native interface that can take DBCA templates as inputs, calls the functionality of DBCA to create databases, and then calls OCI APIs to integrate the database into the cloud automation framework. Customers using DBCA in scripts today can update their existing scripts to call `dbaascli` instead of DBCA. If `dbaascli` cannot be used due to a particular feature of DBCA being unavailable in `dbaascli`, then customers should open a My Oracle Support (MOS) request to add that functionality to `dbaascli`.

To use the `dbaascli` utility, you must be connected to an Exadata Cloud Infrastructure compute node. See [Connecting to an Exadata Cloud Infrastructure Instance](#) for instructions.

Some `dbaascli` commands can be run as the `oracle` or the `opc` user, but many commands require `root` administrator privileges. Refer to each command for specific requirements.

- [dbaascli admin updateAHF](#)
To install or update Autonomous Health Framework (AHF), use the `dbaascli admin updateAHF` command.

- [dbaascli admin updateStack](#)
To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.
- [dbaascli cswlib deleteLocal](#)
To delete the local image, use the `dbaascli cswlib deleteLocal` command.
- [dbaascli cswlib download](#)
To download available software images and make them available in your Exadata Cloud Infrastructure environment, use the `dbaascli cswlib download` command.
- [dbaascli cswlib listLocal](#)
To view the list of locally available Database and Grid Infrastructure images, use the `dbaascli cswlib listLocal` command.
- [dbaascli cswlib showImages](#)
To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.
- [dbaascli database addInstance](#)
To add the database instance on the specified node, use the `dbaascli database addInstance` command.
- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.
- [dbaascli database bounce](#)
To shut down and restart a specified Exadata Cloud Infrastructure database, use the `dbaascli database bounce` command.
- [dbaascli database changepassword](#)
To change the password of a specified Oracle Database user, use the `dbaascli database changePassword` command. When prompted enter the user name for which you want to change the password and then enter the password.
- [dbaascli database convertToPDB](#)
To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.
- [dbaascli database create](#)
To create Oracle Database, use the `dbaascli database create` command. When prompted, enter the `sys` and `tde` passwords.
- [dbaascli database delete](#)
To delete an Oracle Database, use the `dbaascli database delete` command.
- [dbaascli database deleteInstance](#)
To delete the database instance on the specified node, use the `dbaascli database deleteInstance` command.
- [dbaascli database duplicate](#)
To create a database from an active database, use the `dbaascli database duplicate` command.
- [dbaascli database getDetails](#)
This command shows the detailed information of a given database e.g. `dbname`, node information, pluggable databases information etc.
- [dbaascli database getPDBs](#)
To view the list of all pluggable databases in a container database, use the `dbaascli database getPDBs` command.

- [dbascli database modifyParameters](#)
To modify or reset initialization parameters for an Oracle Database, use the `dbascli database modifyParameters` command.
- [dbascli database move](#)
To move the database from one home to another, use the `dbascli database move` command.
- [dbascli database recover](#)
To recover a database, use the `dbascli database recover` command.
- [dbascli database runDatapatch](#)
To patch an Oracle Database, use the `dbascli database runDatapatch` command.
- [dbascli database createTemplate](#)
Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.
- [dbascli database start](#)
To start an Oracle Database, use the `dbascli database start` command.
- [dbascli database status](#)
To check the status of an Oracle Database, use the `dbascli database status` command.
- [dbascli database stop](#)
To stop an Oracle Database, use the `dbascli database stop` command.
- [dbascli database upgrade](#)
To upgrade an Oracle Database, use the `dbascli database upgrade` command.
- [dbascli dataguard prepareStandbyBlob](#)
To generate a blob file containing various files that are required on the standby site in case of a dataguard environment, use the `dbascli dataguard prepareStandbyBlob` command.
- [dbascli dataguard updateDGConfigAttributes](#)
To update Data Guard automation attributes across all the cluster nodes, use the `dbascli dataguard updateDGConfigAttributes` command.
- [dbascli dataguard failover](#)
To perform a manual failover to the standby database, use the `dataguard failover` command.
- [dbascli dataguard reinstate](#)
To reinstate a failed database as a standby database after a failover, use the `dataguard reinstate` command.
- [dbascli dataguard switchover](#)
To perform a switchover to the standby database, use the `dataguard switchover` command.
- [dbascli dataguard prepareForStandby](#)
To create an Oracle Standby Database, use the `dbascli dataguard prepareForStandby` command as the **first step**.
- [dbascli dataguard configureStandby](#)
To create a new Standby Database, use the `dbascli dataguard configureStandby` command as a second step after the `prepareForStandby` step.

- [dbaascli dataguard registerStandby](#)
To register a newly created standby database with all existing standby databases and to primary database, use the `dbaascli dataguard registerStandby` command as a third step after `configureStandby` step.
- [dbaascli dataguard deregisterStandby](#)
During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.
- [dbaascli dataguard configureAWR](#)
To enable or disable Automatic Workload Repository (AWR) configuration on your Active Data Guard standby, use the `dbaascli dataguard configureAWR` command.
- [dbaascli dataguard updateConfiguration](#)
To update the transport mode or protection mode or both the parameters of a Data Guard environment, use the `dbaascli dataguard updateConfiguration` command.
- [dbaascli dbhome create](#)
To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.
- [dbaascli dbHome delete](#)
To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.
- [dbaascli dbhome getDatabases](#)
To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.
- [dbaascli dbHome getDetails](#)
To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.
- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- [dbaascli dbimage purge](#)
The `dbimage purge` command removes the specified software image from your Exadata Cloud Infrastructure environment.
- [dbaascli diag collect](#)
To collect diagnostics, use the `dbaascli diag collect` command.
- [dbaascli diag healthCheck](#)
To run diagnostic health checks, use the `dbaascli diag healthCheck` command.
- [dbaascli gridHome create](#)
To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.
- [dbaascli grid configureTCPS](#)
To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` command.
- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- [dbaascli grid removeTCPSCert](#)
To remove existing TCPS certificates from Grid Infrastructure wallet, use the `dbaascli grid removeTCPSCert` command.

- [dbaascli grid rotateTCPSCert](#)
To rotate TCPS certificates, use the `dbaascli grid rotateTCPSCert` command.
- [dbaascli grid upgrade](#)
To upgrade Oracle Grid Infrastructure from one major version to another, use the `dbaascli grid upgrade` command.
- [dbaascli job getStatus](#)
To view the status of a specified job, use the `dbaascli job getStatus` command.
- [dbaascli patch db apply](#)
- [dbaascli patch db prereq](#)
- [dbaascli pdb backup](#)
To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.
- [dbaascli pdb bounce](#)
To bounce a pluggable database (PDB), use the `dbaascli pdb bounce` command.
- [dbaascli pdb close](#)
To close a pluggable database (PDB), use the `dbaascli pdb close` command.
- [dbaascli pdb getConnectionString](#)
To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectionString` command.
- [dbaascli pdb create](#)
To create a new pluggable database (PDB), use the `dbaascli pdb create` command.
- [dbaascli pdb delete](#)
To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.
- [dbaascli pdb getDetails](#)
To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.
- [dbaascli pdb list](#)
To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.
- [dbaascli pdb localClone](#)
To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.
- [dbaascli pdb open](#)
To open a pluggable database (PDB), use the `dbaascli pdb open` command.
- [dbaascli pdb recover](#)
To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.
- [dbaascli pdb refresh](#)
To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.
- [dbaascli pdb relocate](#)
To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.
- [dbaascli pdb remoteClone](#)
To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.

- [dbaascli system getDBHomes](#)
To view information about all the Oracle homes, use the `dbaascli system getDBHomes` command.
- [dbaascli system getGridHomes](#)
To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.
- [dbaascli tde changePassword](#)
To change TDE keystore password as well as DB wallet password for the alias `tde_ks_passwd`, use the `dbaascli tde changePassword` command.
- [dbaascli tde addSecondaryHsmKey](#)
To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbaascli tde addSecondaryHsmKey` command.
- [dbaascli tde enableWalletRoot](#)
To enable `wallet_root` spfile parameter for the existing database, use the `dbaascli tde enableWalletRoot` command.
- [dbaascli tde encryptTablespacesInPDB](#)
To encrypt all the tablespaces in the specified PDB, use the `dbaascli tde encryptTablespacesInPDB` command.
- [dbaascli tde fileToHsm](#)
To convert FILE based TDE to HSM (KMS/OKV) based TDE, use the `dbaascli tde fileToHsm` command.
- [dbaascli tde getHsmKeys](#)
To get TDE active key details, use the `dbaascli tde getHsmKeys` command.
- [dbaascli tde getMkidForKeyVersionOCID](#)
To get Master Key ID associated with the KMS key version OCID, use the `dbaascli tde getMkidForKeyVersionOCID` command.
- [dbaascli tde getPrimaryHsmKey](#)
To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.
- [dbaascli tde hsmToFile](#)
To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.
- [dbaascli tde listKeys](#)
To list TDE master keys, use the `dbaascli tde listKeys` command.
- [dbaascli tde removeSecondaryHsmKey](#)
To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.
- [dbaascli tde rotateMasterKey](#)
Rotate the master key for database encryption.
- [dbaascli tde setKeyVersion](#)
To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.
- [dbaascli tde setPrimaryHsmKey](#)
To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbaascli tde setPrimaryHsmKey` command.
- [dbaascli tde status](#)
To display information about the keystore for the specified database, use the `dbaascli tde status` command.

dbaascli admin updateAHF

To install or update Autonomous Health Framework (AHF), use the `dbaascli admin updateAHF` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli admin updateAHF
{
  --targetVersion value | --imageTag value
}
[--resume [--sessionID value]] [--executePrereqs]
```

Where:

- `--targetVersion` specifies the target version to update AHF to
- `--imageTag` specifies the image tag of the AHF artifact to be installed
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` runs prerequisite checks and reports the results

dbaascli admin updateStack

To install or update a dbaastools RPM, use the `dbaascli admin updateStack` command.

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud Infrastructure virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli admin updateStack
[--resume]
[--prechecksOnly]
[--nodes]
```

Where:

- `--resume` resumes the previous execution
- `--prechecksOnly` runs only the prechecks for this operation
- `--nodes` specifies a comma-delimited list of nodes to install the RPM on. If you do not pass this argument, then the RPM will be installed on all of the cluster nodes

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib deleteLocal

To delete the local image, use the `dbaascli cswlib deleteLocal` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswLib deleteLocal --imageTag <value>
```

Where:

- `--imageTag` specifies Oracle home image tag

Example 6-4 dbaascli cswlib deletelocal

```
dbaascli cswlib deletelocal --imagnetag 19.15.0.0.0
DBAAS CLI version MAIN
Executing command cswlib deletelocal --imagnetag 19.15.0.0.0
Job id: 8b3e71de-4b81-4832-b49c-7f892179bb4f
Log file location: /var/opt/oracle/log/cswLib/deleteLocal/
dbaastools_2022-07-18_10-00-02-AM_73658.log
dbaascli execution completed
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib download

To download available software images and make them available in your Exadata Cloud Infrastructure environment, use the `dbaascli cswlib download` command.

Prerequisites

Run the command as the `root` user.

To use the utility, you must connect to an Exadata Cloud Infrastructure virtual machine.

See, *Connecting to a Virtual Machine with SSH*.

Syntax

```
dbaascli cswlib download --version | --imageTag
[--product]
```

Where:

- `--version` specifies an Oracle home image version

- `--imageTag` specifies the image tag of the image
- `--product` specifies the image type. Valid values: `database` or `grid`

Example 6-5 dbaascli cswlib download --product --imageTag

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Example 6-6 dbaascli cswlib download --version 19.9.0.0.0

```
dbaascli cswlib download --product database --imageTag 19.14.0.0.0
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib listLocal

To view the list of locally available Database and Grid Infrastructure images, use the `dbaascli cswlib listLocal` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswLib listLocal [--product <value>]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Example 6-7 dbaascli cswlib listlocal

```
dbaascli cswlib listlocal
DBAAS CLI version MAIN
Executing command cswlib listlocal
Job id: bc4f047c-0a34-4d4d-alea-21ddc2a9c627
Log file location: /var/opt/oracle/log/cswLib/listLocal/
dbaastools_2022-07-18_10-29-53-AM_16077.log
##### List of Available Database Images #####
1.IMAGE_TAG=12.2.0.1.220419
  IMAGE_SIZE=5GB
  VERSION=12.2.0.1.220419
  DESCRIPTION=12.2 APR 2022 DB Image
2.IMAGE_TAG=18.16.0.0.0
  IMAGE_SIZE=6GB
  VERSION=18.16.0.0.0
  DESCRIPTION=18c OCT 2021 DB Image
3.IMAGE_TAG=19.14.0.0.0
  IMAGE_SIZE=5GB
  VERSION=19.14.0.0.0
  DESCRIPTION=19c JAN 2022 DB Image
dbaascli execution completed
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli cswlib showImages

To view the list of available Database and Grid Infrastructure images, use the `dbaascli cswlib showImages` command.

Run the command as the `root` user.

Syntax

```
dbaascli cswlib showImages  
[--product]
```

Where:

- `--product` identifies Oracle home product type. Valid values: `database` or `grid`.

Example 6-8 dbaascli cswlib showImages

```
dbaascli cswlib showImages
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli database addInstance

To add the database instance on the specified node, use the `dbaascli database addInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database addInstance --dbname <value> --node <value> [--newNodeSID  
<value>]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for the database instance
- `--newNodeSID` specifies SID for the instance to add in the new node

dbaascli database backup

To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database backup --dbname <value>
{
  --list
  {
    [--backupType <value>]
    | [--json <value>]
  }
  | --start [--level0] [--level1]
  {
    [--archival --tag <value>]
    | [--archiveLog]
  }
  | --delete --backupTag <value>
  | --status --uuid <value> [--json <value>]
  | --getBackupReport
  {
    --tag <value>
    | --latest
  }
  --json <value>
  | --configure
  {
    --configFile <value>
    | --enableRTRT
    | --disableRTRT
    | --disableCatalog
    | --deleteImmutableConfiguration
  }
  | --getConfig
  {
    [--configFile <value>]
    | [--showOldParams]
  }
  | --validate [--untilTime <value>]
  | --showHistory [--all]
  | --getSchedules
}

```

Where:

- `--dbname` specifies Oracle Database name
- `--list` returns database backup information
[`--backupType` | `--json`]

- [--backupType specifies backupType (REGULAR-L0 | REGULAR-L1 | ARCHIVELOG | LONGTERM)]
- [--json specifies file Name for JSON output]
- --start begins database backup
 - [--level0 creates a Level-0 (full) backup]
 - [--level1 creates a Level-1 (incremental) backup]
 - [--archival | --archivelog]
 - [--archival creates an archival full backup]
 - tag specifies backup tag
 - [--archivelog]
- --delete deletes Archival backup
 - backupTag specifies backup tag to delete
- --status displays the details about a backup job process
 - uuid unique identifier of the backup operation. Input format: xxxxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx.
 - [--json specifies file Name for JSON output]
- --getBackupReport returns BackupReport
 - tag | --latest
 - tag specifies backup tag
 - latest returns latest backup report (all types of database backup)
 - json specifies the file name for JSON output
- --configure configures database for backup
 - configFile | --enableRTRT | --disableRTRT | --disableCatalog | --deleteImmutableConfiguration
 - configFile specifies database backup configuration file
 - enableRTRT enables Real Time Redo Transport
 - disableRTRT disables Real Time Redo Transport
 - disableCatalog disables recovery catalog
 - deleteImmutableConfiguration deletes immutable configuration
- --getConfig returns database backup configuration
 - [--configFile | --showOldParams]
 - [--configFile specifies database backup configuration file]
 - [--showOldParams returns old parameter names of backup configuration.]
- --validate validates that backups are complete and corruption-free
 - [--untilTime validates from closest Level-0 (full) backup until time provided. Input format: DD-MON-YYYY HH24:MI:SS]
- --showHistory displays history of backup operations
 - [--all displays all backup operations]
- --getSchedules returns all backup schedules for a given database

Note

The `enableRTRT` and `disableRTRT` parameters are applicable only for ZDLRA backup destination on Exadata Database Service on Cloud@Customer.

dbaascli database bounce

To shut down and restart a specified Exadata Cloud Infrastructure database, use the `dbaascli database bounce` command.

Prerequisites

Run the command as the `oracle` user.

Syntax

```
dbaascli database bounce
[--dbname][--rolling <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--rolling` specifies `true` or `false` to bounce the database in a rolling manner. Default value is `false`.

The command performs a database shutdown in immediate mode. The database is then restarted and opened. In Oracle Database 12c or later, all of the PDBs are also opened.

Example 6-9 dbaascli database bounce

```
dbaascli database bounce --dbname dbname
```

dbaascli database changepassword

To change the password of a specified Oracle Database user, use the `dbaascli database changePassword` command. When prompted enter the user name for which you want to change the password and then enter the password.

Prerequisites

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli database changePassword [--dbname <value>] [--user <value>]
{
  [--prepareStandbyBlob <value> [--blobLocation <value>]] | [--
standbyBlobFromPrimary <value>]
}
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the Oracle Database that you want to act on

- `--user` specifies the user name whose password change is required
- `--prepareStandbyBlob` specifies `true` to generate a blob file containing the artifacts needed to change the password in a Data Guard environment. Valid values: `true|false`
- `--blobLocation` specifies the custom path where blob file will be generated
- `--standbyBlobFromPrimary` specifies the standby blob file, which is prepared from the primary database
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID

Example 6-10 dbaascli database changePassword

```
dbaascli database changepassword --dbname db19
```

dbaascli database convertToPDB

To convert the specified non-CDB database to PDB, use the `dbaascli database convertToPDB` command.

Syntax

```
dbaascli database convertToPDB --dbname <value> [--cdbName <value>] [--executePrereqs]
    {
        [--copyDatafiles [--keepSourceDB]]|[backupPrepared]
    }
    [--targetPDBName <value>] [--waitForCompletion <value>] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of Oracle Database
- `--cdbName` specifies the name of the target CDB in which the PDB will be created. If the CDB does not exist, then it will be created in the same Oracle home as the source non-CDB
- `--executePrereqs` specifies to run only the pre-conversion checks
- `--copyDatafiles` specifies to create a new copy of the data files instead of using the ones from the source database
 - `--keepSourceDB` - to preserve the source database after completing the operation.
- `--backupPrepared` - flag to acknowledge that a proper database backup is in place for the non CDB prior to performing the conversion to PDB.
- `--backupPrepared` flag to acknowledge that a proper database backup is in place for the non-CDB prior to performing the conversion to PDB
- `--targetPDBName` specifies the name of the PDB that will be created as part of the operation
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID

Example 6-11 dbascli database convertToPDB

To run pre-conversion prechecks:

```
dbascli database convertToPDB --dbname ndb19 --cdbname cdb19 --
backupPrepared --executePrereqs
```

To run a full conversion with a copy of the data files from the non-CDB:

```
dbascli database convertToPDB --dbname tst19 --cdbname cdb19 --copyDatafiles
```

dbascli database create

To create Oracle Database, use the `dbascli database create` command. When prompted, enter the `sys` and `tde` passwords.

Use this command to create Oracle Database version 12.1.0.2 or higher with the release update JAN 2021 or higher. For databases with lower versions, it is recommended to use the OCI Console based API.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbascli database create --dbName {--oracleHome | --oracleHomeName}
[--dbUniqueName <value>]
[--dbSID <value>]
[--createAsCDB <value>]
[--pdbName <value>]
[--pdbAdminUserName <value>]
[--dbCharset <value>]
[--dbNCharset <value>]
[--dbLanguage <value>]
[--dbTerritory <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--nodeList <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
{
    [--resume [--sessionID <value>]]
    | [--revert [--sessionID <value>]]
}
[--executePrereqs]
[--honorNodeNumberForInstance <value>]
[--lockPDBAdminAccount <value>]
[--dbcaTemplateFilePath <value>]
[--waitForCompletion]
```

Where:

- `--dbname` specifies the name of the database
- `--oracleHome` specifies the location of the Oracle home
- `--oracleHomeName` specifies the name of the Oracle home
- `--dbUniqueName` specifies database unique name
- `--dbSID` specifies the SID of the database
- `--createAsCDB` specify `true` or `false` to create database as CDB or Non-CDB
- `--pdbName` specify PDB name
- `--pdbAdminUserName` specify PDB admin user name
- `--dbCharset` specifies database character set
- `--dbNCharset` specifies database national character set
- `--dbLanguage` specifies the database language
- `--dbTerritory` specifies the database territory
- `--sgaSizeInMB` specifies the `sga_target` value in megabyte unit
- `--pgaSizeInMB` specifies the `pga_aggregate_target` value in megabyte unit
- `--datafileDestination` specifies the ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database Fast Recovery Area
- `--fraSizeInMB` specifies the Fast Recovery Area size value in megabyte unit
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--tdeConfigMethod` specifies TDE configuration method. Valid values: `FILE`, `KMS`
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--resume` resumes the previous execution
- `--revert` rolls back the previous run
- `--sessionID` to resume or revert a specific session id.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--honorNodeNumberForInstance` specifies `true` or `false` to indicate instance name to be suffixed with the cluster node numbers. Default value: `true`
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`
- `--dbcaTemplateFilePath` specify the absolute path of the dbca template name to create the database.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`

Example 6-12 dbascli database create

```
dbascli database create --dbName db19 --oracleHomeName myhome19 --dbSid db19sid --nodeList node1,node2 --createAsCDB true
```

dbaascli database delete

To delete an Oracle Database, use the `dbaascli database delete` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database delete --dbname <value>
[--deleteArchiveLogs <value>]
[--deleteBackups <value>]
[--precheckOnly <value>]
[--waitForCompletion <value>]
[--force]
[--dbSID <value>]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the database.
- `--deleteArchiveLogs` specifies `true` or `false` to indicate deletion of database archive logs.
- `--deleteBackups` specifies `true` or `false` to indicate deletion of database backups.
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`.
- `--force` flag to force delete database.
- `--dbSID` specify database SID.
- `--resume` to resume the previous execution.
- `--sessionID` to resume a specific session id.

Example 6-13 dbaascli database delete

```
dbaascli database delete --dbname db19
```

dbaascli database deleteInstance

To delete the database instance on the specified node, use the `dbaascli database deleteInstance` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database deleteInstance --dbname <value> --node <value> [--continueOnUnreachableNode]
```

Where:

- `--dbname` specifies Oracle Database name
- `--node` specifies the node name for database instance
- `--continueOnUnreachableNode` specifies to perform the operation even if the node is unreachable

Example 6-14 database deleteinstance

```
database deleteinstance --node test-node
```

dbaascli database duplicate

To create a database from an active database, use the `dbaascli database duplicate` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli database duplicate --dbName <value> --sourceDBConnectionString <value>
{
  --oracleHome <value>
  | --oracleHomeName <value>
}
[--dbSID <value>]
[--dbUniqueName <value>]
[--sgaSizeInMB <value>]
[--pgaSizeInMB <value>]
[--datafileDestination <value>]
[--fraDestination <value>]
[--fraSizeInMB <value>]
[--sourceDBWalletLocation <value>]
[--nodeList <value>]
{
  --resume [--sessionID <value>]
  | [--revert [--sessionID <value>]]
}
[--rmanParallelism <value>]
[--rmanSectionSizeInGB <value>]
[--tdeConfigMethod <value>]
[--kmsKeyOCID <value>]
[--sourceDBTdeConfigMethod <value>]
[--sourceDBKmsKeyOCID <value>]
[--executePrereqs <value>]
```

```
[--waitForCompletion <value>]  
[--skipPDBs <value>]
```

Where:

- `--dbName` specifies Oracle Database name
- `--sourceDBConnectionString` specifies source database connection string in the format of `<scan_name>:<scan_port>/<database_service_name>`
- `--oracleHome` specifies Oracle home location
- `--oracleHomeName` specifies Oracle home name
- `--dbSID` specifies database SID
- `--dbUniqueName` specifies database unique name
- `--sgaSizeInMB` specifies `sga_target` value in mega byte unit
- `--pgaSizeInMB` specifies `pga_aggregate_target` value in mega byte unit
- `--datafileDestination` specifies ASM disk group name to use for database datafiles
- `--fraDestination` specifies ASM disk group name to use for database fast recovery area
- `--fraSizeInMB` specifies fast recovery area size value in mega byte unit
- `--sourceDBWalletLocation` specifies source database TDE wallet file location. This is required to duplicate database from active database
- `--nodeList` specifies a comma-delimited list of nodes for the database
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to rollback the previous execution
 - `--sessionID` specifies to rollback a specific session ID
- `--rmanParallelism` specifies `parallelsim` value
- `--rmanSectionSizeInGB` specifies RMAN section size in GB
- `--tdeConfigMethod` specifies TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--kmsKeyOCID` specifies KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--sourceDBTdeConfigMethod` specifies source database TDE configuration method. Allowed values are `FILE` and `KMS`.
- `--sourceDBKmsKeyOCID` specifies source database KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE.
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`
- `--skipPDBs` specifies a comma-delimited list of source database PDB names, which needs to be excluded for the duplicate database operation. Example: `pdb1,pdb2...`

Example 6-15 dbascli database duplicate

```
dbascli database duplicate --sourceDBConnectionString test-user-  
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/  
mynew.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com --oracleHome /u02/app/  
oracle/product/19.0.0.0/dbhome_2 --dbName newdup --  
sourceDBWalletLocation /var/opt/oracle/dbaas_acfs/tmp/prim_wallet
```

dbascli database getDetails

This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

Prerequisites

Run the command as the `root` user or the `oracle` user

Syntax

```
dbascli database getDetails --dbname <value>
```

Where :

- `--dbname` - Oracle database name.

dbascli database getPDBs

To view the list of all pluggable databases in a container database, use the `dbascli database getPDBs` command.

Run the command as the `root` or `oracleuser`.

Syntax

```
dbascli database getPDBs --dbname <value>
```

Where:

- `--dbname` specifies the name of the container database

Example 6-16 dbascli database getPDBs --dbname

```
dbascli database getPDBs --dbname apr_db1
```

dbascli database modifyParameters

To modify or reset initialization parameters for an Oracle Database, use the `dbascli database modifyParameters` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database modifyParameters --dbname <value> --setParameters <values> |
--resetParameters <values> | --responseFile
[--backupPrepared]
[--instance]
[--allowBounce]
```

Where:

- `--dbname` specifies the name of the database.
- `--setParameters` specifies a comma-delimited list of parameters to modify with new values. For example: `parameter1=valueA,parameter2=valueB`, and so on. For blank values use `parameter1=valueA,parameter2=""`, etc.
- `--resetParameters` specifies a comma-delimited list of parameters to be reset to their corresponding default values. For example, `parameter1,parameter2`, and so on.
- `--responseFile` specifies the absolute location of the response JSON file to modify the database parameters
- `--backupPrepared` acknowledges that a proper database backup is in place prior to modifying critical or sensitive parameters.
- `--instance` specifies the name of the instance on which the parameters will be processed. If not specified, then the operation will be performed at the database level.
- `--allowBounce` grants permission to bounce the database in order to reflect the changes on applicable static parameters.

Example 6-17 dbaascli database modifyParameters

```
dbaascli database modifyParameters --dbname dbname --setParameters
"log_archive_dest_state_17=ENABLE"
```

dbaascli database move

To move the database from one home to another, use the `dbaascli database move` command.

Prerequisites

- Before performing a move operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database move
{
  --oracleHome <value>
  | --oracleHomeName <value>
}
--dbname <value>
{
  [--resume [--sessionID <value>]]
  | [--rollback [--sessionID <value>]]
}
```

```

    }
  [--executePrereqs]
  [--nonRolling]
  [--skipDatapatch]
  [--skipPDBs <value>]
  [--skipClosedPDBs]
  [--continueWithDbDowntime]
  [--drainTimeoutInSeconds <value>]
  [--allowParallelDBMove]
  [--waitForCompletion <value>]
  [--nodeList <value>]

```

Where:

- `--oracleHome` specifies Oracle home path
- `--oracleHomeName` specifies the name of Oracle home
- `--dbname` specifies the name of the database
- `--executePrereqs` runs the prerequisite checks and report the results
- `--nonRolling` moves the database on all nodes in parallel.
Note that this would cause database downtime.
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--rollback` rolls the database back to previous home
 - `--sessionID` specifies to resume a specific session ID
- `--skipDatapatch` skips running the datapatch on the databases
- `--skipPDBs` skips running the datapatch on a specified comma-delimited list of PDBs.
For example: *pdb1,pdb2...*
- `--skipClosedPDBs` skips patching closed PDBs
- `--continueWithDbDowntime` continues patching with database downtime.
This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--drainTimeoutInSeconds` specifies the time (in seconds) to complete the resource draining while stopping the database
- `--allowParallelDBMove` allows database move in parallel.
- `--waitForCompletion` specifies false to run the operation in the background.
Valid values: true|false
- `--nodeList` specifies a comma-delimited list of nodes if operation has to be performed on a subset of nodes

Example 6-18 dbascli database move

```

dbascli database move --dbname testdb1 --oracleHome /u02/app/oracle/product/
12.1.0/dbhome_2

```

dbaascli database recover

To recover a database, use the dbaascli database recover command.

Prerequisite

- Run the command as the `root` user.
- Database must have been configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli database recover --dbname <value>
{
  --start
  {
    --untilTime <value>
    | --untilSCN <value>
    | --latest
    | --tag <value>
  }
  | --status --uuid <value>
}
```

Where:

```
--dbname: Oracle Database name.
--start | --status
--start: Begins database recovery.
--untilTime | --untilSCN | --latest | --tag
--untilTime: Recovers database until time. Input format: DD-MON-YYYY
HH24:MI:SS.
--untilSCN: Recovers database until SCN.
--latest: Recovers database to last known state.
--tag: Recovers database to archival tag.
--status
--uuid <value>
```

Example 6-19 Examples

- To recover the database *myTestDb* to latest:

```
dbaascli database recover --dbname myTestDb --start --latest
```

- To query the status of recovery request submitted with `uuid 2508ea18be2911eb82d0020017075151`:

```
dbaascli database recover --dbname myTestDb --status --uuid
2508ea18be2911eb82d0020017075151
```

dbaascli database runDatapatch

To patch an Oracle Database, use the `dbaascli database runDatapatch` command.

Prerequisites

- Before performing a `runDatapatch` operation, ensure that all of the database instances associated with the database are up and running.
- Run the command as the `root` user.

Syntax

```
dbaascli database runDatapatch --dbname  
[--resume]  
    [--sessionID]  
[--skipPdfs | --pdfs]  
[--executePrereqs]  
[--patchList]  
[--skipClosedPdfs]  
[--rollback]
```

Where:

- `--dbname` specifies the name of the database
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--skipPdfs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: `pdb1,pdb2...`
- `--pdfs` runs the datapatch only on a specified comma-delimited list of PDBs. For example: `pdb1,pdb2...`
- `--executePrereqs` runs prerequisite checks
- `--patchList` applies or rolls back the specified comma-delimited list of patches. For example: `patch1,patch2...`
- `--skipClosedPdfs` skips running the datapatch on closed PDBs
- `--rollback` rolls back the patches applied

```
dbaascli database runDatapatch --dbname db19
```

dbaascli database createTemplate

Use this command to create database templates (DBCA templates) that can subsequently be used to create databases.

Run the command as the `root` or `oracle` user.

Syntax

Create a new DBCA template from the specified database.

```

dbaascli database createTemplate --dbname <value>
{
  --templateLocation <value> | --uploadToObjectStorage --
objectStorageLoginUser <value> --objectStorageBucketName <value> [--
objectStorageUrl <value>]
}
[--templateName <value>] [--rmanParallelism <value>]

```

Where:

- `--dbname` specifies the name of the database
- `--templateLocation` specifies the template name
- `--uploadToObjectStorage`: specifies to upload the template to Object Storage
 - `--objectStorageLoginUser`: specifies the Object Storage login user
 - `--objectStorageBucketName`: specifies the Object Storage bucket name
 - `--objectStorageUrl`: specifies the Object Storage URL
- `--templateName`: specifies the name of the template
- `--rmanParallelism` specifies the parallelism value

dbaascli database start

To start an Oracle Database, use the `dbaascli database start` command.

Prerequisites

Run the command as the `root` user.

Syntax

```

dbaascli database start
[--dbname]
[--mode]

```

Where:

- `--dbname` specifies the name of the database
- `--mode` specifies mount or nomount to start database in the corresponding mode

The command starts and opens the database. In Oracle Database 12c or later, all of the PDBs are also opened.

Example 6-20 dbaascli database start

```

dbaascli database start --dbname dbname --mode mount

```

dbaascli database status

To check the status of an Oracle Database, use the `dbaascli database status` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database status
[--service][--dbname]
[--user]
[--password]
```

Where:

- `--service` specifies the name of the service
- `--dbname` specifies the name of the database
- `--user` specifies the user name of the service
- `--password` specifies the password of the user

Output from the command includes the open mode of the database, the software release and edition of the database, and release version of other software components.

Example 6-21 dbaascli database status

```
dbaascli database status --dbname db19
```

dbaascli database stop

To stop an Oracle Database, use the `dbaascli database stop` command.

Prerequisites

Run the command as the `root` user.

Syntax

```
dbaascli database stop
[--dbname <value>]
[--mode <value>]
```

Where:

- `--dbname` specifies the name of the database that you want to stop
- `--mode` specifies the mode of the database. Valid values: `abort`, `immediate`, `normal`, `transactional`

The command performs a database shutdown in immediate mode. No new connections or new transactions are permitted. Active transactions are rolled back, and all connected users are disconnected.

Example 6-22 dbaascli database stop

```
dbaascli database stop --dbname db19
```

dbaascli database upgrade

To upgrade an Oracle Database, use the `dbaascli database upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli database upgrade --dbname <value>
{--targetHome <value> | --targetHomeName <value>}
{ [--executePrereqs | --postUpgrade | --rollback]}
{[--standBy | --allStandbyPrepared]}
{[--upgradeOptions <value>] | [--standBy]}
[--removeGRP]
[--increaseCompatibleParameter]
[--resume [--sessionID <value>]]
[--waitForCompletion <value>]
```

Where:

- `--dbname` (mandatory) specifies the name of the database.
- `--targetHome` specifies the target Oracle home location
- `--targetHomeName` specifies the name of the target Oracle Database home
- `--standBy` use this option to upgrade standby databases in Data Guard configurations
- `--allStandbyPrepared` required for Data Guard configured primary databases. Flags to acknowledge that all the required operations are performed on the standby databases prior to upgrading primary database
- `--removeGRP` automatically removes the Guaranteed Restore Point (GRP) backup only if the database upgrade was successful
- `--increaseCompatibleParameter` automatically increases the compatible parameter as part of the database upgrade. The parameter will get increased only if the database upgrade was successful
- `--executePrereqs` runs only the preupgrade checks
- `--postUpgrade` use this option if postupgrade fails and needs to rerun the postupgrade steps
- `--rollback` reverts an Oracle Database to its original Oracle home
- `--upgradeOptions` use this option to pass DBUA-specific arguments to perform the Oracle Database upgrade. Refer to the corresponding Oracle documentation for the supported arguments and options.
 - `--standby`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false`.

Example 6-23 dbaascli database upgrade pre-upgrade requisite checks

```
dbaascli database upgrade --dbname dbname --targetHome Target Oracle home  
location --executePrereqs
```

dbaascli dataguard prepareStandbyBlob

To generate a blob file containing various files that are required on the standby site in case of a dataguard environment, use the `dbaascli dataguard prepareStandbyBlob` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli dataguard prepareStandbyBlob --dbname <value> --blobLocation <value>
```

Where:

- `--dbname` specifies the Oracle Database name
- `--blobLocation` specifies the custom directory location where the standby blob file will be generated in a Data Guard environment

dbaascli dataguard updateDGConfigAttributes

To update Data Guard automation attributes across all the cluster nodes, use the `dbaascli dataguard updateDGConfigAttributes` command.

Run the command as the `root` or `oracleuser`.

Syntax

```
dbaascli dataguard updateDGConfigAttributes --attributes <value>
```

Where:

- `--attributes` contains the Data Guard automation attributes that are to be modified. Accepts comma-delimited values in the format `<attribute=value>`. Attributes must be predefined in the Data Guard configuration file.

dbaascli dataguard failover

To perform a manual failover to the standby database, use the `dataguard failover` command.

Run this command as the `oracle` user on the target standby database.

Syntax

```
dbaascli dataguard failover --dbname <value> [--useImmediateFailover] [--  
executePrereqs] [--waitForCompletion <value>] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--useImmediateFailover` use this flag when the Oracle Data Guard configuration in a warning or error state.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- [Performing a Manual Failover Operation Using the dbaascli Utility](#)
To perform a manual failover to the standby database, use the `dataguard failover` command.

Performing a Manual Failover Operation Using the dbaascli Utility

To perform a manual failover to the standby database, use the `dataguard failover` command.

1. Connect to the virtual machine in the Oracle Data Guard configuration that will host the new primary database as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell and then switch to the `oracle` user:

```
$ sudo -s
# su - oracle
$
```

3. Initiate the failover to the standby database.

```
$ dbaascli dataguard failover --dbname <db-name>
```

4. Return to being the `root` user.

```
$ exit
#
```

5. Exit the root-user command shell and disconnect from the virtual machine.

```
# exit
$ exit
```

dbaascli dataguard reinstate

To reinstate a failed database as a standby database after a failover, use the `dataguard reinstate` command.

Run this command as the `oracle` user on where `reinstate` is required (that is failed standby database).

Syntax

```
dbascli dataguard reinstate --dbname <value> [--primaryDBUniqueName <value>]
[--executePrereqs] [--waitForCompletion <value>] [--resume [--sessionID
<value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--primaryDBUniqueName` specifies database unique name of the current primary database in the Oracle Data Guard setup.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.

To determine when a member should be reinstated in a Data Guard (DG) configuration:

Monitor the `dgmgrl show database` output for the following ORA errors:

- On the new primary cluster:
ORA-16661: The standby database must be reinstated
- On the old primary cluster:
ORA-16623: Member detected role change

These messages indicate that a failover has occurred. To restore full synchronization within the Data Guard configuration, the former primary must be reinstated.

- [Reinstating a Failed Primary Database Using the dbascli Utility](#)
To reinstate a failed primary database after a failover, use the `dataguard reinstate` command.

Reinstating a Failed Primary Database Using the dbascli Utility

To reinstate a failed primary database after a failover, use the `dataguard reinstate` command.

1. Connect to one of the virtual machines in the Oracle Data Guard configuration as the `oracle` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Initiate reinstatement of the failed primary database.

```
$ dbascli dataguard reinstate --dbname <db-name>
```

In a multiple standby setup (Data Guard group), it is recommended to specify the `--primaryDBUniqueName` argument.

```
dbascli dataguard reinstate --dbname <db-name> --primaryDBUniqueName
<primary-DB-unique-name>
```

3. Disconnect from the virtual machine.

```
$ exit
```

dbascli dataguard switchover

To perform a switchover to the standby database, use the `dataguard switchover` command.

Run this command as the `oracle` user.

Syntax

```
dbascli dataguard switchover --dbname <value> [--targetStandbyDBUniqueName  
<value>] [--executePrereqs] [--enableDGDebug] [--waitForCompletion <value>]  
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--targetStandbyDBUniqueName` specifies the standby database unique name to change the role from standby to primary database.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--enableDGDebug` enables the traces while performing the operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- [Performing a Switchover Operation Using the dbascli Utility](#)
To perform a switchover to the standby database, use the `dataguard switchover` command.

Performing a Switchover Operation Using the dbascli Utility

To perform a switchover to the standby database, use the `dataguard switchover` command.

1. Connect to the virtual machine in the Oracle Data Guard configuration that will host the new primary database as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell and then switch to the `oracle` user.

```
$ sudo -s  
# su - oracle  
$
```

3. Initiate the switchover to the standby database.

```
$ dbascli dataguard switchover --dbname <db-name>
```

In a multiple standby setup (Data Guard group), it is recommended to specify the `--targetStandbyDBUniqueName` argument.

```
dbaascli dataguard switchover --dbname <db-name> --
targetStandbyDBUniqueName <target-standby-DB-unique-name>
```

4. Return to being the `root` user.

```
$ exit
#
```

5. Exit the root-user command shell and disconnect from the virtual machine.

```
# exit
$ exit
```

dbaascli dataguard prepareForStandby

To create an Oracle Standby Database, use the `dbaascli dataguard prepareForStandby` command as the **first step**.

Run this command as the `root` user on the primary database. At the end of the command run, a standby BLOB file is created. You must copy this file to the standby database system to proceed with the `configureStandby` step.

Note

For Disaster Recovery (DR) configurations on Exadata Cloud@Customer (ExaDB-C@C), you must use the Oracle Cloud Infrastructure (OCI) Console or the OCI SDK to set up Data Guard. The `dbaascli` utility is not supported for this use case and should not be used.

Syntax

```
dbaascli dataguard prepareForStandby --dbname <value> --standbyDBUniqueName
<value> --standbyDBDomain | --noDBDomain --standbyScanIPAddresses <Standby
SCAN IP Addresses> [ --standbyScanPort ] [ --standbyServiceName ] [ --
primaryScanIPAddresses ] [ --primaryScanPort ] [--executePrereqs] [--resume
[--sessionID <value>]] [--revert [--sessionID <value>]] [--waitForCompletion]
[--skipDRConfiguration]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the standby database unique name for which the primary database will be configured.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.
- `--noDBDomain` specifies not to use the database domain name for standby database.

- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbyServiceName` specifies the name of the standby database service for which the primary database will be configured.
- `--primaryScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the primary database SCAN listener, or the SCAN name of the primary database.
- `--primaryScanPort` specifies the corresponding SCAN port number of the primary database.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--skipDRConfiguration` specifies whether to skip Disaster Recovery (DR) configuration as part of the standby database setup. Valid values: `true` (skip DR configuration) or `false` (configure DR).
- [Performing PrepareForStandby Operation Using the dbascli Utility](#)
To prepare the primary database for creating a new standby database, use the `dbascli dataguard prepareForStandby` command.

Performing PrepareForStandby Operation Using the dbascli Utility

To prepare the primary database for creating a new standby database, use the `dbascli dataguard prepareForStandby` command.

1. Connect to the Virtual Machine where you want to host the primary database as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the `prepareForStandby` command. Enter the `SYS` password when prompted.

```
dbascli dataguard prepareForStandby --dbName <db name> --
standbyDBUniqueName <standby db unique name> --standbyDBDomain <standby db
domain> --standbyScanIPAddresses <comma-delimited list of standby SCAN
IP addresses> --primaryScanIPAddresses <comma-delimited list of primary
SCAN IP addresses> --standbyScanPort <standby SCAN listener port>
```

Upon completion, the command displays the location where the standby BLOB file is created.

- Exit the root-user command shell and disconnect from the virtual machine.

```
# exit
```

dbaascli dataguard configureStandby

To create a new Standby Database, use the `dbaascli dataguard configureStandby` command as a second step after the `prepareForStandby` step.

Run this as the `root` user on the standby cluster.

Syntax

```
dbaascli dataguard configureStandby --dbname <value> --oracleHome <value> |
--oracleHomeName <value> --standbyDBUniqueName <value> [--standbyDBDomain
<value>] | [--noDBDomain] --primaryScanIPAddresses <value> --primaryScanPort
<value> --primaryServiceName <value> --protectionMode <value> --transportType
<value> --activeDG <value> [--standbyBlobFromPrimary <value>] | [--
standbyDBInfoJsonLocation <value>] [--standbyScanIPAddresses <value>] [--
standbyScanPort <value>] [--standbySID <value>] [--nodeList <value>] [--
skipAWRConfiguration] [--primaryDBOCID <value>] [--sgaSizeInMB <value>] [--
pgaSizeInMB <value>] [--datafileDestination <value>] [--fraDestination
<value>] [--redoLogDestination <value>] [--fraSizeInMB <value>] [--
tdeKeyStoreType <value> [--tdeKeyOCID <value>]] [--tdeKeyOCID <value>] [--
executePrereqs] [--resume [--sessionID <value>]] | [--revert [--sessionID
<value>]] --waitForCompletion <value>] [--enableFIPS <value>] [--
skipDRConfiguration] [--okvServer <value> --okvAdminUserName <value> [--
okvServerRestPort <value>]] [--okvWalletName <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--oracleHome` specifies the Oracle home path.
- `--oracleHomeName` specifies the Oracle home name.
- `--standbyDBUniqueName` specifies the database unique name for the standby database.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.
- `--noDBDomain` specifies not to use the database domain name for standby database.
- `--primaryScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the primary database SCAN listener, or the SCAN name of the primary database.
- `--primaryScanPort` specifies the corresponding SCAN port number of the primary database service.
- `--primaryServiceName` specifies the name of the primary database service for which the standby database will be configured.
- `--protectionMode` specifies the Data Guard protection mode to be set when configuring the standby database. Valid values: `MAX_PERFORMANCE|MAX_AVAILABILITY`.
- `--transportType` specifies the Data Guard transport type to be set when configuring the standby database. Valid values: `ASYNC|SYNC`.

- `--activeDG` specifies if the Data Guard configuration will be active or not. Valid values: `true|false`.
- `--standbyBlobFromPrimary` specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.
- `--standbyDBInfoJsonLocation` specifies the location of the info file generated from the primary database for exporting additional metadata. This option is required only for standby operations.
- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbySID` specifies the standby database SID for standby configuration.
- `--nodeList` specifies a list of nodes where the standby database is expected to run, including nodes that are already running or configured.
- `--skipAWRConfiguration` specifies whether to skip Oracle AWR configuration as part of the standby database setup. Valid values: `true` (skip AWR configuration) or `false` (configure AWR).
- `--primaryDBOCID` specifies the resource OCID value corresponding to the primary database.
- `--sgaSizeInMB` specifies `sga_target` value in MB.
- `--pgaSizeInMB` specifies the `pga_aggregate_target` value in MB.
- `--datafileDestination` specifies the storage location to use for database datafiles.
- `--fraDestination` specifies the storage location to use for database fast recovery area.
- `--redoLogDestination` specifies the storage location to use for the redo log files.
- `--fraSizeInMB` specifies fast recovery area size value in MB.
- `--tdeKeyStoreType` specifies the TDE keystore type. Valid values: `FILE|KMS|AZURE|GOOGLE|AWS|OKV`
- `--tdeKeyOCID` specifies `KMS/AZURE/GOOGLE/AWS` key OCID to use for TDE. This is applicable only if `KMS/AZURE/GOOGLE/AWS` is selected for TDE keystore type.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- `--enableFIPS` specifies whether to enable FIPS. Set to `false` to disable it. Valid values: `true|false`.
- `--skipDRConfiguration` specifies whether to skip Disaster Recovery (DR) configuration as part of the standby database setup. Valid values: `true` (skip DR configuration) or `false` (configure DR).
- `--okvServer` specifies the Oracle Key Vault server. Comma-delimited list of multiple IP addresses.

- `--okvAdminUserName` specifies the Oracle Key Vault admin user name.
- `--okvServerRestPort` specifies the REST port number for Oracle Key Vault.
- `--okvWalletName` specifies the Oracle Key Vault wallet name.
- [Performing configureStandby Operation Using the dbascli Utility](#)
To create a standby database, use the `dbascli dataguard configureStandby` command.

Performing configureStandby Operation Using the dbascli Utility

To create a standby database, use the `dbascli dataguard configureStandby` command.

1. Connect to the Virtual Machine where you want to host the standby database as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the `configureStandby` command.

```
dbascli dataguard configureStandby --standbySID <standby SID> --activeDG
<true|false> --dbName <db name> --standbyDBUniqueName <standby db unique
name> --standbyScanName <comma-delimited list of standby SCAN IP
addresses> --protectionMode <MAX_PERFORMANCE|MAX_AVAILABILITY> --
standbyScanPort <standby SCAN port> --oracleHome <oracle home path> --
standbyDBDomain <standby db domain name>. --primaryServiceName <primary db
service name> --transportType <ASYN|SYNC> --primaryScanPort <primary db
SCAN port> --primaryScanIPAddresses <comma-delimited list primary db SCAN
IP addresses> --standbyBlobFromPrimary <standby BLOB from primary>
```

4. Enter the primary database SYS, TDE, and AWR passwords when prompted.

Upon successful completion of the command, the standby database starts and becomes operational.

dbascli dataguard registerStandby

To register a newly created standby database with all existing standby databases and to primary database, use the `dbascli dataguard registerStandby` command as a third step after `configureStandby` step.

Run this command as the `root` user on the primary cluster. Additionally, in a multiple-standby setup, run the command on all standby clusters except the newly created standby database cluster.

Syntax

```
dbascli dataguard registerStandby --dbname <value> --standbyDBUniqueName
<value> --standbyDBDomain <value> | --noDBDomain --standbyScanIPAddresses
<value> [--standbyScanPort <value>] [--standbyServiceName <value>] [--
executePrereqs] [--resume [--sessionID <value>]] | [--revert [--sessionID
<value>]] [--waitForCompletion <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the database unique name of the standby database to be registered with the Oracle Data Guard Broker configuration.
- `--standbyDBDomain` specifies the standby database domain for which the primary database will be configured.
- `--noDBDomain` specifies not to use the database domain name for standby database.
- `--standbyScanIPAddresses` specifies a comma-delimited list of IP addresses corresponding to the standby database SCAN listener, or the SCAN name of the standby database.
- `--standbyScanPort` specifies the corresponding SCAN port number of the standby database.
- `--standbyServiceName` specifies the name of the standby database service for which the primary database will be configured.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--revert` rolls back the previous operation.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- [Performing registerStandby Operation Using the dbascli Utility](#)
To register the specified standby database with the Oracle Data Guard Broker configuration, use the `dbascli dataguard registerStandby` command.

Performing registerStandby Operation Using the dbascli Utility

To register the specified standby database with the Oracle Data Guard Broker configuration, use the `dbascli dataguard registerStandby` command.

For single standby use cases, the `registerStandby` command must be run only on the primary cluster, as there is a one-to-one association between the primary and the standby.

However, in configurations with multiple standby databases, you must run the `registerStandby` command on both the primary cluster and all existing standby clusters—excluding the new standby database being added.

For example, consider a setup with two standby databases: `stdby1` and `stdby2`, where `stdby2` is the new standby to be registered. In this case, run the `registerStandby` command on the primary cluster and on `stdby1`, but not on `stdby2`.

In summary, when adding a new standby database to an existing Oracle Data Guard configuration, run the `registerStandby` command on the primary and on all other previously registered standby clusters, except the new standby being added.

1. Connect to the primary cluster of the Data Guard configuration as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the registerStandby command.

```
dbaascli dataguard registerStandby --dbname <db-name> --
standbyDBUniqueName <standby-DB-unique-name> --standbyDBDomain <standby-DB-
domain>
```

Upon successful completion of the command, the specified standby database will be registered with the Oracle Data Guard Broker configuration.

4. Repeat steps 1 through 3, as performed on the primary cluster, on all existing standby clusters in the Oracle Data Guard Broker configuration, except the one being registered.

dbaascli dataguard deregisterStandby

During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

Run this command as the `root` user on the primary cluster. However, in the context of multiple standby databases, this command needs to be executed on all the standby clusters except the target standby.

Syntax

```
dbaascli dataguard deregisterStandby --dbname <value> --standbyDBUniqueName
<value> [--executePrereqs] [--resume [--sessionID <value>]] [--
waitForCompletion <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--standbyDBUniqueName` specifies the database unique name of the standby database to be deregistered from the Oracle Data Guard Broker configuration.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- [Performing deregisterStandby Operation Using the dbaascli Utility](#)
During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

Performing deregisterStandby Operation Using the dbaascli Utility

During standby deletion, run the `dbaascli dataguard deregisterStandby` command before deleting the database on the standby cluster to deregister the standby database from the Oracle Data Guard Broker configuration.

For single standby use cases, the `deregisterStandby` command must be run only on the primary cluster, as there is a one-to-one association between the primary and the standby.

However, in configurations with multiple standby databases, you must run the `deregisterStandby` command on both the primary cluster and all existing standby clusters—excluding the standby database which is being deregistered currently.

For example, consider a setup with two standby databases: `stdby1` and `stdby2`, where `stdby2` is to be deregistered. In this case, run the `deregisterStandby` command on the primary cluster and on `stdby1`, but not on `stdby2`.

In summary, while deleting a standby database from an existing Oracle Data Guard configuration, run the `deregisterStandby` command on the primary and on all other existing standby clusters before the delete database operation on desired standby cluster.

1. Connect to the primary cluster of the Oracle Data Guard configuration as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the `deregisterStandby` command.

```
dbascli dataguard deregisterStandby --dbname <db-name> --
standbyDBUniqueName <standby-DB-unique-name>
```

Upon successful completion of the command, the specified standby database will be deregistered (removed) from the Oracle Data Guard Broker configuration.

4. Repeat steps 1 through 3, as performed on the primary cluster, on all existing standby clusters in the Oracle Data Guard Broker configuration, except the one being deregistered.

dbascli dataguard configureAWR

To enable or disable Automatic Workload Repository (AWR) configuration on your Active Data Guard standby, use the `dbascli dataguard configureAWR` command.

Run this command as the `root` user on the Active Data Guard standby cluster where you want to enable or disable AWR configuration. Use this command if AWR was not configured during the standby addition process.

Syntax

```
dbascli dataguard configureAWR --dbname <value> { --action <value> | --
enable | --disable } [--executePrereqs] [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--action` specifies whether to enable or disable AWR. Use `--action enable` to enable AWR and `--action disable` to disable it.
The `--action` argument is retained for backward compatibility. However, it is recommended to use `--enable` or `--disable`, as they provide the same functionality but are more explicit
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.

- `--sessionID` resumes a specific session by its ID.
- [Performing configureAWR Operation Using the dbaascli Utility](#)
To configure AWR on an ADG standby database, use the `dbaascli dataguard configureAWR` command.

Performing configureAWR Operation Using the dbaascli Utility

To configure AWR on an ADG standby database, use the `dbaascli dataguard configureAWR` command.

1. Connect to the Virtual Machine where your ADG standby database is hosted as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the `configureAWR` command.

To enable AWR, run:

```
$ dbaascli dataguard configureAWR --dbname <db-name> --enable
```

To disable AWR, run:

```
$ dbaascli dataguard configureAWR --dbname <db-name> --disable
```

4. Enter the `SYS` and `AWR` passwords when prompted.

Upon successful completion of the command, the ADG standby database would have been configured with AWR

dbaascli dataguard updateConfiguration

To update the transport mode or protection mode or both the parameters of a Data Guard environment, use the `dbaascli dataguard updateConfiguration` command.

Run this as the `root` user.

When the update transport mode command is run on the primary, only the transport mode of the primary database is updated. To update the transport mode of a standby database, the command must be run separately on that standby.

In contrast, when the update protection mode command is run on the primary, the protection mode is updated for both the primary and standby databases. The protection mode can also be updated from the standby side, in which case both the primary and standby databases are updated.

When updating the transport or protection mode from the primary, the system checks the current modes on both the primary and standby databases and proceeds with the update only if all required conditions are met.

Syntax

```
dbascli dataguard updateConfiguration --dbname <value> [--protectionMode
<value>] [--transportType <value>] [--standbyDGType <value>] [--
executePrereqs] [--resume [--sessionID <value>]] [--waitForCompletion <value>]
```

Where:

- `--dbname` specifies the Oracle Database name.
- `--protectionMode` specifies the Data Guard protection mode to be set when configuring the standby database. Valid values: `MAX_PERFORMANCE|MAX_AVAILABILITY`.
- `--transportType` specifies the Data Guard transport type to be set when configuring the standby database. Valid values: `ASYNC|SYNC`.
- `--standbyDGType` specifies the standby database Data Guard type to be set. Valid values: `ADG|DG`.
- `--executePrereqs` runs the prerequisite checks and report the results.
- `--resume` resumes the previous operation.
- `--sessionID` resumes a specific session by its ID.
- `--waitForCompletion` specifies whether to wait for the operation to complete. Set to `false` to run the operation in the background. Valid values: `true|false`.
- [Performing updateConfiguration Operation Using the dbascli Utility](#)
To update transport mode and protection mode or both the parameters, use the `dbascli dataguard updateConfiguration` command.

Performing updateConfiguration Operation Using the dbascli Utility

To update transport mode and protection mode or both the parameters, use the `dbascli dataguard updateConfiguration` command.

1. Connect to the Virtual Machine where your ADG standby database is hosted as the `opc` user.

For detailed instructions, see [Connecting to a Virtual Machine with SSH](#).

2. Start a root-user command shell.

```
$ sudo -s
```

3. Run the `updateConfiguration` command.

```
$ dbascli dataguard updateConfiguration --dbname <db-name> --
protectionMode MAX_PERFORMANCE|MAX_AVAILABILITY --transportType ASYNC|SYNC
--standbyDGType ADG|DG.
```

Upon successful completion of the command, the specified Data Guard will be configured with specified transport mode or protection mode.

dbaascli dbhome create

To create an Oracle Database home of desired version, use the `dbaascli dbhome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbhome create --version <value>
[--oracleHome <value>]
[--oracleHomeName <value>]
[--enableUnifiedAuditing <value>]
[--imageTag <value>]
[--ImageLocation <value>
```

Where:

- `--version` specifies the version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0
- `--oracleHome` specifies the location of Oracle home
- `--oracleHomeName` specifies user-defined Oracle home name. If not provided, then the default name will be used
- `--enableUnifiedAuditing` specifies `true` or `false` to enable or disable unified auditing link option in Oracle home
- `--imageTag` specifies Oracle home image tag
- `--imageLocation` - path of the image to be used.
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true` or `false`.

Example 6-24 dbaascli dbhome create

```
dbaascli dbhome create --version 19.11.0.0.0
```

Alternatively, `dbaascli dbhome create --version 19.8.0.0.0.0 --imageTag 19.8.0.0.0` for cases where image tags are different from version.

dbaascli dbHome delete

To delete a given Oracle Database home, use the `dbaascli dbHome delete` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome delete
{ --oracleHome <value>
| --oracleHomeName <value> } [--resume [--sessionID <value>]]
```

Where:

- `--oracleHome` specifies the location of the Oracle home
- `--oracleHomeName` specifies the name of the Oracle home
- `--resume` resumes the previous execution
 - `--sessionID` specifies to resume a specific session ID

dbaascli dbhome getDatabases

To view information about all Oracle Databases running from a given database Oracle home, use the `dbaascli dbHome getDatabases` command. Specify either the Oracle home location or Oracle home name.

Run the command as the `root` user.

Syntax

```
dbaascli dbHome getDatabases
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name
- `--oracleHome` specifies the location (path) of Oracle home

Example 6-25 dbaascli dbHome getDatabases --oracleHome

```
dbaascli dbHome getDatabases --oracleHome /u02/app/mar_home/
```

dbaascli dbHome getDetails

To view information about a specific Oracle home, use the `dbaascli dbHome getDetails` command. Specify either the Oracle home location or Oracle home name.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome getDetails
{ --oracleHomeName value | --oracleHome value }
```

Where:

- `--oracleHomeName` specifies user-defined Oracle home name

- `--oracleHome` specifies the location of Oracle home

Example 6-26 dbaascli dbHome getDetails - using Oracle home location

```
dbaascli dbHome getDetails --oracleHome /u02/app/home_db19c/
```

Example 6-27 dbaascli dbHome getDetails - using Oracle home name

```
dbaascli dbHome getDetails --oracleHomeName home_db19c
```

dbaascli dbHome patch

To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli dbHome patch

{
  --oracleHome <value>
  | --oracleHomeName <value>
}

[--imageFilePath <value>] [--executePrereqs] [--nodes <value>]
{
  [--resume [--sessionID <value>]]
  | [--rollback [--sessionID <value>]]
}

[--skipDatapatch]
[--skipClosedPDBs]
[--skipPDBs <value>]
[--continueWithDbDowntime]
[--skipUnreachableNodes]
[--drainTimeoutInSeconds <value>]
[--waitForCompletion <value>]
[--nonRolling]
[--skipDatapatchForDB <value>]
```

Where:

- `--oracleHome` specifies the path of Oracle home
- `--oracleHomeName` specifies the name of Oracle home
- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods, for example, 19.12.0.0.0.
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID

- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only one active instance up and the patching operation can be continued even with a downtime.
- `--skipUnreachableNodes` skips operation on unreachable nodes
- `--nodes` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--executePrereqs` runs prereqs
- `--skipDatapatch` skips running datapatch on the databases
- `--imageFilePath` specifies the absolute path of the image file to be used
- `--skipPDBs` skips running the datapatch on a specified comma-delimited list of PDBs. For example: `cdb1:pdb1,cdb2:pdb2`, and so on
- `--skipClosedPdb` skips running datapatch on closed PDBs
- `--rollback` rolls back patched Oracle home.
- `--waitForCompletion` specifies false to run the operation in background. Valid values : `true|false`
- `--drainTimeoutInSeconds` specifies time (in seconds) to complete the resource draining while stopping the database
- `--skipUnreachableNodes` skips operation on unreachable nodes
- `[--nonRolling]` patches the dbhome on all nodes in parallel. Note that this would cause a downtime in the database.
- `[--skipDatapatchForDB <value>]` skips running datapatch on the specified comma-delimited list of databases. Example: `db1,db2,db3...`

Example 6-28 dbaascli dbhome patch

```
dbaascli dbhome patch --targetVersion 19.10.0.0.0 --oracleHome /u02/app/oracle/product/19.0.0.0/dbhome_2
```

dbaascli dbimage purge

The `dbimage purge` command removes the specified software image from your Exadata Cloud Infrastructure environment.

Connect to the compute node as the `opc` user and execute this command as the `root` user.

```
# dbaascli dbimage purge --version software_version --bp software_bp [--cdb ( yes | no )]
```

In the preceding command:

- `software_version` — specifies the Oracle Database software version. For example, 11204, 12102, 12201, 18000, 19000.
- `software_bp` — identifies the bundle patch release. For example, APR2018, JAN2019, OCT2019, and so on.
- `--cdb` — optionally specifies whether to remove the software image that supports the Oracle multitenant architecture. Default is `yes`. If you specify `--cdb no`, then the software image that contains binaries to support non-container databases (non-CDB) is removed.

If the command will remove a software image that is not currently available in the software image library, and therefore cannot be downloaded again, then the command pauses and prompts for confirmation.

You cannot remove the current default software image for any software version. To avoid this restriction, you must make another software image the current default.

dbaascli diag collect

To collect diagnostics, use the `dbaascli diag collect` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag collect [--components <value>] [--startTime <value>] [--endTime <value>] [--nodes <value>] [--dbNames <value>]
    {
        [--objectStoreBucketUri <value>]
        | [--destLocation <value>]
    }
    [--waitForCompletion <value>]
```

Where:

- `--components` specifies a list of components for log collection.
Valid values:
 - db
 - gi
 - os
 - dbaastools
 - all
- `--startTime` specifies the start time for log collection. Valid date and time format: YYYY-MM-DDTHH24:MM:SS
- `--endTime` specifies the end time for log collection. Valid date and time format: YYYY-MM-DDTHH24:MM:SS
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--dbNames` specifies the database name for which to collect logs. You can specify only one database name.
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.
- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--waitForCompletion` Values: true|false. Default true. Specify false to run in the background.

Related Topics

- [Using Pre-Authenticated Requests](#)
- [Collecting Tooling Log Data Examples](#)

The `dbaascli dbaascli diag collect` command uses the syntax shown below to collect tooling log data:

dbaascli diag healthCheck

To run diagnostic health checks, use the `dbaascli diag healthCheck` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli diag healthCheck  
[--destLocation]  
[--nodes]  
[--objectStoreBucketURI]
```

Where:

- `--destLocation` specifies the location on Guest VM to collect logs. Default: `/var/opt/oracle/dbaas_acfs`
- `--nodes` specifies a comma-delimited list of nodes to collect logs
- `--objectStoreBucketURI` specifies an Object Storage service pre-authenticated request (PAR) URL used to upload collected logs. Logs are collected from Guest VM. For more information, see *Using Pre-Authenticated Requests*.

Related Topics

- [Using Pre-Authenticated Requests](#)

dbaascli gridHome create

To configure Grid Infrastructure home, use the `dbaascli gridHome create` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli gridHome create --version value [--resume [--sessionID value]] [--  
waitForCompletion value]
```

Where:

- `--version` specifies the Grid home version
- `--resume` resumes the previous run
- `--sessionID` specifies to resume a specific session ID

- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`

dbaascli grid configureTCPS

To configure TCPS for the existing cluster, use the `dbaascli grid configureTCPS` command.

Prerequisite

Run the command as the `root` user.

Syntax

① Note

By default, TCPS is enabled for databases on Oracle Exadata Database Service on Dedicated Infrastructure systems.

① Note

TCPS is not enabled for databases on Exadata Database Service on Cloud@Customer systems. To enable TCPS for a given database, update the database specific `sqlnet.ora` file with `WALLET_LOCATION = (SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/grid/tcps_wallets)))` on all database nodes and then bounce the database. This will enable TCPS usage for the database. However, enabling TCPS will cause ZDLRA connection to fail. On Exadata Database Service on Cloud@Customer systems, you can enable either ZDLRA or TCPS configuration. Enabling both ZDLRA and TCPS simultaneously will not work.

```
dbaascli grid configureTCPS
[--pkcs12WalletPath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletPath` specifies the path of the certificate, which is in `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.
- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`

- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: true or false. Default: true.

Example 6-29 dbaascli grid configureTCPS

To configure grid using self-signed certificate:

```
dbaascli grid configureTCPS
```

To configure grid using CA-signed certificate:

```
dbaascli grid configureTCPS --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted_private.key --private_key_password_protected false
```

dbaascli grid patch

To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.

Prerequisites

Run the command as the root user.

Syntax

```
dbaascli grid patch
{
    --targetVersion <value>
    | --targetHome <value>
}
[--executePrereqs] [--nodeList <value>] [--continueWithDbDowntime] [--
drainTimeoutInSeconds <value>] [--containerURL <value>] [--imageFile <value>]
[--patchInParallel]
{
    [--resume [--sessionID <value>]]
    | [--rollback [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--targetVersion` specifies the target version of Oracle Home specified as five numeric segments separated by periods (e.g. 19.12.0.0.0)
- `--targetHome` specifies the fully qualified path of the target Grid Infrastructure home for the out of place patching
- `--containerURL` specifies custom URL for fetching Grid Infrastructure image
- `--executePrereqs` option to run prereqs
- `--nodeList` specifies a comma-delimited list of nodes if patching has to be performed on a subset of nodes
- `--patchInParallel` specifies to perform patching remote nodes in parallel
- `--rollback` specifies to roll back patched Oracle home

- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--continueWithDbDowntime` continues patching with database downtime. This option can be used in environments wherein there is only 1 active instance up and the patching operation can be continued even with a downtime.
- `--drainTimeoutInSeconds` specifies the time (in seconds) to complete the resource draining while stopping the database
- `--createImage` creates an image from a copy of the active Grid home, patched to the specified target version
 - `--createImageDir` specifies fully qualified path of the directory where the image is to be created
- `--imageFile` specifies fully qualified path of the image to be used
- `--patchInParallel` performs the patching of the remote nodes in parallel
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Example 6-30 dbaascli grid patch

```
dbaascli grid patch --targetVersion 19.12.0.0.0
```

dbaascli grid removeTCPSCert

To remove existing TCPS certificates from Grid Infrastructure wallet, use the `dbaascli grid removeTCPSCert` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid removeTCPSCert --subject <value>
{
  --userCert | --trustedCert | --requestedCert
}
[--serialNumber <value>] [--executePrereqs] [--resume [--sessionID <value>]]
[--bounceListeners]
```

Where:

- `--subject` specifies subject of the certificate
- `--userCert` flag to indicate user certificate
- `--trustedCert` flag to indicate trusted certificate
- `--requestedCert` flag to indicate requested certificate
- `--serialNumber` specifies the serial number of the certificate
- `--executePrereqs` runs the prerequisite checks and reports the results
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID

- `--bounceListeners` flag to bounce the Grid Infrastructure listener and scan listener

dbaascli grid rotateTCPSCert

To rotate TCPS certificates, use the `dbaascli grid rotateTCPSCert` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid rotateTCPSCert
[--pkcs12WalletPath]
[--caCertChain]
[--precheckOnly]
[--serverCert]
[--privateKey]
[--certType]
[--privateKeyPasswordProtected]
```

Where:

- `--pkcs12WalletPath` specifies the path of the certificate, which is in `pkcs12` wallet format
- `--caCertChain` concatenated list of certs, containing intermediate CA's and root CA certs
- `--precheckOnly` specifies `yes` to run only the prechecks for this operation. Valid values: `yes` or `no`.
- `--serverCert` specifies the path of PEM certificate to use or rotate for TCPS configuration.
- `--privateKey` specifies the path of the private key file of the certificate.
- `--certType` type of the cert to be added to the Grid Infrastructure wallet. Accepted values are: `SELF_SIGNED_CERT`, `CA_SIGNED_CERT`, or `PKCS12_CERT`. Default: `SELF_SIGNED_CERT`
- `--privateKeyPasswordProtected` specifies if the private key is password protected or not. Valid values: `true` or `false`. Default: `true`.

Example 6-31 dbaascli grid rotateTCPSCert

To rotate cert using self-signed certificate (default option):

```
dbaascli grid rotateTCPSCert
```

To rotate cert using CA-signed certificate:

```
dbaascli grid rotateTCPSCert --cert_type CA_SIGNED_CERT --server_cert /tmp/
certs/server_cert.pem --ca_cert_chain /tmp/certs/ca.pem --private_key /tmp/
certs/encrypted_private.key --privateKeyPasswordProtected true
```

dbaascli grid upgrade

To upgrade Oracle Grid Infrastructure from one major version to another, use the `dbaascli grid upgrade` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli grid upgrade --version  
[--resume]  
[--executePrereqs]  
[--containerURL]  
[--softwareOnly]  
[--targetHome]  
[--revert]
```

Where:

- `--version` specifies the target version
- `--resume` resumes the previous run
- `--executePrereqs` runs prereqs for Grid Infrastructure upgrade
- `--containerUrl` specifies the custom URL for fetching Grid Infrastructure image
- `--softwareOnly` installs only the Grid Infrastructure software
- `--targetHome` specifies the path of existing target Grid home
- `--revert` reverts failed run

Example 6-32 dbaascli grid upgrade

```
daascli grid upgrade --version 19.11.0.0.0 --executePrereqs  
DBAAS CLI version MAIN  
Executing command grid upgrade --version 19.11.0.0.0 --executePrereqs
```

dbaascli job getStatus

To view the status of a specified job, use the `dbaascli job getStatus` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli job getStatus --jobID
```

Where:

- `--jodID` specifies the job ID

Example 6-33 dbascli job getStatus

```

dbascli job getStatus --jobID 13c82031-f202-41b7-9aef-f4a71df0f551
DBAAS CLI version MAIN
Executing command job getStatus --jobID 13c82031-f202-41b7-9aef-f4a71df0f551
{
  "jobId" : "13c82031-f202-41b7-9aef-f4a71df0f551",
  "status" : "Success",
  "message" : "database create job: Success",
  "createTimestamp" : 1628095442431,
  "updatedAtTime" : 1628095633660,
  "description" : "Service job report for operation database create",
  "appMessages" : {
    "schema" : [ ],
    "errorAction" : "SUCCEED_AND_SHOW"
  },
  "resourceList" : [ ],
  "pct_complete" : "100"
}

```

dbascli patch db apply**Note**

dbascli patch db prereq and dbascli patch db apply commands have been deprecated in dbascli release 21.2.1.2.0, and replaced with dbascli grid patch, dbascli dbhome patch, and dbascli database move commands.

For more information, see:

- dbascli grid patch
- dbascli dbhome patch
- dbascli database move
- *Patching Oracle Grid Infrastructure and Oracle Databases Using dbascli*

Related Topics

- [dbascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the dbascli grid patch command.
- [dbascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the dbascli dbHome patch command.
- [dbascli database move](#)
To move the database from one home to another, use the dbascli database move command.
- [Patching Oracle Grid Infrastructure and Oracle Databases Using dbascli](#)
Learn to use the dbascli utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

dbaascli patch db prereq

Note

`dbaascli patch db prereq` and `dbaascli patch db apply` commands have been deprecated in `dbaascli` release 21.2.1.2.0, and replaced with `dbaascli grid patch`, `dbaascli dbhome patch`, and `dbaascli database move` commands.

For more information, see:

- `dbaascli grid patch`
- `dbaascli dbhome patch`
- `dbaascli database move`
- *Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli*

Related Topics

- [dbaascli grid patch](#)
To patch Oracle Grid Infrastructure to the specified minor version, use the `dbaascli grid patch` command.
- [dbaascli dbHome patch](#)
To patch Oracle home from one patch level to another, use the `dbaascli dbHome patch` command.
- [dbaascli database move](#)
To move the database from one home to another, use the `dbaascli database move` command.
- [Patching Oracle Grid Infrastructure and Oracle Databases Using dbaascli](#)
Learn to use the `dbaascli` utility to perform patching operations for Oracle Grid Infrastructure and Oracle Database on an Exadata Cloud Infrastructure system.

dbaascli pdb backup

To backup a pluggable database (PDB), query PDB backups, and delete a PDB backup, use the `dbaascli pdb backup` command.

Prerequisite

- Run the command as the `root` user.

Syntax

```
dbaascli pdb backup --pdbName <value> --dbname <value>
{
  --start
  {
    [--level1]
    | [--archival --tag <value>]
  }
  | --delete --backupTag <value>
  | --status --uuid <value>
```

```

    | --getBackupReport --json <value> --tag <value>
    | --list [--json <value>]
}

```

Where:

```

--pdbName: PDB name.
--dbname: Oracle Database name.
--start | --delete | --status | --getBackupReport | --list
--start: Begins PDB backup.
    [--level1 | --archival]
    [--level1: Creates a Level-1 (incremental) backup.]
    [--archival: Creates an archival full backup.]
    --tag: Specify backup tag.
--delete: Deletes archival backup.
    --backupTag: Specify backup tag to delete.
--status
    --uuid <value>
--getBackupReport: Returns backup report.
    --json: Specify the file name for JSON output.
    --tag: Specify backup tag.
--list: Returns PDB backup information.
    [--json: Specify the file name for JSON output.]

```

Example 6-34 Examples

- To take level1 backup for a PDB *pdb1* in a CDB *myTestDb*:

```
dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --start --level1
```

- To query the status of PDB backup request submitted with uuid *eef16b26361411ecb13800163e8e4fac*:

```
dbaascli pdb backup --dbname myTestDb --pdbName pdb1 --status --uuid
eef16b26361411ecb13800163e8e4fac
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb bounce

To bounce a pluggable database (PDB), use the `dbaascli pdb bounce` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```

dbaascli pdb bounce
{
    --pdbName <value>
    | --pdbUID <value>
}

```

```
--dbname <value> [--openMode <value>] [--startServices <value>] [--
waitForCompletion <value>]
```

Where:

- `--pdbName` specifies the name of the PDB
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB
- `--openMode` specifies the target `OPEN MODE` of PDB
- `--startServices` specifies to start all or a list of all services corresponding to a PDB. Accepted values are `all` or a comma-delimited list of PDB services.
- `--waitForCompletion` specifies to run the operation in the foreground or background. Valid values: `true|false`.

Example 6-35 dbaascli pdb bounce

```
dbaascli pdb bounce --dbname cdb_name --pdbName pdb name associated with the
CDB
```

```
dbaascli pdb bounce --dbname cdb_name --pdbUID con_uid of that pdb
```

Optional:

- `--openMode READ_WRITE`
- `--openMode READ_ONLY`

dbaascli pdb close

To close a pluggable database (PDB), use the `dbaascli pdb close` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb close
{
  --pdbName <value>
  | --pdbUID <value>
}
--dbname <value> [--waitForCompletion <value>]
```

Where:

- `--pdbname` specifies the name of the PDB that you want to close.
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB.
- `--waitForCompletion` specifies to run the operation in the foreground or background. Valid values: `true|false`.

Upon successful completion of running this command, the PDB is closed on all of the container database instances.

Example 6-36 dbaascli pdb close

```
dbaascli pdb close --dbname cdb name --pdbName pdb name associated with the CDB
```

```
dbaascli pdb close --dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb getConnectString

To display Oracle Net connect string information for a pluggable database (PDB) run the `dbaascli pdb getConnectString` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getConnectString --dbname <value>
{
  --pdbName <value>
  | --pdbUID <value>
}
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB for which you want to display connect string information
- `--pdbUID` specifies the identifier of the PDB

Example 6-37 dbaascli pdb getConnectString

```
dbaascli pdb getConnectString --dbname dbname --pdbName pdbName
```

dbaascli pdb create

To create a new pluggable database (PDB), use the `dbaascli pdb create` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb create --pdbName <value> --dbName <value> [--maxCPU <value>] [--maxSize <value>] [--pdbAdminUserName <value>] [--lockPDBAdminAccount <value>] [--resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion <value>]
{
```

```

        [--blobLocation <value>]
        | [--standbyBlobFromPrimary <value>]
    }
    [--pdbTdeKeyVersionOCID <value>]

```

Where:

- `--pdbName` specifies the name of the new PDB that you want to create
- `--dbName` specifies the name of the container database that hosts the new PDB
- `--maxCPU` optionally specifies the maximum number of CPUs that are available to the PDB. Setting this option is effectively the same as setting the `CPU_COUNT` parameter in the PDB
- `--maxSize` optionally specifies the maximum total size of data files and temporary files for tablespaces belonging to the PDB. Setting this option is effectively the same as setting the `MAXSIZE` PDB storage clause in the `CREATE PLUGGABLE DATABASE` SQL command. You can impose a limit by specifying an integer followed by a size unit (K, M, G, or T), or you can specify `UNLIMITED` to explicitly enforce no limit
- `--pdbAdminUserName` specifies the new PDB admin user name
- `--lockPDBAdminAccount` specifies `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations.

Note

the parameters `blobLocation` and `standbyBlobFromPrimary` are mutually exclusive.

- `--pdbTdeKeyVersionOCID`

During the PDB creation process, you are prompted to specify the administration password for the new PDB.

Example 6-38 dbascli pdb create

To create a PDB from seed in a standard database in a non-Data Guard environment:

```
dbascli pdb create --dbName db721 --pdbName new_pdb1 --maxsize 5G --maxcpu 2
```

To create PDB in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1
```

```
dbaascli pdb create --dbName db721 --pdbName new_pdb1 --  
standbyBlobFromPrimary /tmp/send_db721.tar
```

dbaascli pdb delete

To delete a pluggable database (PDB) run the `dbaascli pdb delete` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb delete --dbName <value>  
    {  
        --pdbName <value>  
        | --pdbUID <value>  
    }  
    [--executePrereqs] [--waitForCompletion <value>] [--resume [--  
sessionID <value>]] [--allStandbyPrepared] [--cleanupRelocatedPDB]
```

Where:

- `--dbName` specifies the name of the container database that hosts the PDB
- `--pdbName` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the UID of the PDB that you want to delete
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID
- `--allStandbyPrepared` specifies to confirm that the operation has been successfully run on all the standby databases
- `--cleanupRelocatedPDB` specify to cleanup source database after a PDB has been relocated

Example: dbaascli pdb delete

To delete a PDB from a standard database in a non-Data Guard environment or from Standby database in Data Guard environment.

```
dbaascli pdb delete --dbName db721 --pdbName pdb1
```

To create PDB from Primary database in Data Guard environment:

```
dbaascli pdb create --dbName db721 --pdbName pdb1 --allStandbyPrepared
```

dbaascli pdb getDetails

To view details of a pluggable database (PDB), use the `dbaascli pdb getDetails` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb getDetails --dbname <value>
  {
    --pdbName <value>
  | --pdbUID <value>
  }
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB
- `--pdbname` specifies the name of the PDB that you want to delete
- `--pdbUID` specifies the identifier of the PDB

Example 6-39 dbaascli pdb getDetails

```
dbaascli pdb getDetails--dbname cdb name --pdbName pdb name associated with  
the CDB
```

```
dbaascli pdb getDetails--dbname cdb name --pdbUID con_uid of that pdb
```

dbaascli pdb list

To view the list of pluggable databases (PDB) in a container database, use the `dbaascli pdb list` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb list --dbname
```

Where:

- `--dbname` specifies the name of the container database that hosts the PDB

Example 6-40 dbaascli pdb list

```
dbaascli pdb list --dbname cdb name
```

dbaascli pdb localClone

To create a new pluggable database (PDB) as a clone of an existing PDB in the same container database (CDB), use the `dbaascli pdb localClone` command.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli pdb localClone --pdbName <value> --dbName <value> [--targetPDBName
<value>] [--powerLimit <value>] [--maxCPU <value>] [--maxSize <value>] [--
resume [--sessionID <value>]] [--executePrereqs] [--waitForCompletion <value>]
{
  [--blobLocation <value>]
  | [--standbyBlobFromPrimary <value>]
}
[--excludeUserTablespaces <value>] [--excludePDBData <value>] [--
pdbAdminUserName <value>] [--lockPDBAdminAccount <value>] [--
sourcePDBServiceConvertList <value>]
{
  [--createFromSnapshot
  {
    --snapshotName <value>
    | --snapshotUID <value>
  }
  [--copyDataFiles]]
  | [--snapshot]
}
```

Where:

- `--pdbName` specifies the name of the new PDB that you want to clone
- `--dbName` specifies the name of the database
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run
- `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` OR `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` OR `false`

- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby blob file which is prepared from the primary database. This is required only for standby database PDB operations.

Note

The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example t1,t2,t3.
- `--excludePDBData` specify true/yes to skip user data from source PDB.
- `--pdbAdminUserName` specify new PDB admin user name.
- `--lockPDBAdminAccount` specify true or false to lock the PDB admin user account. Default value is true.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names which need to be converted. Syntax is `source_srv1:new_srv1,source_srv2:new_srv2`.
- `--createFromSnapshot | --snapshot`
 - `--createFromSnapshot` specify to create PDB from PDB snapshot
`--snapshotName | --snapshotUID`
 - * `--snapshotName` specify the snapshot name to create PDB from PDB.
 - * `--snapshotUID` specify the snapshot UID to create PDB from PDB.
 - `--copyDataFiles` specify this option to skip snapshot copy.
 - `--snapshot` specify to create PDB with snapshot copy.

The newly cloned PDB inherits administration passwords from the source PDB.

Example 6-41 dbascli pdb localClone

```
dbascli pdb localClone --dbName db35 --pdbName PDB35 --targetPDBName
local_clone1 --maxCPU 2 --maxSize 15
```

dbascli pdb open

To open a pluggable database (PDB), use the `dbascli pdb open` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbascli pdb open
{
    --pdbName <value>
    | --pdbUID <value>
}
--dbname <value> [--openMode <value>] [--startServices <value>] [--
waitForCompletion <value>] [--setPDBRefreshModeNone [--skipPDBRefresh] [--
```

```
pdbAdminUserName <value>]] [--executePrereqs] [--resume [--sessionID
<value>]] [--blobLocation <value>]
```

Where:

- `--pdbName` specifies the name of the PDB that you want to open
- `--pdbUID` specifies the identifier of the PDB
- `--dbname` specifies the name of the container database that hosts the PDB.
- `--openMode` specifies the target OPEN MODE of PDB
- `--startServices`: specifies to start all or list all services corresponding to a PDB. Accepted values are `all` or a comma-delimited list of PDB services.
- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`
- `--setPDBRefreshModeNone`: specifies to convert a refreshable PDB to non-refreshable PDB
 - `--skipPDBRefresh`: specifies to skip refreshable PDB refresh
 - `--pdbAdminUserName`: specifies new PDB admin user name
- `--executePrereqs` specifies to run the prerequisite checks and report the results
- `--resume` resumes the previous operation
 - `--sessionID` specifies to resume a specific session ID
- `--blobLocation` specifies the custom directory location where the standby blob file will be generated in a Data Guard environment

Upon successful completion, the PDB is opened on all of the container database instances.

Example 6-42 dbaascli pdb open

```
dbaascli pdb open --dbname cdb name --pdbName pdb name associated with the CDB
```

```
dbaascli pdb open --dbname cdb name --pdbUID con_uid of that pdb
```

Optional: `--openMode READ_WRITE/READ_ONLY`

dbaascli pdb recover

To recover a pluggable database (PDB), use the `dbaascli pdb recover` command.

Prerequisite

- Run the command as the `root` user.
- Database must be configured with backup storage destination details where backups are stored.

Syntax

```
dbaascli pdb recover --dbname <value> --pdbName <value>
{
    --start
    {
```

```

        --untilTime <value> [--nonUTC]
        | --untilSCN <value>
        | --latest
        | --tag <value>
    }
    | --status --uuid <value>
}

```

Where:

- --dbname specifies the name of the container database that hosts the PDB
- --pdbName specifies the name of the PDB that you want to recover
- --start begins the recovery of a PDB.
 - --untilTime recovers PDB until time. Input format: DD-MON-YYYY HH24:MI:SS
 - --untilSCN recovers PDB until SCN
 - --latest recovers PDB to the last known state
 - --tag recovers PDB to archival tag
- --status displays the details about a PDB recovery job process
 - --uuid unique identifier of the PDB recovery operation. Input format: xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx

Example 6-43 Examples

- To recover a PDB *pdb1* in a CDB *myTestDb* to latest:

```
dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --start --latest
```

- To query the status of PDB recovery request submitted with *uuid 81a17352362011ecbc3000163e8e4fac*:

```
dbaascli pdb recover --dbname myTestDb --pdbName pdb1 --status --uuid
81a17352362011ecbc3000163e8e4fac
```

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb refresh

To refresh a specified pluggable database (PDB), use the `dbaascli pdb refresh` command.

Run the command as the `root` or `oracle` user.

Syntax

```

dbaascli pdb refresh --dbname <value>
{
  --pdbName <value>
  | --pdbUID <value>
}

```

```

    }
    [--waitForCompletion <value>]

```

Where:

- `--dbname`: specifies the name of the Oracle Database
- `--pdbName`: specifies the name of the pluggable database
- `--pdbUID`: specifies the identifier of the pluggable database
- `--waitForCompletion`: specify `false` to run the operation in the background. Valid values: `true|false`

Related Topics

- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

dbaascli pdb relocate

To relocate the specified PDB from the remote database into local database, use the `dbaascli pdb relocate` command.

Prerequisite

Run the command as the `oracle` user. When prompted, you must supply the SYS user password for the source database.

Syntax

```

dbaascli pdb relocate --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCpu <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--sourcePDBServices <value>] [--
sourcePDBReadOnlyServices <value>] [--relocateAvailabilityType <value>] [--
waitForCompletion <value>]
    {
        [--blobLocation <value>]
        | [--standbyBlobFromPrimary <value>]
    }
    [--upgradePDB] [--updateDBBlockCacheSize]
    {
        [--skipOpenPDB [--enableRefreshMode --refreshMode <value> [--
refreshIntervalInMinutes <value>]]]
        | [--completePDBRelocate]
    }
    {
        [--dbLinkUsername <value> [--honorCaseSensitiveUserName]]
        | [--dbLinkName <value>]
    }

```

Where:

- `--pdbName` specifies the source PDB name to relocate
- `--dbName` specifies the target database name

- `--sourceDBConnectionString` specifies the source database connection string in the format `<scan_name>:<scan_port>/<database_service_name>`
- `--targetPDBName` specifies a name for the target PDB (new relocated PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the relocate operation
- `--maxCpu` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` specifies to resume the previous execution
 - `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes|no`
- `--sourcePDBServices` specifies a list of comma-delimited source PDB services
- `--sourcePDBReadOnlyServices` specifies a comma-delimited list of source PDB read-only services
- `--relocateAvailabilityType` specifies `MAX|NORMAL` availability to enable the listener connection forwarding
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true|false`
- `--blobLocation` specifies a custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specifies the location of the standby BLOB file, which is prepared from the primary database. This is required only for standby operations.

Note

The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--upgradePDB` specify `true` specifies to upgrade the PDB as part of this operation. Valid values: `true | false`.
- `--updateDBBlockCachesize` specifies to enable application to set db block cache size initialization parameters to support data copy with different block size
- `--skipOpenPDB` specifies to indicate that the PDB should not be opened at the end of the current operation.
 - `--enableRefreshMode` specifies to enable PDB refresh support in PDB relocate step-one.
 - * `--refreshMode` specifies to refresh mode for refreshable PDB. Valid values: `AUTO|MANUAL`
 - * `--refreshIntervalInMinutes` specifies to refresh interval for `refreshablePDB` in minutes
- `--completePDBRelocate` specifies to complete the PDB relocation if done as a two-step operation.

Example 6-44 dbaascli pdb relocate

```
dbaascli pdb relocate --sourceDBConnectionString test-
scan.dbaastoolslrgsu.dbaastoolslrgvc.oraclevcn.com:1521/
source_cdb_service_name --pdbName source_pdb --dbName target_db
```

dbaascli pdb remoteClone

To create a new pluggable database (PDB) as a clone of an existing PDB in another container database (CDB), use the `dbaascli pdb remoteClone` command.

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli pdb remoteClone --pdbName <value> --dbName <value> --
sourceDBConnectionString <value> [--targetPDBName <value>] [--powerLimit
<value>] [--maxCPU <value>] [--maxSize <value>] [--resume [--sessionID
<value>]] [--executePrereqs] [--waitForCompletion <value>] [--
sourcePDBExportedTDEKeyFile <value>]
    {
        [--blobLocation <value>]
        | [--standbyBlobFromPrimary <value>]
    }
    [--excludeUserTablespaces <value>] [--excludePDBData <value>] [--
pdbAdminUserName <value>] [--lockPDBAdminAccount <value>] [--
sourcePDBServiceConvertList <value>] [--upgradePDB]
    {
        [--createFromSnapshot --snapshotName <value> [--copyDataFiles]]
        | [--snapshot]
    }
    [--refreshablePDB --refreshMode <value> [--refreshIntervalInMinutes
<value>]] [--updateDBBlockCacheSize]
    {
        [--dblinkUsername <value> [--honorCaseSensitiveUserName]]
        | [--dbLinkName <value>]
    }
    [--keepSourceDBServices <value>]
```

Where:

- `--pdbName` specifies the name of the source PDB that you want to clone
- `--dbname` specifies the name (DB_NAME) of the CDB that hosts the newly cloned PDB
- `--sourceDBConnectionString` specifies the source database connection string in the format `scan_name:scan_port/database_service_name`
- `--targetPDBName` specifies the name for the target PDB (new cloned PDB)
- `--powerLimit` specifies the degree of parallelism to be used for the clone operation. Valid value is between 1 and 128
- `--maxCPU` specifies the maximum number of CPUs to be allocated for the PDB
- `--maxSize` specifies the maximum storage size in GB for the new PDB
- `--resume` resumes the previous run

- `--sessionID` specifies to resume a specific session ID
- `--executePrereqs` specifies `yes` to run only the prereqs for this operation. Valid values: `yes` or `no`
- `--waitForCompletion` specifies `false` to run the operation in the background. Valid values: `true` or `false`
- `--sourcePDBExportedTDEKeyFile` specifies the source PDB exported key file. This variable is applicable to only 12.1 database.
- `--blobLocation` specifies the custom path where the standby blob file will be generated in a Data Guard environment
- `--standbyBlobFromPrimary` specify the location of the standby blob file, which is prepared from the primary database. This is required only for standby database PDB operations

Note

The parameters `--blobLocation` and `--standbyBlobFromPrimary` are mutually exclusive.

- `--excludeUserTablespaces` option to skip user table spaces, example `t1,t2,t3`.
- `--excludePDBData` specify `true/yes` to skip user data from source PDB.
- `--pdbAdminUserName` specifies new PDB admin user name
- `--lockPDBAdminAccount` specify `true` or `false` to lock the PDB admin user account. Default value is `true`.
- `--sourcePDBServiceConvertList` specify a comma-delimited list of source to target service names, which need to be converted. Syntax is `source_srv1:new_srv1, source_srv2:new_srv2`.
- `--upgradePDB` specify to upgrade the PDB as part of this operation
- `--createFromSnapshot | --snapshot`
 - `--createFromSnapshot` specify to create PDB from PDB snapshot
 - * `--snapshotName` specify the snapshot name to create PDB from PDB.
 - `--copyDataFiles` specify this option to skip snapshot copy.
 - `--snapshot` specify to create PDB with snapshot copy.
- `--refreshablePDB` specifies to create refreshable PDB
 - `--refreshMode` specifies refresh mode for refreshable PDB. Valid values: `AUTO|MANUAL`
 - * `--refreshIntervalInMinutes` specifies refresh interval for `refreshablePDB` in minutes
 - `--dblinkUsername` specifies common user of a remote database used for database link to connect to the remote database
 - * `--honorCaseSensitiveUserName` indicates specified username is case sensitive
- `--dbLinkName` specifies an existing database link name in the target database that points to the remote database

When promoted, you must supply the SYS user password for the source PDB. The newly cloned PDB inherits administration passwords from the source PDB. The cloned PDB is named using the following format: `dbname_sourcepdbname`. This command is supported only for

databases that are not in a Data Guard configuration and use Oracle Database version 12.2.0.1, or later.

Example 6-45 dbaascli pdb remoteClone

```
dbaascli pdb remoteClone --sourceDBConnectionString test-  
can.dbaastoolslrgrsu.dbaastoolslrgrvc.oraclevcn.com:1521 --pdbName source_pdb1  
--dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2
```

```
dbaascli pdb remoteClone --sourceDBConnectionString  
orcla.dbaastoolslrgrsu.dbaastoolslrgrvc.oraclevcn.com --pdbName source_pdb1 --  
dbName db9944 --targetPDBName new_pdb1 --maxsize 5 --maxcpu 2
```

dbaascli system getDBHomes

To view information about all the Oracle homes, use the `dbaascli system getDBHomes` command.

Prerequisite

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli system getDBHomes
```

Example 6-46 dbaascli system getDBHomes

```
dbaascli system getDBHomes
```

dbaascli system getGridHomes

To list the details of all Grid homes, use the `dbaascli system getGridHomes` command.

Prerequisite

Run the command as the `root` or `oracle` user.

Syntax

```
dbaascli system getGridHomes
```

dbaascli tde changePassword

To change TDE keystore password as well as DB wallet password for the alias `tde_ks_passwd`, use the `dbaascli tde changePassword` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbascli tde changePassword [--dbname <value>]
  {
    [--prepareStandbyBlob <value> [--blobLocation <value>]]
    | [--standbyBlobFromPrimary <value>]
  }
  [--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the database
- `--prepareStandbyBlob` - specify true to generate a blob file containing the artifacts needed to perform the operation in a DG environment.
- `--blobLocation` - custom path where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` - specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.
- `--resume` - to resume the previous execution
- `--sessionID` - to resume a specific session id.

```
dbascli tde changepassword --dbname
  <dbname>
```

1. Change the TDE password in primary database.

```
dbascli tde changepassword --dbname
  <dbname> --prepareStandbyBlob true --blobLocation
  <Location where blob file has to be generated>
```

2. Copy the created standby blob to standby database environment.
3. Change the TDE password in standby database

```
dbascli tde changepassword --dbname
  <dbname> --standbyBlobFromPrimary <Location of blob generated from
  primary>
```

dbascli tde addSecondaryHsmKey

To add a secondary HSM (KMS) key to the existing HSM (KMS) configuration, use the `dbascli tde addSecondaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbascli tde addSecondaryHsmKey --dbname <value> --secondaryKmsKeyOCID <value>
  [--executePrereqs]
```

Where:

- `--secondaryKmsKeyOCID` specifies the secondary KMS key to add to the existing HSM (KMS) configuration
- `--dbname` specifies the name of the database
- `--executePrereqs` `s`execute the prerequisites checks and report the results.

Example 6-47 dbaascli tde addSecondaryHsmKey

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde addSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --precheckOnly yes
```

dbaascli tde enableWalletRoot

To enable `wallet_root` `s`file parameter for the existing database, use the `dbaascli tde enableWalletRoot` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde enableWalletRoot --dbname <value>
[--dbRestart <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
```

Where:

- `--dbname` specifies the name of the Oracle Database.
- `--dbrestart` specifies the database restart option. Valid values are: `rolling` or `full`.
Default value: `rolling`
If you do not pass the `dbrestart` argument, then the database restarts in a `rolling` manner.
- `--precheckOnly` runs only the precheck for this operation. Valid values are: `yes` or `no`
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Example 6-48 dbascli tde enableWalletRoot

```
dbascli tde enableWalletRoot --dbname db name --dbrestart rolling|full
```

```
dbascli tde enableWalletRoot --dbname orcl
```

```
dbascli tde enableWalletRoot --dbname orcl--dbrestart full
```

dbascli tde encryptTablespacesInPDB

To encrypt all the tablespaces in the specified PDB, use the `dbascli tde encryptTablespacesInPDB` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbascli tde encryptTablespacesInPDB --pdbName  
[--dbname]  
[--precheckOnly]  
[--useSysdbaCredential]
```

Where:

- `--pdbName` specifies the name of the PDB to encrypt all the tablespaces.
- `--dbname` specifies the name of the Oracle Database.
- `--precheckOnly` runs only the precheck for this operation. Valid values: `yes` or `no`
- `--useSysdbaCredential` uses SYSDBA credentials for this operation if passed value is `true`. Valid values: `true` or `false`

Example 6-49 dbascli tde encryptTablespacesInPDB

```
dbascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb
```

```
dbascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --  
precheckOnly yes
```

```
dbascli tde encryptTablespacesInPDB --dbname dbname --pdbName pdb --  
useSysdbaCredential true
```

dbaascli tde fileToHsm

To convert FILE based TDE to HSM (KMS/OKV) based TDE, use the `dbaascli tde fileToHsm` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde fileToHsm --kmsKeyOCID <value> --dbname <value>
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
  [--resume [--sessionID <value>]] | [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyOCID` specifies the KMS key OCID to use for TDE. This is applicable only if KMS is selected for TDE
- `--dbname` specifies the name of the database
- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--primarySuc` specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- `--resume` specifies to resume the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to rollback the previous run
 - `--sessionID` specifies to rollback a specific session ID
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true`|
`false`.

Example 6-50 dbascli tde fileToHsm --kmsKeyOCID

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-  
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa  
invvxa
```

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-  
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa  
invvxa --executePrereqs
```

```
dbascli tde fileToHSM --dbname dbname --kmsKeyOCID ocid1.key.oc1.eu-  
frankfurt-.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5uxa  
invvxa --resume
```

dbascli tde getHsmKeys

To get TDE active key details, use the `dbascli tde getHsmKeys` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbascli tde getHsmKeys  
[--dbname]  
[--infoFile]
```

Where:

- `--dbname` specifies the name of the database
- `--infoFile` specifies the file path where the list of OCIDs will be saved. The output is in JSON format

Example 6-51 dbascli tde getHsmKeys

```
dbascli tde getHsmkeys --dbname dbname
```

```
dbascli tde getHsmkeys --dbname dbname --infoFile infoFilePath
```

dbascli tde getMkidForKeyVersionOCID

To get Master Key ID associated with the KMS key version OCID, use the `dbascli tde getMkidForKeyVersionOCID` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getMkidForKeyVersionOCID --kmsKeyVersionOCID <value>
[--dbname <value>]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set
- `--dbname` specifies the name of the database
- `--waitForCompletion` specify `false` to run the operation in background. Valid values : `true|false`.

Example 6-52 dbaascli tde getMkidForKeyVersionOCID

```
dbaascli tde getMkidForKeyVersionOCID --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

dbaascli tde getPrimaryHsmKey

To get primary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde getPrimaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde getPrimaryHsmKey
[--dbname]
```

Where:

- `--dbname` specifies the name of the database

Example 6-53 dbaascli tde getPrimaryHsmKey

```
dbaascli tde getPrimaryHsmKey --dbname dbname
```

dbaascli tde hsmToFile

To convert HSM (KMS/OKV) based TDE to FILE based TDE, use the `dbaascli tde hsmToFile` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde hsmToFile
[--dbname <value>]
```

```

{
  [--prepareStandbyBlob <value> [--blobLocation <value>]
  | [--standbyBlobFromPrimary <value>]
}
]
[--skipPatchCheck <value>]
[--executePrereqs ]
[--primarySuc <value>]
{
  [--resume [--sessionID <value>]] |
  [--revert [--sessionID <value>]]
}
[--waitForCompletion <value>]

```

Where:

- `--dbname` specifies the name of the database
- `--prepareStandbyBlob` specify `true` to generate a blob file containing the artifacts needed to perform the operation in a DG environment.
- `--blobLocation` custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations.]
- `--skipPatchCheck` skips validation check for required patches if the value passed for this argument is `true`. Valid values: `true` or `false`
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--primarySuc` specify this property in the standby database of the Data Guard environment once the command is successfully run on the primary database
- `--resume` resumes the previous run
 - `--sessionID` specifies to resume a specific session ID
- `--revert` specifies to roll back the previous run
 - `--sessionID` specifies to rollback a specific session ID
- `--waitForCompletion` specifies `false` to run the operation in background. Valid values: `true|false`

Example 6-54 dbascli tde hsmToFile

```
dbascli tde hsmToFile --dbname dbname
```

```
dbascli tde hsmToFile --dbname dbname --executePrereqs
```

```
dbascli tde hsmToFile --dbname dbname --resume
```

dbaascli tde listKeys

To list TDE master keys, use the `dbaascli tde listKeys` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde listKeys  
[--dbname <value>]  
[--infoFilePath <value>]
```

Where:

- `--dbname` specifies the name of the database
- `--infoFilePath` specify the absolute path of the file where the results will be saved.

Example 6-55 dbaascli tde listKeys

```
dbaascli tde listKeys --dbname dbname
```

```
dbaascli tde listKeys --dbname dbname --infoFilePath infoFilePath
```

dbaascli tde removeSecondaryHsmKey

To remove secondary HSM (KMS) key from the existing HSM (KMS) configuration, use the `dbaascli tde removeSecondaryHsmKey` command.

Prerequisite

Run the command as the `root` user.

Syntax

```
dbaascli tde removeSecondaryHsmKey --dbname <value>  
[--confirmDeletion]  
[--secondaryKmsKeyOCID]  
[--executePrereqs]
```

Where:

- `--dbname` specifies the name of the database
- `--confirmDeletion` if not specified the user will be prompted while deleting all existing HSM(KMS) keys.
- `--secondaryKmsKeyOCID` secondary KMS key to be removed from existing HSM(KMS) configuration. If not specified all secondary KMS keys will be removed.
- `--executePrereqs` execute the prerequisites checks and report the results.

Frequently Asked Questions

Q: What is the purpose of the dbaascli tde removeSecondaryHsmKey command?

A: The `dbaascli tde removeSecondaryHsmKey` command is used to remove a secondary Hardware Security Module (HSM) key from the existing HSM (KMS) configuration in an Oracle Database environment.

Q: What are the prerequisites for running the dbaascli tde removeSecondaryHsmKey command?

A: You must:

- Run the command as the `root` user.
- Be connected to an Exadata Cloud Infrastructure virtual machine using SSH.

Q: What does the --force parameter do in the dbaascli tde removeSecondaryHsmKey command?

A: The `--force` parameter allows the removal of the secondary HSM key without prompting the user for confirmation. If not specified, the command will prompt the user before deleting any keys.

Q: What does the --secondaryKmsKeyOCID parameter specify?

A: The `--secondaryKmsKeyOCID` parameter specifies the OCID (Oracle Cloud Identifier) of the secondary KMS key you want to remove from the existing HSM configuration.

Q: What does the --dbname parameter do?

A: The `--dbname` parameter specifies the name of the database for which the secondary HSM key is being removed.

Q: What is the purpose of the --precheckOnly parameter?

A: The `--precheckOnly` parameter, if set to `yes`, will only run the prechecks to validate the readiness for the removal operation without actually removing the secondary HSM key. If set to `no`, the full removal operation is performed.

Q: Is the --force parameter mandatory?

A: No, the `--force` parameter is optional. If it's not specified, the system will prompt the user for confirmation before proceeding with the key removal.

Q: Is the --secondaryKmsKeyOCID parameter mandatory?

A: Yes, you must provide the `--secondaryKmsKeyOCID` to identify the specific secondary HSM key that you want to remove from the configuration.

Q: Is the --dbname parameter mandatory?

A: No, the `--dbname` parameter is optional. If not specified, the command will attempt to remove the secondary HSM key from the default database on the system.

Q: What should I do if I want to remove the secondary HSM key without any user prompts?

A: You should use the `--force` parameter to bypass the confirmation prompt and remove the secondary HSM key directly:

```
dbaascli tde removeSecondaryHsmKey --force --secondaryKmsKeyOCID <value>
```

Q: How can I test whether the system is ready to remove the secondary HSM key without actually removing it?

A: You can use the `--precheckOnly` parameter set to `yes` to perform a precheck:

```
dbaascli tde removeSecondaryHsmKey --precheckOnly yes --secondaryKmsKeyOCID  
<value>
```

Q: What happens if I don't provide a database name with `--dbname`?

A: If the `--dbname` parameter is not specified, the command will attempt to remove the secondary HSM key from the default database configured on the system.

Q: What should I check if the command fails to remove the secondary HSM key?

A: Ensure that:

- You are running the command as the `root` user.
- You are connected to the Exadata Cloud Infrastructure virtual machine.
- The correct `--secondaryKmsKeyOCID` and `--dbname` values are provided. Check the error messages and logs for more details on the failure.

Q: What should I do if the removal operation fails partway through?

A: If the operation fails, review the error logs and try running the command with `--precheckOnly` to ensure the system is ready for the operation. If necessary, correct any issues before retrying.

Q: Can I run the `dbaascli tde removeSecondaryHsmKey` command while the database is running?

A: Yes, the command can be executed while the database is running, as it does not require the database to be stopped.

Q: What is the purpose of removing a secondary HSM key?

A: Removing a secondary HSM key is typically done when the key is no longer needed or when you want to manage the encryption keys used in your TDE (Transparent Data Encryption) configuration.

Q: How do I connect to the Exadata Cloud Infrastructure virtual machine to run the command?

A: You can connect to the virtual machine using SSH. Refer to the Exadata Cloud Infrastructure documentation for instructions on establishing a secure connection.

Example 6-56 dbaascli tde removeSecondaryHsmKey

```
dbaascli tde removeSecondaryHsmKey --dbname dbname
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID  
ocid1.key.oc1.eu-  
frankfurt-1.bjqnwc1vaafak.abtheljsqfxa2xe5prvlzdxtygoiqpm2pu2afgta54krxwllk5ux  
ainvvxa
```

```
dbaascli tde removeSecondaryHsmKey --dbname dbname --secondaryKmsKeyOCID  
ocid1.key.oc1.eu-
```

```
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

dbaascli tde rotateMasterKey

Rotate the master key for database encryption.

Prerequisites:

Run the command as the `root` user.

Syntax

(Optional) <Enter syntax information here.>

```
dbaascli tde rotateMasterKey --dbname <value>
[--rotateMasterKeyOnAllPDBs]
[--pdbName <value>]
[--executePrereqs]
[--resume [--sessionID <value>]]
{
    [--prepareStandbyBlob <value> [--blobLocation <value>]]
    | [--standbyBlobFromPrimary <value>]
}
```

Where:

- `--dbname` - Oracle database name.
- `--rotateMasterKeyOnAllPDBs` - specify true to rotate master key of all PDBs in CDB. Valid values: true|false
- `--pdbName` - specify PDB name.
- `--executePrereqs` - execute the prerequisites checks and report the results.
- `--resume` - to resume the previous execution
- `--sessionID` - to resume a specific session id.
- `--prepareStandbyBlob | --standbyBlobFromPrimary`
- `--prepareStandbyBlob` - specify true to generate a blob file containing the artifacts needed to perform the operation in a DG environment.
- `--blobLocation` - custom directory location where the standby blob file will be generated in a DG environment.
- `--standbyBlobFromPrimary` - specify the location of the standby blob file which is prepared from the primary database. This is required only for standby operations

dbaascli tde setKeyVersion

To set the version of the primary key to be used in DB/CDB or PDB, use the `dbaascli tde setKeyVersion` command.

Run the command as the `root` user.

Syntax

```
dbascli tde setKeyVersion --kmsKeyVersionOCID <value> --dbname <value>
[--pdbName <value>]
[--masterKeyID <value>]
[--standbySuc]
[--executePrereqs]
[--waitForCompletion <value>]
```

Where:

- `--kmsKeyVersionOCID` specifies the KMS key version OCID to set.
- `--dbname` specifies the name of the database.
- `--pdbName` name of the PDB to use the key version OCID.
- `--masterKeyID` specifies the master key ID of the given key version OCID. This is applicable to the Data Guard environment.
- `--standbySuc` specify this property in the primary database of the Data Guard environment once the command is successfully run on the standby database
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--waitForCompletion` specify `false` to run the operation in background. Valid values: `true|false`

Example 6-57 dbascli tde setKeyVersion

```
dbascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

```
dbascli tde setKeyVersion --dbname dbname --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha --executePrereqs
```

```
dbascli tde setKeyVersion --dbname dbname --pdbName pdb --kmsKeyVersionOCID
ocid1.keyversion.oc1.eu-
frankfurt-1.bjqnwclvaafak.bc4hmd3olgaaa.abtheljsyxtgn4vzi2bbpcej6a7abcwvylkd2l
x56lu2s6iwnxwgigu23nha
```

dbascli tde setPrimaryHsmKey

To change the primary HSM (KMS) key for the existing HSM (KMS) configuration, use the `dbascli tde setPrimaryHsmKey` command.

Run the command as the `root` user.

Syntax

```
dbaascli tde setPrimaryHsmKey --primaryKmsKeyOCID <value> --dbname <value>
[--allStandbyPrepared]
[--bounceDatabase]
[--executePrereqs]
[--resume [--sessionID <value>]]
```

Where:

- `--primaryKmsKeyOCID` specifies the primary KMS key to set
- `--dbname` specifies the name of the database
- `--allStandbyPrepared` specify to confirm that the operation has been successfully run on all the standby databases.
- `--bounceDatabase` specify this flag to do rolling database bounce for this operation
- `--executePrereqs` execute the prerequisites checks and report the results.
- `--resume` to resume the previous execution
- `--sessionID` to resume a specific session id.

Example 6-58 dbaascli tde setPrimaryHsmKey

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5ux
ainvvxza
```

```
dbaascli tde setPrimaryHsmKey --dbname dbname --primaryKmsKeyOCID
ocid1.key.oc1.eu-
frankfurt-1.bjqnwclvaafak.abtheljsgfxa2xe5prvlzdxygoiqpm2pu2afgta54krxwllk5ux
ainvvxza --executePrereqs
```

dbaascli tde status

To display information about the keystore for the specified database, use the `dbaascli tde status` command.

Note

`dbaascli tde status` command has been deprecated in dbaascli release 24.4.1.0.0. Oracle recommends using the `dbaascli database getDetails` command instead.

Prerequisite

Run the command as the `oracle` user.

Syntax

```
dbaascli tde status --dbname dbname
```

Where:

- `--dbname` specifies the name of the database that you want to check.

Output from the command includes the type of keystore, and the status of the keystore.

Example 6-59 dbaascli tde status

```
dbaascli tde status --dbname dbname
```

Related Topics

- [dbaascli database getDetails](#)
This command shows the detailed information of a given database e.g. dbname, node information, pluggable databases information etc.

Monitoring and Managing Exadata Storage Servers with ExaCLI

The ExaCLI command line utility allows you to perform monitoring and management functions on Exadata storage servers in an Exadata Cloud Infrastructure instance.

- [About the ExaCLI Command](#)
The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.
- [Exadata Storage Server Username and Password](#)
You need a username and password to connect to the Exadata Storage Server.
- [ExaCLI Command Syntax](#)
For Exadata Storage Server targets, construct your commands using the syntax that follows.
- [Connecting to a Storage Server with ExaCLI](#)
To use ExaCLI on storage servers, you will need to know your target storage server's IP address.
- [ExaCLI Command Reference](#)
You can execute various ExaCLI commands to monitor and manage Exadata Storage Servers associated with your Oracle Cloud Infrastructure Exadata VM cluster. ExaCLI allows you to get up-to-date, real-time information about your Exadata Cloud Service.

About the ExaCLI Command

The ExaCLI command provides a subset of the commands found in the on-premises Exadata command line utility.

ExaCLI offers a subset of the commands found in the on-premises Exadata command line utility *CellCLI utility*. The utility runs on the database virtual machines in the Exadata Cloud Service.

See the *ExaCLI Command* list in this topic to learn what commands are available.

Related Topics

- [Using the CellCLI Utility](#)
- [ExaCLI Command Syntax](#)
For Exadata Storage Server targets, construct your commands using the syntax that follows.

Exadata Storage Server Username and Password

You need a username and password to connect to the Exadata Storage Server.

On Exadata Cloud Infrastructure, the preconfigured user for Exadata Storage Server is `cloud_user_clustername`, where `clustername` is the name of the virtual machine (VM) cluster that is being used.

You can determine the name of the VM cluster by running the following `crsctl` command as the `grid` user on any cluster node:

```
crsctl get cluster name
```

This command returns `CRS-6724: Current cluster name is <cluster_name>`

The password for `cloud_user_clustername` is initially set to a random value, which you can view by running the following command as the `root` user on any cluster node:

```
/opt/exacloud/get_cs_data.py
```

This returns a password `<pwd>`

Then test with ExaCLI as root:

```
exacli -l cloud_user_<clustername> -c 192.168.136.14
Password: *****
exacli cloud_user_<cluster_name>@192.168.136.14>
```

ExaCLI Command Syntax

For Exadata Storage Server targets, construct your commands using the syntax that follows.

Note that the syntax example assumes you are the `opc` user on a compute node.

```
exacli -c [username@]remotehost[:port] [-l username] [--xml] [--cookie-jar
filename] [-e {command | 'command; command' | @batchfile}]
```

NOT_SUPPORTED

This example shows the user on an Exadata compute node issuing the command to log in to ExaCLI start an interactive ExaCLI session on a storage server:

```
[opc@exacs-node1 ~]$ exacli -l cloud_user_clustername -c 192.168.136.7
```

See *Connecting to a Storage Server with ExaCLI* for information on determining your storage server's IP address.

Once logged in, run additional commands as follows:

```
exacli cloud_user_clustername@192.168.136.7> LIST DATABASE
ASM
HRCDB
```

NOT_SUPPORTED

Example 2

This example shows a single command issued on a compute node that does the following:

- Connects to a storage server
- Performs a LIST action
- Exits the session (specified with the "-e" flag)

```
[opc@exacs-nodel ~]$ exacli -l cloud_user_clustername -c 192.168.136.7 --xml
--cookie-jar -e list griddisk detail
```

NOT_SUPPORTED

Option	Description
-c [username@]remotehost or --connect [username@]remotehost[:port]	Specifies the remote node to which you want to connect. ExaCLI prompts for the user name if not specified.
-l username or --login-name username	Specifies the user name to log into the remote node. The preconfigured user is <code>cloud_user_clustername</code> .
--xml	Displays the output in XML format.
--cookie-jar [filename]	Specifies the filename of the cookie jar to use. If filename is not specified, the cookie is stored in a default cookie jar located at <code>HOME/.exacli/cookiejar</code> , where <code>HOME</code> is the home directory of the OS user running the ExaCLI command. The presence of a valid cookie allows the ExaCLI user to execute commands without requiring to login in subsequent ExaCLI sessions.
-e command or -e 'command[; command]' or -e @batchFile	Specifies either the ExaCLI commands to run or a batch file. ExaCLI exits after running the commands. If specifying multiple commands to run, enclose the commands in single quotes to prevent the shell from interpreting the semi-colon. Omit this option to start an interactive ExaCLI session.
--cert-proxy proxy[:port]	Specifies the proxy server to use when downloading certificates. If <code>port</code> is omitted, port 80 is used by default.
-n or --no-prompt	Suppresses prompting for user input.

NOT_SUPPORTED

- Notes for the `--cookie-jar` option:
 - The user name and password are sent to the remote node for authentication. On successful authentication, the remote node issues a cookie (the login credentials) that is stored in the specified `filename` on the database node. If `filename` is not specified, the cookie is stored in a default cookie jar located at `HOME/.exacli/cookiejar`, where

- HOME is the home directory of the operating system user running the ExaCLI command. For the `opc` user, the home is `/home/opc`.
- The operating system user running the ExaCLI command is the owner of the cookie-jar file.
 - A cookie jar can contain multiple cookies from multiple users on multiple nodes in parallel sessions.
 - Cookies are invalidated after 24 hours.
 - If the cookie is not found or is no longer valid, ExaCLI prompts for the password. The new cookie is stored in the cookie jar identified by `filename`, or the default cookie jar if `filename` is not specified.
 - Even without the `--cookie-jar` option, ExaCLI still checks for cookies from the default cookie jar. However, if the cookie does not exist or is no longer valid, the new cookie will not be stored in the default cookie jar if the `--cookie-jar` option is not specified.
- Notes for the `-e` option:
 - ExaCLI exits after running the commands.
 - If specifying multiple commands to run, be sure to enclose the commands in single quotes to prevent the shell from interpreting the semi-colon.
 - The batch file is a text file that contains one or more ExaCLI commands to run.
 - Notes for the `-n` (`--no-prompt`) option:
 - If ExaCLI needs additional information from the user, for example, if ExaCLI needs to prompt the user for a password (possibly because there were no valid cookies in the cookie-jar) or to prompt the user to confirm the remote node's identity, then ExaCLI prints an error message and exits.

Related Topics

- [Connecting to a Storage Server with ExaCLI](#)
To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

Connecting to a Storage Server with ExaCLI

To use ExaCLI on storage servers, you will need to know your target storage server's IP address.

If you do not know the IP address of the node you want to connect to, you can find it by viewing the contents of the `cellip.ora` file.

The following example illustrates how to do so on the UNIX command line for a quarter rack system. (Note that a quarter rack has three storage cells, and each cell has two connections, so a total of six IP addresses are shown.)

```
cat /etc/oracle/cell/  
network-config/cellip.ora  
cell="192.168.136.5;cell="192.168.136.6"  
cell="192.168.136.7;cell="192.168.136.8"  
cell="192.168.136.9;cell="192.168.136.10"
```

If you are connecting to a storage cell for the first time using ExaCLI, you may be prompted to accept an SSL certificate. The ExaCLI output in this case will look like the following:

```
exacli -l cloud_user_clustername -c 192.168.136.7 --cookie-jar
No cookies found for cloud_user_clustername@192.168.136.7
Password: *****
EXA-30016: This connection is not secure. You have asked ExaCLI to connect to
cell 192.168.136.7 securely. The identity of 192.168.136.7 cannot be verified.
Got certificate from server:
C=US,ST=California,L=Redwood City,O=Oracle Corporation,OU=Oracle
Exadata,CN=ed1cl03clu01-priv2.usdc2.oraclecloud.com
Do you want to accept and store this certificate? (Press y/n)
```

Accept the self-signed Oracle certificate by pressing "y" to continue using ExaCLI.

ExaCLI Command Reference

You can execute various ExaCLI commands to monitor and manage Exadata Storage Servers associated with your Oracle Cloud Infrastructure Exadata VM cluster. ExaCLI allows you to get up-to-date, real-time information about your Exadata Cloud Service.

Use the LIST command with the following services and objects:

- **ACTIVEREQUEST:** Lists all active requests that are currently being served by the storage servers.
- **ALERTDEFINITION:** Lists all possible alerts and their sources for storage servers.
- **ALERTHISTORY:** Lists all alerts that have been issued for the storage servers.
- **CELL:** Lists attribute details of the storage servers (cells).

- To display specific attributes:

```
LIST CELL ATTRIBUTES A, B, C
```

Shows the values of the specified attributes.

- To display all attributes:

```
LIST CELL ATTRIBUTES ALL
```

Shows the values of all available attributes.

- **CELLDISK:** Lists the attributes of the cell disks in the storage servers.

```
LIST CELLDISK cell_disk_name DETAIL
```

Displays detailed information for the specified cell disk.

- **DATABASE:** Lists details of the databases.

```
LIST DATABASE
```

Displays a summary of all databases.

```
LIST DATABASE DETAIL
```

Displays detailed information for all databases.

```
LIST DATABASE ATTRIBUTES NAME
```

Displays the specified attribute (in this case, the NAME) for each database.

- **FLASHCACHE:** Lists the details of the Exadata system's flash cache.

```
LIST FLASHCACHE DETAIL
```

Displays detailed information for the flash cache.

```
LIST FLASHCACHE ATTRIBUTES attribute_name
```

Displays the specified attribute for the flash cache.

- **FLASHCACHECONTENT:** Lists the details of all objects in the flash cache, or the details of a specified object ID.

```
LIST FLASHCACHECONTENT DETAIL
```

Displays detailed information for all objects in the flash cache.

```
LIST FLASHCACHECONTENT WHERE objectNumber=12345 DETAIL
```

Displays detailed information for the object with the specified *objectNumber*.

Note

To find the object ID of a specific object, query `user_objects` using the object's name to retrieve the `data_object_id` of a partition or table.

- **FLASHLOG:** Lists the attributes for the Oracle Exadata Smart Flash Log.
- **GRIDDISK:** Lists the details of a particular grid disk. The syntax is similar to the `CELLDISK` command.

```
LIST GRIDDISK grid_disk_name DETAIL
```

Displays all attributes of the specified grid disk.

```
LIST GRIDDISK grid_disk_name ATTRIBUTES size, name
```

Displays only the specified attributes (*size, name*) of the grid disk.

- **IBPORT:** Lists details of the InfiniBand ports.

```
LIST IBPORT DETAIL
```

Displays detailed information for all InfiniBand ports.

- **IORMPROFILE:** Lists any IORM profiles that have been set on the storage servers. You can also refer to the profile attribute on the DATABASE object to see if a database has an associated IORM profile.

```
LIST IORMPROFILE
```

Displays the IORM profiles configured on the storage servers.

- **LUN:** Represents the logical unit numbers (LUNs) of the physical disks in the storage servers.

```
LIST LUN
```

Displays a summary of all LUNs.

```
LIST LUN lun_number DETAIL
```

Displays detailed information for the specified LUN.

- **METRICCURRENT:** Lists the current metrics for a particular object type.

```
LIST METRICCURRENT WHERE objectType = 'CELLDISK'
```

Displays the current metrics for the specified object type (in this case, CELLDISK).

```
LIST METRICCURRENT ATTRIBUTES name, metricObjectName  
ORDER BY metricObjectName ASC, name DESC LIMIT 5
```

Displays selected attributes, sorted by metricObjectName (ascending) and name (descending), limited to the top 5 results.

- **METRICDEFINITION:** Lists the metric definitions available for a given object type. These definitions can then be used to retrieve details for specific metrics.

```
LIST METRICDEFINITION WHERE objectType = cell
```

Displays all metric definitions for the specified object type (cell).

```
LIST METRICDEFINITION WHERE name = IORM_MODE DETAIL
```

Displays detailed information for the specified metric (IORM_MODE).

- **METRICHISTORY:** Lists metrics collected over a specified period of time.

```
LIST METRICHISTORY WHERE ageInMinutes < 30
```

Displays all metrics collected in the past 30 minutes.

```
LIST METRICHISTORY WHERE collectionTime > '2018-04-01T21:12:00-10:00'
```

Displays all metrics collected after the specified timestamp.

```
LIST METRICHISTORY CT_FD_IO_RQ_SM
```

Displays the history of a specific metric by name.

```
LIST METRICHISTORY WHERE name LIKE 'CT_.*'
```

Displays all metrics with names matching the given pattern.

- **OFFLOADGROUP:** Lists the attributes of the offload groups running on the storage servers.

```
LIST OFFLOADGROUP DETAIL
```

Displays detailed information for all offload groups.

```
LIST OFFLOADGROUP offloadgroup4
```

Displays the details of a specific offload group (offloadgroup4).

```
LIST OFFLOADGROUP ATTRIBUTES name
```

Displays only the specified attribute(s), such as name, for all offload groups.

- **PHYSICALDISK:** Lists all physical disks in the storage servers. Use the results to identify a specific disk for further investigation.

```
LIST PHYSICALDISK
```

Displays a summary of all physical disks.

```
LIST PHYSICALDISK 20:10 DETAIL
```

Displays detailed information for a specific disk (20:10).

```
LIST PHYSICALDISK FLASH_1_0 DETAIL
```

Displays detailed information for a specific flash disk (FLASH_1_0).

- **PLUGGABLEDATABASE:** Lists all pluggable databases (PDBs) in the environment.

```
LIST PLUGGABLEDATABASE
```

Displays a summary of all pluggable databases.

```
LIST PLUGGABLEDATABASE pdb_name
```

Displays detailed information for the specified pluggable database (*pdb_name*).

- **QUARANTINE:** Lists all SQL statements that have been prevented from using Smart Scans.

```
LIST QUARANTINE DETAIL
```

Displays detailed information for all quarantined SQL statements.

```
LIST QUARANTINE WHERE attribute = value
```

Filters results by specific attributes using a WHERE clause.

Use the ExaCLI [CREATE](#), [ALTER](#), [DROP](#), and [LIST](#) commands to act on the following Exadata Storage Server objects:

- **DIAGPACK:** Lists diagnostic packages and their status in the Exadata system.

```
LIST DIAGPACK
LIST DIAGPACK DETAIL
```

Lists all diagnostic packages, with DETAIL providing extended information.

```
CREATE DIAGPACK packStartTime=2019_12_15T00_00_00
```

Creates a diagnostic package starting from the specified time. You can also use `now` to capture diagnostics immediately:

```
CREATE DIAGPACK packStartTime=now
```

```
DOWNLOAD DIAGPACK cfclcx2647_diag_2018_06_03T00_44_24_1 /tmp
```

Downloads the specified diagnostic package to the `/tmp` directory (or another local path).

- **IORMPLAN:** Manage I/O Resource Manager (IORM) plans on Exadata storage servers. You can list, create, alter, and drop IORM plans using ExaCLI.

```
LIST IORMPLAN DETAIL
```

Lists all IORM plans with detailed information.

You can also use `CREATE IORMPLAN`, `ALTER IORMPLAN`, or `DROP IORMPLAN` to manage plans, and apply them to storage servers as needed.

```
select object_name, data_object_id from user_objects where object_name =
'BIG_CENSUS';
OBJECT_NAME                DATA_OBJECT_ID
```

BIG_CENSUS

29152

Monitor Metrics for VM Cluster Resources

You can monitor the health, capacity, and performance of your VM clusters and databases with metrics, alarms, and notifications. You can use Oracle Cloud Infrastructure Console, Monitoring APIs, or Database Management APIs to view metrics.

Note: To view metrics you must have the required access as specified in an Oracle Cloud Infrastructure policy (whether you're using the Console, the REST API, or another tool). See [Getting Started with Policies](#) for information on policies.

Warning

Metrics, events, and audit events will not be sent if Cluster Ready Services (CRS) is not running before Autonomous Health Framework (AHF) starts.

- [Prerequisites for Using Metrics](#)
- [View Metrics for VM Cluster](#)
- [View Metrics for a Database](#)
- [View Metrics for VM Clusters in a Compartment](#)
- [View Metrics for Databases in a Compartment](#)
- [Manage Oracle Trace File Analyzer](#)
- [Manage Database Service Agent](#)

Prerequisites for Using Metrics

The following prerequisites are required for the metrics to flow out of the VM cluster.

1. Metrics on the VM clusters depends on Oracle Trace File Analyzer (TFA) agent. Ensure that these components are up and running. AHF version **22.2.4** or higher is required for capturing metrics from the VM clusters. To start, stop, or check the status of TFA, see *Manage Oracle Trace File Analyzer*.
2. To view the metrics on the Oracle Cloud Infrastructure Console, the TFA flag `defaultocimonitoring` must be set to `ON`. This flag is set to `ON` by default and you need not perform any action to set this. If you are not seeing metrics on the Console, then as `root` user on the guest VM, check if the flag is set to `ON`.

```
tfactl get defaultocimonitoring
```

```

-----
|                                     <host name>                                     |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Configuration Parameter                                                    | Value |
-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Send CEF metrics to OCI Monitoring ( defaultOciMonitoring ) | ON    |
-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

If the `defaultocimonitoring` flag is set to `OFF`, then run the `tfactl set defaultocimonitoring=on` Or `tfactl set defaultocimonitoring=ON` command to turn it on:

```
tfactl set defaultocimonitoring=on
Successfully set defaultOciMonitoring=ON
```

<host name>	
Configuration Parameter	Value
Send CEF metrics to OCI Monitoring (defaultOciMonitoring)	ON

3. The following network configurations are required.
 - a. **Egress rules for outgoing traffic:** The default egress rules are sufficient to enable the required network path : For more information, see [Default Security List](#) .If you have blocked the outgoing traffic by modifying the default egress rules on your Virtual Cloud Network(VCN), you will need to revert the settings to allow outgoing traffic. The default egress rule allowing outgoing traffic (as shown in the *Rules Required for both Client and Backup Networks*) is as follows:
 - Stateless: No (all rules must be stateful)
 - Destination Type: CIDR
 - Destination CIDR: All <region> Services in Oracle Services Network
 - IP Protocol: 443 (HTTPS)
 - b. **Public IP or Service Gateway:** The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.
If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send compute instance metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:
 - i. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
 - ii. When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Access to Oracle Services: Service Gateway](#).

Related Topics

- [Manage Oracle Trace File Analyzer](#)
- [Rules Required for Both the Client Network and Backup Network](#)
This topic has several general rules that enable essential connectivity for hosts in the VCN.
- [Rules Required for Monitoring Service](#)
The compute instance must have either a public IP address or a service gateway to be able to send compute instance metrics to the Monitoring service.

View Metrics for VM Cluster

Perform the following steps to view the metrics for Guest VMs using the console.

Note

When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.

Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the **Resources** section, click **Metrics**.
A chart for each metrics is displayed. By default, the metrics for the last one hour are displayed.
You can only select the `oci_database_cluster` namespace from the **Metric namespace** drop-down.
5. If you want to change the interval, select the required start time and end time. Alternatively, you can select the interval from the Quick Selects drop down menu. The metrics are refreshed immediately for the selected interval.
6. For each metric, you can choose the interval and statistic independently.
 - Interval - The time period for which the metric is calculated.
 - Statistic - The mathematical method by which the metric is calculated.
7. For each metric, you can choose the following options from the 'Options' drop down menu.
 - View Query in Metrics Explorer
 - Copy Chart URL
 - Copy Query (MQL)
 - Create an Alarm on this Query
 - Table View

For Detailed information on various options for viewing the metrics chart, see [Viewing Default Metric Charts](#).

View Metrics for a Database

Perform the following steps to view the metrics for a database using the console.

Note

When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.

Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster that contains the database for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the list of databases, click the database for which you want to view the metrics.
5. In the **Resources** section, click **Metrics**.
A chart for each metrics is displayed. By default, the metrics for the last one hour are displayed.
6. Select a namespace from the **Metric namespace** from where you wish to view metrics.

Note

- When Database Management is enabled, you will have an option to choose from `oci_database` or `oracle_oci_database` namespace.
- When Database Management is disabled, then you can view metrics only from the `oci_database` namespace.

7. If you want to change the interval, select the required start time and end time. Alternatively, you can select the interval from the Quick Selects drop down menu. The metrics are refreshed immediately for the selected interval.
8. For each metric, you can choose the interval and statistic independently.
 - Interval - The time period for which the metric is calculated.
 - Statistic - The mathematical method by which the metric is calculated.
9. For each metric, you can choose the following options from the 'Options' drop down menu.
 - View Query in Metrics Explorer
 - Copy Chart URL
 - Copy Query (MQL)
 - Create an Alarm on this Query
 - Table View

For Detailed information on various options for viewing the metrics chart, see [Viewing Default Metric Charts](#).

View Metrics for a PDB

1. Open the navigation menu. Click **Oracle AI Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**. A list of VM clusters is displayed.
3. In the list of VM clusters, click the VM cluster that contains the database for which you want to view the metrics. Details of the VM cluster you selected are displayed.
4. In the list of databases, click the database that contains the PDB for which you want to view the metrics.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of VM clusters, click the PDB that you wish to view metrics.
7. Select a namespace from the **Metric namespace** from where you wish to view metrics.

Note

- When Database Management is enabled, you will have an option to choose from `oracle_oci_database` namespace.
- When Database Management is disabled, then the system will display a banner asking you to enable Database Management to provide metrics.

View Metrics for VM Clusters in a Compartment

Perform the following steps to view the metrics for databases in a compartment using the console.

Note

When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.

Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the Oracle Cloud Infrastructure **Console** by clicking the menu icon next to **Oracle Cloud**.
2. From the left navigation list click **Observability & Management**.
3. Under **Monitoring**, click **Service Metrics**.
4. On the Service Metrics page, under **Compartment** select your compartment.
5. On the Service Metrics page, under **Metric Namespace** select `oci_database_cluster`.
6. If there are multiple VM clusters in the compartment you can show metrics aggregated across the clusters by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).

8. In the **Dimension Name** field select a dimension.
9. In the **Dimension Value** field select a value.
10. Click **Done**.
11. In the **Edit dimensions** dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See *Managing Alarms* for information on setting and using alarms.

Note

If you don't see any metrics, check the network settings and AHF version listed in the prerequisites section.

Related Topics

- [Managing Alarms](#)

View Metrics for Databases in a Compartment

Perform the following steps to view the metrics for databases in a compartment using the console.

Note

When there is a network problem and Oracle Trace File Analyzer (TFA) is unable to post metrics, TFA will wait for one hour before attempting to retry posting the metrics. This is required to avoid creating a backlog of metrics processing on TFA.

Potentially one hour of metrics will be lost between network restore and the first metric posted.

1. Open the Oracle Cloud Infrastructure **Console** by clicking the menu icon next to **Oracle Cloud**.
2. From the left navigation list click **Observability & Management**.
3. Under **Monitoring**, click **Service Metrics**.
4. On the Service Metrics page, under **Compartment** select your compartment.
5. On the Service Metrics page, under **Metric Namespace** select `oci_database`.
6. If there are multiple databases in the compartment you can show metrics aggregated across the databases by selecting **Aggregate Metric Streams**.
7. If you want to limit the metrics you see, next to **Dimensions** click **Add** (click **Edit** if you have already added dimensions).
8. In the **Dimension Name** field select a dimension.
9. In the **Dimension Value** field select a value.
10. Click **Done**.

11. In the **Edit dimensions** dialog click **+Additional Dimension** to add an additional dimension. Click **X** to remove a dimension.
12. To create an alarm on a specific metric, click **Options** and select **Create an Alarm on this Query**. See [Managing Alarms](#) for information on setting and using alarms.

Manage Oracle Trace File Analyzer

The deployment of the cloud-certified Autonomous Health Framework (AHF), which includes Oracle Trace File Analyzer, is managed by Oracle. You shouldn't install this manually on the guest VMs.

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-----
-----
| Host          | Status of TFA | PID   | Port | Version  | Build
ID            | Inventory Status|
+-----+-----+-----+-----+-----+
+-----+-----+
| node1        | RUNNING       | 41312 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE     |
| node2        | RUNNING       | 272300 | 5000 | 22.1.0.0.0 |
22100020220310214615| COMPLETE     |
'-----+-----+-----+-----+-----'
+-----+-----'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands
```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as `root`:

```
# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
```

```
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..
```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```
# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar"|awk '{print $2}'` (code=exited, status=0/SUCCESS)
Main PID: 10055(sudo)
CGroup: /system.slice/dbcsagent.service
        10055sudo -u opc /bin/bash -c umask 077; /bin/java
```

- To start the agent if it is not running, run the `systemctl start` command as the root user:

```
systemctl start dbcsagent.service
```

Metrics for Oracle Exadata Database Service on Dedicated Infrastructure in the Monitoring Service

This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oci_database_cluster` and `oci_database` namespaces for Oracle Databases.

Dimensions

All the metrics discussed in this topic include the following dimensions.

- resourceId:** The OCID of the VM Cluster.
- resourceName:** The name of the VM Cluster.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the VM cluster.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
ASMDiskgroup Utilization	ASM Diskgroup Utilization	percentage	Percentage of usable space used in a Disk Group. Usable space is the space available for growth. DATA disk group stores our Oracle database files. RECO disk group contains database files for recovery such as archives and flashback logs.	10 minutes	deploymentType diskgroupName
CpuUtilization	CPU Utilization	percentage	Percent CPU utilization.	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode
FilesystemUtilization	Filesystem Utilization	percentage	Percent utilization of provisioned filesystem.	1 minute	hostName deploymentType filesystemName resourceId_dbnode resourceName_dbnode
LoadAverage	Load Average	number	System load average over 5 minutes.	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode
MemoryUtilization	Memory Utilization	percentage	Percentage of memory available for starting new applications, without swapping. The available memory can be obtained via the following command: <code>cat /proc/meminfo</code>	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
NodeStatus	Node Status	integer	Indicates whether the host is reachable.	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode
OcpusAllocated	OCPU Allocated	integer	The number of OCPUs allocated.	1 minute	deploymentType
SwapUtilization	Swap Utilization	percentage	Percent utilization of total swap space.	1 minute	hostName deploymentType resourceId_dbnode resourceName_dbnode

NOT_SUPPORTED

The metrics listed in the following table are automatically available for the database.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
CpuUtilization	CPU Utilization	percentage	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
StorageUtilization	Storage Utilization	percentage	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces.	1 hour	deploymentType resourceId_{database pdb} resourceName_{database pdb}
BlockChanges	DB Block Changes	Changes per second	The Average number of blocks changed per second.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
ExecuteCount	Execute Count	Count	The number of user and recursive calls that executed SQL statements during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
CurrentLogons	Current Logons	Count	The number of successful logons during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
TransactionCount	Transaction Count	Count	The combined number of user commits and user rollbacks during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
UserCalls	User Calls	Count	The combined number of logons, parses, and execute calls during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
ParseCount	Parse Count	Count	The number of hard and soft parses during the selected interval.	5 minutes	instanceNumber instanceName hostName deploymentType resourceId_{database pdb} resourceName_{database pdb}
StorageUsed	Storage Space Used	GB	Total amount of storage space used by the database at the collection time.	1 hour	deploymentType resourceId_{database pdb} resourceName_{database pdb}

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
StorageAllocated	Storage Space Allocated	GB	Total amount of storage space allocated to the database at the collection time.	1 hour	deploymentType resourceId_{database pdb} resourceName_{database pdb}
StorageUsedByTablespace	Storage Space Used By Tablespace	GB	Total amount of storage space used by tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName tablespaceType deploymentType resourceId_{database pdb} resourceName_{database pdb}
StorageAllocatedByTablespace	Allocated Storage Space By Tablespace	GB	Total amount of storage space allocated to the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName tablespaceType deploymentType resourceId_{database pdb} resourceName_{database pdb}
StorageUtilizationByTablespace	Storage Space Utilization By Tablespace	percentage	This indicates the percentage of storage space utilized by the tablespace at the collection time. In case of container database, this metric provides root container tablespaces.	1 hour	tablespaceName tablespaceType deploymentType resourceId_{database pdb} resourceName_{database pdb}

Oracle Exadata Database Service on Dedicated Infrastructure Events

Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources.

Note

In addition to the events listed below, Oracle analyzes additional events to provide highest level of service operations and support for ensuring high availability of services.

- [About Event Types on Exadata Cloud Infrastructure](#)
Learn about the event types available for Exadata Cloud Infrastructure resources.
- [Prerequisites for Event Service](#)
The following prerequisites are required for the Events to flow out of the VM Cluster.
- [Oracle Exadata Database Service on Dedicated Infrastructure Event Types](#)
The events in this section are emitted by the cloud Exadata infrastructure resource
- [Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types](#)
The events in this section are emitted by the cloud Exadata infrastructure resource for Maintenance Events
- [Exadata Cloud Infrastructure Critical and Information Event Types](#)
Exadata Cloud Infrastructure infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs attention.
- [Exascale DB Storage Vaults Event Types](#)
The events in this section are emitted by the `exascale-db-storage-vaults` resource.
- [Exadata Cloud Infrastructure VM Cluster Event Types](#)
Review the list of events that can be emitted by VM Cluster
- [VM Node Subsetting Event Types](#)
Review the list of event types that VM Node Subsetting emits.
- [Data Guard Event Types](#)
Review the list of event types that Data Guard group and Data Guard Associations emit.
- [Oracle Database Home Event Types](#)
Review the list of events emitted by Oracle Database Homes.
- [Database Event Types](#)
These are the event types that Oracle Databases in Exadata Cloud Service instances emit.
- [Pluggable Database Event Types](#)
These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.
- [Database Service Events](#)
The Database Service emits events, which are structured messages that indicate changes in resources.

- [Application VIP Event Types](#)
These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.
- [Interim Software Updates Event Types](#)
These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.
- [Serial Console Connection Event Types](#)
Review the list of event types that serial console connection emits.
- [Serial Console History Event Types](#)
Review the list of new event types that serial console history emits.

About Event Types on Exadata Cloud Infrastructure

Learn about the event types available for Exadata Cloud Infrastructure resources.

Exadata Cloud Infrastructure resources emit events, which are structured messages that indicate changes in resources. For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*. You may subscribe to events and be notified when they occur using the Oracle Notification service, see *Notifications Overview*.

Related Topics

- [Overview of Events](#)
- [Notifications Overview](#)

Prerequisites for Event Service

The following prerequisites are required for the Events to flow out of the VM Cluster.

The Event Service requires the following:

1. Events on the VM Cluster depends on Oracle Trace File Analyzer (TFA) agent. Ensure that these components are up and running. AHF version **22.2.2** or higher is required for capturing events from the VM Cluster. To start, stop, or check the status of TFA, see [Incident Logs and Trace Files](#) . To enable AHF Telemetry for the VM Cluster using the dbcli utility, see [AHF Telemetry Commands](#)
2. The following network configurations are required.
 - a. **Egress rules for outgoing traffic:** The default egress rules are sufficient to enable the required network path : For more information, see [Default Security List](#) .If you have blocked the outgoing traffic by modifying the default egress rules on your Virtual Cloud Network(VCN), you will need to revert the settings to allow outgoing traffic. The default egress rule allowing outgoing traffic (as shown in [Security Rules for the Oracle Exadata Database Service on Dedicated Infrastructure](#)) is as follows:
 - Stateless: No (all rules must be stateful)
 - Destination Type: CIDR
 - Destination CIDR: **All <region> Services in Oracle Services Network**
 - IP Protocol: TCP
 - Destination Port: 443 (HTTPS)
 - b. **Public IP or Service Gateway:** The database server host must have either a public IP address or a service gateway to be able to send database server host metrics to the Monitoring service.

If the instance does not have a public IP address, set up a service gateway on the virtual cloud network (VCN). The service gateway lets the instance send database server host metrics to the Monitoring service without the traffic going over the internet. Here are special notes for setting up the service gateway to access the Monitoring service:

- i. When creating the service gateway, enable the service label called **All <region> Services in Oracle Services Network**. It includes the Monitoring service.
- ii. When setting up routing for the subnet that contains the instance, set up a route rule with **Target Type** set to **Service Gateway**, and the **Destination Service** set to **All <region> Services in Oracle Services Network**.

For detailed instructions, see [Access to Oracle Services: Service Gateway](#).

Related Topics

- [Service Gateway for the VCN](#)
Ensure that your VCN can reach the Oracle Services Network—specifically Object Storage for backups, Oracle YUM repositories for OS updates, IAM (Identity and Access Management), and OCI Vault (KMS integration).

Oracle Exadata Database Service on Dedicated Infrastructure Event Types

The events in this section are emitted by the cloud Exadata infrastructure resource

Note

Exadata systems that use the old DB system resource model are deprecated and desupported.

Friendly Name	Event Type
Cloud Exadata Infrastructure - Create Begin	com.oraclecloud.databaseservice.createcloudexadatainfrastructure.begin
Cloud Exadata Infrastructure - Create End	com.oraclecloud.databaseservice.createcloudexadatainfrastructure.end
Cloud Exadata Infrastructure - Change Compartment Begin	com.oraclecloud.databaseservice.changecloudexadatainfrastructurecompartment.begin
Cloud Exadata Infrastructure - Change Compartment End	com.oraclecloud.databaseservice.changecloudexadatainfrastructurecompartment.end
Cloud Exadata Infrastructure - Critical See Exadata Cloud Service Infrastructure Critical and Information Event Types for details	com.oraclecloud.databaseservice.cloudexadatainfrastructure.critical
Cloud Exadata Infrastructure - Delete Begin	com.oraclecloud.databaseservice.deletecloudexadatainfrastructure.begin
Cloud Exadata Infrastructure - Delete End	com.oraclecloud.databaseservice.deletecloudexadatainfrastructure.end
Cloud Exadata Infrastructure - Information See Exadata Cloud Service Infrastructure Critical and Information Event Types for details	com.oraclecloud.databaseservice.cloudexadatainfrastructure.information
Cloud Exadata Infrastructure - Update Begin	com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin

Friendly Name	Event Type
Cloud Exadata Infrastructure - Update End	com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end
Cloud Exadata Infrastructure - Configure Exascale Begin	com.oraclecloud.databaseservice.configureexascalecloudexadatainfrastructure.begin
Cloud Exadata Infrastructure - Configure Exascale End	com.oraclecloud.databaseservice.configureexascalecloudexadatainfrastructure.end

This is a reference event for a Cloud Exadata Infrastructure resource:

```
{
  "cloudEventsVersion": "0.1",
  "eventId": "<unique_ID>",
  "eventType":
"com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end",
  "source": "DatabaseService",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>",
    "compartmentName": "example_name",
    "resourceName": "my_exadata_infrastructure",
    "resourceId": "ocidl.dbssystem.oc1.eu-frankfurt-1.<unique_ID>",
    "availabilityDomain": "tXPJ:EU-FRANKFURT-1-AD-3",
    "freeFormTags": {
      "Department": "Finance"
    },
    "definedTags": {
      "Operations": {
        "CostCenter": "42"
      }
    },
    "additionalDetails": {
      "subnetId": "ocidl.subnet.oc1.eu-frankfurt-1.<unique_ID>",
      "lifecycleState": "MAINTENANCE_IN_PROGRESS",
      "sshPublicKeys": "...",
      "cpuCoreCount": 32,
      "version": "19.2.8.0.0.191119",
      "nsgIds": "null",
      "backupSubnetId": "ocidl.subnet.oc1.eu-frankfurt-1.<unique_ID>",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "dataStoragePercentage": 80,
      "patchHistoryEntries": "null",
      "lifecycleMessage": "The underlying infrastructure of this system (cell
storage) is being updated and this will not impact database
availability.",
      "exadataIormConfig": "ExadataIormConfigCache(lifecycleState=DISABLED,
lifeCycleDetails=null, objective=Auto,
```

```

        dbPlans=[DbIormConfigCache(dbName=default, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database1>,
share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database2>, share=null, flashCacheLimit=null),
        DbIormConfigCache(dbName=<my_database3>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database4>,
share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database5>, share=null, flashCacheLimit=null),
        DbIormConfigCache(dbName=<my_database6>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database7>,
share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database8>, share=null, flashCacheLimit=null),
        DbIormConfigCache(dbName=<my_database9>, share=null,
flashCacheLimit=null), DbIormConfigCache(dbName=<my_database10>,
share=null, flashCacheLimit=null),
DbIormConfigCache(dbName=<my_database11>, share=null, flashCacheLimit=null)],
        undoData=null)"
    }
},
"eventID" : "<unique_ID>",
"extensions" : {
"compartmentId" : "ocidl.compartment.oc1.<unique_ID>"
}
}

```

This is a reference event for Cloud Exadata Infrastructure - Add Storage Capacity Begin:

```

{
  "id":
"ocidl.eventschema.oc1.phx.zlnzw5klc4r7arlvkxunfvfhtwmeaaylr0j5hjnu2j5uozwlie
xa53gwlk4",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Add Storage Capacity Begin",
  "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",

```

```

        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "shape",
        "type": "string"
    },
    {
        "name": "message",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "description",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "timeZone",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "maintenanceMode",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "maintenanceSubType",
        "type": [
            "null",
            "string"
        ]
    }
],
"exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "10274771-3706-4624-99d1-e036805a9ca7",
    "eventTime": "2023-01-06T21:16:04.000Z",
    "contentType": "application/json",
    "data": {

```

```

    "eventGroupingId": "csida87218404b4291914305ec7a5a86/
d53ffb13f83244bbbfb8c7d0a8f0e2eb/FB95D76D5123C152C25DBF288489077F",
    "eventName": "AddStorageCapacityCloudExadataInfrastructure",
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": null,
    "resourceName": "my_cloud_exadata_infrastructure",
    "resourceId": "ocidl.cloudexadatainfrastructure.oc1.....unique_id",
    "resourceVersion": null,
    "availabilityDomain": "",
    "tagSlug": "tag_slug",
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "authZPolicies": null,
      "userGroups": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "7e83c538-28bf-453d-9fb7-125bf70546c4",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,
      "message": null
    },
    "stateChange": {
      "previous": null,
      "current": {
        "lifecycleState": "AVAILABLE",
        "shape": "Exadata.X9M",
        "displayName": "my_display_name",
        "freeTags": {},
        "definedTags": {}
      }
    },
    "additionalDetails": {
      "timeCreated": "2023-01-06T21:16:04.000Z",
      "timeUpdated": "2023-01-06T21:16:04.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleDetails": null,
      "description": null,
      "message": null,
      "shape": "Exadata.X9M",

```

```

        "timeZone": null,
        "maintenanceMode": null,
        "maintenanceSubType": null
    },
    "internalDetails": {
        "attributes": null
    }
}
},
"timeCreated": "2023-01-06T21:16:04.000Z"
}

```

This is a reference event for Cloud Exadata Infrastructure - Add Storage Capacity End:

```

{
  "id":
  "ocidl.eventschema.oc1.phx.4aeklze2colynub2ojmu49shhduq9gh5qg6fvudm7h77w3og8sf
  kau6a3not",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Add Storage Capacity End",
  "eventType":
  "com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
  end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "shape",
      "type": "string"
    },
    {
      "name": "message",
      "type": [
        "null",

```

```

        "string"
      ]
    },
    {
      "name": "description",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "maintenanceMode",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "maintenanceSubType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.addstoragecapacitycloudexadatainfrastructure.
end",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "b12abcc0-110a-9120-aab5-9a34bc799e72",
    "eventTime": "2023-01-06T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId":
"csida2cd1c8442f9b9fc16354a1f0912/95202d41125e4ce18e8dd52fa9f57f5e/
545A43343BC1D5020A85AA2919C06E25",
      "eventName": "AddStorageCapacityCloudExadataInfrastructure",
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "compartmentName": null,
      "resourceName": "my_cloud_exadata_infrastructure",
      "resourceId": "ocid1.cloudexadatainfrastructure.oc1.....unique_id",
      "resourceVersion": null,
      "availabilityDomain": "",
      "tagSlug": "tag_slug",
      "identity": {
        "principalName": null,

```

```

    "principalId": null,
    "authType": null,
    "callerName": null,
    "callerId": null,
    "tenantId": null,
    "ipAddress": null,
    "credentials": null,
    "authZPolicies": null,
    "userGroups": null,
    "userAgent": null,
    "consoleSessionId": null
  },
  "request": {
    "id": "111b9da5-a7a7-4aca-bd05-a51558f7df55",
    "path": null,
    "action": null,
    "parameters": null,
    "headers": null
  },
  "response": {
    "status": null,
    "responseTime": null,
    "headers": null,
    "payload": null,
    "message": null
  },
  "stateChange": {
    "previous": null,
    "current": {
      "lifecycleState": "AVAILABLE",
      "shape": "Exadata.X9M",
      "displayName": "my_display_name",
      "freeTags": {},
      "definedTags": {}
    }
  },
  "additionalDetails": {
    "timeCreated": "2023-01-06T21:16:04.000Z",
    "timeUpdated": "2023-01-06T21:16:04.000Z",
    "lifecycleState": "AVAILABLE",
    "lifecycleDetails": null,
    "description": null,
    "message": null,
    "shape": "Exadata.X9M",
    "timeZone": null,
    "maintenanceMode": null,
    "maintenanceSubType": null
  },
  "internalDetails": {
    "attributes": null
  }
}
}
},
"timeCreated": "2023-01-06T21:16:04.000Z"
}

```

This is a reference event for Cloud Exadata Infrastructure - Update Begin

```
{
  "id":
  "ocid1.eventschema.oc1.phx.jlx9t3z6igwglicpbba6xs1uaewcb8txsegnuykc65n8rxl5tqd
  26ect7i3f",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Update Begin",
  "eventType":
  "com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
```

```

        "string"
      ]
    },
    {
      "name": "availabilityDomain",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "description",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "tenantId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "message",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "shape",
      "type": [
        "null",
        "String"
      ]
    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {

```

```

    "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
    "eventName": "UpdateCloudExadataInfrastructure",
    "compartmentName": "example_compartment",
    "resourceName": "my_container_database",
    "resourceId": "ocidl.cloudexadatainfrastructure.oc1.....unique_id",
    "resourceVersion": null,
    "additionalDetails": {
      "availabilityDomain": "all",
      "compartmentId": "ocidl.compartment.oc1.....unique_id",
      "freeFormTags": {},
      "definedTags": {},
      "lifecycleState": "AVAILABLE"
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}

```

This is a reference event for Cloud Exadata Infrastructure - Update End

```

{
  "id":
"ocidl.eventschema.oc1.phx.aq2fuvh1nh9h71bnyclhmsuj3bky7dr304xj7nejajjzwbh2n4
Ozy3tdand",
  "serviceName": "Database",
  "displayName": "Cloud Exadata Infrastructure - Update End",
  "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    }
  ],
}

```

```
{
  "name": "timeUpdated",
  "type": "string"
},
{
  "name": "lifecycleState",
  "type": "string"
},
{
  "name": "lifecycleDetails",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "compartmentId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "availabilityDomain",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "description",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "tenantId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "message",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "shape",
  "type": [
    "null",
    "String"
  ]
}
```

```

    },
    {
      "name": "timeZone",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType":
"com.oraclecloud.databaseservice.updatecloudexadatainfrastructure.end",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
      "eventGroupingId": "4976b940-2c2d-4380-a669-1d70d071b187",
      "eventName": "UpdateCloudExadataInfrastructure",
      "compartmentName": "example_compartment",
      "resourceName": "my_container_database",
      "resourceId": "ocidl.dbsystem-.....unique_id",
      "resourceVersion": null,
      "additionalDetails": {
        "availabilityDomain": "all",
        "compartmentId": "ocidl.compartment.ocl.....unique_id",
        "freeFormTags": {},
        "definedTags": {},
        "lifecycleState": "AVAILABLE"
      }
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}

```

Oracle Exadata Database Service on Dedicated Infrastructure Maintenance Event Types

The events in this section are emitted by the cloud Exadata infrastructure resource for Maintenance Events

Note

Exadata systems that use the old DB system resource model are deprecated and desupported.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Reminder (ROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder	This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure instance <i><infra_name></i> , OCID <i>infra_ocid</i> in approximately <i>no_of_days_left</i> days <i>schedule_time</i> . The maintenance method for this maintenance is Rolling as selected per the maintenance preferences.
Cloud Exadata Infrastructure - Maintenance Reminder (NONROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancereminder	This is an Oracle Cloud Operations reminder notice. Oracle has scheduled a quarterly maintenance update installation for Cloud Exadata Infrastructure instance <i><infra_name></i> , OCID <i>infra_ocid</i> in approximately <i>no_of_days_left</i> days on <i>schedule_time</i> . The maintenance method for this maintenance is %s as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.
Cloud Exadata Infrastructure - Maintenance Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <i><infra-name></i> , OCID <i><infra-ocid></i> part of Maintenance Run <i><mr-display-name></i> , OCID <i><mr-ocid></i> . Your maintenance update started at <i><start-time></i> . You will get a notification on the completion of the quarterly maintenance update.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance End	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenance.end	<p>Success: This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> part of Maintenance Run <mr-display-name>, OCID <mr-ocid>. Your maintenance update started at <start-time> and was completed successfully at <end-time>. You have successfully completed maintenance updates for this window.</p> <p>Failed: This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> part of Maintenance Run <mr-display-name>, OCID <mr-ocid>. Your maintenance update started at <start-time> and was failed to complete successfully as planned. Our operations team is evaluating the failure and will notify you of the next steps to complete the maintenance update for this quarter.</p> <p>Canceled: This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> part of Maintenance Run <mr-display-name>, OCID <mr-ocid>. Your maintenance update started at <start-time>. Your maintenance has been canceled as requested. And a new window will be created according to the time given.</p> <p>Duration Exceeded: This is an Oracle Cloud Operations notification for quarterly maintenance update of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> part of Maintenance Run <mr-display-name>, OCID <mr-ocid>. Your maintenance update started at <start-time>. Your window was configured for a WINDOW_DURATION duration. Your maintenance is taking longer than the configured window duration. This window has</p>

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - Maintenance Custom action time Begin (ROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.begin	duration enforcement enabled. Oracle automation will reschedule all updates that have not started to a future maintenance window. Please acknowledge the updates rescheduled to run in a future unplanned maintenance window created by Oracle.
Cloud Exadata Infrastructure - Maintenance Custom action time End (ROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.end	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Server <db-server-name>, OCID <db-server-ocid>. Your custom action time started at <start-time>. You will get a notification on the completion of the custom action time for the Database Server.
Cloud Exadata Infrastructure - Maintenance Custom action time Begin (NONROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.begin	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Servers <db-server-name>, OCID <dbserver-ocid> <db-server-name>, OCID <dbserver-ocid>. Your custom action time started at <start-time>. You will get a notification on the completion of the custom action time for the Database Servers.
Cloud Exadata Infrastructure - Maintenance Custom action time End (NONROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancecustomactiontime.end	This is an Oracle Cloud Operations notification for custom action time configured for your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Servers <db-server-name>, OCID <dbserver-ocid> <db-server-name>, OCID <dbserver-ocid>. Your custom action time started at <start-time> ended at <end-time>.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Storage Server Maintenance Begin	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceStorageServers.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Storage Server(s) count <cell-count>. Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the Storage Servers.
Cloud Exadata Infrastructure – Storage Server Maintenance End	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceStorageServers.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Storage Servers of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Storage Server(s) count <cell-count>. Your maintenance update started at <start-time> and was completed successfully at <end-time>.
Cloud Exadata Infrastructure - DB Server Maintenance Begin (ROLLING)	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Server <db-server-name>, OCID <db-server-ocid>. Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the Database Server.
Cloud Exadata Infrastructure - DB Server Maintenance End (ROLLING)	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Server <db-server-name>, OCID <db-server-ocid>. Your maintenance update started at <start-time> and was completed successfully at <end-time>.

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure - DB Server Maintenance Begin (NONROLLING)	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.begin	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Servers <db-server-name>, OCID <dbserver-ocid> <db-server-name>, OCID <dbserver-ocid>. Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the Database Servers.
Cloud Exadata Infrastructure - DB Server Maintenance End (NONROLLING)	com.oraclecloud.DatabaseService.CloudExadataInfrastructureMaintenanceVM.end	This is an Oracle Cloud Operations notification for quarterly maintenance update of the Database Server component of your ExaDB-D Exadata Infrastructure <infra-name>, OCID <infra-ocid> for Database Servers <db-server-name>, OCID <dbserver-ocid> <db-server-name>, OCID <dbserver-ocid>. Your maintenance update started at <start-time> and was completed successfully at <end-time>.
Cloud Exadata Infrastructure - VM Maintenance Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructurevmmaintenancebegin	Your maintenance update started at <start-time>. You will get a notification on the completion of the quarterly maintenance update for the virtual machines.
Cloud Exadata Infrastructure - VM Maintenance End	com.oraclecloud.databaseservice.cloudexadatainfrastructurevmmaintenanceend	Your maintenance update started at <start-time> and was completed successfully at <end-time>.
Cloud Exadata Infrastructure – Maintenance Method Change	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancemethodchange	Oracle Cloud Operations is announcing a change related to quarterly maintenance update for Cloud Exadata Infrastructure. There's a change in maintenance method on your service instance <infra_name>, OCID <infra_ocid> to <i>new_patching_mode</i> .
Cloud Exadata Infrastructure – Maintenance Rescheduled	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduled	Oracle Cloud Operations is announcing reschedule of a quarterly maintenance update for Cloud Infrastructure. A maintenance run has been rescheduled on your service instance <infra_name>, OCID <infra_ocid> to <i>new_schedule_time</i> .

Friendly Name	Event Type	Event Messages
Cloud Exadata Infrastructure – Maintenance Rescheduled With Reason	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancerescheduledwithreason	The operator entered reason for rescheduling is sent as an email.
Cloud Exadata Infrastructure – Maintenance Scheduled (ROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled	Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <infra_name>, OCID <i>infra_ocid</i> on <i>schedule_time</i> . The maintenance method for this maintenance is %s as selected per the maintenance preferences.
Cloud Exadata Infrastructure – Maintenance Scheduled (NONROLLING)	com.oraclecloud.databaseservice.cloudexadatainfrastructuremaintenancescheduled	Oracle Cloud Operations is announcing the availability of a new quarterly maintenance update for Cloud Exadata Infrastructure. Oracle has scheduled the installation of this new update on your service instance <infra_name>, OCID <i>infra_ocid</i> on <i>schedule_time</i> . The maintenance method for this maintenance is %s as selected per the maintenance preferences. Non-rolling maintenance minimizes maintenance time but will result in full system downtime.
Cloud Exadata Infrastructure - IB Switch Maintenance Begin	com.oraclecloud.databaseservice.cloudexadatainfrastructureibswitchmaintenance.begin	This is an Oracle Cloud Operations notice regarding the quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <infra_name>, OCID <i>infra_ocid</i> has started. A follow-up notice will be sent when the network fabric switches maintenance operation has completed.
Cloud Exadata Infrastructure - IB Switch Maintenance End	com.oraclecloud.databaseservice.cloudexadatainfrastructureibswitchmaintenance.end	This is an Oracle Cloud Operations notice that quarterly maintenance update of the network fabric switches component of your Cloud Exadata Infrastructure instance <infra_name>, OCID <i>infra_ocid</i> has completed.

Exadata Cloud Infrastructure Critical and Information Event Types

Exadata Cloud Infrastructure infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs attention.

Exadata Cloud Service infrastructure resources emit "critical" and "information" data plane events that allow you to receive notifications when your infrastructure resource needs urgent attention ("critical" events), or notifications for events that are not critical, but which you may want to monitor ("information" events). The `eventType` values for these events are the following:

- `com.oraclecloud.databaseservice.exadatainfrastructure.critical`
- `com.oraclecloud.databaseservice.exadatainfrastructure.information`

These events use the `additionalDetails` section of the event message to provide specific details about what is happening within the infrastructure resource emitting the event. In the `additionalDetails` section, the `eventName` field provides the name of the critical or information event. (*Note that some fields in the example that follows have been omitted for brevity.*)

```
{
  "eventType" :
  "com.oraclecloud.databaseservice.exadatainfrastructure.critical",
  ....
  "data" : {
    ....
    "additionalDetails" : {
      ....
      "description" : "SQL statement terminated by Oracle Database Resource
Manager due to excessive consumption of CPU and/or I/O.
The execution plan associated with the terminated SQL
stmt is quarantined. Please find the sql identifier in
sqlId field of this JSON payload. This feature protects
an Oracle database from performance degradation.
Please review the SQL statement. You can see the
statement using the following commands: \"set serveroutput off\",
\"select sql_id, sql_text from v$sqltext where sql_id
=<sqlId>\", \"set serveroutput on\"",
      "component" : "storage",
      "infrastructureType" : "exadata",
      "eventName" : "HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE",
      "quarantineMode" : "\"FULL Quarantine\""
      ....
    }
  },
  "eventID" : "<unique_ID>",
  ....
}
```

In the tables below, you can read about the conditions and operations that trigger "critical" and "information" events. Each condition or operation is identified by a unique `eventName` value.

Critical events for Exadata Cloud Service infrastructure:

Critical Event - EventName	Description
HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE	<p>SQL statement terminated by Oracle Database Resource Manager due to excessive consumption of CPU and/or I/O. The execution plan associated with the terminated SQL stmt is quarantined. Please find the sql identifier in sqlId field of this JSON payload. This feature protects an Oracle database from performance degradation. Please review the SQL statement. You can see the statement using the following commands:</p> <ul style="list-style-type: none"> • <code>\set serveroutput off\</code> • <code>\select sql_id, sql_text from v\$sqltext where sql_id =<sqlId>\</code> • <code>\set serveroutput on\</code>

Informational events for Exadata Cloud Service infrastructure:

Information Event - EventName	Description
HEALTH.INFRASTRUCTURE.CELL.FLASH_DISK_FAILURE	Flash Disk Failure has been detected. This is being investigated by Oracle Exadata team and the disk will be replaced if needed. No action needed from the customer.

NOT_SUPPORTED

In the following example of a "critical" event, you can see within the `additionalDetails` section of the event message that this particular message concerns an SQL statement that was terminated by Oracle Database Resource Manager because it was consuming excessive CPU or I/O resources. The `eventName` and `description` fields within the `additionalDetails` section provide information regarding the critical situation:

```
{
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.critical",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-07-30T04:53:18Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>",
    "compartmentName" : "example_name",
    "resourceName" : "my_exadata_resource",
    "resourceId" : "ocid1.dbssystem.oc1.phx.<unique_ID>",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "sqlID" : "gnwfm1jgqcfuu",
      "systemId" : "ocid1.dbssystem.oc1.eu-frankfurt-1.<unique_ID>",
      "creationTime" : "2021-05-14T13:29:28+00:00",
      "dbUniqueID" : "1558836122",
      "quarantineType" : "SQLID",
      "dbUniqueName" : "AB0503_FRA1S6",

```

"description" : "SQL statement terminated by Oracle Database Resource Manager due to excessive consumption of CPU and/or I/O.

The execution plan associated with the terminated SQL stmt is quarantined. Please find the sql identifier in sqlId field of this JSON payload. This feature protects an Oracle database from performance degradation.

Please review the SQL statement. You can see the statement using the following commands: \ "set serveroutput off\" ,

\ "select sql_id, sql_text from v\$sqltext where sql_id =<sqlId>\", \ "set serveroutput on\" ,

```

    "quarantineReason" : "Manual",
    "asmClusterName" : "None",
    "component" : "storage",
    "infrastructureType" : "exadata",
    "name" : "143",
    "eventName" : "HEALTH.INFRASTRUCTURE.CELL.SQL_QUARANTINE",
    "comment" : "None",
    "quarantineMode" : "\"FULL Quarantine\"",
    "rpmVersion" : "OSS_20.1.8.0.0_LINUX.X64_210317",
    "cellsrvChecksum" : "14f73eb107dc1be0bde757267e931991",
    "quarantinePlan" : "SYSTEM"
  }
},
"eventID" : "<unique_ID>",
"extensions" : {
  "compartmentId" : "ocid1.tenancy.oc1.<unique_ID>"
}
}

```

NOT_SUPPORTED

In the following example of an "information" event, you can see within the additionalDetails section of the event message that this particular message concerns a flash disk failure that is being investigated by the Oracle Exadata operations team. The eventName and description fields within the additionalDetails section provide information regarding the event:

```

{
  "eventType" :
"com.oraclecloud.databaseservice.exadatainfrastructure.information",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "Exadata Storage",
  "eventTime" : "2021-12-17T19:14:42Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" :
"ocid1.tenancy.oc1..aaaaaaaao3lj36x6lwxyvc4wausjouca7pwyjfw5ebsq5emrpqlq12gj5iq",
    "compartmentName" : "intexadatateam",
    "resourceId" :
"ocid1.dbsystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcctw45ykkaxyzc5isxoula",
    "availabilityDomain" : "phx-ad-2",
    "additionalDetails" : {
      "serviceType" : "exacs",
      "component" : "storage",

```

```

    "systemId" :
"ocid1.dbssystem.oc1.phx.abyhqljt5y3taezn7ug445fzwlngjfszbedxlcctw45ykkaxyzc5i
sxoula",
    "infrastructureType" : "exadata",
    "description" : "Flash Disk Failure has been detected. This is being
investigated by Oracle Exadata team and the disk will be
replaced if needed. No action needed from the
customer.",
    "eventName" : "HEALTH.INFRASTRUCTURE.CELL.FLASH_DISK_FAILURE",
    "FLASH_1_1" : "S2T7NA0HC01251 failed",
    "otto-ingestion-time" : "2021-12-17T19:14:43.205Z",
    "otto-send-EventService-time" : "2021-12-17T19:14:44.198Z"
  }
},
"eventID" : "30130ab4-42fa-4285-93a7-47e49522c698",
"extensions" : {
  "compartmentId" :
"ocid1.tenancy.oc1..aaaaaaaa03lj36x6lwxyc4wausjouca7pwyjfw5ebsq5emrpqlql2gj5
iq"
}
}

```

Exascale DB Storage Vaults Event Types

The events in this section are emitted by the `exascale-db-storage-vaults` resource.

Friendly Name	Event Type
Exascale Database Storage Vault - Create Begin	com.oraclecloud.databaseservice.createexascaledbstoragevault.begin
Exascale Database Storage Vault - Create End	com.oraclecloud.databaseservice.createexascaledbstoragevault.end
Exascale Database Storage Vault - Update Begin	com.oraclecloud.databaseservice.updateexascaledbstoragevault.begin
Exascale Database Storage Vault - Update End	com.oraclecloud.databaseservice.updateexascaledbstoragevault.end
Exascale Database Storage Vault - Delete Begin	com.oraclecloud.databaseservice.deleteexascaledbstoragevault.begin
Exascale Database Storage Vault - Delete End	com.oraclecloud.databaseservice.deleteexascaledbstoragevault.end
Exascale Database Storage Vault - Change Compartment Begin	com.oraclecloud.databaseservice.changeexascaledbstoragevaultcompartment.begin
Exascale Database Storage Vault - Change Compartment End	com.oraclecloud.databaseservice.changeexascaledbstoragevaultcompartment.end

Exadata Cloud Infrastructure VM Cluster Event Types

Review the list of events that can be emitted by VM Cluster

Friendly Name	Event Type
Cloud VM Cluster - Change Compartment Begin	com.oraclecloud.databaseservice.changecloudvmclustercompartment.begin

Friendly Name	Event Type
Cloud VM Cluster - Change Compartment End	com.oraclecloud.databaseservice.changecloudvmclustercompartment.end
Cloud VM Cluster - Create Begin	com.oraclecloud.databaseservice.createcloudvmcluster.begin
Cloud VM Cluster - Create End	com.oraclecloud.databaseservice.createcloudvmcluster.end
Cloud VM Cluster - Delete Begin	com.oraclecloud.databaseservice.deletecloudvmcluster.begin
Cloud VM Cluster - Delete End	com.oraclecloud.databaseservice.deletecloudvmcluster.end
Cloud VM Cluster - Update Begin	com.oraclecloud.databaseservice.updatecloudvmcluster.begin
Cloud VM Cluster - Update End	com.oraclecloud.databaseservice.updatecloudvmcluster.end
Cloud VM Cluster - Update IORM Configuration Begin	com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.begin
Cloud VM Cluster - Update IORM Configuration End	com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.end
Cloud VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin
Cloud VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end

NOT_SUPPORTED

This is a reference event for a cloud VM cluster resource:

```
{
  "cloudEventsVersion": "0.1",
  "eventID": "<unique_ID>",
  "eventType":
"com.oraclecloud.databaseservice.updatecloudvmclusteriormconfig.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "<unique_ID>",
    "eventName": "UpdateCloudVmClusterIormConfig",
    "compartmentName": "example_compartment",
    "resourceName": "my_container_database",
    "resourceId": "ocid1.cloudvmcluster.oc1.<unique_ID>",
    "resourceVersion": null,
    "additionalDetails": {
      "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.<unique_ID>",
      "freeFormTags": {},
      "definedTags": {},
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "lifecycleState": "AVAILABLE",
      "giVersion": "19.0.0.0.0",
      "cpuCoreCount": 16
    }
  }
}
```

```

    }
  },
  "timeCreated": "2022-06-15T16:31:31.979Z"
}

```

This is a reference event for Add Virtual Machine Begin:

```

{
  "id":
"ocidl.eventschema.oc1.phx.n2p4ijm0jyuia5p6lzhps0axtqft2d2ueywaq4oxcr3ywlzt9jd
689kvxazo",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Add Virtual Machine Begin",
  "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudExadataInfrastructureId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cpuCoreCount",
      "type": [
        "null",
        "Integer"
      ]
    }
  ]
}

```

```

        "name": "ocpuCountFractional",
        "type": [
            "null",
            "Float"
        ]
    },
    {
        "name": "dataStorageSizeInTBs",
        "type": [
            "null",
            "Integer"
        ]
    },
    {
        "name": "dataStorageSizeInGBs",
        "type": [
            "null",
            "Integer"
        ]
    },
    {
        "name": "licenseType",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "giVersion",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "dbNodeIds",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "timeZone",
        "type": [
            "null",
            "string"
        ]
    }
],
"exampleEvent": {
    "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.begin",
    "cloudEventsVersion": "0.1",
    "eventTypeVersion": "2.0",
    "source": "databaseservice",
    "eventID": "bc78609a-783a-9034-ccd1-12ab908df913",

```

```

"eventTime": "2023-01-06T23:18:04.000Z",
"contentType": "application/json",
"data": {
  "eventGroupingId": "csid201fe4f3443a853d76e9cec3ef4a/
3200918f142a44adb715d8aaf4f5ba99/DC62865A826A6E98699590E7F33C5064",
  "eventName": "CloudVmClusterAddVirtualMachine",
  "compartmentId": "ocidl.compartment.oc1.....unique_id",
  "compartmentName": null,
  "resourceName": "my_cloud_vm_cluster",
  "resourceId": "ocidl.cloudvmcluster.oc1.....unique_id",
  "resourceVersion": null,
  "availabilityDomain": "",
  "tagSlug": "tag_slug",
  "identity": {
    "principalName": null,
    "principalId": null,
    "authType": null,
    "callerName": null,
    "callerId": null,
    "tenantId": null,
    "ipAddress": null,
    "credentials": null,
    "authZPolicies": null,
    "userGroups": null,
    "userAgent": null,
    "consoleSessionId": null
  },
  "request": {
    "id": "01858321-0045-4bc5-b0d9-a917a6a40901",
    "path": null,
    "action": null,
    "parameters": null,
    "headers": null
  },
  "response": {
    "status": null,
    "responseTime": null,
    "headers": null,
    "payload": null,
    "message": null
  },
  "stateChange": {
    "previous": null,
    "current": {
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "dataStorageSizeGb": 60,
      "lifecycleState": "AVAILABLE",
      "sshPublicKeys": "...",
      "displayName": "my_cloud_vm_cluster",
      "cpuCoreCount": 16,
      "freeTags": {},
      "definedTags": {},
      "ocpuCountFractional": 16.0
    }
  },
  "additionalDetails": {

```

```

        "timeCreated": "2023-01-06T22:18:04.000Z",
        "timeUpdated": "2023-01-06T22:20:04.000Z",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": null,
        "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
        "cpuCoreCount": 16,
        "ocpuCountFractional": 16.0,
        "dataStorageSizeInTBs": 4,
        "dataStorageSizeInGBs": 60,
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "giVersion": "19.0.0.0.0",
        "dbNodeIds": "[ocid1.dbnode.oc1.....unique_id,...]",
        "timeZone": "UTC"
    },
    "internalDetails": {
        "attributes": null
    }
}
},
"timeCreated": "2023-01-06T23:18:04.000Z"
}

```

This is a reference event for Add Virtual Machine End:

```

{
  "id":
"ocid1.eventschema.oc1.phx.v87pkelz9k9u6xaqo51taf6bunf0gc2wyhrbmjzjbh3h1pjwakav
mf2borxgb",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Add Virtual Machine End",
  "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2023-01-06T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    }
  ]
}

```

```
]
},
{
  "name": "cloudExadataInfrastructureId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "cpuCoreCount",
  "type": [
    "null",
    "Integer"
  ]
},
{
  "name": "ocpuCountFractional",
  "type": [
    "null",
    "Float"
  ]
},
{
  "name": "dataStorageSizeInTBs",
  "type": [
    "null",
    "Integer"
  ]
},
{
  "name": "dataStorageSizeInGBs",
  "type": [
    "null",
    "Integer"
  ]
},
{
  "name": "licenseType",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "giVersion",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "dbNodeIds",
  "type": [
    "null",
    "string"
  ]
}
```

```

    ]
  },
  {
    "name": "timeZone",
    "type": [
      "null",
      "string"
    ]
  }
],
"exampleEvent": {
  "eventType":
"com.oraclecloud.databaseservice.cloudvmclusteraddvirtualmachine.end",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "databaseservice",
  "eventID": "ced78bb7-3903-acd8-ff78-5567aa01a912",
  "eventTime": "2023-01-06T23:18:04.000Z",
  "contentType": "application/json",
  "data": {
    "eventGroupingId": "csid89a04ef74ccb8b48340f56e656cf/
729c99d3e5a34d548ddc31c054810454/634F086E8618E0A660946A6862C82A68",
    "eventName": "CloudVmClusterAddVirtualMachine",
    "compartmentId": "ocid1.compartment.oc1.....unique_id",
    "compartmentName": null,
    "resourceName": "my_cloud_vm_cluster",
    "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
    "resourceVersion": null,
    "availabilityDomain": "",
    "tagSlug": "tag_slug",
    "identity": {
      "principalName": null,
      "principalId": null,
      "authType": null,
      "callerName": null,
      "callerId": null,
      "tenantId": null,
      "ipAddress": null,
      "credentials": null,
      "authZPolicies": null,
      "userGroups": null,
      "userAgent": null,
      "consoleSessionId": null
    },
    "request": {
      "id": "07197e12-b680-475e-851e-bb89fcd8376d",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,

```

```

    "message": null
  },
  "stateChange": {
    "previous": null,
    "current": {
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "dataStorageSizeGb": 60,
      "lifecycleState": "AVAILABLE",
      "sshPublicKeys": "...",
      "displayName": "my_cloud_vm_cluster",
      "cpuCoreCount": 16,
      "freeTags": {},
      "definedTags": {},
      "ocpuCountFractional": 16.0
    }
  },
  "additionalDetails": {
    "timeCreated": "2023-01-06T22:18:04.000Z",
    "timeUpdated": "2023-01-06T22:20:04.000Z",
    "lifecycleState": "AVAILABLE",
    "lifecycleDetails": null,
    "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
    "cpuCoreCount": 16,
    "ocpuCountFractional": 16.0,
    "dataStorageSizeInTBs": 4,
    "dataStorageSizeInGBs": 60,
    "licenseType": "BRING_YOUR_OWN_LICENSE",
    "giVersion": "19.0.0.0.0",
    "dbNodeIds": "[ocid1.dbnode.oc1.....unique_id,...]",
    "timeZone": "UTC"
  },
  "internalDetails": {
    "attributes": null
  }
}
},
"timeCreated": "2023-01-06T23:18:04.000Z"
}

```

This is a reference event for Cloud VM Cluster - Update Begin:

```

{
  "id":
"ocid1.eventschema.oc1.phx.ekmz1phzp4b1lk7m7tbygulbnakmjnr99eqjops3zvpt337pn
nfmj6r79j",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Update Begin",
  "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {

```

```
    "name": "id",
    "type": "string"
  },
  {
    "name": "defineTags",
    "type": [
      "null",
      "Map<String, Map<String, Object>>"
    ]
  },
  {
    "name": "freeFormTags",
    "type": [
      "null",
      "Map<String, String>"
    ]
  },
  {
    "name": "timeCreated",
    "type": "string"
  },
  {
    "name": "timeUpdated",
    "type": "string"
  },
  {
    "name": "lifecycleState",
    "type": "string"
  },
  {
    "name": "lifecycleDetails",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "cloudExadataInfrastructureId",
    "type": "string"
  },
  {
    "name": "cpuCoreCount",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "dataStorageSizeInGBs",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "licenseType",
```

```

        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "giVersion",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "dbNodeIds",
        "type": [
            "null",
            "string"
        ]
    },
    {
        "name": "timeZone",
        "type": [
            "null",
            "string"
        ]
    }
],
"exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
    "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.begin",
    "source": "databaseservice",
    "eventTypeVersion": "2.0",
    "eventTime": "2019-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "data": {
        "eventGroupId": "4976b940-2c2d-4380-a669-1d70d071b187",
        "eventName": "UpdateCloudVmCluster",
        "compartmentName": "example_compartment",
        "resourceName": "my_container_database",
        "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",
        "resourceVersion": null,
        "additionalDetails": {
            "cloudExadataInfrastructureId":
"ocid1.cloudexadatainfrastructure.oc1.....unique_id",
            "freeFormTags": {},
            "definedTags": {},
            "licenseType": "BRING_YOUR_OWN_LICENSE",
            "lifecycleState": "AVAILABLE",
            "giVersion": "19.0.0.0.0",
            "cpuCoreCount": 16
        }
    }
},
"timeCreated": "2020-06-15T16:31:31.979Z"
}

```

This is a reference event for Cloud VM Cluster - Update End:

```
{
  "id":
  "ocid1.eventschema.oc1.phx.svwkildsx63clplq6phba7d6lns1r192yc3uyc2ea5utjprqcu
  hbgvht4we",
  "serviceName": "Database",
  "displayName": "Cloud VM Cluster - Update End",
  "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "defineTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudExadataInfrastructureId",
      "type": "string"
    }
  ],
}
```

```

    "name": "cpuCoreCount",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "dataStorageSizeInGBs",
    "type": [
      "null",
      "Integer"
    ]
  },
  {
    "name": "licenseType",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "giVersion",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "dbNodeIds",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "timeZone",
    "type": [
      "null",
      "string"
    ]
  }
],
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "b28fcda6-3d7b-4044-aa8e-7c21cde84b44",
  "eventType": "com.oraclecloud.databaseservice.updatecloudvmcluster.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "data": {
    "eventGroupId": "4976b940-2c2d-4380-a669-1d70d071b187",
    "eventName": "UpdateCloudVmCluster",
    "compartmentName": "example_compartment",
    "resourceName": "my_container_database",
    "resourceId": "ocid1.cloudvmcluster.oc1.....unique_id",

```

```

    "resourceVersion": null,
    "additionalDetails": {
      "cloudExadataInfrastructureId":
"ocidl.cloudexadatainfrastructure.ocl.....unique_id",
      "freeFormTags": {},
      "definedTags": {},
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "lifecycleState": "AVAILABLE",
      "giVersion": "19.0.0.0.0",
      "cpuCoreCount": 16
    }
  },
  "timeCreated": "2020-06-15T16:31:31.979Z"
}

```

VM Node Subsetting Event Types

Review the list of event types that VM Node Subsetting emits.

Table 6-1 VM Node Subsetting Events

Friendly Name	Event Type
VM Cluster - Add Virtual Machine Begin	com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.begin
VM Cluster - Add Virtual Machine End	com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.end
VM Cluster - Terminate Virtual Machine Begin	com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.begin
VM Cluster - Terminate Virtual Machine End	com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.end

Example 6-60 VM Node Subsetting Examples

This is a reference event for VM Cluster - Add Virtual Machine Begin:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.ocl..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.ocl..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",

```

```

    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..ocl...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbNodeStorageSizeInGBs": 60,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocidl.dbnode.1, ocidl.dbnode.2,...]",
      "dbServerIds": "[ocidl.dbserver.1, ocidl.dbserver.2,...]",
      "timeZone": "US/Pacific"
    }
  }
}

```

This is a reference event for VM Cluster - Add Virtual Machine End:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusteraddvirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocidl.compartment.ocl..unique_ID"
  },
  "data": {
"compartmentId": "ocidl.compartment.ocl..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..ocl...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-09-03T12:00:00.000Z",
      "timeUpdated": "2019-09-03T12:30:00.000Z",
      "displayName": "testDisplayName",
      "lifecycleDetails": "detail message",
    }
  }
}

```

```

        "exadataInfrastructureId": "ExatraInfra-unique_ID",
        "vmClusterNetworkId": "VmCluster-unique_ID",
        "cpuCoreCount": 2,
        "dataStorageSizeInTBs": 4,
        "memorySizeInGBs": 30,
        "dbNodeStorageSizeInGBs": 60,
        "dbVersion": "19.0.0.0",
        "licenseType": "BRING_YOUR_OWN_LICENSE",
        "giVersion": "19.0.0.0",
        "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
        "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
        "timeZone": "US/Pacific"
    }
}
}

```

This is a reference event for VM Cluster - Terminate Virtual Machine Begin:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
  },
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_database",
    "resourceId": "Vmcluster-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
"dbNodeStorageSizeInGBs": 60,
      "exadataInfrastructureId": "ExatraInfra-unique_ID",
      "vmClusterNetworkId": "VmCluster-unique_ID",
      "cpuCoreCount": 2,
      "dataStorageSizeInTBs": 4,
      "memorySizeInGBs": 30,
      "dbVersion": "19.0.0.0",
      "licenseType": "BRING_YOUR_OWN_LICENSE",
      "giVersion": "19.0.0.0",
      "dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
    }
  }
}

```

```

        "dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
        "timeZone": "US/Pacific"
    }
}
}

```

This is a reference event for VM Cluster - Terminate Virtual Machine End:

```

"exampleEvent": {
"cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.vmclusterterminatevirtualmachine.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
"compartmentId": "ocid1.compartment.oc1..unique_ID"
},
  "data": {
"compartmentId": "ocid1.compartment.oc1..unique_ID",
  "compartmentName": "example_name",
  "resourceName": "my_database",
  "resourceId": "Vmcluster-unique_ID",
  "availabilityDomain": "all",
  "freeFormTags": {},
  "definedTags": {},
  "additionalDetails": {
"dbNodeStorageSizeInGBs": 60,
"dbVersion": "19.0.0.0",
"licenseType": "BRING_YOUR_OWN_LICENSE",
"giVersion": "19.0.0.0",
"dbNodeIds": "[ocid1.dbnode.1, ocid1.dbnode.2,...]",
"dbServerIds": "[ocid1.dbserver.1, ocid1.dbserver.2,...]",
"timeZone": "US/Pacific"
}
}
}
}

```

Data Guard Event Types

Review the list of event types that Data Guard group and Data Guard Associations emit.

Note

To receive events related to Data Guard actions on multiple standby databases, subscribe to the Data Guard group resource events. If you have not switched to the new model, you can continue to subscribe to the Data Guard Associations resource events. However, after switching to the new model, you will need to explicitly subscribe to the new Data Guard Group resource events.

Data Guard Event Types (Data Guard Group resource)

Review the list of event types that Data Guard Group emit.

Friendly Name	Event Type
Change Protection Mode Begin	com.oraclecloud.databaseservice.changeprotectionmode.begin
Change Protection Mode End	com.oraclecloud.databaseservice.changeprotectionmode.end
Data Guard Create Standby Database - Create Begin	com.oraclecloud.databaseservice.createtandbydatabase.begin
Data Guard Create Standby Database - Create End	com.oraclecloud.databaseservice.createtandbydatabase.end
Data Guard Switchover - Begin	com.oraclecloud.databaseservice.dataguardswitchover.begin
Data Guard Switchover - End	com.oraclecloud.databaseservice.dataguardswitchover.end
Data Guard Failover - Begin	com.oraclecloud.databaseservice.dataguardfailover.begin
Data Guard Failover - End	com.oraclecloud.databaseservice.dataguardfailover.end
Data Guard Reinstate - Begin	com.oraclecloud.databaseservice.dataguardreinstate.begin
Data Guard Reinstate - End	com.oraclecloud.databaseservice.dataguardreinstate.end
Data Guard Update Config - Begin	com.oraclecloud.databaseservice.updatedataguardconfig.begin
Data Guard Update Config - End	com.oraclecloud.databaseservice.updatedataguardconfig.end
Refresh Data Guard Health Status - Begin	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin
Refresh Data Guard Health Status - End	com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.end

Example 6-61 Data Guard Health Status

This is a reference event for Refresh Data Guard Health Status - Begin

```
"exampleEvent": {
  "eventID": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
  "eventTime": "2024-12-04T10:06:21.000Z",
  "extensions": {
    "compartmentId": "ocid1.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "displayName": "Refresh Data Guard Health Status - Begin",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocid1.compartment.oc1..unique_id",
    "compartmentName": "example_name",
    "resourceName": "my_standby_database",
    "resourceId": "ocid1.database.oc1.phx.unique_id",
    "availabilityDomain": "AD1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocid1.database.oc1.phx.unique_id",
      "timeCreated": "2025-10-04T10:06:21.000Z",
      "timeUpdated": "2025-10-04T10:06:21.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleMessage": null,
      "dbSystemId": "ocid1.dbsystem.oc1.phx.unique_id",
      "databaseId": "ocid1.database.oc1.phx.unique_id",
      "dbHomeId": "ocid1.dbhome.oc1.phx.unique_id",
      "lastSyncedTime": "2025-10-04T10:06:21.000Z",
      "applyLag": "0 seconds",
      "syncState": "SYNCED",
      "switchoverReadiness": "HEALTHY",
      "switchoverReadinessMessage": null,
      "failoverReadiness": "HEALTHY",
      "failoverReadinessMessage": null,
      "dataLossExposure": "0 seconds",
      "memberTimeUpdated": "2025-10-04T10:06:21.000Z",
      "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-f45c89b1cb17",
      "currentDataGuardRole": "STANDBY",
      "noOfStandbyDatabases": 3,
      "primaryDatabaseId": "ocid1.database.oc1.iad.unique_id",
      "protectionMode": "MAXIMUM_AVAILABILITY",
      "transportType": "SYNC"
    }
  }
}
```

This is a reference event for Refresh Data Guard Health Status - End

```
"exampleEvent": {
  "eventID": "022a63a4-ff77-11e9-a0af-f45c89b1cb17",
  "eventTime": "2024-12-04T10:06:21.000Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.dataguardrefreshhealthstatus.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "displayName": "Refresh Data Guard Health Status - End",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_id",
    "compartmentName": "example_name",
    "resourceName": "my_standby_database",
    "resourceId": "ocidl.database.oc1.phx.unique_id",
    "availabilityDomain": "AD1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.database.oc1.phx.unique_id",
      "timeCreated": "2025-10-04T10:06:21.000Z",
      "timeUpdated": "2025-10-05T10:10:21.000Z",
      "lifecycleState": "AVAILABLE",
      "lifecycleMessage": null,
      "dbSystemId": "ocidl.dbsystem.oc1.phx.unique_id",
      "databaseId": "ocidl.database.oc1.phx.unique_id",
      "dbHomeId": "ocidl.dbhome.oc1.phx.unique_id",
      "lastSyncedTime": "2025-10-04T10:06:21.000Z",
      "applyLag": "0 seconds",
      "syncState": "SYNCED",
      "switchoverReadiness": "HEALTHY",
      "switchoverReadinessMessage": null,
      "failoverReadiness": "ALERT",
      "failoverReadinessMessage": "The Oracle Data Guard broker indicates
that the member 'DGHS1DB_s9b_sea' is not ready for failover.",
      "dataLossExposure": "12 seconds",
      "memberTimeUpdated": "2025-10-05T10:10:21.000Z",
      "lastUpdatedIdentifier": "022a6912-ff77-11e9-9e77-f45c89b1cb17",
      "currentDataGuardRole": "STANDBY",
      "noOfStandbyDatabases": 3,
      "primaryDatabaseId": "ocidl.database.oc1.iad.unique_id",
      "protectionMode": "MAXIMUM_AVAILABILITY",
      "transportType": "SYNC"
    }
  }
}
```

Data Guard Event Types (Data Guard Associations resource)

Review the list of event types that Data Guard Associations emit.

Table 6-2 Data Guard Associations Events

Friendly Name	Event Type
Change Protection Mode Begin	com.oraclecloud.databaseservice.change protectionmode.begin
Change Protection Mode End	com.oraclecloud.databaseservice.change protectionmode.end
Data Guard Association - Create Begin	com.oraclecloud.databaseservice.created ataguardassociation.begin
Data Guard Association - Create End	com.oraclecloud.databaseservice.created ataguardassociation.end
Data Guard Association - Failover Begin	com.oraclecloud.databaseservice.failover dataguardassociation.begin
Data Guard Association - Failover End	com.oraclecloud.databaseservice.failover dataguardassociation.end
Data Guard Association - Reinstatement Begin	com.oraclecloud.databaseservice.reinstatement dataguardassociation.begin
Data Guard Association - Reinstatement End	com.oraclecloud.databaseservice.reinstatement dataguardassociation.end
Data Guard Association - Switchover Begin	com.oraclecloud.databaseservice.switchover dataguardassociation.begin
Data Guard Association - Switchover End	com.oraclecloud.databaseservice.switchover dataguardassociation.end

Oracle Database Home Event Types

Review the list of events emitted by Oracle Database Homes.

Friendly Name	Event Type
DB Home - Create Begin	com.oraclecloud.databaseservice.created bhome.begin
DB Home - Create End	com.oraclecloud.databaseservice.created bhome.end
DB Home - Patch Begin	com.oraclecloud.databaseservice.patchdb home.begin
DB Home - Patch End	com.oraclecloud.databaseservice.patchdb home.end
DB Home - Terminate Begin	com.oraclecloud.databaseservice.deleted bhome.begin
DB Home - Terminate End	com.oraclecloud.databaseservice.deleted bhome.end
DB Home - Update Begin	com.oraclecloud.databaseservice.updated bhome.begin
DB Home - Update End	com.oraclecloud.databaseservice.updated bhome.end

NOT_SUPPORTED

This is a reference event for Database Homes:

```
{
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.createdbhome.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2019-08-29T21:16:04Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>",
    "compartmentName": "example_compartment",
    "resourceName": "my_dbhome",
    "resourceId": "DbHome-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id.oc1.<unique_ID>",
      "lifecycleState": "PROVISIONING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbSystemId": "DbSystem-unique_ID",
      "dbVersion": "19.0.0.0",
      "recordVersion": 4,
      "displayName": "example_display_name"
    }
  }
}
```

Database Event Types

These are the event types that Oracle Databases in Exadata Cloud Service instances emit.

Friendly Name	Event Type
Database - Automatic Backup Begin	com.oraclecloud.databaseservice.automat icbackupdatabase.begin
Database - Automatic Backup End	com.oraclecloud.databaseservice.automat icbackupdatabase.end
Database - Create Backup Begin	com.oraclecloud.databaseservice.backupd atabase.begin
Database - Create Backup End	com.oraclecloud.databaseservice.backupd atabase.end

Friendly Name	Event Type
Database - Critical (see Database Service Event Types for more information)	com.oraclecloud.databaseservice.databases.e.critical
Database - Information	com.oraclecloud.databaseservice.databases.e.information
Database - Delete Backup Begin	com.oraclecloud.databaseservice.deletebackup.begin
Database - Delete Backup End	com.oraclecloud.databaseservice.deletebackup.end
Database - Migrate to KMS Key Begin	com.oraclecloud.databaseservice.migrate.databasekmskey.begin
Database - Migrate to KMS Key End	com.oraclecloud.databaseservice.migrate.databasekmskey.end
Database - Move Begin	com.oraclecloud.databaseservice.movedatabase.begin
Database - Move End	com.oraclecloud.databaseservice.movedatabase.end
Database - Restore Begin	com.oraclecloud.databaseservice.restore.database.begin
Database - Restore End	com.oraclecloud.databaseservice.restore.database.end
Database - Rotate KMS Key Begin	com.oraclecloud.databaseservice.rotatedatabasekmskey.begin
Database - Rotate KMS Key End	com.oraclecloud.databaseservice.rotatedatabasekmskey.end
Database - Terminate Begin	com.oraclecloud.databaseservice.databases.e.terminate.begin
Database - Terminate End	com.oraclecloud.databaseservice.databases.e.terminate.end
Database - Update Begin	com.oraclecloud.databaseservice.updatedatabase.begin
Database - Update End	com.oraclecloud.databaseservice.updatedatabase.end
Database - Upgrade Begin	com.oraclecloud.databaseservice.upgrade.database.begin
Database - Upgrade End	com.oraclecloud.databaseservice.upgrade.database.end

NOT_SUPPORTED

This is a reference event for databases:

```
{
  "eventType" : "com.oraclecloud.databaseservice.backupdatabase.begin",
  "udEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "DatabaseService",
  "eventTime" : "2020-01-08T17:31:43.666Z",
  "contentType" : "application/json",
  "data" : {
    "compartmentId" : "ocidl.compartment.oc1.<unique_ID>",

```

```

"compartmentName": "example_compartment_name",
"resourceName": "my_backup",
"resourceId": "ocidl.dbbackup.oc1.<unique_ID>",
"availabilityDomain": "<availability_domain>",
"additionalDetails" : {
"timeCreated" : "2020-01-08T17:31:44Z",
"lifecycleState" : "CREATING",
"dbSystemId" : "ocidl.dbssystem.oc1.<unique_ID>",
"dbHomeId" : ocidl.dbhome.oc1.<unique_ID>",
"dbUniqueName" : DB1115_iad1dv",
"dbVersion" : "11.2.0.4.190716",
"databaseEdition" : "ENTERPRISE_EDITION_HIGH_PERFORMANCE",
"autoBackupsEnabled" : "false",
"backupType" : "FULL",
"databaseId" : "ocidl.database.oc1.<unique_ID>",
},
"definedTags" : {
  "My_example_tag_name" :
    { "Example_key" : "Example_value" }
},
"eventID": "<unique_ID>",
"extensions" : {
  "compartmentId": "ocidl.compartment.oc1.<unique_ID>"
}
}

```

Pluggable Database Event Types

These are the event types that Oracle pluggable databases in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Pluggable Database - Create Begin	com.oraclecloud.databaseservice.createpluggabledatabase.begin
Pluggable Database - Create End	com.oraclecloud.databaseservice.createpluggabledatabase.end
Pluggable Database - Delete Begin	com.oraclecloud.databaseservice.deletepluggabledatabase.begin
Pluggable Database - Delete End	com.oraclecloud.databaseservice.deletepluggabledatabase.end
Pluggable Database - Local Clone Begin	com.oraclecloud.databaseservice.localclonepluggabledatabase.begin
Pluggable Database - Local Clone End	com.oraclecloud.databaseservice.localclonepluggabledatabase.end
Pluggable Database - Remote Clone Begin	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin
Pluggable Database - Remote Clone End	com.oraclecloud.databaseservice.remoteclonepluggabledatabase.end
Start Pluggable Database - Begin	com.oraclecloud.databaseservice.startpluggabledatabase.begin
Start Pluggable Database - End	com.oraclecloud.databaseservice.startpluggabledatabase.end
Stop Pluggable Database - Begin	com.oraclecloud.databaseservice.stoppluggabledatabase.begin

Friendly Name	Event Type
Stop Pluggable Database - End	com.oraclecloud.databaseservice.stoppluggabledatabase.end
Pluggable Database - Convert to Regular Begin	com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.begin
Pluggable Database - Convert to Regular End	com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end
Pluggable Database - Inplace Restore Begin	com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin
Pluggable Database - Inplace Restore End	com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end
Pluggable Database - Refresh Begin	com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin
Pluggable Database - Refresh End	com.oraclecloud.databaseservice.pluggabledatabase.refresh.end
Pluggable Database - Relocate Begin	com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin
Pluggable Database - Relocate End	com.oraclecloud.databaseservice.pluggabledatabase.relocate.end

NOT_SUPPORTED

This is a reference event for pluggable databases (PDBs):

```
{
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>"
  },
  "eventType":
"com.oraclecloud.databaseservice.remoteclonepluggabledatabase.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.<unique_ID>",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.<unique_ID>",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.pluggabledatabases.oc1.phx.<unique_ID>",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocidl.database.oc1.<unique_ID>",
      "lifecycleState": "AVAILABLE",
      "lifecycleDetails": "Pluggable Database is available",
      "displayName": "Pluggable Database - Remote Clone Begin"
    }
  }
}
```

```

    }
  }
}

```

This is a reference event for Pluggable Database - Convert to Regular Begin:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
      "isRefreshableClone": true,
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocidl.database.oc1....unique_id",
      "lifecycleState": "UPDATING",
      "displayName": "Pluggable Database - Convert to Regular Begin"
    }
  }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Convert to Regular End:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.converttoregular.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",

```

```

"contentType": "application/json",
"definedTags": {},
"data": {
  "compartmentId": "ocidl.compartment.oc1.....unique_id",
  "compartmentName": "MyCompartment",
  "resourceName": "11092020_PKS_PDB1",
  "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
  "availabilityDomain": "XXIT:PHX-AD-1",
  "freeFormTags": {},
  "definedTags": {},
  "additionalDetails": {
    "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "isRefreshableClone": false,
    "timeCreated": "2021-03-13T21:15:59.000Z",
    "timeUpdated": "2021-03-13T21:15:59.000Z",
    "databaseId": "ocidl.database.oc1.....unique_id",
    "lifecycleState": "AVAILABLE",
    "displayName": "Pluggable Database - Convert to Regular End"
  }
}
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Inplace Restore Begin:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.begin",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocidl.database.oc1.....unique_id",
      "lifecycleState": "RESTORE_IN_PROGRESS",
      "isRefreshableClone": false,

```

```

        "displayName": "Pluggable Database - Inplace Restore Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Inplace Restore End:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.inplacerestore.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocidl.database.oc1....unique_id",
      "lifecycleState": "AVAILABLE",
      "isRefreshableClone": false,
      "lifecycleDetails": "Pluggable Database is available",
      "displayName": "Pluggable Database - Inplace Restore End"
    }
  }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Refresh Begin:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },

```

```

    "eventType":
    "com.oraclecloud.databaseservice.pluggabledatabase.refresh.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
      "compartmentId": "ocidl.compartment.oc1.....unique_id",
      "compartmentName": "MyCompartment",
      "resourceName": "11092020_PKS_PDB1",
      "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
      "availabilityDomain": "XXIT:PHX-AD-1",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "isRefreshableClone": true,
        "databaseId": "ocidl.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh Begin"
      }
    }
  },
  "activationTime": "2021-03-23T15:00:00.000Z",
  "eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Refresh End:

```

"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
  "com.oraclecloud.databaseservice.pluggabledatabase.refresh.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": "MyCompartment",
    "resourceName": "11092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {

```

```

        "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "timeCreated": "2021-03-13T21:15:59.000Z",
        "timeUpdated": "2021-03-13T21:15:59.000Z",
        "databaseId": "ocid1.database.oc1.....unique_id",
        "lifecycleState": "AVAILABLE",
        "isRefreshableClone": true,
        "lifecycleDetails": "Pluggable Database is available",
        "displayName": "Pluggable Database - Refresh End"
    }
}
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Relocate Begin:

```

"exampleEvent": {
    "eventID": "unique_id",
    "eventTime": "2021-03-23T00:49:14.123Z",
    "extensions": {
        "compartmentId": "ocid1.compartment.oc1..unique_id"
    },
    "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.begin",
    "eventTypeVersion": "2.0",
    "cloudEventsVersion": "0.1",
    "source": "databaseservice",
    "contentType": "application/json",
    "definedTags": {},
    "data": {
        "compartmentId": "ocid1.compartment.oc1.....unique_id",
        "compartmentName": "MyCompartment",
        "resourceName": "11092020_PKS_PDB1",
        "resourceId": "ocid1.pluggabledatabases.oc1.phx.unique_id",
        "availabilityDomain": "XXIT:PHX-AD-1",
        "freeFormTags": {},
        "definedTags": {},
        "additionalDetails": {
            "id": "ocid1.pluggabledatabases.oc1.phx.unique_id",
            "timeCreated": "2021-03-13T21:15:59.000Z",
            "timeUpdated": "2021-03-13T21:15:59.000Z",
            "databaseId": "ocid1.database.oc1.....unique_id",
            "lifecycleState": "AVAILABLE",
            "isRefreshableClone": false,
            "lifecycleDetails": "Pluggable Database is available",
            "displayName": "Pluggable Database - Relocate Begin"
        }
    }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}

```

This is a reference event for Pluggable Database - Relocate End:

```
"exampleEvent": {
  "eventID": "unique_id",
  "eventTime": "2021-03-23T00:49:14.123Z",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_id"
  },
  "eventType":
"com.oraclecloud.databaseservice.pluggabledatabase.relocate.end",
  "eventTypeVersion": "2.0",
  "cloudEventsVersion": "0.1",
  "source": "databaseservice",
  "contentType": "application/json",
  "definedTags": {},
  "data": {
    "compartmentId": "ocidl.compartment.oc1.....unique_id",
    "compartmentName": "MyCompartment",
    "resourceName": "l1092020_PKS_PDB1",
    "resourceId": "ocidl.pluggabledatabases.oc1.phx.unique_id",
    "availabilityDomain": "XXIT:PHX-AD-1",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.pluggabledatabases.oc1.phx.unique_id",
      "timeCreated": "2021-03-13T21:15:59.000Z",
      "timeUpdated": "2021-03-13T21:15:59.000Z",
      "databaseId": "ocidl.database.oc1.....unique_id",
      "lifecycleState": "AVAILABLE",
      "lifecycleDetails": "Pluggable Database is available",
      "displayName": "Pluggable Database - Relocate End"
    }
  }
},
"activationTime": "2021-03-23T15:00:00.000Z",
"eventTypeVersion": "2.0"
}
```

Database Service Events

The Database Service emits events, which are structured messages that indicate changes in resources.

- [Overview of Database Service Events](#)
The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.
- [Database Service Event Types](#)
Review the list of event types that the Database Service emits.
- [Receive Notifications about Database Service Events](#)
Subscribe to the Database Service Events and get notified.
- [Temporarily Restrict Automatic Diagnostic Collections for Specific Events](#)
Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.

- [Remediation](#)
These topics cover some common issues you might run into and how to address them.

Overview of Database Service Events

The Database Service Events feature implementation enables you to be notified about health issues with your Oracle Databases, or with other components on the Guest VM.

It is possible that Oracle Database or Clusterware may not be healthy or various system components may be running out of space in the Guest VM. You are not notified of this situation, unless you opt-in.

Note

You are opting in with the understanding that the list of events can change in the future. You can opt-out of this feature at any time

Database Service Events feature implementation generates events for Guest VM operations and conditions, as well as Notifications for customers by leveraging the existing OCI Events service and Notification mechanisms in their tenancy. Customers can then create topics and subscribe to these topics through email, functions, or streams.

Note

Events flow on Exadata Cloud Infrastructure depends on the following components: Oracle Trace File Analyzer (TFA), sysLens, and Oracle Database Cloud Service (DBCS) agent. Ensure that these components are up and running.

Manage Oracle Trace File Analyzer

- To check the run status of Oracle Trace File Analyzer, run the `tfactl status` command as `root` or a non-root user:

```
# tfactl status
.-----
.
| Host      | Status of TFA | PID      | Port | Version  | Build ID
| Inventory Status|
+-----+-----+-----+-----+-----+-----+
| node1     | RUNNING      | 41312   | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE     |
| node2     | RUNNING      | 272300  | 5000 | 22.1.0.0.0 |
22100020220310214615 | COMPLETE     |
'-----+-----+-----+-----+-----+-----'
```

- To start the Oracle Trace File Analyzer daemon on the local node, run the `tfactl start` command as `root`:

```
# tfactl start
Starting TFA..
```

```

Waiting up to 100 seconds for TFA to be started..
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
. . . . .
Successfully started TFA Process..
. . . . .
TFA Started and listening for commands

```

- To stop the Oracle Trace File Analyzer daemon on the local node, run the `tfactl stop` command as root:

```

# tfactl stop
Stopping TFA from the Command Line
Nothing to do !
Please wait while TFA stops
Please wait while TFA stops
TFA-00002 Oracle Trace File Analyzer (TFA) is not running
TFA Stopped Successfully
Successfully stopped TFA..

```

Manage sysLens

- If sysLens is running, then once every 15 minutes data is collected in the local domU to discover the events to be reported. To check if sysLens is running, run the `systemctl status syslens` command as root in the domU:

```

# systemctl status syslens
? syslens.service
Loaded: loaded (/etc/systemd/system/syslens.service; enabled; vendor
preset: disabled)
Active: active (running) since Wed 2025-03-19 20:23:00 UTC; 44min ago
Process: 137603 ExecStopPost=/var/opt/oracle/syslens/bin/syslens --stop
(code=exited, status=0/SUCCESS)
Main PID: 137794 (python3)
Tasks: 7 (limit: 319999)
Memory: 194.7M
CGroup: /system.slice/syslens.service
??137794 /usr/bin/python3 /var/opt/oracle/syslens/bin/syslens_main.py --
level DRIFT=0 --daemon --service

Mar 19 20:53:12 scaqar07dv0201 su[336679]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:13 scaqar07dv0201 su[336728]: (to oracle) root on none
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:13 scaqar07dv0201 su[336728]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:14 scaqar07dv0201 su[336808]: (to oracle) root on none
Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)

```

```

Mar 19 20:53:14 scaqar07dv0201 su[336808]: pam_unix(su-l:session): session
closed for user oracle
Mar 19 20:53:15 scaqar07dv0201 su[336869]: (to oracle) root on none
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
opened for user oracle by (uid=0)
Mar 19 20:53:15 scaqar07dv0201 su[336869]: pam_unix(su-l:session): session
closed for user oracle
[root@scaqar07dv0201 opc]#

```

- If the sysLens is enabled, when there is a reboot of the domU, then sysLens starts automatically. To validate if sysLens is enabled to collect telemetry, run the `systemctl is-enabled syslens` command as root in the domU:

```

# systemctl is-enabled syslens
enabled

```

- To validate if sysLens is able to run as daemon:

```

# /var/opt/oracle/syslens/bin/applin_ctl config_file get enable --
file /etc/oracle/syslens/config/syslens.config
true

```

- To validate if sysLens is configured to notify events, run the `tfactl get customerDiagnosticsNotifications` command as root in the domU:

```

# tfactl get customerDiagnosticsNotifications
-----
|                               Hostname                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Configuration Parameter                                             | Value |
+-----+-----+-----+-----+-----+-----+-----+
| Send CEF notifications ( customerDiagnosticsNotifications ) | ON    |
+-----+-----+-----+-----+-----+-----+-----+

```

Manage Database Service Agent

View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

- To check the status of the Database Service Agent, run the `systemctl status` command:

```

# systemctl status dbcsagent.service
dbcsagent.service
Loaded: loaded (/usr/lib/systemd/system/dbcsagent.service; enabled; vendor
preset: disabled)
Active: active (running) since Fri 2022-04-01 13:40:19 UTC; 6min ago
Process: 9603 ExecStopPost=/bin/bash -c kill `ps -fu opc |grep "java.*dbcs-
agent.*jar" |awk '{print $2}' ` (code=exited, status=0/SUCCESS)
Main PID: 10055 (sudo)
CGroup: /system.slice/dbcsagent.service
        10055 sudo -u opc /bin/bash -c umask 077; /bin/java -
Doracle.security.jps.config=/opt/oracle/...

```

- To start the agent if it is not running, run the `systemctl start` command as the root user:

```

systemctl start dbcsagent.service

```

Related Topics

- [To Enable, Partially Enable, or Disable Diagnostics Collection](#)
You can enable, partially enable, or disable diagnostics collection for your Guest VMs after provisioning the VM cluster. Enabling diagnostics collection at the VM cluster level applies the configuration to all the resources such as DB home, Database, and so on under the VM cluster.
- [Overview of Events](#)
- [Notifications Overview](#)

Database Service Event Types

Review the list of event types that the Database Service emits.

Note

- Critical events are triggered due to several types of critical conditions and errors that cause disruption to the database and other critical components. For example, database hang errors, and availability errors for databases, database nodes, and database systems to let you know if a resource becomes unavailable.
- Information events are triggered when the database and other critical components work as expected. For example, a clean shutdown of CRS, CDB, client, or scan listener, or a startup of these components will create an event with the severity of INFORMATION.
- Threshold limits reduce the number of notifications customers will receive for similar incident events whilst at the same time ensuring they receive the incident events and are reminded in a timely fashion.

Table 6-3 Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Resource Utilization - Disk Usage	HEALTH.DB_GUEST_FILESYSTEM.FREE_SPACE This event is reported when VM guest file system free space falls below 10% free, as determined by the operating system <code>df(1)</code> command, for the following file systems: <ul style="list-style-type: none"> • / • /u01 • /u02 • /var (X8M and later only) • /tmp (X8M and later only) 	HEALTH.DB_GUEST_FILESYSTEM.FREE_SPACE	com.oraclecloud.databaseservice.dbnode.critical	Critical threshold: 90%

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CRS status Up/ Down	AVAILABILITY.DB _GUEST.CRS_INST ANCE.DOWN.	AVAILABILITY.DB_ GUEST.CRS_INST ANCE.DOWN	com.oraclecloud .databaseservic e.dbnode.critic al (if .DOWN and NOT "user_action")	N/A
	An event of type CRITICAL is created when the Cluster Ready Service (CRS) is detected to be down.			
	AVAILABILITY.DB _GUEST.CRS_INST ANCE.DOWN_CLEAR ED	N/A	com.oraclecloud .databaseservic e.dbnode.inform ation (if .DOWN_CLEAR ED)	
	An event of type INFORMATION is created once it is determined that the event for CRS down has cleared.			

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
SCAN Listener Up/ Down	AVAILABILITY.DB _CLUSTER.SCAN_L ISTENER.DOWN A DOWN event is created when a SCAN listener goes down. The event is of type INFORMATION when a SCAN listener is shutdown due to user action, such as with the Server Control Utility (srvctl) or Listener Control (lsnrctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a SCAN listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a SCAN listener is started. There are three SCAN listeners per cluster called LISTENER_SCAN[1,2,3].	AVAILABILITY.DB_ CLUSTER.SCAN_ LISTENER.DOWN	com.oraclecloud .databaseservic e.dbnode.critic al (if .DOWN and NOT "user_action")	N/A
	AVAILABILITY.DB _CLUSTER.SCAN_L ISTENER.DOWN_CL EARED An event of type INFORMATION is created once it is determined that the event for SCAN Listener down has cleared.	N/A	com.oraclecloud .databaseservic e.dbnode.inform ation (if .DOWN_CLEAR ED)	

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Net Listener Up/ Down	AVAILABILITY.DB _GUEST.CLIENT_L ISTENER.DOWN A DOWN event is created when a client listener goes down. The event is of type INFORMATION when a client listener is shutdown due to user action, such as with the Server Control Utility (srvctl) or Listener Control (lsnrctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a grid infrastructure software update. The event is of type CRITICAL when a client listener goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a client listener is started. There is one client listener per node, each called LISTENER.	AVAILABILITY.DB_ GUEST.CLIENT_LI STENER.DOWN	com.oraclecloud .databaseservic e.database.crit ical (if .DOWN and NOT "user_action")	N/A
	AVAILABILITY.DB _GUEST.CLIENT_L ISTENER.DOWN_CL EARED An event of type INFORMATION is created once it is determined that the event for Client Listener down has cleared.	N/A	com.oraclecloud .databaseservic e.database.info rmation (if .DOWN_CLEAR ED)	

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CDB Up/Down	AVAILABILITY.DB _GUEST.CDB_INST ANCE.DOWN	AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN	com.oraclecloud .databaseservice.database.critical (if .DOWN and NOT "user_action")	N/A
	A DOWN event is created when a database instance goes down. The event is of type INFORMATION when a database instance is shutdown due to user action, such as with the SQL*Plus (sqlplus) or Server Control Utility (srvctl) commands, or any Oracle Cloud maintenance action that uses those commands, such as performing a database home software update. The event is of type CRITICAL when a database instance goes down unexpectedly. A corresponding DOWN_CLEARED event is created when a database instance is started.			
	AVAILABILITY.DB _GUEST.CDB_INST ANCE.DOWN_CLEARED	N/A	com.oraclecloud .databaseservice.database.information (if .DOWN_CLEARED)	
	An event of type INFORMATION is created once it is determined that the event for the CDB down has cleared.			

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
CRS Eviction	AVAILABILITY.DB _GUEST.CRS_INST ANCE.EVICTION An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster. The CRS alert.log is parsed for the CRS-1632 error indicating that a node is being removed from the cluster.	AVAILABILITY.DB_ GUEST.CRS_INST ANCE.EVICTION	An event of type CRITICAL is created when the Cluster Ready Service (CRS) evicts a node from the cluster. The CRS alert.log is parsed for the CRS-1632 error indicating that a node is being removed from the cluster.	N/A
Critical DB Errors	HEALTH.DB_CLUST ER.CDB.CORRUPTI ON Database corruption has been detected on your primary or standby database. The database alert.log is parsed for any specific errors that are indicative of physical block corruptions, logical block corruptions, or logical block corruptions caused by lost writes.	HEALTH.DB_CLUS TER.CDB.CORRU PTION	com.oraclecloud .databaseservic e.database.crit ical	N/A
Other DB Errors	HEALTH.DB_CLUST ER.CDB.ARCHIVER _HANG An event of type CRITICAL is created if a CDB is either unable to archive the active online redo log or unable to archive the active online redo log fast enough to the log archive destinations.	HEALTH.DB_CLUS TER.CDB.ARCHIV ER_HANG	com.oraclecloud .databaseservic e.database.crit ical	N/A

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG	HEALTH.DB_CLUSTER.CDB.DATABASE_HANG		
	An event of type CRITICAL is created when a process/session hang is detected in the CDB.			
Backup Failures	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE	com.oraclecloud .databaseservice.database.critical	N/A
	An event of type CRITICAL is created if there is a CDB backup with a FAILED status reported in the v\$rman_status view.			
Disk Group Usage	HEALTH.DB_CLUSTER.DISK_GROUP_FREE_SPACE	HEALTH.DB_CLUSTER.DISK_GROUP_FREE_SPACE	com.oraclecloud .databaseservice.dbsystem.critical com.oraclecloud .databaseservice.dbsystem.information (if < 90%)	Critical threshold: 90%
	An event of type CRITICAL is created when an ASM disk group reaches space usage of 90% or higher. An event of type INFORMATION is created when the ASM disk group space usage drops below 90%.			
Memory Usage	CONFIGURATION.DB_GUEST.MEMORY.HUGE_PAGES_TOO_LARGE	CONFIGURATION.DB_GUEST.MEMORY.HUGE_PAGES_TOO_LARGE	com.oraclecloud .databaseservice.dbnode.critical	90%
	An event of type CRITICAL is created when the amount of memory in the VM configured for HugePages is 90% or more of the total VM memory.			

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
sshd Configuration	CONFIGURATION.D B_GUEST.SSHD.IN VALID	CONFIGURATION. DB_GUEST.SSHD. INVALID	com.oraclecloud .databaseservic e.dbnode.critic al	N/A
	An event of type CRITICAL is created if unexpected values are set in the <code>/etc/ssh/sshd_config</code> file.			
Disk Issues	HEALTH.DB_GUEST .FILESYSTEM.COR RUPTION	HEALTH.DB_GUE ST.FILESYSTEM.C ORRUPTION	com.oraclecloud .databaseservic e.dbnode.critic al	N/A
	A Write-then-Read operation with a dummy file has failed for a file system, typically indicating the operating system had detected an I/O error or inconsistency (i.e. corruption) with the file system and changed the file system mount mode from read-write to read-only. The following file systems are tested:			
	<ul style="list-style-type: none"> • / • /u01 • /u02 			

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
Oracle EXAchk Reported Issues	HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT Oracle EXAchk is Exadata database platform's holistic health check that includes software, infrastructure and database configuration checks. CRITICAL check alerts should be addressed in 24 hours to maintain the maximum stability and availability of your system. This database service event alerts every 24 hours whenever there are any CRITICAL checks that are flagged in the most recent Oracle EXAchk report. The event will point to the latest Oracle EXAchk zip report.	HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT	com.oraclecloud .databaseservice.dbnode.critical	N/A
DB I/O Latency	HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY	HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY	com.oraclecloud .databaseservice.dbnode.warnings	Threshold: > 16ms
DB IO Latency Distribution	HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM	HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM	com.oraclecloud .databaseservice.dbnode.warnings	Threshold: >32ms
Unusual CPU Waits	HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS	HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS	com.oraclecloud .databaseservice.dbnode.warnings	Threshold: >5 minutes
DB Waits	HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS	HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS	com.oraclecloud .databaseservice.dbnode.warnings	Threshold: >20% > 35% "Exadata User I/O Wait-event"
CPU Throttling	HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING	HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING	com.oraclecloud .databaseservice.dbnode.warnings	Threshold: >0.25 * CPU_COUNT

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
FAST RECOVERY AREA (FRA) usage	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1	com.oraclecloud.databaseservice.database.warnings	Threshold: VM cluster RECO disk group > 90%. Threshold: Database FAST RECOVERY AREA (FRA) > 90%.
	When the VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 90%, Oracle preserves archive logs locally to simplify the fully cloud-automated conversion of a Snapshot Standby to a Physical Standby. If the RECO disk group or FRA usage exceeds 95%, an additional alert is triggered, and archived logs that have already been backed up are purged. At this stage, converting to a Physical Standby will require manual intervention. Reverting to a Physical Standby before reaching this threshold is recommended to avoid manual intervention.			

Table 6-3 (Cont.) Database Service Events

Friendly Name	Event Name	Remediation	Event Type	Threshold
FAST RECOVERY AREA (FRA) space usage	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2	HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2	com.oraclecloud.databaseservice.database.critical	Threshold: VM cluster RECO disk group > 95%. Threshold: Database FAST RECOVERY AREA (FRA) > 95%.
	VM Cluster's RECO disk group or database's RECOVERY AREA space exceeds 95% used, archives that have been backed up will be purged to reduce space utilization. To convert the Snapshot Standby to Physical Standby will require following these manual steps described in <i>How to Roll Forward a Standby Database Using Recover Database From Service (Doc ID 2850185.1)</i> .			
DB Blocker Detector	HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK	HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK	com.oraclecloud.databaseservice.dbnode.warning	Threshold: >= 10 number of Session waiting Threshold: >= 600 sec wait time for the chain
Exadata Cache Efficiency	HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO	HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO	com.oraclecloud.databaseservice.dbnode.warning	Threshold: < 90%
Redo Log Wait	HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS	HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS	com.oraclecloud.databaseservice.dbnode.warning	Threshold: > 50% of Overall Background Processes Time

Example 6-62 Database Service DB Node Critical Events Examples

DB node critical reference events:

```
{
  "eventType" : "com.oraclecloud.databaseservice.dbnode.critical",
  "cloudEventsVersion" : "0.1",
  "eventTypeVersion" : "2.0",
  "source" : "SYSENS/host_Name/DomU",
  "eventTime" : "2022-03-04T18:19:42Z",
  "contentType" : "application/json",
```

```

"data" : {
  "compartmentId" : "compartment_ID",
  "compartmentName" : "compartment_Name",
  "resourceName" : "resource_Name",
  "resourceId" : "resource_ID",
  "additionalDetails" : {
    "serviceType" : "EXACS",
    "hostName" : "host_Name",
    "description" : "The '/' filesystem is over 90% used.",
    "eventName" : "HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE",
    "status" : "online"
  }
},
"eventID" : "a9752630-9be7-11ec-a203-00163eb980bb",
"extensions" : {
  "compartmentId" : "compartment_ID"
}
}

```

Receive Notifications about Database Service Events

Subscribe to the Database Service Events and get notified.

To receive notifications, subscribe to Database Service Events and get notified using the Oracle Notification service, see *Notifications Overview*. For more information about Oracle Cloud Infrastructure Events, see *Overview of Events*.

Events Service - Event Types:

- Database - Critical
- DB Node - Critical
- DB Node - Error
- DB Node - Warning
- DB Node - Information

Related Topics

- [Overview of Events](#)
- [Notifications Overview](#)

Temporarily Restrict Automatic Diagnostic Collections for Specific Events

Use the `tfactl blackout` command to temporarily suppress automatic diagnostic collections.

If you set blackout for a target, then Oracle Trace File Analyzer stops automatic diagnostic collections if it finds events in the alert logs for that target while scanning. By default, blackout will be in effect for 24 hours.

You can also restrict automatic diagnostic collection at a granular level, for example, only for **ORA-00600** or even only **ORA-00600** with specific arguments.

Syntax

```

tfactl blackout add|remove|print
-targettype host|crs|asm|asmdg|database|dbbackup|db_dataguard|db_tablespace|

```


Table 6-4 (Cont.) tfactl blackout Command Parameters

Parameter	Description
events <i>all</i> " <i>str1, str2</i> "	Limits blackout only to the availability events, or event strings, which should not trigger auto collections, or be marked as blacked out in telemetry JSON. <i>all</i> : Blackout everything for the target specified. <i>string</i> : Blackout for incidents where any part of the line contains the strings specified. Specify a comma-delimited list of strings.
timeout <i>nh</i> <i>nd</i> none	Specify the duration for blackout in number of hours or days before timing out. By default, the timeout is set to 24 hours (24h).
c local	Specify if blackout should be set to cluster-wide or local. By default, blackout is set to local.
reason <i>comment</i>	Specify a descriptive reason for the blackout.
docollection	Use this option to do an automatic diagnostic collection even if a blackout is set for this target.

Example 6-63 tfactl blackout

- To blackout **event**: ORA-00600 on **target type**: database, **target**: mydb

```
tfactl blackout add -targettype database -target mydb -event "ORA-00600"
```

- To blackout **event**: ORA-04031 on **target type**: database, **target**: all

```
tfactl blackout add -targettype database -target all -event "ORA-04031" -  
timeout 1h
```

- To blackout **db backup events** on **target type**: dbbackup, **target**: mydb

```
tfactl blackout add -targettype dbbackup -target mydb
```

- To blackout **db dataguard events** on **target type**: db_dataguard, **target**: mydb

```
tfactl blackout add -targettype db_dataguard -target mydb -timeout 30m
```

- To blackout **db tablespace events** on **target type**: db_tablespace, **target**: system, **container**: mydb

```
tfactl blackout add -targettype db_tablespace -target system -container  
mydb -timeout 30m
```

- To blackout **ALL events** on **target type**: host, **target**: all

```
tfactl blackout add -targettype host -event all -target all -timeout 1h -  
reason "Disabling all events during patching"
```

- To print blackout details

```
tfactl blackout print
```

```

-----
-----
-----
|
myhostname
|
+-----+-----+-----+
+-----+-----+-----+
+-----+-----+-----+
| Target Type | Target | Events | Start
Time | End Time | Status | Do
Collection | Reason |
+-----+-----+-----+
+-----+-----+-----+
| HOST | ALL | ALL | Thu Mar 24 16:48:39
UTC 2022 | Thu Mar 24 17:48:39 UTC 2022 | ACTIVE | false |
Disabling all events during patching |
| DATABASE | MYDB | ORA-00600 | Thu Mar 24 16:39:03
UTC 2022 | Fri Mar 25 16:39:03 UTC 2022 | ACTIVE | false |
NA |
| DATABASE | ALL | ORA-04031 | Thu Mar 24 16:39:54
UTC 2022 | Thu Mar 24 17:39:54 UTC 2022 | ACTIVE | false |
NA |
| DB_DATAGUARD | MYDB | ALL | Thu Mar 24 16:41:38
UTC 2022 | Thu Mar 24 17:11:38 UTC 2022 | ACTIVE | false |
NA |
| DBBACKUP | MYDB | ALL | Thu Mar 24 16:40:47
UTC 2022 | Fri Mar 25 16:40:47 UTC 2022 | ACTIVE | false |
NA |
| DB_TABLESPACE | SYSTEM_CDBNAME_MYDB | ALL | Thu Mar 24 16:45:56
UTC 2022 | Thu Mar 24 17:15:56 UTC 2022 | ACTIVE | false |
NA |
'-----+-----+-----+
+-----+-----+-----+
+-----+-----+-----+'

```

- To remove blackout for **event:** ORA-00600 on **target type:** database, **target:** mydb

```
tfactl blackout remove -targettype database -event "ORA-00600" -target mydb
```

- To remove blackout for **db backup events** on **target type:** dbbackup, **target:** mydb

```
tfactl blackout remove -targettype dbbackup -target mydb
```

- To remove blackout for **db tablespace events** on **target type:** db_tablespace, **target:** system, **container:** mydb

```
tfactl blackout remove -targettype db_tablespace -target system -container mydb
```

- To remove blackout for **host events** on **target type:** host, **target:** all

```
tfactl blackout remove -targettype host -event all -target all
```

Remediation

These topics cover some common issues you might run into and how to address them.

- [HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE](#)
- [AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN](#)
- [AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN](#)
- [AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN](#)
- [AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN](#)
- [AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION](#)
- [HEALTH.DB_CLUSTER.CDB.CORRUPTION](#)
- [HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG](#)
- [HEALTH.DB_CLUSTER.CDB.DATABASE_HANG](#)
- [HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE](#)
- [HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE](#)
- [Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure](#)
- [CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE](#)
- [CONFIGURATION.DB_GUEST.SSHD.INVALID](#)
- [HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION](#)
- [HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT](#)
- [HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY](#)
- [HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM](#)
- [HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS](#)
- [HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS](#)
- [HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING](#)
- [HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1](#)
- [HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2](#)
- [HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK](#)
- [HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO](#)
- [HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS](#)

HEALTH.DB_GUEST.FILESYSTEM.FREE_SPACE

Problem Statement: One or more VM guest file systems has free space below 10% free.

Risk: Insufficient VM guest file system free space can cause disk space allocation failure, which can result in wide-ranging errors and failures in Oracle software (Database, Clusterware, Cloud, Exadata).

Action:

Oracle Cloud and Exadata utilities run automatically to purge old log files and trace files created by Oracle software to reclaim file system space.

If the automatic file system space reclamation utilities cannot sufficiently purge old files to clear this event, then perform the following actions:

1. Remove unneeded files and/or directories created manually or by customer-installed applications or utilities. Files created by customer-installed software are outside the scope of Oracle's automatic file system space reclamation utilities. The following operating system command, run as the `opc` user, is useful for identifying directories consuming excessive disk space:

```
$ sudo du -hx file-system-mount-point | sort -hr
```

Only remove files or directories you are certain can be safely removed.

2. Reclaim `/u02` file system disk space by removing Database Homes that have no databases. For more information about managing Database Homes, see *Manage Oracle Database Homes on Exadata Database Service on Exadata Cloud Infrastructure Instance*.
3. Open service request to receive additional guidance about reducing file system space use.

Related Topics

- [Managing Oracle Database Homes on an Exadata Cloud Infrastructure Instance](#)
You can delete or view information about Oracle Database Homes (referred to as "Database Homes" in Oracle Cloud Infrastructure) by using the Oracle Cloud Infrastructure Console, the API, or the CLI.

AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN

Problem Statement: The Cluster Ready Stack is in an offline state or has failed.

Risk: If the Cluster Ready Service is offline on a node, then the node cannot provide database services for the application.

Action:

1. Check if CRS was stopped by your administrator, as part of a planned maintenance event, or a scale up or down of local storage.
 - a. The following patching events will stop CRS:
 - i. GRID Patching
 - ii. Exadata VM patching of Guest
 - iii. Exadata VM Patching of Host

2. If CRS has stopped unexpectedly, then the current status can be checked by issuing the `crsctl check crs` command.
 - a. If the node is not responding, then the VM node may be rebooting. Wait for the node reboot to finish, CRS will normally be started through the `init` process.
3. If CRS is still down, then investigate the cause of the failure by referring to the `alert.log` found in `/u01/app/grid/diag/crs/<node_name>/crs/trace`. Review the log entries corresponding to the date/time of the down event. Act on any potential remediation.
4. Restart the CRS, by issuing the `crsctl start crs` command.
5. A successful restart of CRS will generate the clearing event:
`AVAILABILITY.DB_GUEST.CRS_INSTANCE.DOWN_CLEARED`.

AVAILABILITY.DB_CLUSTER.SCAN_LISTENER.DOWN

Problem Statement: A SCAN listener is down and unable to accept application connections.

Risk: If all SCAN listeners are down, then application connections to the database through the SCAN listener will fail.

Action:

Start the SCAN listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected SCAN listener will automatically failover to an available instance.
2. If the event was caused by user action, then start the SCAN listener at the next opportunity.

DOWN event of type CRITICAL

Check SCAN status and restart the SCAN listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

```
sudo su - grid
```

2. Check the SCAN listener status on any node:

```
srvctl status scan_listener
```

3. Start the SCAN listener:

```
srvctl start scan_listener
```

4. Recheck the SCAN listeners status on any node:

If the `scan_listener` is still down, then investigate the cause of the scan listener failure:

- a. Collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the `<hostName>` indicated in the log. Note the time in the event payload is always

provided in UTC. For `tfactl` collection, adjust the time to the timezone of the VM Cluster. As the `grid` user:

```
tfactl diagcollect -crs -os -node <hostName> -from "<eventTime
adjusted for local vm timezone> - 30 minute " -to "<eventTime adjusted
for local vm timezone> + 10 minutes"
```

- b. Review the SCAN listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CLIENT_LISTENER.DOWN

Problem Statement: A client listener is down and unable to accept application connections.

Risk:

- If the node's client listener is down, then the database instances on the node cannot provide services for the application.
- If the client listener is down on all nodes, then any application that connects to any database using the SCAN or VIP will fail.

Action:

Start the client listener to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Grid Infrastructure software update, then no action is required. The affected client listener will automatically restart when maintenance affecting the grid instance is complete.
2. If the event was caused by user action, then start the client listener at the next opportunity.

DOWN event of type CRITICAL

Check the client listener status and then restart the client listener.

1. Login to the VM as `opc` user and `sudo` to the `grid` user:

```
[opc@vm ~] sudo su - grid
```

2. Check the client listener status on any node:

```
[grid@vm ~] srvctl status listener
```

3. Start the client listener:

```
[grid@vm ~] srvctl start listener
```

4. Recheck the client listener status on any node:

If the client listener is still down, then investigate the cause of the client listener failure:

- a. Use `tfactl` to collect both the CRS and operating system logs 30 minutes prior and 10 minutes for the `<hostName>` indicated in the log. Note the time in the event payload is

always provided in UTC. For tfactl collection, adjust the time to the timezone of the VM Cluster.

```
[grid@vm ~] tfactl diagcollect -crs -os -node <hostName> -from
"<eventTime adjusted for local vm timezone> - 30 minute " -to
"<eventTime adjusted for local vm timezone> + 10 minutes"
```

- b. Review the listener log located under `/u01/app/grid/diag/tnslsnr/<hostName>/<listenerName>/trace`

AVAILABILITY.DB_GUEST.CDB_INSTANCE.DOWN

Problem Statement: A database instance has gone down.

Risk: A database instance has gone down, which may result in reduced performance if database instances are available on other nodes in the cluster, or complete downtime if database instances on all nodes are down.

Action:

Start the database instance to receive the `DOWN_CLEARED` event.

DOWN event of type INFORMATION

1. If the event was caused by an Oracle Cloud maintenance action, such as performing a Database Home software update, then no action is required. The affected database instance will automatically restart when maintenance affecting the instance is complete.
2. If the event was caused by user action, then start the affected database instance at the next opportunity.

DOWN event of type CRITICAL

1. Check database status and restart the down database instance.

- a. Login to the VM as `oracle` user:
- b. Set the environment:

```
[oracle@vm ~] . <dbName>.env
```

- c. Check the database status:

```
[oracle@vm ~] srvctl status database -db <dbName>
```

- d. Start the database instance:

```
[oracle@vm ~] srvctl start instance -db <dbName> -instance
<instanceName>
```

2. Investigate the cause of the database instance failure.

- a. Review Trace File Analyzer (TFA) events for the database:

```
[oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
```

- b. Review the database alert log located at `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`

AVAILABILITY.DB_GUEST.CRS_INSTANCE.EVICTION

Problem Statement: The Oracle Clusterware is designed to perform a node eviction by removing one or more nodes from the cluster if some critical problem is detected. A critical problem could be a node not responding via a network heartbeat, a node not responding via a disk heartbeat, a hung or severely degraded machine, or a hung ocssd.bin process. The purpose of this node eviction is to maintain the overall health of the cluster by removing impaired members.

Risks: During the time it takes to restart the evicted node, the node cannot provide database services for the application.

Action: CRS node eviction could be caused by OCSSD (CSS daemon), CSSDAGENT, or CSSDMONITOR processes. This requires determining which process was responsible for the node eviction and reviewing the relevant log files. Common causes of OCSSD eviction are network failures/latencies, IO issues with CSS voting disks, a member kill escalation. CSSDAGENT or CSSDMONITOR evictions could be OS scheduler problem or a hung thread within CSS daemon.

Log files to review include:

- clusterware alert log
- cssdagent log
- cssdmonitor log
- ocssd log
- lastgasp log
- /var/log/messages
- CHM/OS Watcher data
- `opatch lsinventory detail`

For more information on collecting together most of the files, see *Autonomous Health Framework (AHF) - Including TFA and ORAchk/EXAchk (Doc ID 2550798.1)*.

For more information on troubleshooting CRS node eviction, see *Troubleshooting Clusterware Node Evictions (Reboots) (Doc ID 1050693.1)*.

Related Topics

- [Autonomous Health Framework \(AHF\) - Including TFA and ORAchk/EXAchk \(Doc ID 2550798.1\)](#)
- [Troubleshooting Clusterware Node Evictions \(Reboots\) \(Doc ID 1050693.1\)](#)

HEALTH.DB_CLUSTER.CDB.CORRUPTION

Problem Statement: Corruptions can lead to application or database errors and in worse case result in significant data loss if not addressed promptly.

A corrupt block is a block that was changed so that it differs from what Oracle Database expects to find. Block corruptions can be categorized as physical or logical:

- In a physical block corruption, which is also called a media corruption, the database does not recognize the block at all; the checksum is invalid or the block contains all zeros. An

example of a more sophisticated block corruption is when the block header and footer do not match.

- In a logical block corruption, the contents of the block are physically sound and pass the physical block checks; however, the block can be logically inconsistent. Examples of logical block corruption include incorrect block type, incorrect data or redo block sequence number, corruption of a row piece or index entry, or data dictionary corruptions.

For more information, see *Physical and Logical Block Corruptions. All you wanted to know about it.* (Doc ID 840978.1).

Block corruptions can also be divided into interblock corruption and intrablock corruption:

- In an intrablock corruption, the corruption occurs in the block itself and can be either a physical or a logical block corruption.
- In an interblock corruption, the corruption occurs between blocks and can only be a logical block corruption.

Oracle checks for the following errors in the `alert.log`:

- ORA-01578
- ORA-00752
- ORA-00753
- ORA-00600 [3020]
- ORA-00600 [kdsgrp1]
- ORA-00600 [kclchkblk_3]
- ORA-00600 [13013]
- ORA-00600 [5463]

Risk: A data corruption outage occurs when a hardware, software, or network component causes corrupt data to be read or written. The service-level impact of a data corruption outage may vary, from a small portion of the application or database (down to a single database block) to a large portion of the application or database (making it essentially unusable). If remediation action is not taken promptly, then potential downtime and data loss can increase.

Action:

The current event notification currently triggers on physical block corruptions (ORA-01578), lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020), and logical corruptions (typical detected from ORA-00600 with first argument of kdsgrp1, kdsgrp1, kclchkblk_3, 13013 OR 5463).

Oracle recommends the following steps:

1. Confirm that these corruptions were reported in the alert.log trace file. Log a Service Request (SR) with latest EXAchk report, excerpt of the alert.log and trace file containing the corruption errors, any history of recent application, database or software changes and any system, clusterware and database logs for the same time period. For all these cases, a TFA collection should be available and should be attached to the SR.
2. For repair recommendations, refer to *Handling Oracle Database Corruption Issues* (Doc ID 1088018.1).

For physical corruptions or ORA-1578 errors, the following notes will be helpful:

- Doc ID 1578.1 : OERR: ORA-1578 "ORACLE data block corrupted (file # %s, block # %s)" Primary Note

- Doc ID 472231.1 : How to identify all the Corrupted Objects in the Database reported by RMAN
- Doc ID 819533.1 : How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY
- Depending on the object that has the corruption, follow the guidance in Doc ID 1088018.1. Note RMAN can be used to recover one or many data block that are physically corrupted. Also using Active Data Guard with real time apply, auto block repair of physical data corruptions would have occurred automatically.

For logical corruptions caused by lost writes (ORA-00752, ORA-00753 and ORA-00600 with first argument 3020) on the primary or standby databases, they will be detected on the primary or with standby's redo apply process. The following notes will be helpful:

- Follow the guidance, follow Doc ID 1088018.1.
- If you have a standby and lost write corruption on the primary or standby, refer to Resolving ORA-00752 or ORA-00600 [3020] During Standby Recovery (Doc ID 1265884.1)

For logical corruptions (typical detected from ORA-00600 with arguments of kdsgrp1, kclchblk_3, 13013 OR 5463):

- Follow the guidance, follow Doc ID 1088018.1 for specific guidance on the error that was detected.
- If you have a standby and logical corruption on the primary, refer to Resolving Logical Block Corruption Errors in a Physical Standby Database (Doc ID 2821699.1)

Related Topics

- [Physical and Logical Block Corruptions. All you wanted to know about it. \(Doc ID 840978.1\)](#)
- [OERR: ORA-1578 "ORACLE data block corrupted \(file # %s, block # %s\)" Primary Note \(Doc ID 1578.1\)](#)
- [How to identify all the Corrupted Objects in the Database with RMAN \(Doc ID 472231.1\)](#)
- [How to identify the corrupt Object reported by ORA-1578 / RMAN / DBVERIFY \(Doc ID 819533.1\)](#)
- [Resolving ORA-00752 or ORA-600 \[3020\] During Standby Recovery \(Doc ID 1265884.1\)](#)
- [Resolving Logical Block Corruption Errors in a Physical Standby Database \(Doc ID 2821699.1\)](#)

HEALTH.DB_CLUSTER.CDB.ARCHIVER_HANG

Problem Statement: CDB RAC Instance may temporarily or permanently stall due to the log writer's (LGWR) inability to write the log buffers to an online redo log. This occurs because all online logs need archiving. Once the archiver (ARC) can archive at least one online redo log, LGWR will be able to resume writing the log buffers to online redo logs and the application impact will be alleviated.

Risk: If the archiver hang is temporary, then this can result in a small application brown out or stall for application processes attempting to commit their database changes. If the archiver is not unblocked, applications can experience extended delay in processing.

Action:

- See, [Script To Find Redo log Switch History And Find Archivelog Size For Each instance In RAC \(Doc ID 2373477.1\)](#) to determine the hourly frequency for each thread/instance.

- If any hourly bucket is greater than 12, then consider resizing the online redo logs. See item 2 below for resizing steps.
- If the database hangs are temporary, then the archiver may be unable to keep up with the redo log generated. Check the `alert.log`, `$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/alert_<instanceName>.log`, for "All online logs need archiving", multiple events in a short period can indicate 2 possible solutions.
 - If the number of redo logs groups per thread is less than 4, then consider adding additional logs groups to reach 4, see item 1 below for add redo log steps.
 - The other possible solution is to resize the redo logs, see item 2 below for resizing steps.
- For Data Guard and Non Data Guard review the *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices for sizing guidelines.

1. Add a redo log group for each thread. The additional redo log should equal the current log size.

- a. Use the following query:

```
select max(group#) Ending_group_number, thread#, count(*)
number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
by thread#,bytes
```

- b. Add one new group per thread using the same size as the current redo logs.

```
alter database add logfile thread <thread_number> Group <max group + 1>
('<DATA_DISKGROUP>') size <redo_size_in_bytes>
```

2. Resize the online redo logs by adding larger redo logs and dropping the current smaller redo logs.

- a. Use the following query:

```
select max(group#) Ending_group_number, thread#, count(*)
number_of_groups_per_thread, bytes redo_size_in_bytes from v$log group
by thread#,bytes
```

- b. Add the same number of redo logs for each thread `<number_of_groups_per_thread>` that currently exist. The `<new_redo_size_in_bytes>` should be based on *Configure Online Redo Logs Appropriately* of section Oracle Database High Availability Overview and Best Practices.

- i.

```
alter database add logfile thread <thread_number> Group <max group
+ 1> ('<DATA_DISKGROUP>') size <new_redo_size_in_bytes>
```

- ii. The original smaller redo logs should be deleted. A redo log can only be deleted if its status is inactive. To determine the status of a redo logs issue:

```
select group#, thread#, status, bytes from v$log order by bytes,
group#, thread#;
```

To delete the original smaller redo logs:

```
alter database drop logfile <group#>
```

- If the database is hung, the primary log archive destination and alternate may be full. Review the `HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE` for details on freeing space in RECO and DATA disk groups.

Related Topics

- [Script To Find Redolog Switch History And Find Archivelog Size For Each Instances In RAC \(Doc ID 2373477.1\)](#)
- [Configure Online Redo Logs Appropriately](#)
- [HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE](#)

HEALTH.DB_CLUSTER.CDB.DATABASE_HANG

Problem Statement: Hang management detected a process hang and generated a ORA-32701 error message. Additionally, this event may be raised if Diagnostic Process (DIA0) process detects a hang in a critical database process.

Risk: A hang can indicate resource, operating system, or application coding related issues.

Action:

Investigate the cause of the session hang.

1. Review TFA events for the database for the following message patterns corresponding to the date/time of the event: ORA-32701, "DIA0 Critical Database Process Blocked" or "DIA0 Critical Database Process As Root".

```
[oracle@vm ~] tfactl events -database <dbName> -instance <instanceName>
```

2. Review the `alert.log` file.

```
$ORACLE_BASE/diag/rdbms/<dbName>/<instanceName>/trace/  
alert_<instanceName>.log
```

3. **For ora-32701:** An overloaded system can cause slow progress, which can be interpreted as a hang. The hang manager may attempt to resolve the hang by terminating the final blocker process.
4. **For DIA0 Critical Database Process messages:** Review the related diagnostic lines indicating the process and the reason for the hang.

HEALTH.DB_CLUSTER.CDB.BACKUP_FAILURE

Problem Statement: A daily incremental BACKUP of the CDB failed.

Risk: A failure of the backup can compromise the ability to use the backups for restore/recoverability of the database. Recoverability Point Object (RPO) and the Recoverability Time Object (RTO) can be impacted.

Action:

Review the RMAN logs corresponding to the date/time of the event. Note the event time stamp `<eventTime>` is in UTC, adjust as necessary for the VM's timezone.

- For Exadata Cloud Infrastructure Oracle Managed Backups or User Configured Backups under `dbaascli`:
 - RMAN output can be found at `/var/opt/oracle/log/<DB_NAME>/obkup`. Daily incremental logs have the following format `obkup_YYYY-mm-dd_24h:mm:ss.zzzzzzzzzzzz.log` within the `obkup` directory. The logs are located on the lowest active node/instance of the database when the backup was initiated.
 - Review the log for any failures:
 - * If the failure is due to an external event outside of RMAN, for example the backup location was full or a networking issue, resolve the external issue.
 - * For other RMAN script errors, collect the diagnostic logs and open a Service Request. See *DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check*.
 - If the issue is transient or is resolved, take a new incremental backup: See *dbaascli database backup*.
- For Customer owned and managed backup taken through RMAN:
 - Review the RMAN logs for the backup.

Related Topics

- [DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check](#)
- [dbaascli database backup](#)
To configure Oracle Database with a backup storage destination, take database backups, query backups, and delete a backup, use the `dbaascli database backup` command.

HEALTH.DB_CLUSTER.DISK_GROUP.FREE_SPACE

Problem Statement: ASM disk group space usage is at or exceeds 90%.

Risk: Insufficient ASM disk group space can cause database creation failure, tablespace and data file creation failure, automatic data file extension failure, or ASM rebalance failure.

Action:

ASM disk group used space is determined by the running the following query while connected to the ASM instance.

```
[opc@node ~] sudo su - grid
[grid@node ~] sqlplus / as sysasm
```

```
SQL> select 'ora.'||name||'.dg', total_mb, free_mb, round ((1-(free_mb/
total_mb))*100,2) pct_used from v$asm_diskgroup;
```

NAME	TOTAL_MB	FREE_MB	PCT_USED
ora.DATA1.dg	75497472	7408292	90.19
ora.RECO1.dg	18874368	17720208	6.11

ASM disk group capacity can be increased in the following ways:

1. Scale Exadata VM Cluster storage to add more ASM disk group capacity. See *Scaling an Exadata Cloud Infrastructure Instance*.
2. Scale Exadata Infrastructure storage to add more ASM disk group capacity. See *Scaling Exadata X8M and X9M Compute and Storage*.

DATA disk group space use can be reduced in the following ways:

1. Drop unused data files and temp files from databases. See *Dropping Data Files*.
2. Terminate unused databases (e.g. test databases). See *Using the Console to Terminate a Database*.

RECO disk group space use can be reduced in the following ways:

1. Drop unnecessary Guaranteed Restore Points. See *Using Normal and Guaranteed Restore Points*.
2. Delete archived redo logs or database backups already backed up outside the Flash Recovery Area (FRA). See *Maintaining the Fast Recovery Area*.

SPARSE disk group space use can be reduced in the following ways:

1. Move full copy test master databases to another disk group (e.g. DATA).
2. Drop unused snapshot databases or test master databases. See *Managing Exadata Snapshots*.

For more information about managing the log and diagnostic files, see *Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure*.

Related Topics

- [Scaling Exadata X8M, X9M, and X11M Compute and Storage](#)
The flexible X8M, X9M, and X11M system model is designed to be easily scaled in place, with no need to migrate the database using a backup or Data Guard.
- [Dropping Data Files](#)
- [To terminate a database](#)
- [Using Normal and Guaranteed Restore Points](#)
- [Maintaining the Fast Recovery Area](#)
- [Managing Exadata Snapshots](#)
- [Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure](#)

Managing the Log and Diagnostic Files on Oracle Exadata Database Service on Dedicated Infrastructure

The software components in Oracle Exadata Database Service on Dedicated Infrastructure generate a variety of log and diagnostic files, and not all these files are automatically archived and purged. Thus, managing the identification and removal of these files to avoid running out of file storage space is an important administrative task.

Database deployments on ExaDB-D include the `cleandblogs` script to simplify this administrative task. The script runs daily as a `cron` job on each compute node to archive key files and remove old log and diagnostic files.

The `cleandblogs` script operates by using the `adrci` (Automatic Diagnostic Repository [ADR] Command Interpreter) tool to identify and purge target diagnostic folders and files for each

Oracle Home listed in `/etc/oratab`. It also targets Oracle Net Listener logs, audit files, and core dumps.

On ExaDB-D, the script is run separately as the `oracle` user to clean log and diagnostic files that are associated with Oracle Database, and as the `grid` user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.

The `cleandblogs` script uses a configuration file to determine how long to retain each type of log or diagnostic file. You can edit the file to change the default retention periods. The file is located at `/var/opt/oracle/cleandb/cleandblogs.cfg` on each compute node.

Note

Configure an optimal retention period for each type of log or diagnostic file. An insufficient retention period will hinder root cause analysis and problem investigation.

Parameter	Description and Default Value
AlertRetention	Alert log (<code>alert_instance.log</code>) retention value in days. Default value: 14
ListenerRetention	Listener log (<code>listener.log</code>) retention value in days. Default value: 14
AuditRetentionDB	Database audit (<code>*.aud</code>) retention value in days. Default value: 1
CoreRetention	Core dump/files (<code>*.cmdp*</code>) retention value in days. Default value: 7
TraceRetention	Trace file (<code>*.tr*</code> and <code>*.prf</code>) retention value in days. Default value: 7
longpRetention	Data designated in the Automatic Diagnostic Repository (ADR) as having a long life (the <code>LONGP_POLICY</code> attribute). For information about ADR, see <i>Automatic Diagnostic Repository (ADR)</i> in the <i>Oracle Database Administrator's Guide</i> . Default value: 14
shortpRetention	Data designated in the Automatic Diagnostic Repository (ADR) as having a short life (the <code>SHORTP_POLICY</code> attribute). For information about ADR, see <i>Automatic Diagnostic Repository (ADR)</i> in the <i>Oracle Database Administrator's Guide</i> . Default value: 7
LogRetention	Log file retention in days for files under <code>/var/opt/oracle/log</code> and log files in ACFS under <code>/var/opt/oracle/dbaas_acfs/log</code> . Default value: 14
LogDirRetention	<code>cleandblogs</code> logfile retention in days. Default value: 14

Parameter	Description and Default Value
ScratchRetention	Temporary file retention in days for files under / scratch. Default value: 7

Archiving Alert Logs and Listener Logs

When cleaning up alert and listener logs, `cleandblogs` first archives and compresses the logs, operating as follows:

1. The current log file is copied to an archive file that ends with a date stamp.
2. The current log file is emptied.
3. The archive file is compressed using `gzip`.
4. Any existing compressed archive files older than the retention period are deleted.

Running the `cleandblogs` Script Manually

The `cleandblogs` script automatically runs daily on each compute node, but you can also run the script manually if the need arises.

1. Connect to the compute node as the `oracle` user to clean log and diagnostic files that are associated with Oracle Database, or connect as the `grid` user to clean log and diagnostic files that are associated with Oracle Grid Infrastructure.
For detailed instructions, see *Connecting to a Virtual Machine with SSH*.

Change to the directory containing the `cleandblogs` script:

```
$ cd /var/opt/oracle/cleandb
```

2. Run the `cleandblogs` script:

```
$ ./cleandblogs.pl
```

When running the script manually, you can specify an alternate configuration file to use instead of `cleandblogs.cfg` by using the `--pfile` option:

```
$ ./cleandblogs.pl --pfile config-file-name
```

3. Close your connection to the compute node:

```
$ exit
```

Related Topics

- [Automatic Diagnostic Repository \(ADR\)](#)
- [Connecting to a Virtual Machine with SSH](#)
You can connect to the virtual machines in an Exadata Cloud Infrastructure system by using a Secure Shell (SSH) connection.

CONFIGURATION.DB_GUEST.MEMORY.HUGEPAGES_TOO_LARGE

Problem Statement: Too much VM memory is allocated for HugePages use.

Risk: Excessive memory allocated to HugePages may result in poor database performance, or the system running out of memory, experiencing excessive swapping, or having crucial system services fail, causing system crash or node eviction.

Action:

1. Reduce HugePages memory use. To determine the proper setting for operating system parameter `vm.nr_hugepages`, see My Oracle Support document 361323.1.
2. Scale up VM memory. For more information about scaling VM memory, see *Introduction to Scale Up or Scale Down Operations*.

Related Topics

- <https://support.oracle.com/rs?type=doc&id=361323.1>
- [Introduction to Scale Up or Scale Down Operations](#)
With the Multiple VMs per Exadata system (MultiVM) feature release, you can scale up or scale down your VM cluster resources.

CONFIGURATION.DB_GUEST.SSHD.INVALID

Problem Statement: SSHD configuration is unexpected.

SSHD Configuration Setting	Expected Value
<code>PubkeyAuthentication</code>	yes
<code>AuthorizedKeysFile</code>	<code>.ssh/authorized_keys</code> This file must exist in root user home directory.
<code>HostbasedAuthentication</code>	no
<code>IgnoreUserKnownHosts</code>	yes
<code>IgnoreRhosts</code>	yes
<code>PermitEmptyPasswords</code>	no
<code>PasswordAuthentication</code>	no
<code>ChallengeResponseAuthentication</code>	no
<code>GSSAPIAuthentication</code>	no
<code>UsePAM</code>	yes
<code>PrintMotd</code>	no
<code>UsePrivilegeSeparation</code>	yes
<code>PermitUserEnvironment</code>	no
<code>Compression</code>	delayed
<code>MaxStartups</code>	100

SSHD Configuration Setting	Expected Value
AcceptEnv	Must contain one of the following: <ul style="list-style-type: none"> • LANG • LC_CTYPE • LC_NUMERIC • LC_TIME • LC_COLLATE • LC_MONETARY • LC_MESSAGES • LC_PAPER • LC_NAME • LC_ADDRESS • LC_TELEPHONE • LC_MEASUREMENT • LC_IDENTIFICATION • LC_ALL • LANGUAGE • XMODIFIERS
Subsystem	sftp /usr/libexec/openssh/sftp-server
Protocol	2
AddressFamily	inet

Risk: SSHD configuration is unexpected which may cause Oracle Cloud automation failure or prevent customer SSH access to the VM.

Action: Change SSHD to match expected configuration.

1. Verify SSHD service is active.

```
$ sudo systemctl is-active sshd.service
active
```

If SSHD service is inactive, then start it.

```
$ sudo systemctl start sshd.service
```

2. Verify SSHD service is enabled.

```
$ sudo /opt/oracle.cellos/host_access_control ssh-service --status
[INFO] [IMG-SEC-1201] Service sshd is enabled {1}
```

If SSHD service is disabled, then enable it.

```
$ sudo /opt/oracle.cellos/host_access_control ssh-service --enable
```

3. Change SSHD configuration to match the expected values according to the table shown in the Problem Statement section above.

SSHD Configuration Setting	How to Change Current setting
Ciphers	/opt/oracle.cellos/host_access_control sshciphers --help
MACs	/opt/oracle.cellos/host_access_control ssh-macs --help
PermitRootLogin	/opt/oracle.cellos/host_access_control rootssh --help
ClientAliveInterval	/opt/oracle.cellos/host_access_control idle-timeout --help
ClientAliveCountMax	/opt/oracle.cellos/host_access_control idle-timeout --help
ListenAddress	/opt/oracle.cellos/host_access_control ssh-listen --help
ALL OTHER PARAMETERS	<ul style="list-style-type: none"> Edit /etc/ssh/sshd_config. Restart sshd.service.\$ sudo systemctl restart sshd.service

HEALTH.DB_GUEST.FILESYSTEM.CORRUPTION

Problem Statement: A file system that is expected to be read-write can no longer be written to.

Risk: Oracle software (Linux, Database, Clusterware, Cloud, Exadata) requires write access to file systems to operate correctly.

Action:

/u01 and /u02 file systems:

1. Stop running services, if any, that are using the file system, such as Oracle Clusterware, Trace File Analyzer (TFA), and Enterprise Manager (EM) agent.
2. Unmount the file system.
3. Run file system check and repair.
 - **ext4:** Refer to *Checking and Repairing a File System*.
 - **xfs:** Refer to *Checking and Repairing an XFS File System*.
 - If the file system cannot be repaired then open a service request with Oracle Support for assistance.
4. Mount the file system.
5. Start the services.

/ (root) file system:

Open a service request with Oracle Support for assistance.

- If there is VM access, then collect full `dmesg(1)` command output and provide it to Oracle Support.
- Note that / (root) file system repair is possible only with console access.

Related Topics

- [Checking and Repairing a File System](#)

- [Checking and Repairing an XFS File System](#)

HEALTH.DB_CLUSTER.EXACHK.CRITICAL_ALERT

Problem Statement: A CRITICAL Exachk check failed and should be reviewed and addressed as soon as possible.

Risk: A CRITICAL check is expected to impact a large number of customers AND should be addressed immediately (for example, within 24 hours) AND meets one or more of the following criteria:

1. On-disk corruption or data loss
2. Intermittent wrong results with Exadata feature usage (e.g. smart scan)
3. System wide availability impact
4. Severe system wide performance impact seriously affecting application service Service Level Agreements (SLAs)
5. Compromised redundancy and inability to restore redundancy
6. Inability to update software in a rolling manner
7. Configuration error that could lead to an unexpected or unknown impact

Action:

Recommend that you bring up the EXAchk HTML report from the latest EXAchk zip file and click "**view**" on each CRITICAL check and follow the recommendation guidance that contains: Benefit/Impact, Risk, and Action/Repair guidance. Once the CRITICAL check is addressed, the next EXAchk run will pass that check. For more information about Oracle EXAchk, see *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

As the `root` user, you can re-run EXAchk command by issuing:

```
/usr/bin/exachk -profile exatier1 -noupgrade -dball
```

If the check results are returning false data, then log a Service Request.

If there is a CRITICAL check that needs to be temporarily excluded, then follow the "**Skipping Specific Best Practice Checks in Exadata Cloud**" section of *Oracle Exadata Database Machine Exachk (Doc ID 1070954.1)*.

Related Topics

- [Oracle Exadata Database Machine Exachk \(Doc ID 1070954.1\)](#)

HEALTH.DB_GUEST.PDB.HIGH_IO_LATENCY

Problem Statement: Higher I/O latency can lead to significant IO issues.

Risk: Increased I/O latency directly translates to slower database operations and SQL query response times. High latency can cause significant database performance degradation, especially for OLTP workloads.

Action: To investigate high cell single block physical read waits on Exadata, follow My Oracle Support (MOS) notes: [2119510.1](#) and [2530864.1](#).

HEALTH.DB_GUEST.CDB.HIGH_IO_LATENCY_HISTOGRAM

Problem Statement: High wait count values pertaining to IO wait class exceeding 32 ms, can indicate a potential problem with underlying Exadata Infrastructure.

Risk: Increased I/O latency directly translates to slower database operations and SQL query response times. High latency can cause significant.

Action: To investigate high cell single block physical read waits on Exadata, follow My Oracle Support (MOS) notes: [2119510.1](#) and [2530864.1](#).

HEALTH.DB_GUEST.PDB.HIGH_CPU_WAITS

Problem Statement: CPU intensive CDB/PDB sessions waiting on latch/cursor related wait-events for more than 300 seconds and CPU intensive CDB/PDB sessions waiting on Library cache related wait-events for more than 900 seconds.

Risk: These CPU intensive Wait events generally impact the overall system performance. Since mutexes are CPU intensive resource, in the event of mutex contention, CPU usage can rise and can impact user sessions.

Action: To investigate high CPU intensive waits in the Database, follow My Oracle Support (MOS) notes: [1349387.1](#), [1357946.1](#), [1377998.1](#) and [444560.1](#), which provide specific troubleshooting guidance for various scenarios.

HEALTH.DB_GUEST.PDB.EXCESSIVE_WAITS

Problem Statement: The Foreground Wait Event is consuming over 20% of total DB time, excluding the 'IDLE' wait-event class and Exadata User I/O wait-events. The Exadata User I/O wait-event alone is accounting for more than 35% of total DB time.

Risk: High wait event times indicate that the database is spending a significant amount of time waiting for resources, processes, or I/O operations, leading to performance bottlenecks and potentially impacting application responsiveness.

Action: To investigate high waits in the Database, follow My Oracle Support (MOS) notes: [1377446.1](#).

HEALTH.DB_GUEST.PDB.EXCESSIVE_CPU_THROTTLING

Problem Statement: Excessive CPU throttling observed in the Database.

Risk: Excessive CPU throttling in Database can occur when the resource manager limits do not match the consumer group utilization, potentially leading to performance issues. If the throttling occurs frequently, it can lead to performance problems, such as slow queries or application response times.

Action: To troubleshoot excessive CPU throttling in the Database, follow My Oracle Support (MOS) notes: [1339769.1](#).

HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_1

Problem Statement: VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 90%.

Action: Convert snapshot back to physical standby using `dbaascli`:

```
dbaascli dataguard convertStandby --standbyType physical --dbname <dbname>
```

HEALTH.DB_CLUSTER.CDB.SNAPSHOT_STBY_FRA_SPACE_ALERT_2

Problem Statement: VM cluster's RECO disk group or the database's FAST RECOVERY AREA (FRA) usage exceeds 95%.

Action: The standby database must be rolled forward using the SQL*PLUS `RECOVER DATABASE FROM SERVICE` statement.

Convert back to physical database

```
dbaascli dataguard convertStandby --dbname <db name> --standbyType physical
```

Stop managed recovery and redo transport

Stop transport and apply of redologs on the standby database.

```
DGMGRL> edit database <primary> set state=TRANSPORT-OFF;
DGMGRL> edit database <standby> set state=APPLY-OFF;
```

Save the existing logfile names

The roll forward process will create new online redo logs and standby redo logs, leaving the existing logs as orphans that consume space. Use the following query to save the names of these logfiles so they can be removed later in the process.

```
SQL> set heading off linesize 999 pagesize 0 feedback off trimspool on
SQL> spool /tmp/delete_logfiles.log
SQL> select member from v$logfile;
SQL> spool off;
```

Save the current RMAN configuration settings

Save the standby database RMAN configuration so it can be reapplied after restoring the control file.

```
$ rman target / nocatalog log=/tmp/RMAN_settings.log <<EOF
show all;
$ grep ^CONFIGURE /tmp/RMAN_settings.log | grep -v 'RETENTION POLICY' >/tmp/
RMAN_settings.rman
$ rm /tmp/RMAN_settings.log
```

Refresh the standby control file from the primary database

1. Before deleting backups, determine whether AL backups taken by the snapshot database need to be retained.
 - For OSS, NAS, and disk backups, these backups will no longer be tracked by the database and therefore will not be included in retention-based deletion.
 - Such backups must be manually deleted directly from their backup destination if no longer needed.

2. If backups are no longer needed, they can be deleted using RMAN. This avoids the need for manual deletion in the future.

```
RMAN> delete force noprompt backup device type 'sbt_tape';
RMAN> delete force noprompt archivelog all device type 'sbt_tape';
RMAN> delete force noprompt datafilecopy all device type 'sbt_tape';
```

3. Restore the control file from the primary database.

```
$ srvctl stop database -d <db> -o immediate
$ rman target / nocatalog
RMAN> startup nomount
RMAN> restore standby controlfile from service <tns alias for primary
database>;
RMAN> alter database mount;
```

4. Replace the RMAN configuration settings and remove the orphaned online and standby redo log files.

```
RMAN> @/tmp/RMAN_settings.rman
<output from CONFIGURE commands in the script>
RMAN> exit
$ rm /tmp/RMAN_settings.rman
```

5. Restore the datafiles:

- a. Switch to the primary database's current incarnation. From the primary database, run:

```
RMAN> list incarnation;
list incarnation;
using target database control file instead of recovery catalog
```

List of Database Incarnations

DB Key	Inc Key	DB Name	DB ID	STATUS	Reset SCN	Reset Time
1	1	OGG_DG	1153531161	PARENT	1	21-JAN-25
2	2	OGG_DG	1153531161	CURRENT	1514702	01-JUL-25

- b. From the standby database, run:

```
RMAN> reset database to incarnation <>;
```

6. Roll the Standby Database Forward: The standby database is now ready to perform the incremental roll forward using the recover-from-service operation. Restart all instances to mount.

```
$ srvctl stop database -db <dbname> -o immediate
$ srvctl start database -db <dbname> -o mount
```

7. Re-Enable Redo Transport: Archived logs generated during the roll forward are required to bring the database to a consistent state. It is more efficient to allow the primary database to ship redo to the standby while the roll forward is in progress, rather than waiting for the logs to be transported after the roll forward completes.

- a. Check the log archive destination for this standby database and note the destination number.

```
SQL> show parameter LOG_ARCHIVE_DEST_1;
NAME                                TYPE        VALUE
-----
log_archive_dest_1                   string      service="snby002", ASYNC NOAFF
                                        IRM delay=0 optional compressi
                                        on=disable max_failure=0 reope
                                        n=300 db_unique_name="snby002"
                                        net_timeout=30, valid_for=(on
                                        line_logfile,all_roles)
```

- b. Enable redo transport.

```
DGMGRL> ALTER SYSTEM SET LOG_ARCHIVE_DEST_STATE_2=ENABLE;
```

- c. Using Data Guard Broker, perform this operation from either the primary or the standby database.

```
DGMGRL> edit database <primary> set state=TRANSPORT-ON;
```

8. Run Recover From Service: The recovery command can use all available standby instances during the recovery process.

```
$ rman target sys/<password> <- It is necessary to connect with the
password
RMAN > run {

allocate channel c1 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c2 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c3 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c4 type disk connect '/@<standby instance 1 SID_NAME>';
allocate channel c5 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c6 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c7 type disk connect '/@<standby instance 2 SID_NAME>';
allocate channel c8 type disk connect '/@<standby instance 2 SID_NAME>';
restore database from service '<primary unique name>' section size 64G;
recover database from service '<primary unique name>' section size 64G;
}
```

9. Recover Until Consistent: At the completion of the RECOVER DATABASE FROM SERVICE command, additional recovery is required to make the standby database consistent before you can re-enable Flashback Database and open it in read-only mode.

A database is considered consistent when the control file and all datafiles are at the same SCN. The amount of additional recovery needed depends on how long the initial recovery took and the level of activity on the primary during that period.

As a first step, switch a log on the primary database to ensure that the last of the redo generated during the recovery process is archived.

From the primary database, run:

```
SQL> alter system archive log current;
```

10. Re-Enable Flashback Database: Restoring the standby control file automatically disables Flashback Database. Once the standby is brought to a consistent state, you can re-enable Flashback Database.

```
SQL> alter database flashback on;
```

11. Restart the standby database and re-enable managed recovery.

```
$ srvctl stop database -db <dbname> -o immediate
$ srvctl start database -db <dbname> -o 'read only'
```

12. Enable redo apply.

```
DGMGRL> edit database <standby> set state=APPLY-ON;
```

HEALTH.DB_GUEST.PDB.DB_BLOCKER_STALL_CHECK

Problem Statement: Identify Stalled DB Sessions Blocker

Risk: A high number of blocked or hanging sessions can lead to a database hang if critical operations wait indefinitely, and can also negatively impact application performance.

Action: To troubleshoot DB Blocker, follow My Oracle Support (MOS) notes: [1378583.1](#).

HEALTH.DB_GUEST.PDB.TOTAL_EXADATA_CACHE_HIT_RATIO

Problem Statement: Identify Exadata Cache Efficiency

Risk: If the Consolidated Exadata Cache Hit Ratio (%) drops below 90%, it suggests suboptimal cache utilization, potentially contributing to increased physical I/O and overall database performance degradation.

Action: To troubleshoot the low Exadata Cache Hit Ratio, refer to [Monitoring Smart Flash Cache – Common Issues](#).

HEALTH.DB_GUEST.CDB.HIGH_REDO_LOG_WAITS

Problem Statement: Elevated redo log write waits observed

Risk: Prolonged redo log write waits can significantly delay transaction commits, degrade throughput, and lead to overall database performance degradation.

Action: To troubleshoot high redo log waits, refer to the following My Oracle Support (MOS) notes: 223117.1, 2634755.1, and 1376916.1.

Application VIP Event Types

These are the event types that Application VIPs in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Application Virtual IP (VIP) - Create Begin	com.oraclecloud.databaseservice.createapplicationvip.begin

Friendly Name	Event Type
Application Virtual IP (VIP) - Create End	com.oraclecloud.databaseservice.createapplicationvip.end
Application Virtual IP (VIP) - Delete Begin	com.oraclecloud.databaseservice.deleteapplicationvip.begin
Application Virtual IP (VIP) - Delete End	com.oraclecloud.databaseservice.deleteapplicationvip.end

Application VIP Event Types Examples:

This is a reference event for Application Virtual IP (VIP) - Create Begin:

```
{
  "id":
  "ocid1.eventschema.oc1.phx.5ur5er8bddumnu9r84rtt2c3282s5no31vsthibyqvvsisotnwp
  csg9idv6q",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    }
  ],
}
```

```
{
  "name": "lifecycleDetails",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "hostnameLabel",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "cloudVmClusterId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "compartmentId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "vcnIpId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "ipAddress",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "subnetId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "networkType",
  "type": [
    "null",
    "string"
  ]
}
}
```

```

],
"exampleEvent": {
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "databaseservice",
  "contentType": "application/json",
  "eventID": "ab2ac219-b435-1045-aaf3-13cd909ec106",
  "eventTime": "2022-12-16T21:16:04.000Z",
  "data": {
    "resourceId": "ocid1.applicationvip.oc1.....unique_id",
    "resourceName": "my_application_vip",
    "tagSlug": null,
    "compartmentId": "ocid1.compartment.oc1.....unique_id",
    "request": {
      "id": "4260c9fd-d36b-4bc8-866e-c2dd53f34b2f",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,
      "message": ""
    },
    "stateChange": {
      "previous": null,
      "current": {
        "lifecycleState": "PROVISIONING",
        "hostnameLabel": "my_application_vip",
        "freeTags": {},
        "definedTags": {}
      }
    },
    "eventGroupingId": "csid74237ee84398b60cf1b834c81602/
f43a881dc99542318d46fa9285bdf2c5/6AC9F7641E1A5AD5C27D1650CB17E822",
    "eventName": "CreateApplicationVip",
    "availabilityDomain": "",
    "resourceVersion": null,
    "additionalDetails": {
      "id": "ocid1.applicationvip.oc1.....unique_id",
      "freeformTags": {},
      "definedTags": {},
      "timeCreated": "2022-12-15T21:17:59.000Z",
      "timeUpdated": "2022-12-15T21:18:04.389Z",
      "lifecycleState": "PROVISIONING",
      "lifecycleDetails": "",
      "hostnameLabel": "my_application_vip",
      "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
      "compartmentId": "ocid1.compartment.oc1.....unique_id",
      "vcnIpId": "ocid1.privateip.oc1.....unique_id",
      "ipAddress": "10.0.0.0",
      "subnetId": "ocid1.subnet.oc1.....unique_id",

```

```

        "networkType": "CLIENT"
      }
    }
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}

```

This is a reference event for Application Virtual IP (VIP) - Create End:

```

{
  "id":
"ocid1.eventschema.oc1.phx.clok1948lwge4il6m85ta4jdlbnhlyjrjltrabujyv52calb0el
p263oyqrm",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Create End",
  "eventType": "com.oraclecloud.databaseservice.createapplicationvip.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",

```

```

        "string"
      ]
    },
    {
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "cloudVmClusterId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "compartmentId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "vcnIpId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "ipAddress",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "subnetId",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "networkType",
      "type": [
        "null",
        "string"
      ]
    }
  ],
  "exampleEvent": {
    "eventType": "com.oraclecloud.databaseservice.createapplicationvip.end",
    "cloudEventsVersion": "0.1",

```

```

"eventTypeVersion": "2.0",
"source": "databaseservice",
"contentType": "application/json",
"eventID": "bc122d87-ac42-8731-ccd1-09ab320eef11",
"eventTime": "2022-12-16T21:16:04.000Z",
"data": {
  "resourceId": "ocid1.applicationvip.oc1.....unique_id",
  "resourceName": "my_application_vip",
  "tagSlug": null,
  "compartmentId": "ocid1.compartment.oc1.....unique_id",
  "request": {
    "id": "195eb9b5-b5a0-474d-a1c3-86189d8eeb2c",
    "path": null,
    "action": null,
    "parameters": null,
    "headers": null
  },
  "response": {
    "status": null,
    "responseTime": null,
    "headers": null,
    "payload": null,
    "message": ""
  },
  "stateChange": {
    "previous": null,
    "current": {
      "lifecycleState": "AVAILABLE",
      "hostnameLabel": "my_application_vip",
      "freeTags": {},
      "definedTags": {}
    }
  },
  "eventGroupingId":
  "6CEB05B6C81E4B19855AD716E90F5BC3/070ECF4976BDD89B16849A92B95564A6/1418EDD7590
  B8D5DDFF947FC3161F358",
  "eventName": "CreateApplicationVip",
  "availabilityDomain": "",
  "resourceVersion": null,
  "additionalDetails": {
    "id": "ocid1.applicationvip.oc1.....unique_id",
    "freeformTags": {},
    "definedTags": {},
    "timeCreated": "2022-12-15T21:17:59.000Z",
    "timeUpdated": "2022-12-15T21:18:04.389Z",
    "lifecycleState": "AVAILABLE",
    "lifecycleDetails": "",
    "hostnameLabel": "my_application_vip",
    "cloudVmClusterId": "ocid1.cloudvmcluster.oc1.....unique_id",
    "compartmentId": "ocid1.compartment.oc1.....unique_id",
    "vcnIpId": "ocid1.privateip.oc1.....unique_id",
    "ipAddress": "10.0.0.0",
    "subnetId": "ocid1.subnet.oc1.....unique_id",
    "networkType": "CLIENT"
  }
}
}

```

```

    },
    "timeCreated": "2022-12-15T16:31:31.979Z"
  }
}

```

This is a reference event for Application Virtual IP (VIP) - Delete Begin:

```

{
  "id":
  "ocid1.eventschema.oc1.phx.m2gheil6f1nfzb9ggpkv17wdomdks8zin9nntqlghui6bckh17
  yu0m5jcqt",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Delete Begin",
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.begin",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    }
  ],
}

```

```
{
  "name": "hostnameLabel",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "cloudVmClusterId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "compartmentId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "vcnIpId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "ipAddress",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "subnetId",
  "type": [
    "null",
    "string"
  ]
},
{
  "name": "networkType",
  "type": [
    "null",
    "string"
  ]
}
],
"exampleEvent": {
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "databaseservice",
  "contentType": "application/json",
```

```

"eventID": "e32cb1fe-123d-8341-de13-2be5f18ab31e",
"eventTime": "2022-12-16T21:16:04.000Z",
"data": {
  "resourceId": "ocidl.applicationvip.oc1....unique_id",
  "resourceName": "my_application_vip",
  "tagSlug": null,
  "compartmentId": "ocidl.compartment.oc1....unique_id",
  "request": {
    "id": "23a08e08-6b1e-40f0-a027-f2601dfd44ea",
    "path": null,
    "action": null,
    "parameters": null,
    "headers": null
  },
  "response": {
    "status": null,
    "responseTime": null,
    "headers": null,
    "payload": null,
    "message": ""
  },
  "stateChange": {
    "previous": null,
    "current": {
      "lifecycleState": "TERMINATING",
      "hostnameLabel": "my_application_vip",
      "freeTags": {},
      "definedTags": {}
    }
  },
  "eventGroupingId": "csidb3f42d234534bc8bc8849b892e84/
fbd51970d2a2486f94671614b5ea0571/9DFE1BEB5433FF69BABCCB7E34F2EAF4",
  "eventName": "DeleteApplicationVip",
  "availabilityDomain": "",
  "resourceVersion": null,
  "additionalDetails": {
    "id": "ocidl.applicationvip.oc1....unique_id",
    "freeformTags": {},
    "definedTags": {},
    "timeCreated": "2022-12-15T21:17:59.000Z",
    "timeUpdated": "2022-12-15T21:18:04.389Z",
    "lifecycleState": "TERMINATING",
    "lifecycleDetails": "",
    "hostnameLabel": "my_application_vip",
    "cloudVmClusterId": "ocidl.cloudvmcluster.oc1....unique_id",
    "compartmentId": "ocidl.compartment.oc1....unique_id",
    "vcnIpId": "ocidl.privateip.oc1....unique_id",
    "ipAddress": "10.0.0.0",
    "subnetId": "ocidl.subnet.oc1....unique_id",
    "networkType": "CLIENT"
  }
},
"timeCreated": "2022-12-15T16:31:31.979Z"
}

```

This is a reference event for Application Virtual IP (VIP) - Delete End:

```
{
  "id":
  "ocid1.eventschema.oc1.phx.9d1tjgkavhn0rq4qdlmofrjro9npvugu73dp07uht0igxs9732x
  6yarlm5l5",
  "serviceName": "Database",
  "displayName": "Application Virtual IP (VIP) - Delete End",
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.end",
  "source": "databaseservice",
  "eventTypeVersion": "2.0",
  "eventTime": "2022-12-15T21:16:04.000Z",
  "contentType": "application/json",
  "additionalDetails": [
    {
      "name": "id",
      "type": "string"
    },
    {
      "name": "definedTags",
      "type": [
        "null",
        "Map<String, Map<String, Object>>"
      ]
    },
    {
      "name": "freeFormTags",
      "type": [
        "null",
        "Map<String, String>"
      ]
    },
    {
      "name": "timeCreated",
      "type": "string"
    },
    {
      "name": "timeUpdated",
      "type": "string"
    },
    {
      "name": "lifecycleState",
      "type": "string"
    },
    {
      "name": "lifecycleDetails",
      "type": [
        "null",
        "string"
      ]
    },
    {
      "name": "hostnameLabel",
      "type": [
        "null",
        "string"
      ]
    }
  ]
}
```

```

    ]
  },
  {
    "name": "cloudVmClusterId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "compartmentId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "vcnIpId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "ipAddress",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "subnetId",
    "type": [
      "null",
      "string"
    ]
  },
  {
    "name": "networkType",
    "type": [
      "null",
      "string"
    ]
  }
],
"exampleEvent": {
  "eventType": "com.oraclecloud.databaseservice.deleteapplicationvip.end",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "databaseservice",
  "contentType": "application/json",
  "eventID": "17619ca1-07ae-4e2d-a818-5b5f1fcd4f70",
  "eventTime": "2022-12-16T21:16:04.000Z",
  "data": {
    "resourceId": "ocidl.applicationvip.oc1.....unique_id",
    "resourceName": "my_application_vip",

```

```

    "tagSlug": null,
    "compartmentId": "ocidl.compartment.oc1....unique_id",
    "request": {
      "id": "1b0d242b-b3cd-4d61-9779-2de23e0e6742",
      "path": null,
      "action": null,
      "parameters": null,
      "headers": null
    },
    "response": {
      "status": null,
      "responseTime": null,
      "headers": null,
      "payload": null,
      "message": ""
    },
    "stateChange": {
      "previous": null,
      "current": {
        "lifecycleState": "TERMINATED",
        "hostnameLabel": "my_application_vip",
        "freeTags": {},
        "definedTags": {}
      }
    },
    "eventGroupingId": "csid80b16d4d459eaaa60ad25a9829d8/
b3e19f76a81549e6b7bffd8619f7c191/C683214FCB0BF3CEC1C8B23C2FEE983E",
    "eventName": "DeleteApplicationVip",
    "availabilityDomain": "",
    "resourceVersion": null,
    "additionalDetails": {
      "id": "ocidl.applicationvip.oc1....unique_id",
      "freeformTags": {},
      "definedTags": {},
      "timeCreated": "2022-12-15T21:17:59.000Z",
      "timeUpdated": "2022-12-15T21:18:04.389Z",
      "lifecycleState": "TERMINATED",
      "lifecycleDetails": "",
      "hostnameLabel": "my_application_vip",
      "cloudVmClusterId": "ocidl.cloudvmcluster.oc1....unique_id",
      "compartmentId": "ocidl.compartment.oc1....unique_id",
      "vcnIpId": "ocidl.privateip.oc1....unique_id",
      "ipAddress": "10.0.0.0",
      "subnetId": "ocidl.subnet.oc1....unique_id",
      "networkType": "CLIENT"
    }
  },
  "timeCreated": "2022-12-15T16:31:31.979Z"
}

```

Interim Software Updates Event Types

These are the event types that Interim Software Updates in Oracle Cloud Infrastructure emit.

Friendly Name	Event Type
Oneoff Patch - Create Begin	com.oraclecloud.databaseservice.createoneoffpatch.begin
Oneoff Patch - Create End	com.oraclecloud.databaseservice.createoneoffpatch.end
Oneoff Patch - Delete Begin	com.oraclecloud.databaseservice.deleteoneoffpatch.begin
Oneoff Patch - Delete End	com.oraclecloud.databaseservice.deleteoneoffpatch.end
Oneoff Patch - Download Begin	com.oraclecloud.databaseservice.downloadoneoffpatch.begin
Oneoff Patch - Download End	com.oraclecloud.databaseservice.downloadoneoffpatch.end

Interim Software Updates Event Types Examples:

This is a reference event for This is a reference event for Oneoff Patch - Create Begin:

```
{
  "id":
  "ocidl.eventschema.oc1.phx.abyhqljrsl1p7rfnea9gq2knxbqopwux24za7qzoe3mfj2bzfxtnwqcxpbcq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    },
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_oneoffpatch",
    "resourceId": "OneOffPatch-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2020-08-26T12:00:00.000Z",
      "displayName": "testDisplayName",
      "databaseVersion": "19.6.0.0",
      "patchSet": "test_patch_set"
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create Begin",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.begin",
}
```

```

"additionalDetails": [
  { "name": "id", "type": "string" },
  { "name": "lifecycleState", "type": "string" },
  { "name": "timeCreated", "type": "string" },
  { "name": "displayName", "type": "string" },
  { "name": "dbVersion", "type": "string" },
  { "name": "patchType", "type": "string" },
  { "name": "patchShapeFamily", "type": "string" },
  { "name": "releaseUpdate", "type": "string" }
],
"timeCreated": "2020-06-26T13:31:31.979Z"
}

```

This is a reference event for Oneoff Patch - Create End:

```

{
  "id":
"ocidl.eventschema.oc1.phx.abyhqljrj4vvuph4qvj5eateeel6axblhkq3caqndgmjvw13sld
pgb255j2q",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocidl.id.oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Create End",
  "eventType": "com.oraclecloud.databaseservice.createoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },

```

```

    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}

```

This is a reference event for Oneoff Patch - Delete Begin:

```

{
  "id":
  "ocidl.eventschema.oc1.phx.abyhqljrdripga5rryplwmv4ws6hqzr3pjyl7wfvoaquvtvg2ey2
  vtycn5onq",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    },
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_oneoffpatch",
    "resourceId": "OneOffPatch-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id.oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2020-08-26T12:00:00.000Z",
      "displayName": "testDisplayName",
      "databaseVersion": "19.6.0.0",
      "patchSet": "test_patch_set"
    }
  }
},
"serviceName": "Database",
"displayName": "Oneoff Patch - Delete Begin",
"eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.begin",
"additionalDetails": [
  { "name": "id", "type": "string" },
  { "name": "lifecycleState", "type": "string" },
  { "name": "timeCreated", "type": "string" },
  { "name": "displayName", "type": "string" },
  { "name": "dbVersion", "type": "string" },
  { "name": "patchType", "type": "string" },
  { "name": "patchShapeFamily", "type": "string" },

```

```

    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}

```

This is a reference event for Oneoff Patch - Delete End:

```

{
  "id":
  "ocidl.eventschema.oc1.phx.abyhqljrgwk2gvx5lmx6fiwotgdy32mdmrnkyzznz37dpb4mmeh
  gzt37v17a",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    },
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_oneoffpatch",
    "resourceId": "OneOffPatch-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id.oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2020-08-26T12:00:00.000Z",
      "displayName": "testDisplayName",
      "databaseVersion": "19.6.0.0",
      "patchSet": "test_patch_set"
    }
  }
},
"serviceName": "Database",
"displayName": "Oneoff Patch - Delete End",
"eventType": "com.oraclecloud.databaseservice.deleteoneoffpatch.end",
"additionalDetails": [
  { "name": "id", "type": "string" },
  { "name": "lifecycleState", "type": "string" },
  { "name": "timeCreated", "type": "string" },
  { "name": "displayName", "type": "string" },
  { "name": "dbVersion", "type": "string" },
  { "name": "patchType", "type": "string" },
  { "name": "patchShapeFamily", "type": "string" },
  { "name": "releaseUpdate", "type": "string" }
],
"timeCreated": "2020-06-26T13:31:31.979Z"
}

```

This is a reference event for Oneoff Patch - Download Begin:

```
{
  "id":
  "ocidl.eventschema.oc1.phx.abyhqljr3vkb7klt5hkbsnqzjaxmszsqomanlbqmr2tsrcq7xaf
  cv2c7412q",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    },
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_oneoffpatch",
    "resourceId": "OneOffPatch-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..oc1...unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2020-08-26T12:00:00.000Z",
      "displayName": "testDisplayName",
      "databaseVersion": "19.6.0.0",
      "patchSet": "test_patch_set"
    }
  }
},
"serviceName": "Database",
"displayName": "Oneoff Patch - Download Begin",
"eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.begin",
"additionalDetails": [
  { "name": "id", "type": "string" },
  { "name": "lifecycleState", "type": "string" },
  { "name": "timeCreated", "type": "string" },
  { "name": "displayName", "type": "string" },
  { "name": "dbVersion", "type": "string" },
  { "name": "patchType", "type": "string" },
  { "name": "patchShapeFamily", "type": "string" },
  { "name": "releaseUpdate", "type": "string" }
],
"timeCreated": "2020-06-26T13:31:31.979Z"
}
```

This is a reference event for Oneoff Patch - Download End:

```
{
  "id":
```

```

"ocid1.eventschema.oc1.phx.abyhqljrn2lruez55ah56kqksi5qfg6m7igvven7o2qkahlk5tk
wrj5ll3oa",
  "exampleEvent": {
    "cloudEventsVersion": "0.1",
    "eventID": "60600c06-d6a7-4e85-b59a-1de3e6042f57",
    "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
    "source": "databaseservice",
    "eventTypeVersion": "1.0",
    "eventTime": "2020-06-27T21:16:04.000Z",
    "contentType": "application/json",
    "extensions": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID"
    },
    "data": {
      "compartmentId": "ocid1.compartment.oc1..unique_ID",
      "compartmentName": "example_name",
      "resourceName": "my_oneoffpatch",
      "resourceId": "OneOffPatch-unique_ID",
      "availabilityDomain": "all",
      "freeFormTags": {},
      "definedTags": {},
      "additionalDetails": {
        "id": "ocid1.id..oc1...unique_ID",
        "lifecycleState": "AVAILABLE",
        "timeCreated": "2020-08-26T12:00:00.000Z",
        "displayName": "testDisplayName",
        "databaseVersion": "19.6.0.0",
        "patchSet": "test_patch_set"
      }
    }
  },
  "serviceName": "Database",
  "displayName": "Oneoff Patch - Download End",
  "eventType": "com.oraclecloud.databaseservice.downloadoneoffpatch.end",
  "additionalDetails": [
    { "name": "id", "type": "string" },
    { "name": "lifecycleState", "type": "string" },
    { "name": "timeCreated", "type": "string" },
    { "name": "displayName", "type": "string" },
    { "name": "dbVersion", "type": "string" },
    { "name": "patchType", "type": "string" },
    { "name": "patchShapeFamily", "type": "string" },
    { "name": "releaseUpdate", "type": "string" }
  ],
  "timeCreated": "2020-06-26T13:31:31.979Z"
}

```

Serial Console Connection Event Types

Review the list of event types that serial console connection emits.

Table 6-5 Serial Console Connection Events

Friendly Name	Event Type
DB Node Console Connection - Create Begin	com.oraclecloud.databaseservice.createdbnodeconsoleconnection.begin
DB Node Console Connection - Create End	com.oraclecloud.databaseservice.createdbnodeconsoleconnection.end
DB Node Console Connection - Delete Begin	com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.begin
DB Node Console Connection - Delete End	com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.end
DB Node Console Connection - Update	com.oraclecloud.databaseservice.updatedbnodeconsoleconnection
DB Node - Update	com.oraclecloud.databaseservice.updatedbnode

Example 6-64 Serial Console Connection Event Types Examples

This is a reference event for DB Node Console Connection - Create Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "CREATING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
```

This is a reference event for DB Node Console Connection - Create End:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
```

This is a reference event for DB Node Console Connection - Delete Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETING",
      "timeCreated": "2019-08-29T12:00:00.000Z",

```

```

        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
        "tenantId": "ocidl.tenant.oc1..unique_ID",
        "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
}
}

```

This is a reference event for DB Node Console Connection - Delete End:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETED",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
}

```

This is a reference event for DB Node Console Connection - Update:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.updatedbnodeconsoleconnection",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },

```

```

"data": {
  "compartmentId": "ocidl.compartment.oc1..unique_ID",
  "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
  "freeFormTags": {},
  "definedTags": {},
  "additionalDetails": {
    "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "lifecycleState": "ACTIVE",
    "timeCreated": "2019-08-29T12:00:00.000Z",
    "timeUpdated": "2019-08-29T12:30:00.000Z",
    "lifecycleDetails": "detail message",
    "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
    "tenantId": "ocidl.tenant.oc1..unique_ID",
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  }
}
}
}

```

This is a reference event for DB Node - Update:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.updatedbnode",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_dbnode",
    "resourceId": "DbNode-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..oc1..unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-26T12:00:00.000Z",
      "timeUpdated": "2019-08-26T12:30:00.000Z",
      "dbSystemId": "ocidl.dbssystem.oc1.phx.unique_ID",
      "lifecycleDetails": "detail message",
      "vmClusterId": "VmCluster-unique_ID",
      "dbHostId": "dbHost-unique_ID",
      "nodeNumber": 2,
      "powerAction": "HardReset",
      "hostName": "testHostName"
    }
  }
}
}
}

```

- [Viewing Audit Log Events](#)
Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

Viewing Audit Log Events

Oracle Cloud Infrastructure Audit service provides records of API operations performed against supported services as a list of log events.

An audit event is generated when you connect to the serial console using a Secure Shell (SSH) connection. Navigate to Audit in the Console and search for `VmConsoleConnected`. When you navigate to Audit in the Console, a list of results is generated for the current compartment. Audit logs are organized by compartment, so if you are looking for a particular event, you must know which compartment the event occurred in. You can filter the list in the following ways:

- Date and time
- Request Action Types (operations)
- Keywords

For more information, see *Viewing Audit Log Events*.

Example 6-65 Serial Console Connection Audit Event Example

This is a reference event for Serial Console Connection:

```
{
  "eventType": "VmConsoleConnected",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "VmConsoleConnectionAPI",
  "eventId": "2367d627-cff8-11ed-bfd3-02001714f979",
  "eventTime": "2023-03-31T19:13:37.120Z",
  "contentType": "application/json",

  "data": {
    "eventGroupId": "2367d62d-cff8-11ed-bfd3-02001714f979",
    "eventName": "VmConsoleConnected",
    "compartmentId": "ocidl.compartment.oc1..<TRUNCATED>aaaaxxxxx",
    "compartmentName": "exacc-dev",
    "resourceName": "",
    "resourceId":
"ocidl.dbnodeconsoleconnection.oc1.iad.<TRUNCATED>aaaaaxxxxx",
    "availabilityDomain": null,
    "freeformTags": null,
    "definedTags": null,

    "identity": {
      "principalName": "dsaes",
      "principalId": "ocidl.user.oc1..<TRUNCATED>aaaaaaaaaxxxxxxxxx",
      "authType": "Native",
      "callerName": null,
      "callerId": null,
      "tenantId": "ocidl.tenancy.oc1..<TRUNCATED>aaaaaxxxxx",
      "ipAddress": null,
    }
  }
}
```

```

    "credentials": null,
    "userAgent": null,
    "consoleSessionId": null
  },

  "request": {
    "id": "",
    "path": "",
    "action": "",
    "parameters": null,
    "headers": null
  },

  "response": {
    "status": "",
    "responseTime": "0001-01-01T00:00:00.000Z",
    "headers": null,
    "payload": null,
    "message": ""
  },

  "stateChange": null,

  "additionalDetails": {
    "DBNodeId": "ocidl.dbnode.oc1.iad.<TRUNCATED>aaaaaxxxxxxx"
  }
}

```

Related Topics

- [Overview of Audit](#)
- [Viewing Audit Log Events](#)
- [Setting Audit Log Retention Period](#)

Serial Console History Event Types

Review the list of new event types that serial console history emits.

Table 6-6 Serial Console History Events

User Action	Event Type	Friendly Name	Event Type
Create Console History	Async	DB Node Console History - Create Begin	com.oraclecloud.dat abaseservice.create dbnodeconsolehistor y.beginn
Create Console History	Async	DB Node Console History - Create End	com.oraclecloud.dat abaseservice.create dbnodeconsolehistor y.end

Table 6-6 (Cont.) Serial Console History Events

User Action	Event Type	Friendly Name	Event Type
Terminate Console History	Async	DB Node Console History - Delete Begin	com.oraclecloud.dat abaseservice.delete dbnodeconsolehistor y.begin
Terminate Console History	Async	DB Node Console History - Delete End	com.oraclecloud.dat abaseservice.delete dbnodeconsolehistor y.end
Update Console History	Sync	DB Node Console History - Update	com.oraclecloud.dat abaseservice.update dbnodeconsolehistor y
Get Console History Content	Sync	DB Node Console History - Get Content	com.oraclecloud.dat abaseservice.getdbn odeconsolehistoryco ntent

Example 6-66 Serial Console Connection Event Types Examples

This is a reference event for DB Node Console Connection - Create Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "CREATING",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
```

This is a reference event for DB Node Console Connection - Create End:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.createdbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "ACTIVE",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
```

This is a reference event for DB Node Console Connection - Delete Begin:

```
"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.begin",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETING",
      "timeCreated": "2019-08-29T12:00:00.000Z",

```

```

        "timeUpdated": "2019-08-29T12:30:00.000Z",
        "lifecycleDetails": "detail message",
        "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
        "tenantId": "ocidl.tenant.oc1..unique_ID",
        "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
}
}

```

This is a reference event for DB Node Console Connection - Delete End:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.deletedbnodeconsoleconnection.end",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
      "lifecycleState": "DELETED",
      "timeCreated": "2019-08-29T12:00:00.000Z",
      "timeUpdated": "2019-08-29T12:30:00.000Z",
      "lifecycleDetails": "detail message",
      "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
      "tenantId": "ocidl.tenant.oc1..unique_ID",
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}
}

```

This is a reference event for DB Node Console Connection - Update:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventID": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType":
"com.oraclecloud.databaseservice.updatedbnodeconsoleconnection",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-08-29T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
}

```

```

"data": {
  "compartmentId": "ocidl.compartment.oc1..unique_ID",
  "resourceId": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
  "freeFormTags": {},
  "definedTags": {},
  "additionalDetails": {
    "id": "ocidl.dbnodeconsoleconnection.oc1..unique_ID",
    "lifecycleState": "ACTIVE",
    "timeCreated": "2019-08-29T12:00:00.000Z",
    "timeUpdated": "2019-08-29T12:30:00.000Z",
    "lifecycleDetails": "detail message",
    "dbnodeId": "ocidl.dbnode.oc1..unique_ID",
    "tenantId": "ocidl.tenant.oc1..unique_ID",
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  }
}
}
}

```

This is a reference event for DB Node - Update:

```

"exampleEvent": {
  "cloudEventsVersion": "0.1",
  "eventId": "60600c06-d6a7-4e85-b56a-1de3e6042f57",
  "eventType": "com.oraclecloud.databaseservice.updatedbnode",
  "source": "databaseservice",
  "eventTypeVersion": "1.0",
  "eventTime": "2019-06-27T21:16:04.000Z",
  "contentType": "application/json",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  },
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_name",
    "resourceName": "my_dbnode",
    "resourceId": "DbNode-unique_ID",
    "availabilityDomain": "all",
    "freeFormTags": {},
    "definedTags": {},
    "additionalDetails": {
      "id": "ocidl.id..oc1..unique_ID",
      "lifecycleState": "AVAILABLE",
      "timeCreated": "2019-08-26T12:00:00.000Z",
      "timeUpdated": "2019-08-26T12:30:00.000Z",
      "dbSystemId": "ocidl.dbsystem.oc1.phx.unique_ID",
      "lifecycleDetails": "detail message",
      "vmClusterId": "VmCluster-unique_ID",
      "dbHostId": "dbHost-unique_ID",
      "nodeNumber": 2,
      "powerAction": "HardReset",
      "hostName": "testHostName"
    }
  }
}
}
}

```

Policy Details for Exadata Cloud Infrastructure

This topic covers details for writing policies to control access to Exadata Cloud Infrastructure resources.

Note

For more information on Policies, see "How Policies Work".

For a sample policy, see "Let database admins manage Exadata Cloud Infrastructure instances".

- [About Resource-Types](#)
Learn about resource-types you can use in your policies.
- [Resource-Types for Exadata Cloud Service Instances](#)
- [Supported Variables](#)
Use variables when adding conditions to a policy.
- [Details for Verb + Resource-Type Combinations](#)
Review the list of permissions and API operations covered by each verb.

Related Topics

- [How Policies Work](#)
- [Let database admins manage Exadata Cloud Infrastructure instances](#)

About Resource-Types

Learn about resource-types you can use in your policies.

An aggregate resource-type covers the list of individual resource-types that directly follow. For example, writing one policy to allow a group to have access to the `database-family` is equivalent to writing separate policies for the group that would grant access to the `cloud-exadata-infrastructures`, `cloud-vmclusters`, `db-nodes`, `db-homes`, `databases`, `database-software-image`, and `backups` resource-types. For more information, see [Resource-Types](#).

Resource-Types for Exadata Cloud Service Instances

Aggregate Resource-Type

- `database-family`
- `db-multi-cloud-family`

Individual Resource-Types:

- `database-family`
 - `cloud-exadata-infrastructures`
 - `cloud-vmclusters`
 - `db-nodes`
 - `db-homes`

- databases
- pluggable-databases
- exascale-db-storage-vaults
- db-backups
- application-vips
- dbnode-console-connection
- dbnode-console-history
- scheduling-policies
- scheduling-windows
- scheduling-plan
- scheduling-action
- execution-windows
- execution-action
- db-multi-cloud-family
 - oracle-db-azure-connectors
 - oracle-db-azure-blob-mounts
 - oracle-db-azure-blob-containers
 - oracle-db-mci-work-requests
 - multi-cloud-resource-discoveries
 - oracle-db-azure-vaults
 - oracle-db-azure-keys
 - oracle-db-azure-vault-associations
 - oracle-db-gcp-connectors
 - oracle-db-gcp-keyrings
 - oracle-db-gcp-keys
 - oracle-db-aws-connectors
 - oracle-db-aws-keys

Note

The following API calls require additional permissions:

- oracle-db-azure-connectors
- multi-cloud-resource-discoveries
- oracle-db-gcp-connectors
- oracle-db-gcp-keyrings

Additional Required Permissions

To use these APIs, make sure the following permissions are granted:

- CLOUD_VM_CLUSTER_UPDATE
- CLOUD_VM_CLUSTER_INSPECT

Without these permissions, API requests may fail or return incomplete results.

Supported Variables

Use variables when adding conditions to a policy.

Exadata Cloud Infrastructure supports only the general variables. For more information, see "General Variables for All Requests".

Related Topics

- [General Variables for All Requests](#)

Details for Verb + Resource-Type Combinations

Review the list of permissions and API operations covered by each verb.

For more information, see "Permissions", "Verbs", and "Resource-Types".

- [Database-Family Resource Types](#)
- [Permissions and API operation details for Cloud Exadata Infrastructures](#)
- [Permissions and API operation details for VM Clusters](#)
- [Permissions and API operation details for Exascale DB Storage Vaults](#)
- [Permissions and API operation details for DB Nodes](#)
- [Permissions and API operation details for DB Node Console Connection](#)
- [Permissions and API operation details for DB Node Console History](#)
- [Permissions and API operation details for DB Homes](#)
- [Permissions and API operation details for DB Servers](#)
- [Permissions and API operation details for Database Software Image](#)
- [Permissions and API operation details for Pluggable Databases \(PDBs\)](#)
- [Permissions and API operation details for Databases \(CDBs\)](#)
- [Permissions and API operation details for DB Backups](#)
- [Permissions and API operation details for Data Guard Association](#)

- [Permissions and API operation details for Data Guard Group](#)
- [Permissions and API operation details for Key Stores](#)
- [Permissions and API operation details for Application VIPs](#)
- [Permissions and API operation details for Interim Software Updates](#)
- [Permissions and API operation details for Scheduling Policies](#)
- [Permissions and API operation details for Scheduling Windows](#)
- [Permissions and API operation details for Scheduling Plan](#)
- [Permissions and API operation details for Scheduled Action](#)
- [Permissions and API operation details for Execution Windows](#)
- [Permissions and API operation details for Execution Action](#)
- [Permissions and API operation details for Oracle DB Azure Connectors](#)
- [Permissions and API operation details for Oracle DB Azure Blob Mounts](#)
- [Permissions and API operation details for Oracle DB Azure Blob Containers](#)
- [Permissions and API operation details for Oracle DB MCI Work Requests](#)
- [Permissions and API operation details for Multicloud Resource Discoveries](#)
- [Permissions and API operation details for Oracle DB Azure Vaults](#)
- [Permissions and API operation details for Oracle DB Azure Keys](#)
- [Permissions and API operation details for Oracle DB Azure Vault Associations](#)
- [Permissions and API operation details for Oracle DB GCP Connectors](#)
- [Permissions and API operation details for Oracle DB GCP Keyrings](#)
- [Permissions and API operation details for Oracle DB GCP Keys](#)
- [Permissions and API operation details for Oracle DB AWS Identity Connectors](#)
- [Permissions and API operation details for Oracle DB AWS Key](#)
- [Permissions Required for Each API Operation](#)

Related Topics

- [Permissions](#)
- [Verbs](#)
- [Resource-Types](#)

Database-Family Resource Types

The level of access is cumulative as you go from `inspect` > `read` > `use` > `manage`. A plus sign (+) in a table cell indicates incremental access compared to the cell directly above it, whereas "no extra" indicates no incremental access.

For example, the `read` verb for the `vmclusters` resource-type covers no extra permissions or API operations compared to the `inspect` verb. However, the `use` verb includes one more permission, fully covers one more operation, and partially covers another additional operation.

Permissions and API operation details for Cloud Exadata Infrastructures

The table below lists permissions and API operations for `cloud-exadata-infrastructures`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	CLOUD_EXADATA_INFRA STRUCTURE_INSPECT	ListCloudExadataInf rastructures GetCloudExadataInfr astructures	none
read	<i>no extra</i>	<i>no extra</i>	none
use	CLOUD_EXADATA_INFRA STRUCTURE_UPDATE	ConfigureExascaleCl oudExadataInfrastru cture	ChangeCloudExadataI nfrastructureCompar tment (also needs use cloud-vmclusters, use db-homes, use databases, and inspect db-backups)
manage	USE + CLOUD_EXADATA_INFRA STRUCTURE_CREATE CLOUD_EXADATA_INFRA STRUCTURE_DELETE	UpdateCloudExadataI nfrastructure	CreateCloudExadataI nfrastructure, DeleteCloudExadataI nfrastructure, AddStorageCapacityC loudExadataInfrastr ucture (also needs use cloud-vmclusters)

Permissions and API operation details for VM Clusters

The table below lists permissions and API operations for `cloud-vmclusters`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	CLOUD_VM_CLUSTER_IN SPECT	ListCloudVmClusters GetCloudVmCluster ListCloudVmClusterU pdates ListCloudVmClusterU pdateHistoryEntries GetCloudVmClusterUp date GetCloudVmClusterUp dateHistoryEntry	none
read	<i>no extra</i>	<i>no extra</i>	none

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	READ + CLOUD_VM_CLUSTER_UPDATE CLOUD_VM_CLUSTER_UPDATE_TAGS CLOUD_VM_CLUSTER_UPDATE_COMPARTMENT CLOUD_VM_CLUSTER_UPDATE_SSH_KEY CLOUD_VM_CLUSTER_UPDATE_LICENSE CLOUD_VM_CLUSTER_UPDATE_CPU CLOUD_VM_CLUSTER_UPDATE_MEMORY CLOUD_VM_CLUSTER_UPDATE_LOCAL_STORAGE CLOUD_VM_CLUSTER_UPDATE_EXADATA_STORAGE CLOUD_VM_CLUSTER_UPDATE_GI_SOFTWARE CLOUD_VM_CLUSTER_UPDATE_GUEST_OS_SOFTWARE CLOUD_VM_CLUSTER_UPDATE_FILE_SYSTEM CLOUD_VM_CLUSTER_UPDATE_DIAGNOSTIC_LOGS CLOUD_VM_CLUSTER_UPDATE_IORM	<i>no extra</i>	ChangeCloudVmClusterCompartments (also needs use db-homes, use databases, and inspect db-backups)
manage	USE + CLOUD_VM_CLUSTER_CREATE CLOUD_VM_CLUSTER_DELETE	UpdateCloudVmCluster	CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage db-homes, manage databases, use vnics, and use subnets); RemoveVmFromCloudVmCluster, AddVmToCloudVmCluster (both also need use cloud_exadata_infrastructure_update)

Note

The `CLOUD_VM_CLUSTER_UPDATE_SSH_KEY` permission is a highly privileged permission that allows the user to be a `root` user on the guest VM and gives them the ability to run other cluster update operations on the guest VM using `dbaascli`.

Using fine-grained permissions, you can write policies as follows:

- To allow any update operations:

```
allow group abc to use cloud-vmclusters in compartment comp1
```

- To allow only scale CPU:

```
allow group abc to use cloud-vmclusters in compartment comp1 where
request.permission = 'CLOUD_VM_CLUSTER_UPDATE_CPU'
```

- To allow GI update and any scale operations:

```
allow group abc to use cloud-vmclusters in compartment comp1
where any
  { request.permission = 'CLOUD_VM_CLUSTER_UPDATE_CPU',
    request.permission = 'CLOUD_VM_CLUSTER_UPDATE_EXADATA_STORAGE',
      request.permission = 'CLOUD_VM_CLUSTER_UPDATE_MEMORY',
    request.permission = 'CLOUD_VM_CLUSTER_UPDATE_LOCAL_STORAGE',
    request.permission = 'CLOUD_VM_CLUSTER_UPDATE_GI_SOFTWARE' }
```

- To allow any operations except add SSH key:

```
allow group abc to use cloud-vmclusters in compartment comp1 where all
{ request.permission != 'CLOUD_VM_CLUSTER_UPDATE_SSH_KEY' ,
  request.permission != 'CLOUD_VM_CLUSTER_UPDATE' }
```

Permissions and API operation details for Exascale DB Storage Vaults

The table below lists permissions and API operations for `exascale-db-storage-vaults`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	EXASCALE_DB_STORAGE _VAULT_INSPECT	ListExascaleDbStora geVaults GetExascaleDbStorag eVault	<i>none</i>
read	<i>no extra</i>	<i>no extra</i>	<i>none</i>
use	READ + EXASCALE_DB_STORAGE _VAULT_UPDATE	ChangeExascaleDbSto rageVaultCompartment UpdateExascaleDbSto rageVault	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	USE + EXASCALE_DB_STORAGE _VAULT_CREATE EXASCALE_DB_STORAGE _VAULT_DELETE	CreateExascaleDbSto rageVault DeleteExascaleDbSto rageVault	<i>none</i>

Permissions and API operation details for DB Nodes

Note

For Exadata Cloud Infrastructure VM clusters, the database node is sometimes referred to as a virtual machine.

The table below lists permissions and API operations for db-nodes.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_NODE_INSPECT DB_NODE_QUERY	GetDbNode	<i>none</i>
read	<i>no extra</i>	<i>no extra</i>	<i>none</i>
use	DB_NODE_UPDATE	UpdateDbNode	<i>none</i>
manage	USE + DB_NODE_POWER_ACTIO NS	DbNodeAction	<i>none</i>

Permissions and API operation details for DB Node Console Connection

The table below lists permissions and API operations for dbnode-console-connection.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DBNODE_CONSOLE_CONN ECTION_INSPECT	GetDbNodeConsoleCon nection ListDbNodeConsoleCo nnections	<i>none</i>
read	<i>no extra</i>	<i>no extra</i>	<i>none</i>
use	READ + DBNODE_CONSOLE_CONN ECTION_UPDATE PLUGGABLE_DATABASE_ UPDATE	UpdateDbNodeConsole Connection	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	<i>USE</i> + DBNODE_CONSOLE_CONN ECTION_CREATE DBNODE_CONSOLE_CONN ECTION_DELETE	CreateDbNodeConsole Connection DeleteDbNodeConsole Connection	<i>none</i>

Permissions and API operation details for DB Node Console History

The table below lists permissions and API operations for dbnode-console-history.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DBNODE_CONSOLE_HIST ORY_INSPECT	GetDbNodeConsoleHis tory ListDbNodeConsoleHi stories	<i>none</i>
read	<i>INSPECT</i> + DBNODE_CONSOLE_HIST ORY_CONTENT_READ	getDbNodeConsoleHis toryContent	<i>none</i>
use	<i>READ</i> + DBNODE_CONSOLE_HIST ORY_UPDATE PLUGGABLE_DATABASE_ UPDATE	UpdateDbNodeConsole History	<i>none</i>
manage	<i>USE</i> + DBNODE_CONSOLE_HIST ORY_CREATE DBNODE_CONSOLE_HIST ORY_DELETE	CreateDbNodeConsole History DeleteDbNodeConsole History	<i>none</i>

Permissions and API operation details for DB Homes

The table below lists permissions and API operations for db-homes.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_HOME_INSPECT	ListDBHome GetDBHome ListDbHomePatches ListDbHomePatchHist oryEntries GetDbHomePatch GetDbHomePatchHisto ryEntry	<i>none</i>
read	<i>no extra</i>	<i>no extra</i>	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
use	DB_HOME_UPDATE	UpdateDBHome	ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use databases, and inspect backups)
manage	USE + DB_HOME_CREATE DB_HOME_DELETE	no extra	CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage cloud-vmclusters, manage databases, use vnics, and use subnets). If automatic backups are enabled on the default database, also needs manage backups CreateDbHome, (also needs use cloud-vmclusters and manage databases). If creating the Database Home by restoring from a backup, also needs read backups DeleteDbHome, (also needs use cloud-vmclusters and manage databases). If automatic backups are enabled on the default database, also needs manage backups. If the performFinalBackup option is selected, also needs manage backups and read databases.

Permissions and API operation details for DB Servers

The table below lists permissions and API operations for dbServers.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	EXADATA_INFRASTRUCTURE_INSPECT	none	GetDbServer ListDbServers
READ	no extra	no extra	none

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
USE	<i>READ +</i> VM_CLUSTER_UPDATE EXADATA_INFRASTRUCTURE_UPDATE	<i>none</i>	AddVirtualMachineToVmCluster, RemoveVirtualMachineFromVmCluster
MANAGE	<i>No extra</i>	<i>No extra</i>	<i>none</i>

Permissions and API operation details for Database Software Image

The table below lists permissions and API operations for database-software-image.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_SOFTWARE_IMG_INSPECT	ListDatabaseSoftwareImages GetDatabaseSoftwareImage	<i>none</i>
read	<i>no extra</i>	<i>none</i>	<i>none</i>
use	<i>READ +</i> DB_SOFTWARE_IMG_UPDATE	UpdateDatabaseSoftwareImage ChangeDatabaseSoftwareImageCompartment	<i>none</i>
manage	<i>USE +</i> DB_SOFTWARE_IMG_CREATE DB_SOFTWARE_IMG_DELETE	CreateDatabaseSoftwareImage DeleteDatabaseSoftwareImage	<i>none</i>

Permissions and API operation details for Pluggable Databases (PDBs)

The table below lists permissions and API operations for pluggable-databases.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	PLUGGABLE_DATABASE_INSPECT	ListPluggableDatabases GetPluggableDatabase	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
	DATABASE_INSPECT	<i>no extra</i>	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase
read	<i>INSPECT +</i> PLUGGABLE_DATABASE_CONTENT_READ	<i>no extra</i>	CreatePluggableDatabase (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.) UpdatePluggableDatabase (Additional permissions are required if auto-backups are enabled on the CDB and includes this PDB.) LocalClonePluggableDatabase RemoteClonePluggableDatabase
use	<i>READ +</i> PLUGGABLE_DATABASE_CONTENT_WRITE	<i>no extra</i>	LocalClonePluggableDatabase RemoteClonePluggableDatabase
	PLUGGABLE_DATABASE_UPDATE	<i>no extra</i>	UpdatePluggableDatabase StartPluggableDatabase StopPluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase RefreshPluggableDatabase ConvertRefreshablePluggableDatabase
	DATABASE_UPDATE	<i>no extra</i>	CreatePluggableDatabase DeletePluggableDatabase LocalClonePluggableDatabase RemoteClonePluggableDatabase

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	<i>USE +</i> PLUGGABLE_DATABASE_ CREATE	<i>no extra</i>	CreatePluggableData base LocalClonePluggable Database RemoteClonePluggabl eDatabase
	PLUGGABLE_DATABASE_ DELETE	<i>no extra</i>	DeletePluggableData base

Permissions and API operation details for Databases (CDBs)

The table below lists permissions and API operations for databases.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DATABASE_INSPECT	ListDatabases GetDatabase ListDataGuardAssoci ations GetDataGuardAssocia tion	enableDatabaseManag ement disableDatabaseMana gement updateDatabaseManag ement
read	<i>INSPECT+</i> DATABASE_CONTENT_RE AD	<i>no extra</i>	<i>no extra</i>
use	<i>READ +</i> DATABASE_CONTENT_WR ITE DATABASE_UPDATE	UpdateDatabase SwitchoverDataGuard Association FailoverDataGuardAs sociation ReinstateDataGuardA ssociation	CreateDataGuardAsso ciation ChangeCloudVmCluste rCompartment (also needs use cloud- vmclusters, use db- homes, and inspect db-backups) enableDatabaseManag ement disableDatabaseMana gement updateDatabaseManag ement

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
manage	<i>USE</i> + DATABASE_CREATE DATABASE_DELETE	<i>no extra</i>	CreateDatabase (also needs use cloud-vmclusters, use db-homes, and if automatic backups to be enabled, also needs manage backups) DeleteDatabase (also needs use cloud-vmclusters, use db-homes, and if automatic backups to be enabled, also needs manage backups) CreateCloudVmCluster, DeleteCloudVmCluster (both also need manage cloud-vmclusters, manage db-homes, use vnics, and use subnets)

Permissions and API operation details for DB Backups

The table below lists permissions and API operations for db-backups.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	DB_BACKUP_INSPECT	GetBackup ListBackups	ChangeCloudVmClusterCompartment (also needs use cloud-vmclusters, use db-homes, and use databases)
read	<i>INSPECT</i> + DB_BACKUP_CONTENT_READ	<i>none</i>	RestoreDatabase (also needs use databases)
use	<i>no extra</i>	<i>no extra</i>	<i>none</i>
manage	<i>USE</i> + DB_BACKUP_CREATE DB_BACKUP_DELETE	DeleteBackup	CreateBackup (also needs read databases)

Permissions and API operation details for Data Guard Association

The table below lists permissions and API operations for data-guard-association.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	DATABASE_INSPECT	ListDataGuardAssociations, GetDataGuardAssociation	CreateDataGuardAssociation
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ + VM_CLUSTER_UPDATE + DB_HOME_UPDATE DATABASE_UPDATE	DeleteDatabase SwitchoverDataGuardAssociation,FailoverDataGuardAssociation, ReinstateDataGuardAssociation	CreateDataGuardAssociation
MANAGE	USE + DATABASE_DELETE	DeleteDatabase	<i>none</i>

Permissions and API operation details for Data Guard Group

The table below lists permissions and API operations for Data Guard with multiple standby databases.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	DATABASE_INSPECT	ListDataGuardAssociations, GetDatabase	CreateDatabase
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ + Standby: CLOUD_VM_CLUSTER_INSPECT + DB_HOME_INSPECT + DATABASE_CREATE Primary: DATABASE_INSPECT + DATABASE_CONTENT_READ + DATABASE_UPDATE	CreateDatabase (Standby database) DeleteDatabase	DataguardAction (Switchover, Failover, Reinstate,, UpdateDatabase) Needs only DATABASE_INSPECT + DATABASE_UPDATE
MANAGE	USE + DATABASE_DELETE	DeleteDatabase	<i>none</i>

Permissions and API operation details for Key Stores

The table below lists permissions and API operations for key-stores.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	KEY_STORE_INPSECT	GetKeyStore	ChangeKeyStoreCompartment
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
USE	READ + KEY_STORE_UPDATE	UpdateKeyStore	ChangeKeyStoreCompartment
MANAGE	USE + KEY_STORE_CREATE KEY_STORE_DELETE	CreateKeyStore DeleteKeyStore	<i>none</i>

Permissions and API operation details for Application VIPs

The table below lists permissions and API operations for application-vips.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	APPLICATION_VIP_INSPECT	ListApplicationVips GetApplicationVips	<i>none</i>
read	INSPECT +	<i>no extra</i>	<i>none</i>
use	READ +	<i>no extra</i>	<i>none</i>
manage	USE + APPLICATION_VIP_CREATE APPLICATION_VIP_DELETE	CreateApplicationVips DeleteApplicationVips	<i>none</i>

Permissions and API operation details for Interim Software Updates

The table below lists permissions and API operations for oneoffPatch.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
inspect	ONEOFF_PATCH_INSPECT	DownloadOneoffPatch GetOneoffPatch ListOneoffPatches	CreateOneoffPatch DeleteOneoffPatch UpdateOneoffPatch ChangeOneoffPatchCompartment
read	INSPECT + <i>no extra</i>	DownloadOneoffPatch	<i>none</i>
use	READ + ONEOFF_PATCH_UPDATE	<i>no extra</i>	UpdateOneoffPatch ChangeOneoffPatchCompartment
manage	USE + ONEOFF_PATCH_CREATE ONEOFF_PATCH_DELETE	<i>no extra</i>	CreateOneoffPatch DeleteOneoffPatch

Related Topics

- [OneoffPatch Reference](#)

Permissions and API operation details for Scheduling Policies

The table below lists permissions and API operations for `scheduling-policies`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	SCHEDULING_POLICY_INSPECT	ListSchedulingPolicies	CreateSchedulingPolicy
		GetSchedulingPolicy	UpdateSchedulingPolicy
		ListRecommendedScheduledActions	DeleteSchedulingPolicy
			ChangeSchedulingPolicyCompartment
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ+ SCHEDULING_POLICY_UPDATE	<i>no extra</i>	UpdateSchedulingPolicy
			ChangeSchedulingPolicyCompartment
MANAGE	USE + SCHEDULING_POLICY_CREATE SCHEDULING_POLICY_DELETE	<i>no extra</i>	CreateSchedulingPolicy
			DeleteSchedulingPolicy

Permissions and API operation details for Scheduling Windows

The table below lists permissions and API operations for `scheduling-windows`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	SCHEDULING_WINDOW_INSPECT	ListSchedulingWindows	CreateSchedulingWindow
		GetSchedulingWindow	UpdateSchedulingWindow
			DeleteSchedulingWindow
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ+ SCHEDULING_WINDOW_UPDATE	<i>no extra</i>	UpdateSchedulingWindow
MANAGE	USE + SCHEDULING_WINDOW_CREATE SCHEDULING_WINDOW_DELETE	<i>no extra</i>	CreateSchedulingWindow
			DeleteSchedulingWindow

Permissions and API operation details for Scheduling Plan

The table below lists permissions and API operations for `scheduling-plan`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	CLOUD_EXADATA_INFRA STRUCTURE_INSPECT	ListCloudExadataInfrastructures GetCloudExadataInfrastructures ListSchedulingPlans GetSchedulingPlan	<i>none</i>
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ+ CLOUD_EXADATA_INFRA STRUCTURE_UPDATE	ChangeSchedulingPlanCompartment CascadingDeleteSchedulingPlan	ChangeCloudExadataInfrastructureCompartment (also needs use <code>cloud-vmclusters</code> , use <code>db-homes</code> , use <code>databases</code> , and inspect <code>db-backups</code>)
MANAGE	USE + CLOUD_EXADATA_INFRA STRUCTURE_CREATE CLOUD_EXADATA_INFRA STRUCTURE_DELETE	UpdateCloudExadataInfrastructure CreateSchedulingPlan DeleteSchedulingPlan	CreateCloudExadataInfrastructure DeleteCloudExadataInfrastructure AddStorageCapacityCloudExadataInfrastructure (also needs use <code>cloud-vmclusters</code>)

Permissions and API operation details for Scheduled Action

The table below lists permissions and API operations for `scheduled-action`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	CLOUD_EXADATA_INFRA STRUCTURE_INSPECT	ListCloudExadataInfrastructures GetCloudExadataInfrastructures ListScheduledActions GetScheduledAction ListParamsForActionType	<i>none</i>
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
USE	READ+ CLOUD_EXADATA_INFRA STRUCTURE_UPDATE	UpdateScheduledActi on ReorderScheduledAct ions	ChangeCloudExadataI nfrastructureCompar tment (also needs use cloud-vmclusters, use db-homes, use databases, and inspect db-backups)
MANAGE	USE + CLOUD_EXADATA_INFRA STRUCTURE_CREATE CLOUD_EXADATA_INFRA STRUCTURE_DELETE	UpdateCloudExadataI nfrastructure CreateScheduledActi on DeleteScheduledActi on	CreateCloudExadataI nfrastructure DeleteCloudExadataI nfrastructure AddStorageCapacityC loudExadataInfrastr ucture (also needs use cloud-vmclusters)

Permissions and API operation details for Execution Windows

The table below lists permissions and API operations for execution-windows.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	CLOUD_EXADATA_INFRA STRUCTURE_INSPECT	ListCloudExadataInf rastructures GetCloudExadataInfr astructures ListExecutionWindow s GetExecutionWindow	<i>none</i>
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ+ CLOUD_EXADATA_INFRA STRUCTURE_UPDATE	UpdateExecutionWind ow CancelExecutionWind ow	ChangeCloudExadataI nfrastructureCompar tment (also needs use cloud-vmclusters, use db-homes, use databases, and inspect db-backups)
MANAGE	USE + CLOUD_EXADATA_INFRA STRUCTURE_CREATE CLOUD_EXADATA_INFRA STRUCTURE_DELETE	UpdateCloudExadataI nfrastructure CreateExecutionWind ow DeleteExecutionWind ow	CreateCloudExadataI nfrastructure DeleteCloudExadataI nfrastructure AddStorageCapacityC loudExadataInfrastr ucture (also needs use cloud-vmclusters)

Permissions and API operation details for Execution Action

The table below lists permissions and API operations for execution-action.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	CLOUD_EXADATA_INFRA STRUCTURE_INSPECT	ListCloudExadataInfrastructures GetCloudExadataInfrastructures ListExecutionActions GetExecutionAction	<i>none</i>
READ	<i>no extra</i>	<i>no extra</i>	<i>none</i>
USE	READ+ CLOUD_EXADATA_INFRA STRUCTURE_UPDATE	UpdateExecutionAction MoveExecutionActionMember	ChangeCloudExadataInfrastructureCompartment (also needs use cloud-vmclusters, use db-homes, use databases, and inspect db-backups)
MANAGE	USE + CLOUD_EXADATA_INFRA STRUCTURE_CREATE CLOUD_EXADATA_INFRA STRUCTURE_DELETE	UpdateCloudExadataInfrastructure CreateExecutionAction DeleteExecutionAction	CreateCloudExadataInfrastructure DeleteCloudExadataInfrastructure AddStorageCapacityCloudExadataInfrastructure (also needs use cloud-vmclusters)

Permissions and API operation details for Oracle DB Azure Connectors

The table below lists permissions and API operations for `oracle-db-azure-connectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_CONNECTOR_INSPECT	ListOracleDbAzureConnectors	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_CONNECTOR_READ	GetOracleDbAzureConnector	<i>none</i>
USE	READ+ ORACLE_DB_AZURE_CONNECTOR_UPDATE	UpdateOracleDbAzureConnector	<i>none</i>
MANAGE	USE + ORACLE_DB_AZURE_CONNECTOR_CREATE ORACLE_DB_AZURE_CONNECTOR_MOVE ORACLE_DB_AZURE_CONNECTOR_DELETE	<i>no extra</i>	CreateOracleDbAzureConnector ChangeOracleDbAzureConnectorCompartment DeleteOracleDbAzureConnector

Permissions and API operation details for Oracle DB Azure Blob Mounts

The table below lists permissions and API operations for `oracle-db-azure-blob-mounts`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_BLOB_MOUNT_INSPECT	ListOracleDbAzureBlobMounts	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_BLOB_MOUNT_READ	GetOracleDbAzureBlobMount	<i>none</i>
USE	READ+ ORACLE_DB_AZURE_BLOB_MOUNT_UPDATE	UpdateOracleDbAzureBlobMount	<i>none</i>
MANAGE	USE + ORACLE_DB_AZURE_BLOB_MOUNT_CREATE ORACLE_DB_AZURE_BLOB_MOUNT_MOVE ORACLE_DB_AZURE_BLOB_MOUNT_DELETE	<i>no extra</i>	CreateOracleDbAzureBlobMount ChangeOracleDbAzureBlobMountCompartment DeleteOracleDbAzureBlobMount

Permissions and API operation details for Oracle DB Azure Blob Containers

The table below lists permissions and API operations for `oracle-db-azure-blob-containers`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_BLOB_CONTAINER_INSPECT	ListOracleDbAzureBlobContainers	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_BLOB_CONTAINER_READ	GetOracleDbAzureBlobContainer	<i>none</i>
USE	READ+ ORACLE_DB_AZURE_BLOB_CONTAINER_UPDATE	UpdateOracleDbAzureBlobContainer	<i>none</i>
MANAGE	USE + ORACLE_DB_AZURE_BLOB_CONTAINER_CREATE ORACLE_DB_AZURE_BLOB_CONTAINER_MOVE ORACLE_DB_AZURE_BLOB_CONTAINER_DELETE	<i>no extra</i>	CreateOracleDbAzureBlobContainer ChangeOracleDbAzureBlobContainerCompartment DeleteOracleDbAzureBlobContainer

Permissions and API operation details for Oracle DB MCI Work Requests

The table below lists permissions and API operations for `oracle-db-mci-work-requests`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_INSPECT	<i>no extra</i>	ListWorkRequests ListWorkRequestErrors ListWorkRequestLogs
READ	INSPECT+ ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_READ	GetWorkRequest	<i>none</i>
USE	<i>no extra</i>	<i>no extra</i>	<i>none</i>
MANAGE	USE + ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_CANCEL	CancelWorkRequest	<i>none</i>

Permissions and API operation details for Multicloud Resource Discoveries

The table below lists permissions and API operations for multi-cloud-resource-discoveries.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	MULTICLOUD_DISCOVERY_INSPECT	ListMultiCloudResourceDiscoveries	<i>none</i>
READ	INSPECT+ MULTICLOUD_DISCOVERY_READ	GetMultiCloudResourceDiscovery	<i>none</i>
USE	READ+ MULTICLOUD_DISCOVERY_UPDATE	UpdateMultiCloudResourceDiscovery	<i>none</i>
MANAGE	USE + MULTICLOUD_DISCOVERY_CREATE MULTICLOUD_DISCOVERY_MOVE MULTICLOUD_DISCOVERY_DELETE	<i>no extra</i>	CreateMultiCloudResourceDiscovery ChangeMultiCloudResourceDiscoveryCompartment DeleteMultiCloudResourceDiscovery

Permissions and API operation details for Oracle DB Azure Vaults

The table below lists permissions and API operations for oracle-db-azure-vaults.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_VAULTS_INSPECT	ListOracleDbAzureVaults	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_VAULTS_READ	GetOracleDbAzureVault	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
USE	READ+ ORACLE_DB_AZURE_VAU LT_UPDATE ORACLE_DB_AZURE_VAU LT_REFRESH	<i>no extra</i>	UpdateOracleDbAzure Vault RefreshOracleDbAzur eVault
MANAGE	USE + ORACLE_DB_AZURE_VAU LT_CREATE ORACLE_DB_AZURE_VAU LT_MOVE ORACLE_DB_AZURE_VAU LT_DELETE	<i>no extra</i>	CreateOracleDbAzure Vault ChangeOracleDbAzure VaultCompartment DeleteOracleDbAzure Vault

Permissions and API operation details for Oracle DB Azure Keys

The table below lists permissions and API operations for `oracle-db-azure-keys`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_KEY _INSPECT	ListOracleDbAzureKe ys	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_KEY _READ	GetOracleDbAzureKey	<i>none</i>
USE	<i>no extra</i>	<i>no extra</i>	<i>none</i>
MANAGE	<i>no extra</i>	<i>no extra</i>	<i>none</i>

Permissions and API operation details for Oracle DB Azure Vault Associations

The table below lists permissions and API operations for `oracle-db-azure-vault-associations`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AZURE_ASS OCIATION_INSPECT	ListOracleDbAzureVa ultAssociations	<i>none</i>
READ	INSPECT+ ORACLE_DB_AZURE_ASS OCIATION_READ	GetOracleDbAzureVau ltAssociation	<i>none</i>
USE	READ+ ORACLE_DB_AZURE_ASS OCIATION_UPDATE	UpdateOracleDbAzure VaultAssociation	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	ORACLE_DB_AZURE_ASSOCIATION_CREATE ORACLE_DB_AZURE_ASSOCIATION_MOVE ORACLE_DB_AZURE_ASSOCIATION_DELETE	<i>no extra</i>	CreateOracleDbAzureVaultAssociation ChangeOracleDbAzureVaultAssociationCompartments DeleteOracleDbAzureVaultAssociation CascadingDeleteOracleDbAzureVaultAssociation

Permissions and API operation details for Oracle DB GCP Connectors

The table below lists permissions and API operations for `oracle-db-gcp-connectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_GCP_IDENTITY_CONNECTOR_INSPECT	ListOracleDbGcpIdentityConnectors	<i>none</i>
READ	INSPECT+ ORACLE_DB_GCP_IDENTITY_CONNECTOR_READ	GetOracleDbGcpIdentityConnector	<i>none</i>
USE	READ+ ORACLE_DB_GCP_IDENTITY_CONNECTOR_UPDATE	UpdateOracleDbGcpIdentityConnector	<i>none</i>
MANAGE	USE+ ORACLE_DB_GCP_KEYRING_CREATE ORACLE_DB_GCP_KEYRING_MOVE ORACLE_DB_GCP_KEYRING_DELETE	<i>no extra</i>	CreateOracleDbGcpIdentityConnector ChangeOracleDbGcpIdentityConnectorCompartments DeleteOracleDbGcpIdentityConnector

Permissions and API operation details for Oracle DB GCP Keyrings

The table below lists permissions and API operations for `oracle-db-gcp-keyrings`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_GCP_KEYRING_INSPECT	ListOracleDbGcpKeyRings	<i>none</i>
READ	INSPECT+ ORACLE_DB_GCP_KEYRING_READ	GetOracleDbGcpKeyRing	<i>none</i>

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
USE	READ+	UpdateOracleDbGcpKe	<i>none</i>
	ORACLE_DB_GCP_KEY_R	yRing	
	ING_UPDATE	RefreshOracleDbGcpK	
	ORACLE_DB_GCP_KEY_R	eyRing	
MANAGE	USE+	<i>no extra</i>	CreateOracleDbGcpKe
	ORACLE_DB_GCP_KEY_R		yRing
	ING_CREATE		ChangeOracleDbGcpKe
	ORACLE_DB_GCP_KEY_R		yRingCompartment
	ING_MOVE		DeleteOracleDbGcpKe
	ORACLE_DB_GCP_KEY_R		yRing
	ING_DELETE		

Permissions and API operation details for Oracle DB GCP Keys

The table below lists permissions and API operations for `oracle-db-gcp-keys`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_GCP_KEY_I	ListOracleDbGcpKeys	<i>none</i>
	NSPECT		
READ	INSPECT+	GetOracleDbGcpKey	<i>none</i>
	ORACLE_DB_GCP_KEY_R		
	EAD		
USE	READ+	<i>no extra</i>	<i>none</i>
	<i>no extra</i>		
MANAGE	USE+	<i>no extra</i>	<i>none</i>
	<i>no extra</i>		

Permissions and API operation details for Oracle DB AWS Identity Connectors

The table below lists permissions and API operations for `OracleDbAwsIdentityConnectors`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AWS_IDENT	ListOracleDbAwsIden	<i>none</i>
	ITY_CONNECTOR_INSPE	tityConnectors	
	CT		
READ	INSPECT+	GetOracleDbAwsIdent	<i>none</i>
	ORACLE_DB_AWS_IDENT	ityConnector	
	ITY_CONNECTOR_READ		
USE	READ+	<i>no extra</i>	UpdateOracleDbAwsId
	ORACLE_DB_AWS_IDENT		entityConnector
	ITY_CONNECTOR_UPDAT		RefreshOracleDbAwsI
	E		entityConnector

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
MANAGE	USE+ ORACLE_DB_AWS_IDENT ITY_CONNECTOR_CREAT E ORACLE_DB_AWS_IDENT ITY_CONNECTOR_MOVE ORACLE_DB_AWS_IDENT ITY_CONNECTOR_DELET E	<i>no extra</i>	CreateOracleDbAwsId entityConnector ChangeOracleDbAwsId entityConnectorComp artment DeleteOracleDbAwsId entityConnector

Permissions and API operation details for Oracle DB AWS Key

The table below lists permissions and API operations for `OracleDbAwsKey`.

Verbs	Permissions	APIs Fully Covered	APIs Partially Covered
INSPECT	ORACLE_DB_AWS_KEY_I NSPECT	ListOracleDbAwsKeys	<i>none</i>
READ	INSPECT+ ORACLE_DB_AWS_KEY_R EAD	GetOracleDbAwsKey	<i>none</i>
USE	READ+ ORACLE_DB_AWS_KEY_U PDATE	<i>no extra</i>	UpdateOracleDbAwsKe Y RefreshOracleDbAwsK ey
MANAGE	USE+ ORACLE_DB_AWS_KEY_C REATE ORACLE_DB_AWS_KEY_M OVE ORACLE_DB_AWS_KEY_D ELETE	<i>no extra</i>	CreateOracleDbAwsKe Y ChangeOracleDbAwsKe yCompartment DeleteOracleDbAwsKe Y

Permissions Required for Each API Operation

Database API Operations

For information about permissions, see:

[Permissions](#).

The following tables list of API operations and permissions by API operation.

Table 6-7 Cloud Exadata Infrastructure Resource

API Operation	Permissions Required to Use the Operation
ListCloudExadataInfrastructures	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT

Table 6-7 (Cont.) Cloud Exadata Infrastructure Resource

API Operation	Permissions Required to Use the Operation
CreateCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_CREATE
UpdateCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
ChangeCloudExadataInfrastructureCompartment	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
DeleteCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_DELETE
AddStorageCapacityCloudExadataInfrastructure	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
ListSchedulingPolicies	SCHEDULING_POLICY_INSPECT
GetSchedulingPolicy	SCHEDULING_POLICY_INSPECT
UpdateSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
ChangeSchedulingPolicyCompartment	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
ListRecommendedScheduledActions	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_UPDATE
CreateSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_CREATE
DeleteSchedulingPolicy	SCHEDULING_POLICY_INSPECT & SCHEDULING_POLICY_DELETE
ListSchedulingWindows	SCHEDULING_WINDOW_INSPECT
GetSchedulingWindow	SCHEDULING_WINDOW_INSPECT
UpdateSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_UPDATE
CreateSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_CREATE
DeleteSchedulingWindow	SCHEDULING_WINDOW_INSPECT & SCHEDULING_WINDOW_DELETE
ListSchedulingPlans	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetSchedulingPlan	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
CreateSchedulingPlan	CLOUD_EXADATA_INFRASTRUCTURE_CREATE
ChangeSchedulingPlanCompartment	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
ReorderScheduledActions	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
CascadingDeleteSchedulingPlan	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
DeleteSchedulingPlan	CLOUD_EXADATA_INFRASTRUCTURE_DELETE
ListScheduledActions	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetScheduledAction	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
ListParamsForActionType	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
ReorderScheduledActions	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
CreateScheduledAction	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
DeleteScheduledAction	CLOUD_EXADATA_INFRASTRUCTURE_DELETE
ListExecutionWindows	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetExecutionWindow	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
UpdateExecutionWindow	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE

Table 6-7 (Cont.) Cloud Exadata Infrastructure Resource

API Operation	Permissions Required to Use the Operation
ReorderExecutionActions	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
CancelExecutionWindow	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
CreateExecutionWindow	CLOUD_EXADATA_INFRASTRUCTURE_CREATE
DeleteExecutionWindow	CLOUD_EXADATA_INFRASTRUCTURE_DELETE
ListExecutionActions	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
GetExecutionAction	CLOUD_EXADATA_INFRASTRUCTURE_INSPECT
UpdateExecutionAction	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
MoveExecutionActionMember	CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
CreateExecutionAction	CLOUD_EXADATA_INFRASTRUCTURE_CREATE
DeleteExecutionAction	CLOUD_EXADATA_INFRASTRUCTURE_DELETE

Table 6-8 Cloud VM Cluster

API Operation	Permissions Required to Use the Operation
ListCloudVmClusters	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmCluster	CLOUD_VM_CLUSTER_INSPECT
CreateCloudVmCluster	CLOUD_VM_CLUSTER_CREATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and VNIC_CREATE and VNIC_ATTACH and SUBNET_ATTACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE DNS_VIEW_INSPECT)
ChangeCloudVmClusterCompartment	CLOUD_VM_CLUSTER_UPDATE
UpdateCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE
GetCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_INSPECT
UpdateCloudVmClusterIormConfig	CLOUD_VM_CLUSTER_UPDATE
DeleteCloudVmCluster	CLOUD_VM_CLUSTER_DELETE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and DB_HOME_DELETE and VNIC_DELETE and SUBNET_DETACH and VNIC_DETACH and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)
AddVmToCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_CREATE, DNS_VIEW_INSPECT)
RemoveVmFromCloudVmCluster	CLOUD_VM_CLUSTER_UPDATE and CLOUD_EXADATA_INFRASTRUCTURE_UPDATE and (needed if Private DNS is used: DNS_ZONE_READ, DNS_RECORD_UPDATE, DNS_ZONE_DELETE)

Table 6-9 Cloud VM Cluster Maintenance Updates and Update History

API Operation	Permissions Required to Use the Operation
ListCloudVmClusterUpdates	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmClusterUpdate	CLOUD_VM_CLUSTER_INSPECT
ListCloudVmClusterUpdateHistoryEntries	CLOUD_VM_CLUSTER_INSPECT
GetCloudVmClusterUpdateHistoryEntry	CLOUD_VM_CLUSTER_INSPECT

Table 6-10 Virtual Machines / Nodes

API Operation	Permissions Required to Use the Operation
ListDbNodes	DB_NODE_INSPECT
GetDbNode	DB_NODE_INSPECT
DbNodeAction	DB_NODE_POWER_ACTIONS

Table 6-11 Database Homes

API Operation	Permissions Required to Use the Operation
ListDbHomes	DB_HOME_INSPECT
GetDbHome	DB_HOME_INSPECT
ListDbHomePatches	DB_HOME_INSPECT
ListDbHomePatchHistoryEntries	DB_HOME_INSPECT
GetDbHomePatch	DB_HOME_INSPECT
GetDbHomePatchHistoryEntry	DB_HOME_INSPECT
CreateDbHome	DB_SYSTEM_INSPECT and DB_SYSTEM_UPDATE and DB_HOME_CREATE and DATABASE_CREATE To enable automatic backups for the database, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
UpdateDbHome	DB_HOME_UPDATE
DeleteDbHome	DB_SYSTEM_UPDATE and DB_HOME_DELETE and DATABASE_DELETE If automatic backups are enabled, also need DELETE_BACKUP If performing a final backup on termination, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ

Table 6-12 Databases (CDB)

API Operation	Permissions Required to Use the Operation
ListDatabases	DATABASE_INSPECT
GetDatabase	DATABASE_INSPECT

Table 6-12 (Cont.) Databases (CDB)

API Operation	Permissions Required to Use the Operation
CreateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
UpdateDatabase	DATABASE_UPDATE To enable automatic backups, also need DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteDatabase	For new resource model using VM cluster resource: CLOUD_VM_CLUSTER_INSPECT and DB_HOME_UPDATE and DATABASE_DELETE
enableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE

Table 6-13 Pluggable Databases (PDBs)

API Operation	Permissions Required to Use the Operation
ListPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
GetPluggableDatabase	PLUGGABLE_DATABASE_INSPECT
CreatePluggableDatabase	PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
UpdatePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StartPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
StopPluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE
DeletePluggableDatabase	PLUGGABLE_DATABASE_DELETE and DATABASE_INSPECT and DATABASE_UPDATE
LocalClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
RemoteClonePluggableDatabase	PLUGGABLE_DATABASE_INSPECT and PLUGGABLE_DATABASE_UPDATE and PLUGGABLE_DATABASE_CONTENT_READ and PLUGGABLE_DATABASE_CONTENT_WRITE and PLUGGABLE_DATABASE_CREATE and DATABASE_INSPECT and DATABASE_UPDATE
enableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE
disableDatabaseManagement	DATABASE_INSPECT and DATABASE_UPDATE

Table 6-14 System Shapes and Database Versions

API Operation	Permissions Required to Use the Operation
ListDbSystemShapes	(no permissions required; available to anyone)
ListDbVersions	(no permissions required; available to anyone)

Table 6-15 Oracle Data Guard Associations

API Operation	Permissions Required to Use the Operation
GetDataGuardAssociation	DATABASE_INSPECT
ListDataGuardAssociations	DATABASE_INSPECT
CreateDataGuardAssociation	DB_SYSTEM_UPDATE and DB_HOME_CREATE and DB_HOME_UPDATE and DATABASE_CREATE and DATABASE_UPDATE
SwitchoverDataGuardAssociation	DATABASE_UPDATE
FailoverDataGuardAssociation	DATABASE_UPDATE
ReinstateDataGuardAssociation	DATABASE_UPDATE

Table 6-16 Backups and Database Restore

API Operation	Permissions Required to Use the Operation
GetBackup	DB_BACKUP_INSPECT
ListBackups	DB_BACKUP_INSPECT
CreateBackup	DB_BACKUP_CREATE and DATABASE_CONTENT_READ
DeleteBackup	DB_BACKUP_DELETE and DB_BACKUP_INSPECT
RestoreDatabase	DB_BACKUP_INSPECT and DB_BACKUP_CONTENT_READ and DATABASE_CONTENT_WRITE

Table 6-17 Application VIP

API Operation	Permissions Required to Use the Operation
CreateApplicationVip	APPLICATION_VIP_CREATE and CLOUD_VM_CLUSTER_UPDATE and PRIVATE_IP_CREATE and PRIVATE_IP_ASSIGN and VNIC_ASSIGN and SUBNET_ATTACH
DeleteApplicationVip	APPLICATION_VIP_DELETE and CLOUD_VM_CLUSTER_UPDATE and PRIVATE_IP_DELETE and PRIVATE_IP_UNASSIGN and VNIC_UNASSIGN and SUBNET_DETACH
ListApplicationVips	APPLICATION_VIP_INSPECT
ListApplicationVips	APPLICATION_VIP_INSPECT

Table 6-18 Serial Console Access to VM

API Operation	Permissions Required to Use the Operation
AddVirtualMachineToVmCluster	VM_CLUSTER_UPDATE and EXADATA_INFRASTRUCTURE_UPDATE
RemoveVirtualMachineFromVmCluster	VM_CLUSTER_UPDATE and EXADATA_INFRASTRUCTURE_UPDATE
CreateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_CREATE and DBNODE_CONSOLE_CONNECTION_INSPECT
GetDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_INSPECT
ListDbNodeConsoleConnections	DBNODE_CONSOLE_CONNECTION_INSPECT
DeleteDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_DELETE
UpdateDbNodeConsoleConnection	DBNODE_CONSOLE_CONNECTION_UPDATE
UpdateDbNode	DB_NODE_UPDATE

Table 6-19 Oracle DB Azure Connector Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureConnectors	ORACLE_DB_AZURE_CONNECTOR_INSPECT
GetOracleDbAzureConnector	ORACLE_DB_AZURE_CONNECTOR_READ
CreateOracleDbAzureConnector	ORACLE_DB_AZURE_CONNECTOR_CREATE
UpdateOracleDbAzureConnector	ORACLE_DB_AZURE_CONNECTOR_UPDATE
ChangeOracleDbAzureConnectorCompartment	ORACLE_DB_AZURE_CONNECTOR_MOVE
DeleteOracleDbAzureConnector	ORACLE_DB_AZURE_CONNECTOR_DELETE

Table 6-20 Oracle DB Azure Blob Container Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureBlobContainers	ORACLE_DB_AZURE_BLOB_CONTAINER_INSPECT
CreateOracleDbAzureBlobContainer	ORACLE_DB_AZURE_BLOB_CONTAINER_CREATE
ChangeOracleDbAzureBlobContainerCompartment	ORACLE_DB_AZURE_BLOB_CONTAINER_MOVE
GetOracleDbAzureBlobContainer	ORACLE_DB_AZURE_BLOB_CONTAINER_READ
UpdateOracleDbAzureBlobContainer	ORACLE_DB_AZURE_BLOB_CONTAINER_UPDATE
DeleteOracleDbAzureBlobContainer	ORACLE_DB_AZURE_BLOB_CONTAINER_DELETE

Table 6-21 Oracle DB Azure Blob Mount Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureBlobMounts	ORACLE_DB_AZURE_BLOB_MOUNT_INSPECT
CreateOracleDbAzureBlobMount	ORACLE_DB_AZURE_BLOB_MOUNT_CREATE
ChangeOracleDbAzureBlobMountCompartment	ORACLE_DB_AZURE_BLOB_MOUNT_MOVE
GetOracleDbAzureBlobMount	ORACLE_DB_AZURE_BLOB_MOUNT_READ
UpdateOracleDbAzureBlobMount	ORACLE_DB_AZURE_BLOB_MOUNT_UPDATE
DeleteOracleDbAzureBlobMount	ORACLE_DB_AZURE_BLOB_MOUNT_DELETE

Table 6-22 Work Request Resource

API Operation	Permissions Required to Use the Operation
ListWorkRequests	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_INSPECT
GetWorkRequest	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_READ
CancelWorkRequest	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_CANCEL
ListWorkRequestErrors	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_INSPECT
ListWorkRequestLogs	ORACLE_DB_MULTI_CLOUD_WORK_REQUEST_INSPECT

Table 6-23 MultiCloudResourceDiscovery Resource

API Operation	Permissions Required to Use the Operation
ListMultiCloudResourceDiscoveries	MULTICLOUD_DISCOVERY_INSPECT
CreateMultiCloudResourceDiscovery	MULTICLOUD_DISCOVERY_CREATE
ChangeMultiCloudResourceDiscoveryCompartment	MULTICLOUD_DISCOVERY_MOVE
GetMultiCloudResourceDiscovery	MULTICLOUD_DISCOVERY_READ
UpdateMultiCloudResourceDiscovery	MULTICLOUD_DISCOVERY_UPDATE
DeleteMultiCloudResourceDiscovery	MULTICLOUD_DISCOVERY_DELETE

Table 6-24 OracleDbAzureVault Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureVaults	ORACLE_DB_AZURE_VAULT_INSPECT
CreateOracleDbAzureVault	ORACLE_DB_AZURE_VAULT_CREATE
ChangeOracleDbAzureVaultCompartment	ORACLE_DB_AZURE_VAULT_MOVE
RefreshOracleDbAzureVault	ORACLE_DB_AZURE_VAULT_REFRESH
GetOracleDbAzureVault	ORACLE_DB_AZURE_VAULT_READ
UpdateOracleDbAzureVault	ORACLE_DB_AZURE_VAULT_UPDATE
DeleteOracleDbAzureVault	ORACLE_DB_AZURE_VAULT_DELETE

Table 6-25 OracleDbAzureKey Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureKeys	ORACLE_DB_AZURE_KEY_INSPECT
GetOracleDbAzureKey	ORACLE_DB_AZURE_KEY_READ

Table 6-26 OracleDbAzureVaultAssociation Resource

API Operation	Permissions Required to Use the Operation
ListOracleDbAzureVaultAssociations	ORACLE_DB_AZURE_ASSOCIATION_INSPECT

Table 6-26 (Cont.) OracleDbAzureVaultAssociation Resource

API Operation	Permissions Required to Use the Operation
CreateOracleDbAzureVaultAssociation	ORACLE_DB_AZURE_ASSOCIATION_CREATE
ChangeOracleDbAzureVaultAssociationCompartment	ORACLE_DB_AZURE_ASSOCIATION_MOVE
GetOracleDbAzureVaultAssociation	ORACLE_DB_AZURE_ASSOCIATION_READ
UpdateOracleDbAzureVaultAssociation	ORACLE_DB_AZURE_ASSOCIATION_UPDATE
DeleteOracleDbAzureVaultAssociation	ORACLE_DB_AZURE_ASSOCIATION_DELETE
CascadingDeleteOracleDbAzureVaultAssociation	ORACLE_DB_AZURE_ASSOCIATION_DELETE
createDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_CREATE &DBNODE_CONSOLE_HISTORY_INSPECT
getDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_INSPECT
getDbNodeConsoleHistoryContent	DBNODE_CONSOLE_HISTORY_CONTENT_READ
listDbNodeConsoleHistories	DBNODE_CONSOLE_HISTORY_INSPECT
updateDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_UPDATE
deleteDbNodeConsoleHistory	DBNODE_CONSOLE_HISTORY_DELETE

Table 6-27 OracleDbGcpIdentityConnectors

API Operation	Permissions Required to Use the Operation
ListOracleDbGcpIdentityConnectors	ORACLE_DB_GCP_IDENTITY_CONNECTOR_INSPECT
GetOracleDbGcpIdentityConnector	ORACLE_DB_GCP_IDENTITY_CONNECTOR_READ
CreateOracleDbGcpIdentityConnector	ORACLE_DB_GCP_IDENTITY_CONNECTOR_CREATE
UpdateOracleDbGcpIdentityConnector	ORACLE_DB_GCP_IDENTITY_CONNECTOR_UPDATE
ChangeOracleDbGcpIdentityConnectorCompartment	ORACLE_DB_GCP_IDENTITY_CONNECTOR_MOVE
DeleteOracleDbGcpIdentityConnector	ORACLE_DB_GCP_IDENTITY_CONNECTOR_DELETE

Table 6-28 OracleDbGcpKeyRings

API Operation	Permissions Required to Use the Operation
ListOracleDbGcpKeyRings	ORACLE_DB_GCP_KEY_RING_INSPECT
CreateOracleDbGcpKeyRing	ORACLE_DB_GCP_KEY_RING_CREATE
ChangeOracleDbGcpKeyRingCompartment	ORACLE_DB_GCP_KEY_RING_MOVE
RefreshOracleDbGcpKeyRing	ORACLE_DB_GCP_KEY_RING_REFRESH
GetOracleDbGcpKeyRing	ORACLE_DB_GCP_KEY_RING_READ
UpdateOracleDbGcpKeyRing	ORACLE_DB_GCP_KEY_RING_UPDATE
DeleteOracleDbGcpKeyRing	ORACLE_DB_GCP_KEY_RING_DELETE

Table 6-29 OracleDbGcpKeyKeys

API Operation	Permissions Required to Use the Operation
ListOracleDbGcpKeys	ORACLE_DB_GCP_KEY_INSPECT
GetOracleDbGcpKey	ORACLE_DB_GCP_KEY_READ

Table 6-30 OracleDbAwsIdentityConnectors

API Operation	Permissions Required to Use the Operation
ListOracleDbAwsIdentityConnectors	ORACLE_DB_AWS_IDENTITY_CONNECTOR_INSPECT
GetOracleDbAwsIdentityConnector	ORACLE_DB_AWS_IDENTITY_CONNECTOR_READ
CreateOracleDbAwsIdentityConnector	ORACLE_DB_AWS_IDENTITY_CONNECTOR_CREATE
UpdateOracleDbAwsIdentityConnector	ORACLE_DB_AWS_IDENTITY_CONNECTOR_UPDATE
ChangeOracleDbAwsIdentityConnectorCompartment	ORACLE_DB_AWS_IDENTITY_CONNECTOR_MOVE
DeleteOracleDbAwsIdentityConnector	ORACLE_DB_AWS_IDENTITY_CONNECTOR_DELETE
RefreshOracleDbAwsIdentityConnector	ORACLE_DB_AWS_IDENTITY_CONNECTOR_UPDATE

Table 6-31 OracleDbAwsKey

API Operation	Permissions Required to Use the Operation
ListOracleDbAwsKeys	ORACLE_DB_AWS_KEY_INSPECT
CreateOracleDbAwsKey	ORACLE_DB_AWS_KEY_CREATE
ChangeOracleDbAwsKeyCompartment	ORACLE_DB_AWS_KEY_MOVE
GetOracleDbAwsKey	ORACLE_DB_AWS_KEY_READ
UpdateOracleDbAwsKey	ORACLE_DB_AWS_KEY_UPDATE
DeleteOracleDbAwsKey	ORACLE_DB_AWS_KEY_DELETE
RefreshOracleDbAwsKey	ORACLE_DB_AWS_KEY_UPDATE

Managing Exadata Resources with Oracle Enterprise Manager Cloud Control

To manage and monitor Exadata Cloud Infrastructure and Exadata Database Service on Cloud@Customer resources, use Oracle Enterprise Manager Cloud Control.

For complete documentation and Oracle By Example tutorials, see the following documentation resources: *Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud* and *Setting Up Oracle Enterprise Manager 13.4 on Oracle Cloud Infrastructure*.

- [Overview of Oracle Enterprise Manager Cloud Control](#)
Oracle Enterprise Manager Cloud Control provides a complete lifecycle management solution for Oracle Cloud Infrastructure's Exadata Cloud Infrastructure (ExaDB-D) and Exadata Database Service on Cloud@Customer (ExaDB-C@C) services.
- [Features of Enterprise Manager Cloud Control](#)
Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Exadata Cloud and Exadata Cloud@Customer resources.

- [Analyzing Exadata Database Service Database Performance](#)
This topic describes how to use Database Metrics and Performance Hub to monitor, analyze, and tune the performance of OCI user-managed databases, including Oracle Exadata Database Service on Dedicated Infrastructure databases.

Related Topics

- [Oracle Enterprise Manager Cloud Control for Oracle Exadata Cloud](#)
- [Setting Up Oracle Enterprise Manager 13.4 on Oracle Cloud Infrastructure](#)

Overview of Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control provides a complete lifecycle management solution for Oracle Cloud Infrastructure's Exadata Cloud Infrastructure (ExaDB-D) and Exadata Database Service on Cloud@Customer (ExaDB-C@C) services.

Enterprise Manager Cloud Control discovers ExaDB-D and ExaDB-C@C services as a single target and automatically identifies and organizes all dependent components. Using Enterprise Manager Cloud Control you can then:

- Monitor and manage all Exadata, ExaDB-D and ExaDB-C@C systems, along with any other targets, from a single interface
- Visualize storage and compute data
- View performance metrics of your Exadata components

Features of Enterprise Manager Cloud Control

Familiarize yourself with the features of Enterprise Manager Cloud Control to manage and monitor Exadata Cloud and Exadata Cloud@Customer resources.

Enterprise Manager Target for Exadata Cloud

The target for Oracle Cloud Infrastructure Exadata resources, which covers both Exadata Cloud and Exadata Cloud@Customer does the following:

- Automatically identifies and organizes related targets.
- Provides a high-level integration point for Enterprise Manager framework features such as incident rules, groups, notifications, and monitoring templates.

Improved Performance Monitoring

Enterprise Manager Cloud Control enhances performance monitoring in the following ways:

- Adds Exadata Storage Server and Exadata Storage Grid targets.
- Offers visualization of storage and compute performance for your Exadata Cloud and Exadata Cloud@Customer resources.
- Enables use of the same Maximum Availability Architecture (MAA) key performance indicators (KPI) developed for Oracle Exadata Database Machine.

Scripted CLI-based Discovery

Enterprise Manager Cloud Control uses scripts to discover Oracle Cloud Infrastructure Exadata resources. Scripts search the existing hosts, clusters, ASM, databases and related targets, and add the storage server targets.

"Single Pane of Glass" View of On-Premises and Oracle Cloud Infrastructure Exadata Resources

Enterprise Manager Cloud Control 's use of a single Exadata target type provides a consistent Enterprise Manager experience across on-premises, Exadata Cloud, and Exadata Cloud@Customer resources. The common Exadata target menu allows you to easily navigate to, monitor and manage all of your Exadata systems.

Visualization

Enterprise Manager Cloud Control allows you to visualize the database and related targets associated with each Exadata Cloud and Exadata Cloud@Customer system.

Analyzing Exadata Database Service Database Performance

This topic describes how to use Database Metrics and Performance Hub to monitor, analyze, and tune the performance of OCI user-managed databases, including Oracle Exadata Database Service on Dedicated Infrastructure databases.

With this tool, you can view real-time and historical performance data. For information about using Performance Hub, see [Using Performance Hub to Analyze Database Performance](#).

To use Database Metrics and Performance Hub for Exadata Cloud Infrastructure, Database Management must be enabled for the database. When enabling a database, the database administrator can choose from two database management options: Basic Management and Full Management.

Note

Using Identity and Access Management (IAM), you can create a policy that grants users access to Performance Hub while limiting actions they can take on Autonomous Databases, Oracle Database Cloud Service, and external databases. For information about IAM policies and Exadata Cloud Infrastructure databases, see *Required IAM Policy*. For information about policies and how to use them, see [How Policies Work](#).

Related Topics

- [Required IAM Policy for Exadata Cloud Infrastructure](#)
Review the identity access management (IAM) policy for provisioning Oracle Exadata Database Service on Dedicated Infrastructure systems.

Observability and Management for Exadata Database Service on Dedicated Infrastructure

- [Metrics for Exadata Cloud Infrastructure in the Database Management Service](#)
Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.
- [Oracle Cloud Infrastructure Operations Insights](#)
Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Exadata Cloud Infrastructure.

- [Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases](#)
Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.

Metrics for Exadata Cloud Infrastructure in the Database Management Service

Database Management provides comprehensive database performance diagnostics and management capabilities for Oracle Databases.

This article describes the metrics emitted by the Exadata Cloud Infrastructure Database service in the `oracle_oci_database` namespace for Oracle Databases.

To use database metrics for these Oracle Databases in Exadata Cloud Infrastructure Database Service, you must enable Database Management for the database you want to monitor. You can enable either Basic Management or Full Management for your database. See [Enable Database Management](#) for instructions.

Dimensions

All the metrics discussed in this topic include the following dimensions.

- RESOURCEID - The OCID of the database.
- RESOURCENAME - The name of the database.
- DEPLOYMENTTYPE - The deployment type of the database.

Note: Valid alarm intervals are 5 minutes or greater due to the frequency at which these metrics are emitted. See [To create an alarm](#) for details on creating alarms.

The database metrics can be provided for the basic and full Database Management options.

NOT_SUPPORTED

The metrics listed in the following table are automatically available for Oracle Databases when the **Basic Database Management** option is enabled.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
BlockChanges	DB Block Changes	changes per second	The average number of blocks changed per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
CpuUtilization	CPU Utilization	percent	The CPU utilization expressed as a percentage, aggregated across all consumer groups. The utilization percentage is reported with respect to the number of CPUs the database is allowed to use, which is two times the number of OCPUs. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
CurrentLogons	Current Logons	count	The number of successful logons during the selected interval. Statistics: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
ExecuteCount	Execute Count	count	The number of user and recursive calls that executed SQL statements during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
OcpusAllocated	OCPU Allocated	integer	The actual number of OCPUs allocated by the service during the selected interval of time. Statistic: Count Interval: 1 minute	5 minutes	N/A

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
ParseCount	Parse Count (Total)	count	The number of hard and soft parses during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
StorageAllocated	Allocated Storage Space	GB	The maximum amount of space allocated by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	N/A
StorageAllocatedByTablespace	Allocated Storage Space By Tablespace	GB	The maximum amount of space allocated by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
StorageUsed	Storage Space Used	GB	The maximum amount of space used during the interval. Statistic: Max Interval: 30 minutes	30 minutes	N/A

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
StorageUsedByTablespace	Storage Space Used By Tablespace	GB	The maximum amount of space used by tablespace during the interval. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
StorageUtilization	Storage Utilization	percent	The percentage of provisioned storage capacity currently in use. Represents the total allocated space for all tablespaces. Statistic: Mean Interval: 30 minutes	30 minutes	N/A
StorageUtilizationByTablespace	Storage Space Utilization By Tablespace	percent	The percentage of the space utilized, by tablespace. For container databases, this metric provides data for root container tablespaces. Statistic: mean Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
TransactionCount	Transaction Count	count	The combined number of user commits and user rollbacks during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
UserCalls	User Calls	count	The combined number of logons, parses, and execute calls during the selected interval. Statistic: Sum Interval: 1 minute	5 minutes	instanceNumber instanceName hostName

NOT_SUPPORTED

The metrics listed in the following table are automatically available for Oracle Databases when the **Full Database Management** option is enabled.

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
AllocatedStorageUtilizationByTablespace	Allocated Space Utilization By Tablespace	percent	The percentage of space used by tablespace, out of all allocated. For container databases, this metric provides data for root container tablespaces. Statistic: Mean Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
AvgGCCRBlockReceiveTime	Average GC CR Block Receive Time	milliseconds	The average global cache CR (consistent-read) block receive time. Statistic: Mean Interval: 5 minutes <i>For RAC / cluster databases only.</i>	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
BlockingSessions	Blocking Sessions	count	Current blocking sessions. Statistic: Max Interval: 15 minutes <i>Not applicable for container databases.</i>	15 minutes	N/A
CPUTime	CPU Time	seconds per second	The average rate of accumulation of CPU time by foreground sessions in the database instance over the time interval. The CPU time component of Average Active Sessions. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
DbmgmtJobExecutionsCount	??	??	The number of SQL job executions on a single managed database or a database group, and their status. Status dimensions can be the following values: "Succeeded," "Failed," "InProgress." Statistic: Sum Interval: 1 minute	??	managedDatabaseId managedDatabaseGroupId jobId status

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
DBTime	DB Time	seconds per second	The average rate of accumulation of database time (CPU + Wait) by foreground sessions in the database instance over the time interval. Also known as Average Active Sessions. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
FRASpaceLimit	Flash Recovery Area Limit	GB	The flash recovery area space limit. Statistic: Max Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
FRAUtilization	Flash Recovery Area Utilization	percent	The flash recovery area utilization. Statistic: Mean Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
GCCRBlocksReceived	GC CR Blocks Received	blocks per second	The global cache CR (consistent-read) blocks received per second. Statistic: Mean Interval: 5 minutes <i>For RAC / cluster databases only.</i>	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
GCCurrentBlocksReceived	GC Current Blocks Received	blocks per second	Represents global cache current blocks received per second. Statistic reports the mean value. Statistic: Mean Interval: 5 minutes <i>For Real Application Cluster (RAC) databases only.</i>	5 minutes	instanceNumber instanceName hostName
InterconnectTraffic	Average Interconnect Traffic	MB per second	The average internode data transfer rate. Statistic: Mean Interval: 5 minutes <i>For RAC / cluster databases only.</i>	5 minutes	instanceNumber instanceName hostName
InvalidObjects	Invalid Objects	count	Invalid database objects count. Statistic: Max Interval: 24 hours <i>Not applicable for container databases.</i>	24 hours	N/A
IOPS	IOPS	operations per second	The average number of input-output operations per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName ioType (Read, Write)
IOThroughput	IO Throughput	MB per second	The average throughput in MB per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName ioType (Read, Write)

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
LogicalBlocksRead	Logical Reads	reads per second	The average number of blocks read from SGA/Memory (buffer cache) per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
MaxTablespaceSize	Max Tablespace Size	GB	The maximum possible tablespace size. For container databases, this metric provides data for root container tablespaces. Statistic: Max Interval: 30 minutes	30 minutes	tablespaceName tablespaceType
MemoryUsage	Memory Usage	MB	Memory pool total size in MB. Statistic: Mean Interval: 15 minutes	15 minutes	instanceNumber instanceName hostName memoryType (SGA, PGA) memoryPool (AllocatedPGA, Buffercache, FixedSGA, JavaPool, LargePool, LogBuffer, OtherPools, SharedPool, StreamsPool)
MonitoringStatus	Monitoring Status	not applicable	The monitoring status of the resource. If a metric collection fails, error information is captured in this metric. Statistic: Mean Interval: 5 minutes	5 minutes	collectionName errorSeverity errorCode

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
NonReclaimableFRA	Non Reclaimable Fast Recovery Area	percent	The Non-reclaimable fast recovery area. Statistic: Mean Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
ParsesByType	Parses By Type	parses per second	The number of hard or soft parses per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName parseType (HardParse, SoftParse)
ProblematicScheduledDBMSJobs	Problematic Scheduled DBMS Jobs	count	The problematic scheduled database jobs count. Statistic: Max Interval: 15 minutes <i>Not applicable for container databases.</i>	15 minutes	type (Broken, Failed)
Processes	Process Count	count	The database processes count. Statistic: Max Interval: 1 minute <i>Not applicable for pluggable databases.</i>	5 minutes	instanceNumber instanceName hostName
ProcessLimitUtilization	Process Limit Utilization	percent	The process limit utilization. Statistic: Mean Interval: 1 minute <i>Not applicable for pluggable databases.</i>	5 minutes	instanceNumber instanceName hostName

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
ReclaimableFRA	Reclaimable Fast Recovery Area	percent	The reclaimable fast recovery area. Statistic: Mean Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
ReclaimableFRASpace	Flash Recovery Area Reclaimable Space	GB	The flash recovery area reclaimable space. Statistic: Mean Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
RedoSize	Redo Generated	MB per second	The average amount of redo generated, in MB per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
SessionLimitUtilization	Session Limit Utilization	percent	The session limit utilization. Statistic: Mean Interval: 1 minute <i>Not applicable for pluggable databases.</i>	5 minutes	instanceNumber instanceName hostName
Sessions	Sessions	count	The number of sessions in the database. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName
TransactionsByStatus	Transactions By Status	transactions per second	The number of committed or rolled back transactions per second. Statistic: Mean Interval: 1 minute	5 minutes	instanceNumber instanceName hostName transactionStatus (Committed, RolledBack)

Metric Name	Metric Display Name	Unit	Description and Metric Chart Defaults	Collection Frequency	Dimensions
UnusableIndexes	Unusable Indexes	count	Unusable indexes count in database schema. Statistic: Max Interval: 24 hours <i>Not applicable for container databases.</i>	24 hours	schemaName
UsableFRA	Usable Fast Recovery Area	percent	The useable fast recovery area. Statistic: Mean Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
UsedFRASpace	Flash Recovery Area Usage	GB	The flash recovery area space usage. Statistic: Max Interval: 15 minutes <i>Not applicable for pluggable databases.</i>	15 minutes	N/A
WaitTime	Wait Time	seconds per second	The average rate of accumulation of non-idle wait time by foreground sessions in the database instance over the time interval. The wait time component of Average Active Sessions. Statistic: Mean Interval: 5 minutes	5 minutes	instanceNumber instanceName hostName waitClass

Oracle Cloud Infrastructure Operations Insights

Oracle Cloud Infrastructure Operations Insights allows you to use the Capacity Planning and SQL Warehouse functionality to gain insight into Oracle Databases deployed in Exadata Cloud Infrastructure.

Using Operations Insights on Oracle Cloud Databases allows you to:

- Analyze resource usage of databases across cloud databases
- Forecast future demand for database resources such as CPU, memory, and storage based on historical trends
- Improve resource utilization by identifying under and over utilized resources
- Identify Exadata systems projected to reach high utilization
- Identify total lead time to expand capacity using machine learning based forecast based on long term historic data to project future resource growth
- Compare SQL performance across databases and identify common patterns

Related Topics

- [Enabling Database Cloud Service Databases](#)

Monitor Metrics to Diagnose and Troubleshoot Problems with Pluggable Databases

Enable Database Management service to view metrics to diagnose and troubleshoot problems with pluggable databases.

- [About Database Management](#)
- [Using the Console to Enable Database Management for a Container Database \(CDB\)](#)
To enable Database Management for a container database (CDB), use this procedure.
- [Using the Console to Enable Database Management for a Pluggable Database \(PDB\)](#)
To enable Database Management for a pluggable database (PDB), use this procedure.
- [Using the Console to Edit Database Management for a Pluggable Database \(PDB\)](#)
To edit the Database Management configuration for a pluggable database (PDB), use this procedure.
- [Using the Console to Disable Database Management for a Pluggable Database \(PDB\)](#)
To disable Database Management for a pluggable database (PDB), use this procedure.
- [Using the Console to View Performance Hub for a Container Database \(CDB\)](#)
To view Performance Hub for a container database (CDB), use this procedure. You must first enable Database Management to view the performance report.
- [Using the Console to View Performance Hub for a Pluggable Database \(PDB\)](#)
To view Performance Hub for a pluggable database (PDB), use this procedure. You must first enable Database Management to view the performance report.
- [Using the API to Enable, Disable, or Update Database Management Service](#)
- [Oracle Cloud Database Metrics](#)
Use the metrics to diagnose and troubleshoot issues.

About Database Management

As a Database Administrator, you can use the Oracle Cloud Infrastructure Database Management service to monitor and manage Oracle Databases. For more information, see *About Database Management*.

Performance Hub provides a visual representation of diagnostic data that you can leverage to fix performance issues or tune the database to improve performance. For more information about Performance Hub, see *Performance Hub*.

Related Topics

- [About Database Management](#)
- [Performance Hub](#)

Using the Console to Enable Database Management for a Container Database (CDB)

To enable Database Management for a container database (CDB), use this procedure.

Note

You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to enable Database Management.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to enable Database Management.
Database Details page is displayed.
5. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
If the Database Management is displayed as **Not Enabled**, perform the following steps:

Enable Database Management

1. Click **Enable**.
Enable Database Management window is displayed.
2. In the **Database information** section, provide the following details:
 - **Database type**: Read-only. Type of the database.

- **Exadata VM Cluster:** Read-only. Compartment in which the database is located.
- **Database home:** Read-only. Database home of the database.
- **Database name:** Read-only. Name of the database.
- **Service name:** The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol:** Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

Note

- If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
- Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port:** Specify the port number.
If TCP is selected in the **Protocol** field, then the port number 1521 is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret:** This field is only displayed if TCPS is selected in the **Protocol** field.
 - Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
The Create database wallet secret panel is displayed and you can create a new secret.

For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.
 - If the Database Management (dpm) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.
- 3. In the **Specify credentials for the connection** section, provide the following details:
 - **Database user name:** Enter the database user name.
 - **Database user password secret:**
 - Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

- b. If the Database Management (dpm) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.
4. In the **Private endpoint information** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the Oracle Cloud Database can be accessed.

You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available.

Here are the two types of Database Management private endpoints:

- Private endpoint for single instance Databases.
- Private endpoint for Oracle RAC Databases.

If a Database Management private endpoint is not available, then you must create one.

For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:
 - **Full management:** This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
 - **Basic management:** This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.
6. Click **Enable Database Management**.
7. A confirmation message with a link to the Oracle Cloud Database's **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.
8. In the **Database Information** section, under the **Associated Services**, verify if the status of **Database Management** is **Enabled**.

The **Disable** option is also displayed, which you can click to disable Database Management.

If you encounter issues when enabling Database Management, see *Issues Encountered When Enabling Database Management for Oracle Cloud Databases* for likely causes and solutions.

Related Topics

- [Permissions Required to Enable Database Management for Oracle Cloud Databases](#)
- [Oracle Cloud Database-related Prerequisite Tasks](#)
- [Enable Database Management for Oracle Cloud Databases](#)
- [Issues Encountered When Enabling Database Management for Oracle Cloud Databases](#)

Using the Console to Enable Database Management for a Pluggable Database (PDB)

To enable Database Management for a pluggable database (PDB), use this procedure.

Note

You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

Prerequisite

To enable the Database Management for a pluggable database, enable Database Management for the associated database with the **Full Management** option.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to enable Database Management.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to enable Database Management.
Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to enable Database Management.
Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
If the Database Management is displayed as **Not Enabled**, perform the following steps:

Enable Database Management

1. Click **Enable**.
Enable Database Management window is displayed.
2. In the **Database information** section, provide the following details:
 - **Database type**: Read-only. Type of the database.
 - **Exadata VM Cluster**: Read-only. Compartment in which the database is located.
 - **Database home**: Read-only. Database home of the database.
 - **Pluggable Database name**: Read-only. Name of the database.

- **Service name:** The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol:** Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

Note

- If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
- Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port:** Specify the port number.
If TCP is selected in the **Protocol** field, then the port number 1521 is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret:** This field is only displayed if TCPS is selected in the **Protocol** field.
 - a. Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
The Create database wallet secret panel is displayed and you can create a new secret.

For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.
 - b. If the Database Management (dpd) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.
- 3. In the **Specify credentials for the connection** section, provide the following details:
 - **Database user name:** Enter the database user name.
 - **Database user password secret:**
 - a. Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

- b. If the Database Management (dpm) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy. For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.
4. In the **Private endpoint information** section, select the private endpoint that will act as a representation of Database Management in the VCN in which the Oracle Cloud Database can be accessed.

You can choose the private endpoint from a different compartment as well. You must ensure that the appropriate Database Management private endpoint is available.

Here are the two types of Database Management private endpoints:

- Private endpoint for single instance Databases.
- Private endpoint for Oracle RAC Databases.

If a Database Management private endpoint is not available, then you must create one.

For information on how to create a private endpoint, see *Create a Database Management Private Endpoint*.

5. In the **Management options** section, choose between the following options:
 - **Full management:** This includes fleet management, advanced Performance Hub, and other SKU features along with basic management capabilities.
 - **Basic management:** This includes basic monitoring metrics and the ASH Analytics and SQL Monitoring features in Performance Hub for container databases. For more information on management options, see *About Management Options*.
6. Click **Enable Database Management**.
7. A confirmation message with a link to the Oracle Cloud Database's **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.
8. In the **Database Information** section, under the **Associated Services**, verify if the status of **Database Management** is **Enabled**.

The **Disable** option is also displayed, which you can click to disable Database Management.

If you encounter issues when enabling Database Management, see *Issues Encountered When Enabling Database Management for Oracle Cloud Databases* for likely causes and solutions.

Related Topics

- [Permissions Required to Enable Database Management for Oracle Cloud Databases](#)
- [Oracle Cloud Database-related Prerequisite Tasks](#)
- [Enable Database Management for Oracle Cloud Databases](#)
- [Issues Encountered When Enabling Database Management for Oracle Cloud Databases](#)

Using the Console to Edit Database Management for a Pluggable Database (PDB)

To edit the Database Management configuration for a pluggable database (PDB), use this procedure.

Note

You can also enable Database Management for a database from the Database Management Administration page. For more information, see *Enable Database Management for Oracle Cloud Databases*.

Prerequisite

To enable the Database Management for a pluggable database, enable Database Management for the associated database with the **Full Management** option.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to edit Database Management.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to edit Database Management.
Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to edit Database Management.
Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
If the Database Management is displayed as **Enabled**, perform the following steps to edit Database Management:

Edit Database Management

1. Click **Enable**.
Edit Database Management window is displayed.
2. In the **Database information** section, provide the following details:
 - **Database type**: Read-only. Type of the database.
 - **Exadata VM Cluster**: Read-only. Compartment in which the database is located.
 - **Database home**: Read-only. Database home of the database.
 - **Pluggable Database name**: Read-only. Name of the database.

- **Service name:** The unique service name of the database. A default unique name is displayed, which can be changed if required.
- **Protocol:** Select either TCP or TCPS to connect to the Oracle Cloud Database. TCP is selected by default.

Note

- If Oracle Data Guard is enabled after Database Management was enabled for an Exadata VM Cluster using the TCPS protocol, then TCPS will have to be reconfigured. Enabling Oracle Data Guard is causing TCPS configuration to be overwritten, and it's recommended that TCPS is configured on an Exadata VM Cluster after enabling Oracle Data Guard.
- Database Management currently does not support Oracle Data Guard configuration and Database Management features are not available for standby databases.

- **Port:** Specify the port number.
If TCP is selected in the **Protocol** field, then the port number 1521 is displayed by default. You can change it if required. You can select the port number from a range of 1 to 65535.
- **Database wallet secret:** This field is only displayed if TCPS is selected in the **Protocol** field.
 - Select the secret that contains the database wallet from the drop-down list. If an existing database wallet secret is not available, then select **Create new secret...** from the drop-down list.
The Create database wallet secret panel is displayed and you can create a new secret.

For information on database wallets and creating a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.
 - If the Database Management (dpd) service policy that grants Database Management permission to read the secret that contains the database wallet is not created, then the `System policies are required...` message is displayed. You can click **Add policy** to view and automatically create the service policy.
For information on Vault service permissions required to use existing secrets or create new secrets, see *Permissions Required to Enable Database Management for Oracle Cloud Databases*.
- 3. In the **Specify credentials for the connection** section, provide the following details:
 - **Database user name:** Enter the database user name.
 - **Database user password secret:**
 - Select the secret that contains the database user password from the drop-down list. If the compartment in which the secret resides is different from the compartment displayed, then click **Change compartment** and select another compartment. If an existing secret with the database user password is not available, then select **Create new secret...** from the drop-down list.
The Create password secret panel is displayed and you can create a new secret.

For information on database monitoring user credentials and saving the database user password as a secret in the Vault service, see *Oracle Cloud Database-related Prerequisite Tasks*.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to disable Database Management.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database for which you want to disable Database Management.
Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database for which you want to disable Database Management.
Pluggable Database Details page is displayed.
7. In the **Database Information** section, under the **Associated Services**, check the status of Database Management.
8. If the Database Management is displayed as **Enabled**, perform the following steps to disable Database Management:
 - a. Click **Disable**.
 - b. A confirmation message with a link to the **Work requests** section on the **Database information** page is displayed. Click the link to monitor the progress of the work request.
 - c. In the **Database Information** section, under the **Associated Services**, verify if the status of Database Management is **Disabled**.

Using the Console to View Performance Hub for a Container Database (CDB)

To view Performance Hub for a container database (CDB), use this procedure. You must first enable Database Management to view the performance report.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to view Performance Hub.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to view Performance Hub.
Database Details page is displayed.
5. Click **Performance Hub**.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Performance Hub allows you to download reports for your managed databases. For more information about downloading reports, see *Automatic Workload Repository (AWR) Report*, *Active Sessions History (ASH) Report*, and *Performance Hub Report*.

Related Topics

- [Automatic Workload Repository \(AWR\) Report](#)
- [Active Sessions History \(ASH\) Report](#)
- [Performance Hub Report](#)

Using the Console to View Performance Hub for a Pluggable Database (PDB)

To view Performance Hub for a pluggable database (PDB), use this procedure. You must first enable Database Management to view the performance report.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to view Performance Hub.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database that contains the pluggable database.
Database Details page is displayed.
5. Under **Resources**, click **Pluggable Databases**.
6. In the list of pluggable databases, click the pluggable database that you're interested in.
Pluggable Database Details page is displayed.
7. Click **Performance Hub**.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Performance Hub allows you to download reports for your managed databases. For more information about downloading reports, see *Automatic Workload Repository (AWR) Report*, *Active Sessions History (ASH) Report*, and *Performance Hub Report*.

Related Topics

- [Automatic Workload Repository \(AWR\) Report](#)
- [Active Sessions History \(ASH\) Report](#)
- [Performance Hub Report](#)

Using the API to Enable, Disable, or Update Database Management Service

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

Use these API operations to configure the Database Management service.

- Enable Database Management service for an Oracle Database located in Oracle Cloud Infrastructure to access tools including Metrics and Performance hub:
`enableDatabaseManagement`
- Disable Database Management service: `disableDatabaseManagement`
- Update Database Management configuration: `updateDatabaseManagement`

Oracle Cloud Database Metrics

Use the metrics to diagnose and troubleshoot issues.

The metrics for Oracle Cloud Databases help measure useful quantitative data, such as CPU and storage utilization, the number of successful and failed database logon and connection attempts, database operations, SQL queries, transactions, and so on.

For more information, see *Oracle Cloud Database Metrics*.

- [Using the Console View Metrics for a Container Database \(CDB\)](#)
To view metrics for a container database (CDB), you must first enable Database Management with the **Full Management** option.
- [Using the Console to View Metrics for a Pluggable Database \(PDB\)](#)
To view metrics for a Pluggable Database (PDB), the following prerequisites must be met:

Related Topics

- [Oracle Cloud Database Metrics](#)

Using the Console View Metrics for a Container Database (CDB)

To view metrics for a container database (CDB), you must first enable Database Management with the **Full Management** option.

To enable Database Management for databases, see *Using the Console to Enable Database Management for a Database*.

1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the database for which you want to view the metrics.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
4. In the list of databases, click the database for which you want to view the metrics.
Database Details page is displayed.
5. Under **Resources**, click **Metrics**.

Related Topics

- [Using the Console to Enable Database Management for a Container Database \(CDB\)](#)
To enable Database Management for a container database (CDB), use this procedure.

Using the Console to View Metrics for a Pluggable Database (PDB)

To view metrics for a Pluggable Database (PDB), the following prerequisites must be met:

- Enable Database Management for databases with the **Full Management** option.
- Enable Database Management for pluggable databases.
- 1. Open the navigation menu. Click **Oracle Database**, then click **Oracle Exadata Database Service on Dedicated Infrastructure**.
- 2. Choose your **Compartment**.
A list of Exadata VM Clusters is displayed.
- 3. In the list of Exadata VM Clusters, click the Exadata VM Cluster that contains the pluggable database for which you want to view the metrics.
Exadata VM Cluster Details page is displayed.
Under **Resources**, **Databases** is selected by default.
- 4. In the list of databases, click the database that contains the pluggable database.
Database Details page is displayed.
- 5. Under **Resources**, click **Pluggable Databases**.
- 6. In the list of pluggable databases, click the pluggable database for which you want to view the metrics.
Pluggable Database Details page is displayed.
- 7. Under **Resources**, click **Metrics**.
- 8. Select a namespace from the **Metric namespace** from where you wish to view metrics.

Note

- When Database Management is enabled, then you can view metrics only from the `oracle_oci_database` namespace.
- When Database Management is disabled, then a banner, "Database management must be enabled to provide data for metrics." is displayed.

With Basic Management, Performance Hub provides **ASH Analytics** and **SQL Monitoring**. Advanced Management will additionally provide **ADDM**, **Workload**, and **Blocking Sessions**.

Related Topics

- [Using the Console to Enable Database Management for a Container Database \(CDB\)](#)
To enable Database Management for a container database (CDB), use this procedure.
- [Using the Console to Enable Database Management for a Pluggable Database \(PDB\)](#)
To enable Database Management for a pluggable database (PDB), use this procedure.

Security Guide for Oracle Exadata Database Service on Dedicated Infrastructure

This guide describes security for an Exadata Cloud Infrastructure. It includes information about the best practices for securing the Exadata Cloud Infrastructure.

- [Part 1: Security Configurations and Default Enabled Features](#)
- [Part 2: Additional Procedures for Updating Security Posture](#)

Part 1: Security Configurations and Default Enabled Features

- [Responsibilities](#)
Exadata Cloud Infrastructure is jointly managed by the customer and Oracle, and is divided into two areas of responsibility.
- [Infrastructure Security](#)
Security features offered by Exadata Cloud Infrastructure.
- [Guiding Principles Followed for Security Configuration Defaults](#)
- [Security Features](#)
- [Guest VM Default Fixed Users](#)
Several user accounts regularly manage the components of Exadata Cloud Infrastructure. These users are required and may not be modified.
- [Default Security Settings: Customer VM](#)
- [Default Processes on Customer VM](#)
A list of the processes that run by default on the customer VM, also called DOMU, or Guest VM and Guest OS
- [Default Database Security Configuration](#)
- [Default Backup Security Configuration](#)
- [Operator Access to Customer System and Customer Data](#)
Only automated tooling is permitted to access guest VM for purposes of lifecycle automation.
- [Compliance Requirements](#)
- [Break Glass Procedure for Accessing Customer's Guest VM](#)
There are situations where some problems can only be resolved by Oracle logging into the customer guest VM.

Responsibilities

Exadata Cloud Infrastructure is jointly managed by the customer and Oracle, and is divided into two areas of responsibility.

These areas are as follows:

1. Customer accessible services: components that the customer can access as part of their subscription to Exadata Cloud Service:
 - Customer accessible virtual machines (VM)
 - Customer accessible database services

2. Oracle Managed Infrastructure: components that are owned and operated by Oracle to run customer accessible services
 - Power Distribution Units (PDUs)
 - Out of band (OOB) management switches
 - Storage network switches
 - Exascale Storage Servers
 - Physical Exadata database servers
 - Data center security

Customers control and monitor access to customer services, including network access to their VMs (through OCI Virtual Cloud Networks and OCI Security Lists), authentication to access the VM, and authentication to access databases running in the VMs. Oracle controls and monitors access to Oracle Managed Infrastructure components and physical server security. Oracle staff are not authorized to access customer services, including customer VMs and databases except where customers are unable to access the customer VM.

Customers access Oracle Databases (DB) running on Exadata Cloud Infrastructure via client and backup VCNs to the databases running in the customer VM using standard Oracle database connection methods, such as Oracle Net on port 1521. Customer's access the VM running the Oracle databases via standard Oracle Linux methods, such as token based ssh on port 22.

Infrastructure Security

Security features offered by Exadata Cloud Infrastructure.

- **Oracle Cloud Physical Security**

Oracle Cloud data centers align with Uptime Institute and Telecommunications Industry Association (TIA) ANSI/TIA-942-A Tier 3 (99.982% Availability) or Tier 4 (99.995% Availability) standards and follow a N2 ('N' stands for Need) redundancy methodology for critical equipment operation. Data centers housing Oracle Cloud Infrastructure services use redundant power sources and maintain generator backups in case of widespread electrical outage. Server rooms are closely monitored for air temperature and humidity, and fire-suppression systems are in place. Data center staff are trained in incident response and escalation procedures to address security and availability events that may arise. For more information see: [Oracle Cloud Infrastructure Security Guide](#). For further details on Oracle Cloud Infrastructure Data Center compliance, see [Oracle Cloud Compliance](#). Also see: [Oracle Corporate Security Practices: Data Security: Physical and Environmental Controls](#).
- **Oracle Data Center Access Controls**

To provide secure systems, Oracle access protocols to physical data centers include the following:

 - Physical access to facilities to data centers is limited to certain Oracle employees, contractors, and authorized visitors.
 - Oracle employees, subcontractors, and authorized visitors are issued identification cards that must be worn while on Oracle premises.
 - Visitors are required to sign a visitor's register, be escorted and/or observed when they are on Oracle premises, and/or be bound by the terms of a confidentiality agreement with Oracle.

- Security monitors the possession of keys/access cards, and monitors the ability to access facilities. Staff leaving Oracle's employment must return keys/cards, and key/cards are deactivated upon termination.
- Security authorizes all repairs and modifications to the physical security barriers or entry controls at service locations.
- Oracle use a mixture of 24/7 onsite security officers or patrol officers, depending on the risk/protection level of the facility. In all cases, officers are responsible for patrols, alarm response, and recording of security incidents.
- Oracle has implemented centrally managed electronic access control systems with integrated intruder alarm capability. The access logs are kept for a minimum of six months. Furthermore, the retention period for CCTV monitoring and recording ranges from 30-90 days minimum, depending on the facility's functions and risk level.

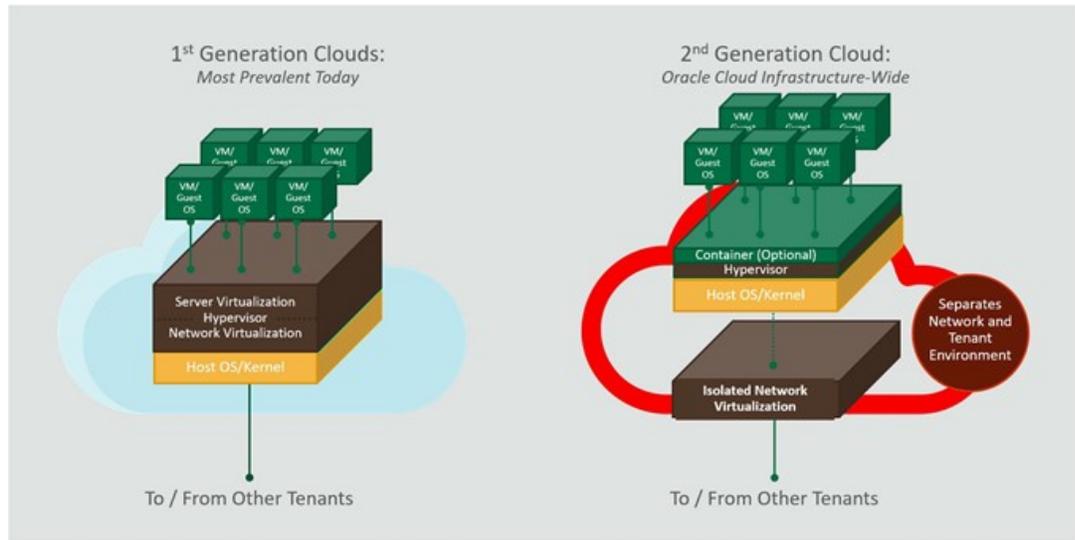
For more details about Oracle site access controls, see: [Oracle Corporate Security Practices: Oracle Access Control](#)

- **Hypervisor Customer Isolation**

The hypervisor is the software that manages virtual devices in a cloud environment, handling server and network virtualization. In traditional virtualization environments, the hypervisor manages network traffic, enabling it to flow between VM instances and between VM instances and physical hosts. This adds considerable complexity and computational overhead in the hypervisor. Proof-of concept computer security attacks, such as virtual machine escape attacks, have highlighted the substantial risk that can come with this design. These attacks exploit hypervisor complexity by enabling an attacker to “breakout” of a VM instance, access the underlying operating system, and gain control of the hypervisor. The attacker can then access other hosts, sometimes undetected. Oracle Cloud Infrastructure reduces this risk by decoupling network virtualization from the hypervisor.

To address potential attacks, Oracle has implemented a security-first architecture using Isolated network virtualization, which is foundational element of Oracle Cloud infrastructure's architecture. This architecture stops malware in its tracks with a custom-designed SmartNIC to isolate and virtualize the network. Isolated network virtualization reduces risk by limiting the attack surface. Even if a malicious actor succeeds with a VM escape attack on a single host, the virtual architecture is designed so that actor can't reach other hosts in the cloud infrastructure. The attack is effectively contained to the one host. Secure Isolated network virtualization architecture is implemented in every data center in every region. Every Oracle Cloud Infrastructure tenant is provided with the benefits of this security-first architecture.

Figure 6-1 Isolated Network Virtualization Reduces Risk in Oracle Generation 2 Cloud



- **Multitenant Security**

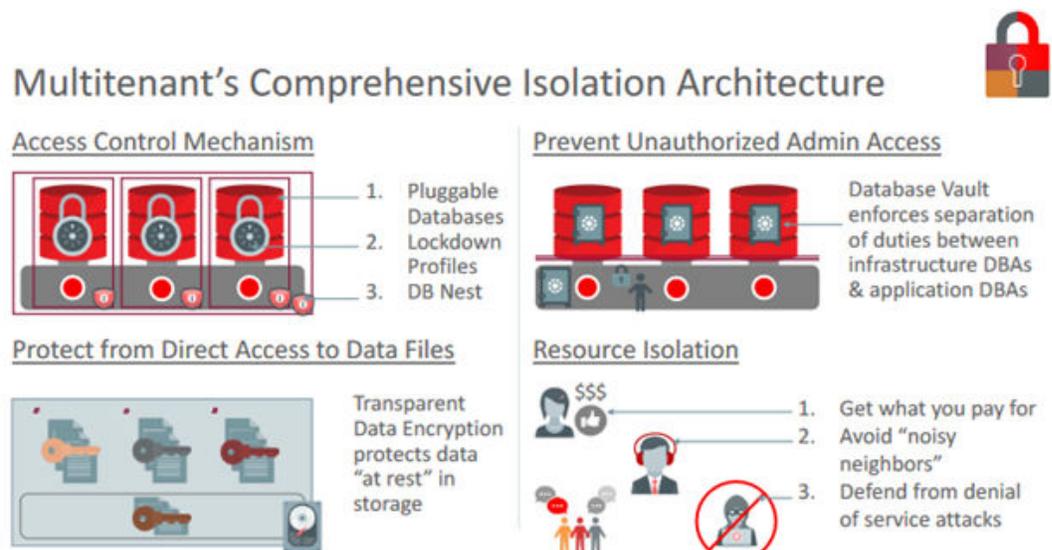
Consistent with our security philosophy of Defense in Depth, Multitenant has a comprehensive isolation architecture.

There are four major categories to this, with several important features in each category.

1. Access Control Mechanism
2. Prevent Unauthorized Admin Access
3. Protect from direct access to Data Files
4. Resource Isolation

e

Figure 6-2 Multitenant’s Comprehensive Isolation Architecture



Related Topics

- [Oracle Cloud Infrastructure Security Architecture](#)
- [Oracle Cloud Infrastructure Security Guide](#)
- [Data Security: Physical and Environmental Controls](#)
- [Oracle Multitenant with Oracle Database 19c](#)
- [Oracle Cloud Compliance](#)

Guiding Principles Followed for Security Configuration Defaults

- **Defense in Depth** Exadata Cloud Infrastructure offers a number of controls to ensure confidentiality, integrity, and availability throughout the service. First, Exadata Cloud Infrastructure is built from the hardened operating system image provided by Exadata Database Machine (<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadata-security-overview.html>). This secures the core operating environment by restricting the installation image to only the required software packages, disabling unnecessary services, and implementing secure configuration parameters throughout the system.

In addition to inheriting all the strength of Exadata Database Machine's mature platform, because Exadata Cloud Infrastructure provisions systems for customers, additional secure default configuration choices are implemented in the service instances. For example, all database tablespaces require transparent data encryption (TDE), strong password enforcement for initial database users and superusers, and enhanced audit and event rules.

Exadata Cloud Infrastructure also constitutes a complete deployment and service, so it is subjected to industry-standard external audits such as PCI, HIPPA and ISO27001. These external audit requirements impose additional value-added service features such as antivirus scanning, automated alerting for unexpected changes to the system, and daily vulnerability scans for all Oracle-managed infrastructure systems in the fleet.

- **Least privilege** [Oracle Secure Coding Standards](#) require software processes run at the minimum privilege level to implement their functionality.

Each process and daemon, must run as a normal, unprivileged user unless it can prove a requirement for a higher level of privilege. This helps contain any unforeseen issues or vulnerabilities to unprivileged user space and not compromise an entire system.

[This principle also applies to Oracle operations team members](#) who use individual named accounts to access the Exadata Cloud Infrastructure for maintenance or troubleshooting. Only when necessary will they use the audited access to higher levels of privilege to solve or resolve an issue. Most issues are resolved through automation, so we also employ least privilege by not permitting human operators to access a system unless the automation is unable to resolve the issue.

- **Auditing and accountability** When required, access to the system is allowed, but all access and actions are logged and tracked for accountability.

Exadata Cloud Infrastructure audit logs are controlled by Oracle and used for security monitoring and compliance purposes. Oracle can share relevant audit logs with customers per [Oracle Incident Response Practices](#) and the [Oracle Data Processing Agreement](#).

Auditing capabilities are provided at all infrastructure components to ensure all actions are captured. Customers also have ability to configure auditing for their database and guest

VM configuration and may choose to integrate those with other enterprise auditing systems.

- **Automating cloud operations**

By eliminating manual operations required to provision, patch, maintain, troubleshoot, and configure systems, the opportunity for error is reduced.

Security Features

- **Hardened OS image**

- Minimal package installation:

Only the necessary packages required to run an efficient system are installed. By installing a smaller set of packages, the attack surface of the operating system is reduced and the system remains more secure.

- Secure configuration:

Many non-default configuration parameters are set during installation to enhance the security posture of the system and its content. For example, SSH is configured to only listen on certain network interfaces, sendmail is configured to only accept localhost connections, and many other similar restrictions are implemented during installation.

- Run only necessary services:

Any services that may be installed on the system, but not required for normal operation, are disabled by default. For example, while NFS is a service often configured by customers for various application purposes, it is disabled by default as it is not required for normal database operations. Customers may choose to optionally configure services per their requirements.

- **Minimized attack surface**

As part of the hardened image, attack surface is reduced installing and running only the software required to deliver the service.

- **Additional security features enabled (grub passwords, secure boot)**

- Leveraging Exadata image capabilities, ExaDB-D enjoys the features integrated into the base image such as grub passwords and secure boot in addition to many others.

- Through customization, customers may wish to further enhance their security posture with additional configurations.

- **Secure access methods**

- Accessing database servers via SSH using strong cryptographic ciphers. Weak ciphers are disabled by default.

- Accessing databases via encrypted Oracle Net connections. By default, our services are available using encrypted channels and a default configured Oracle Net client will use encrypted sessions.

- Accessing diagnostics via Exadata MS web interface (https).

- **Auditing and logging**

- By default, auditing is enabled for administrative operations and those audit records are communicated to OCI internal systems for automated review and alerting when required.

Guest VM Default Fixed Users

Several user accounts regularly manage the components of Exadata Cloud Infrastructure. These users are required and may not be modified.

In all ExaDB-D machines, Oracle uses and recommends token-based SSH login.

There are three classes of users:

- [Default Users: No Logon Privileges](#)
- [Default Users WITH RESTRICTED SHELL Access](#)
These users are used for accomplishing a defined task with a restricted shell login. These users should not be removed as the defined task will fail in case these users are deleted.
- [Default Users with Login Permissions](#)
These privileged users are used for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have significant impact on the running system.

Default Users: No Logon Privileges

This user list consists of default operating system users along with some specialized users like exawatch and dbmsvc. These users should not be altered. These users cannot login to the system as all are set to /sbin/nologin.

In the list of users below, most are either standard Linux OS users or related to standard Linux packages except for the exawatch and dbmsvc users.

- exawatch: The exawatch user is responsible for collecting and archiving system statistics on both the database servers and the storage servers
- dbmsvc: User is used for Management Server on the database node service in Oracle Exadata System

NOLOGIN Users

```
bin:x:1:1:bin:/bin:/sbin/nologin
Daemon:x:2:2:daemon:/sbin:/sbin/nologin
adm:x:3:4:adm:/dev/null:/sbin/nologin
mail:x:8:12:mail:/var/spool/mail:/sbin/nologin
nobody:x:99:99:Nobody:/:/sbin/nologin
systemd-network:x:192:192:systemd Network Management:/:/sbin/nologin
dbus:x:81:81:System message bus:/:/sbin/nologin
rpm:x:37:37:/:/var/lib/rpm:/sbin/nologin
sshd:x:74:74:Privilege-separated SSH:/var/empty/ssh:/sbin/nologin
rpc:x:32:32:Rpcbind Daemon:/var/lib/rpcbind:/sbin/nologin
unbound:x:999:997:Unbound DNS resolver:/etc/unbound:/sbin/nologin
nscd:x:28:28:NSCD Daemon:/:/sbin/nologin
tss:x:59:59:Account used by the trousers package to sandbox the tcsd
daemon:/dev/null:/sbin/nologin
saslauth:x:998:76:Saslauthd user:/run/saslauthd:/sbin/nologin
mailnull:x:47:47:/:/var/spool/mqueue:/sbin/nologin
smmsp:x:51:51:/:/var/spool/mqueue:/sbin/nologin
chrony:x:997:996:/:/var/lib/chrony:/sbin/nologin
rpcuser:x:29:29:RPC Service User:/var/lib/nfs:/sbin/nologin
nslcd:x:65:55:LDAP Client User:/:/sbin/nologin
uucp:x:10:14:Uucp user:/var/spool/uucp:/sbin/nologin
```

```

tcpdump:x:72:72:::/sbin/nologin
exawatch:x:1010:510::/opt/oracle.ExaWatcher:/sbin/nologin
sssd:x:996:508:User forsssd:/sbin/nologin
dbmsvc:x:2001:2001:::/sbin/nologin
clamupdate:x:995:504:Clamav database update user:/var/lib/clamav:/sbin/nologin

```

Default Users WITH RESTRICTED SHELL Access

These users are used for accomplishing a defined task with a restricted shell login. These users should not be removed as the defined task will fail in case these users are deleted.

`dbmmonitor` password is set to a random string during deployment, which must change on first use.

- `dbmmonitor`: The `dbmmonitor` user is used for DBMCLI Utility. For more details refer to [Using the DBMCLI Utility](#)

Restricted Shell Users

```
dbmmonitor:x:2003:2003::/home/dbmmonitor:/bin/rbash
```

Default Users with Login Permissions

These privileged users are used for accomplishing most of the tasks in the system. These users should never be altered or deleted as it would have significant impact on the running system.

SSH keys are used for login by customer staff and cloud automation software.

Customer-added SSH keys may be added by the `UpdateVmCluster` method, or by customers directly accessing the customer VM and managing SSH keys inside of the customer VM. Customers are responsible for adding comments to keys to make them identifiable. When a customer adds the SSH key by the `UpdateVmCluster` method, the key is only added to the `authorized_keys` file of the `opc` user.

Cloud automation keys are temporary, specific to a given cloud automation task, for example, VM Cluster Memory resize, and unique. Cloud automation access keys can be identified by the following comments: `OEDA_PUB`, `EXACLOUD_KEY`, `ControlPlane`. Cloud automation keys are removed after the cloud automation task completes so the `authorized_keys` files of the `root`, `opc`, `oracle`, and `grid` accounts should only contain cloud automation keys while the cloud automation actions are running.

Privileged Users

```

root:x:0:0:root:/root:/bin/bash
oracle:x:1001:1001::/home/oracle:/bin/bash
grid:x:1000:1001::/home/grid:/bin/bash
opc:x:2000:2000::/home/opc:/bin/bash
dbmadmin:x:2002:2002::/home/dbmadmin:/bin/bash

```

- `root`: Linux requirement, used sparingly to run local privileged commands. `root` is also used for some processes like Oracle Trace File Analyzer Agent and `ExaWatcher`.
- `grid`: Owns Oracle Grid Infrastructure software installation and runs Grid Infrastructure processes.
- `oracle`: Owns Oracle database software installation and runs Oracle Database processes.

- `opc:`
 - Used by Oracle Cloud automation for automation tasks.
 - Has the ability to run certain privileged commands without further authentication (to support automation functions).
 - Runs the local agent, also known as “DCS Agent” that performs lifecycle operations for Oracle Database and Oracle Grid Infrastructure software (patching, create database, and so on).
- `dbmadmin:`
 - The `dbmadmin` user is used for Oracle Exadata Database Machine Command-Line Interface (DBMCLI) utility.
 - The `dbmadmin` user should be used to run all services on the database server. For more information, see [Using the DBMCLI Utility](#).

Related Topics

- [Using the DBMCLI Utility](#)

Default Security Settings: Customer VM

In addition to all of the Exadata features explained in [Security Features of Oracle Exadata Database Machine](#), the following security settings are also applicable, to Exadata Cloud Infrastructure instances.

- Custom database deployment with non-default parameters.
The command `host_access_control` is to configure Exadata security settings:
 - Implementing password aging and complexity policies.
 - Defining account lockout and session timeout policies.
 - Restricting remote root access.
 - Restricting network access to certain accounts.
 - Implementing login warning banner.
- `account-disable`: Disables a user account when certain configured conditions are met.
- `pam-auth`: Various PAM settings for password changes and password authentication.
- `rootssh`: Adjusts the `PermitRootLogin` value in `/etc/ssh/sshd_config`, which permits or denies the `root` user to login through SSH.
 - By default, `PermitRootLogin` is set to `no`.
 - `PermitRootLogin=without-password` is required for the cloud automation to perform some lifecycle management operations, disabling root login will cause that service functionality to fail.
- `session-limit`: Sets the `* hard maxlogins` parameter in `/etc/security/limits.conf`, which is the maximum number of logins for all users. This limit does not apply to a user with `uid=0`.
Defaults to `* hard maxlogins 10` and it is the recommended secure value.
- `ssh-macs`: Specifies the available Message Authentication Code (MAC) algorithms.
 - The MAC algorithm is used in protocol version 2 for data integrity protection.
 - Defaults to `hmac-sha1, hmac-sha2-256, hmac-sha2-512` for both server and client.

- Secure recommended values: `hmac-sha2-256`, `hmac-sha2-512` for both server and client.
- `password-aging`: Sets or displays the current password aging for interactive user accounts.
 - `-M`: Maximum number of days a password may be used.
 - `-m`: Minimum number of days allowed between password changes.
 - `-W`: Number of days warning given before a password expires.
 - Defaults to `-M 99999`, `-m 0`, `-W 7`
 - `--strict_compliance_only-M 60`, `-m 1`, `-W 7`
 - Secure recommended values: `-M 60`, `-m 1`, `-W 7`

Related Topics

- [Security Features of Oracle Exadata Database Machine](#)

Default Processes on Customer VM

A list of the processes that run by default on the customer VM, also called DOMU, or Guest VM and Guest OS

- **Exadata Cloud Infrastructure VM agent:**
Cloud agent for handling database lifecycle operations.
 - Runs as `opc` user
 - Process table shows it running as a Java process with `jar` names `-dbcs-agent-VersionNumber-SNAPSHOT.jar` and `dbcs-admin-VersionNumver-SNAPSHOT.jar`.
- **Oracle Trace File Analyzer agent:**
Oracle Trace File Analyzer (TFA) provides a number of diagnostic tools in a single bundle, making it easy to gather diagnostic information about the Oracle database and clusterware, which in turn helps with problem resolution when dealing with Oracle Support
 - Runs as `root` user
 - Runs as `initd` demon (`/etc/init.d/init.tfa`)
 - Process tables show a Java application (`oracle.rat.tfa.TFAMain`)
 - Runs as `root` and `exawatch` users.
 - Runs as background script, `ExaWatcher.sh` and all its child process run as a Perl process.
 - Process table shows as multiple Perl applications.`ExaWatcher`:
- **Database and GI (clusterware):**
 - Runs as `dbmsvc` and `grid` users
 - Process table shows following applications:
 - * `oraagent.bin`, `apx_*` and `ams_*` as `grid` user
 - * `dbrsMain`, and Java applications `-derbyclient.jar`, `weblogic.Server` as `oracle` user.
- **Management Server (MS):**
Part of Exadata image software for managing and monitoring the image functions.
 - Runs as `dbmadmin`.

- Process table shows it running as a Java process.
- [Guest VM Network Security](#)
- [Compliance Requirements](#)

Guest VM Network Security

Table 6-32 Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
Bridge on client VLAN	bondeth0	22	sshd
		1521	Oracle TNS listener
		5000	Oracle Trace File Analyzer Collector
	bondeth0:1	7879	Jetty Management Server
		1521	Oracle TNS listener
		1521	Oracle TNS listener
Bridge on backup VLAN	bondeth1	7879	Jetty Management Server
		7879	Jetty Management Server
Oracle Clusterware running on each cluster node communicates through these interfaces.	clib0/clre0	1525	Oracle TNS listener
		3260	Synology DSM iSCSI
		5054	Oracle Grid Interprocess Communication
		7879	Jetty Management Server
		Dynamic Port: 9000-65500	System Monitor service (osysmond)
		Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	

Table 6-32 (Cont.) Default Port Matrix for Guest VM Services

Type of interface	Name of interface	Port	Process running
		Dynamic Port: 9000-65500	Cluster Logger service (ologgerd)
		Ports are controlled by the configured ephemeral range in the operating system and are dynamic.	
	clib1/clre1	5054	Oracle Grid Interprocess communication
		7879	Jetty Management Server
Cluster nodes use these interfaces to access storage cells (ASM disks). However, the IP/ports 7060/7070 attached to the storage interfaces are used to access DBCS agent from the Control Plane server.	stib0/stre0	7060	dbcs-admin
		7070	dbcs-agent
	stib1/stre1	7060	dbcs-admin
		7070	dbcs-agent
Control Plane server to domU	eth0	22	sshd
Loopback	lo	22	sshd
		2016	Oracle Grid Infrastructure
		6100	Oracle Notification Service (ONS), part of Oracle Grid Infrastructure
		7879	Jetty Management Server
		Dynamic Port 9000-65500	Oracle Trace File Analyzer

Note

TNS listener opens dynamic ports after initial contact to well known ports (1521, 1525).

Default iptables rules for Guest VM:

The default iptables are setup to ACCEPT connections on input, forward, and output chains.

```
#iptables -L -n -v
Chain INPUT (policy ACCEPT 0 packets, 0 bytes)
pkts bytes target      prot opt in      out     source
```

```
destination
```

```
Chain FORWARD (policy ACCEPT 0 packets, 0 bytes)
  pkts bytes target    prot opt in     out    source
destination
```

```
Chain OUTPUT (policy ACCEPT 0 packets, 0 bytes)
  pkts bytes target    prot opt in     out    source
destination
```

Compliance Requirements

PII (Personally Identifiable Information) This information is considered confidential and sensitive, and must be protected to prevent unauthorized use of personal information for the purposes of legal regulation, financial liability, and personal reputation.

You must configure a set of explicit rules to prevent Personally Identifiable Information (PII) from being displayed in your data.

The default Application Performance Monitoring rules hide PII in URLs by recognizing monetary values, bank-account numbers, and dates. However, the default rules only catch obvious PII and are not exhaustive. You must evaluate the default rules and further configure rules to ensure correct reporting in your environment and ensure that PII is not displayed in your data.

For more information, see [Hide Personally Identifiable Information](#) and [Security and Personally Identifiable Information](#)

Backup Retention

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

For more information, see [Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure](#)

Audit Log Retention Period

The OCI Audit service provides records of API operations performed against supported services as a list of log events. By default, Audit service records are retained for 365 days.

For more information, see [Audit Log Retention Period](#)

Service Log Retention

Oracle Cloud Infrastructure services, such as API Gateway, Events, Functions, Load Balancing, Object Storage, and VCN Flow Logs emit service logs. Each of these supported services has a Logs resource that allows you to enable or disable logging for that service. By default, Log retention is 1 month, but it can be set until 6 months.

Logs groups can be used to limit access to sensitive logs generated by services using IAM policy. You don't have to rely on complex compartment hierarchies to secure your logs. For example, say the default log group in a single compartment is where you store logs for the

entire tenancy. You grant access to the compartment for log administrators with IAM policy as you normally would. However, let's say some projects contain personally identifiable information (PII) and those logs can only be viewed by a select group of log administrators. Log groups allow you to put logs that contain PII into a separate log group, and then use IAM policy to restrict access to all but a few log administrators.

For more information, see [Service Logs](#) and [Managing Logs and Log Groups](#)

Default Database Security Configuration

Default database security features enabled and used:

- Transparent Database Encryption (TDE) is used for database tablespaces created by Oracle Database Cloud tools.
 - CDB\$ROOT: users tablespace is encrypted
 - PDBs: all tablespaces encrypted
 - Wallet password is provided during initial DB creation. Wallet passwords may be changed using `dbaascli`. Customers should change this password periodically.
- Users in the database
 - No additional users are created in the database.
 - After DB creation, all DB users are locked except for SYS, SYSTEM and DBSNMP.
 - Auditing is enabled for the following operations:
 - * DATABASE LINK
 - * PUBLIC DATABASE LINK
 - * PUBLIC SYNONYM
 - * DROP ANY PROCEDURE
 - * PROCEDURE
 - * ALTER SYSTEM
 - * TRIGGER
 - * CREATE DATABASE LINK
 - * ALTER DATABASE LINK
 - * CREATE PROCEDURE
 - * ALTER SYSTEM
 - * CREATE TRIGGER
 - * CREATE ANY TRIGGER
 - * SELECT ANY DICTIONARY
 - * **DB VERSION_11_2: EXEMPT REDACTION POLICY**
 - * **DB VERSION_12_1 or DB VERSION_12_2: BECOME USER**
 - * **DB VERSION_12_1: SESSION**
 - * DBAASecure profile is created and it is set as default profile for database user account.

- Native SQL*Net encryption for all network connections - Relevant `sqlnet.ora` parameters set in Exadata Cloud Infrastructure by default are:
 - `SQLNET.ENCRYPTION_TYPES_SERVER = (AES256, AES192, AES128)`
 - `SQLNET.ENCRYPTION_SERVER = requested`
 - `SQLNET.CRYPTO_CHECKSUM_SERVER = accepted`
 - `SQLNET.CRYPTO_CHECKSUM_TYPES_SERVER = (SHA256, SHA384, SHA512)`
- TCPS protocol offered for network connection to the database on port 2484 (wallet configured at `/var/opt/oracle/dbaas_acfs/grid/tcps_wallets`). Relevant `sqlnet.ora` parameters set in Exadata Cloud Infrastructure by default are:
 - `SSL_CIPHER_SUITES = (SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256, SSL_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384, SSL_ECDHE_RSA_WITH_AES_128_GCM_SHA256, SSL_ECDHE_RSA_WITH_AES_256_GCM_SHA384)`
 - `WALLET_LOCATION = (SOURCE=(METHOD=FILE)(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/grid/tcps_wallets)))`
 - `SQLNET.IGNORE_ANO_ENCRYPTION_FOR_TCPS = TRUE`
 - `SSL_VERSION = 1.2`
- Remote listener registration - listeners run from GI home. Exadata Cloud Infrastructure deployments the Grid Infrastructure version specified in *Oracle Support Document 2333222.1 (Exadata Cloud Service Software Versions)*. Exadata Cloud Infrastructure default configuration includes `listener.ora` parameter `VALID_NODE_CHECKING_REGISTRATION_LISTENER=SUBNET` combined with `REMOTE_REGISTRATION_ADDRESS_<SCANLISTENER>=<value>` to restrict remote listener registrations for security purposes.
- OCI Vault integration - TDE encryption key may be stored in OCI Vault (a Key Management System). For more information and instructions to configure principals, vaults, etc. see *Customer-Managed Keys in Exadata Cloud Infrastructure*. Both private vs shared vault types are supported for Exadata Cloud Infrastructure OCI Vault integration. DB user authentication is not integrated with OCI Vault.

Related Topics

- [Securing Database](#)
- [Customer-Managed Keys in Exadata Cloud Infrastructure](#)
Customer-managed keys for Exadata Cloud Infrastructure is a feature of Oracle Cloud Infrastructure (OCI) Vault service that enables you to encrypt your data using encryption keys that you control.
- [Oracle Support Document 2333222.1 \(Exadata Cloud Service Software Versions\)](#)

Default Backup Security Configuration

OS/VM backups:

Oracle does a full backup of guest VM weekly and maintains one or more backup copies. These backups are full disk snapshots of the guest VM (local OS filesystems, not ASM disk groups which reside on Exadata storage). This backup is triggered at a preset time every week. The backups are stored locally in the dom0. Customers can request Oracle to restore the guest VM image from the most recent backup by filing a My Oracle Support (MOS) Service Request (SR). Oracle cannot restore specific files from the image backup. Customers should perform file level backups in the guest VM if they require the ability to perform single-file restore.

Managed DB backups:

- Weekly full backup (level 0)
- Daily rolling incremental backup (level 1) on seven day cycle
- Automatic backups daily at a specific time set during the database deployment creation process

Retention period for backups vary from 30 days (on Object Storage) to 7 days (on local storage)

Encryption:

- Both Object Storage and local storage: All backups to cloud storage are encrypted.
- Object Storage only: All backups to cloud storage are encrypted.

All backups can be configured via CP UI or CP API.

All backups are encrypted with the same master key used for Transparent Data Encryption (TDE) wallet encryption.

Operator Access to Customer System and Customer Data

Only automated tooling is permitted to access guest VM for purposes of lifecycle automation.

One specific use case is when guest VM is unable to boot. In this case, customers must provide permission to access the guest VM for recovery purposes. Details to handle this scenario are described in section "Exception Workflows" of [Exadata Cloud Service Security Controls](#).

Customers control and monitor access to customer services, including network access to their guest VMs (through layer 2 VLANs and firewalls implemented in the guest VM), authentication to access the guest VM, and authentication to access databases running in the guest VMs. Oracle controls and monitors access to Oracle-managed infrastructure components. Oracle staff are not authorized to access customer services, including guest VMs and databases.

Compliance Requirements

PII (Personally Identifiable Information) This information is considered confidential and sensitive, and must be protected to prevent unauthorized use of personal information for the purposes of legal regulation, financial liability, and personal reputation.

You must configure a set of explicit rules to prevent Personally Identifiable Information (PII) from being displayed in your data.

The default Application Performance Monitoring rules hide PII in URLs by recognizing monetary values, bank-account numbers, and dates. However, the default rules only catch obvious PII and are not exhaustive. You must evaluate the default rules and further configure

rules to ensure correct reporting in your environment and ensure that PII is not displayed in your data.

For more information, see [Hide Personally Identifiable Information](#) and [Security and Personally Identifiable Information](#)

Backup Retention

When you enable the Automatic Backup feature, the service creates daily incremental backups of the database to Object Storage. The first backup created is a level 0 backup. Then, level 1 backups are created every day until the next weekend. Every weekend, the cycle repeats, starting with a new level 0 backup.

If you choose to enable automatic backups, you can choose one of the following preset retention periods: 7 days, 15 days, 30 days, 45 days, or 60 days. The system automatically deletes your incremental backups at the end of your chosen retention period.

For more information, see [Manage Database Backup and Recovery on Oracle Exadata Database Service on Dedicated Infrastructure](#)

Audit Log Retention Period

The OCI Audit service provides records of API operations performed against supported services as a list of log events. By default, Audit service records are retained for 365 days.

For more information, see [Audit Log Retention Period](#)

Service Log Retention

Oracle Cloud Infrastructure services, such as API Gateway, Events, Functions, Load Balancing, Object Storage, and VCN Flow Logs emit service logs. Each of these supported services has a Logs resource that allows you to enable or disable logging for that service. By default, Log retention is 1 month, but it can be set until 6 months.

Logs groups can be used to limit access to sensitive logs generated by services using IAM policy. You don't have to rely on complex compartment hierarchies to secure your logs. For example, say the default log group in a single compartment is where you store logs for the entire tenancy. You grant access to the compartment for log administrators with IAM policy as you normally would. However, let's say some projects contain personally identifiable information (PII) and those logs can only be viewed by a select group of log administrators. Log groups allow you to put logs that contain PII into a separate log group, and then use IAM policy to restrict access to all but a few log administrators.

For more information, see [Service Logs](#) and [Managing Logs and Log Groups](#)

Break Glass Procedure for Accessing Customer's Guest VM

There are situations where some problems can only be resolved by Oracle logging into the customer guest VM.

Below are situations where customer's guest VM access is require and recommended procedures for accessing guest VM:

1. Situations where the **starter database is not yet created and customer do not have ssh access to their guest VM yet**. An example would be SR opened by customer to troubleshoot why customer is unable to create a starter database. In this situation, customer never had access to guest VM and no database have yet been created and hence no customer data exists in guest VM.

As per the security policy associated with ExaDB-D service, Oracle personnel are prohibited to access customer guest VM without customer's explicit permission. To comply

with this policy, Oracle requires to get Customer permission to access guest VM by asking the following question.

"In order for Oracle to resolve the issue described in this SR, we need customer's explicit permission allowing us to login to customer guest VM. By giving us explicit permission to access guest VM, you are confirming that there is no confidential data that is stored in customer guest VM or associated databases and customer security team is authorizing Oracle to have access to customer guest VM in order for Oracle to help fix this issue. Do I have your explicit permission to access guest VM?"

After affirmative response by customer, Oracle support staff can login to customer guest VM to resolve the issue.

2. Situations where a number of databases exist in customer system and customer have access to guest VM but now support needs to login to guest VM to resolve one of the many situations

We have encountered (Nodes doesn't start because of changes on guest VM, eg. Non-existing mounts in fstab, need to run fsck, Hugepage / sysctl conf modification or lvm backup not completed successfully, fstab has wrong entries for non-existing mounts, customer changed the sshd configurations or permissions in /etc/ssh/sshd_config file, etc.) or simply because customer wants Oracle to help resolve the issue they are facing.

This case is more serious than the first one as there could be some sensitive data in customer guest VM file system or database. In this case, our support staff will be required to ask the customer to open a new explicit SR specifically to get this permission with the following SR title and content.

As per the security policy associated with ExaDB-D service, Oracle personnel are prohibited to access customer guest VM without customer's explicit permission. For Oracle to comply with this policy, We are required to ask you to open a new SR with exact language as shown below granting Oracle an explicit permission to access guest VM. Please note any modification to the language below may delay resolution of your SR.

New SR Title: SR granting Oracle explicit permission to access DomU of ExaDB-C@C with AK serial number AK99999999

New SR Content: We are opening this SR to grant explicit permission to Oracle to access our DomU in order for support to help resolve issue described in SR# 1-xxxxxxx.

We acknowledge that by providing this permission, we understand that Oracle will have access to ALL FILES in DomU and agree that there are no confidential

files stored in any of the file systems in DomU. In addition, we also agree that customer security team has authorized Oracle to have access to customer DomU

in order to resolve the issue described in the above SR.

After affirmative response by customer in the above SR, Oracle support staff can login to customer guest VM to resolve the issue.

Part 2: Additional Procedures for Updating Security Posture

- [Customer Responsibilities](#)
A list of Oracle Cloud Operations responsibilities and customer responsibilities for various operations by components
- [Enabling Additional Security Capabilities](#)

Customer Responsibilities

A list of Oracle Cloud Operations responsibilities and customer responsibilities for various operations by components

Table 6-33 Oracle Cloud Ops and Customer Responsibilities for various operations

Operations	Oracle Cloud Ops responsibilities for ORACLE CLOUD PLAFTRM	Customer responsibilities for ORACLE CLOUD PLAFTRM	Oracle Cloud Ops responsibilities for CUSTOMER / TENANT INSTANCES	Customer responsibilities for CUSTOMER / TENANT INSTANCES
DATABASE DEPLOYMENT	Software infrastructure and guidance for ExaCS deployment	Network Admin: Configure cloud network infrastructure (VCN, Backup/Client subnet, Gateway, etc) Database Admin: Setup database requirements (Memory, Storage, Computation, Database version, Database type, etc)	Install Operating System, Database and Grid Infrastructure	Database Admin: Maintain customer hardware requirements based on workloads
MONITORING	Physical Security, Infrastructure, Control Plane, Hardware Faults, Availability, Capacity	Nothing required	Infrastructure availability to support customer monitoring of customer services	Database Admin: Monitoring of Customer Operating System, Databases, Apps and Grid Infrastructure
INCIDENT MANAGEMENT & RESOLUTION	Incident Management and Remediation Spare parts and field dispatch	Nothing required	Support for any incidents related to the underlying platform	Database Admin: Incident Management and resolution for Customer's apps
PATCH MANAGEMENT	Proactive patching of hardware, IaaS/PaaS control stack	Nothing required	Staging of available patches, for example, Oracle Database patch set	Database Admin: Patching of tenant instances Testing
BACKUP & RESTORATION	Infrastructure and Control Plane backup and recovery, recreate customer VMs	Nothing required	Provide running and customer accessible VM	Database Admin: Snapshots / backup and recovery of customer's IaaS and PaaS data using Oracle native or third-party capability

Enabling Additional Security Capabilities

- [KMS Integration \(HSM keys\)](#)
Oracle Exadata Database Service on Dedicated Infrastructure has integration with the OCI Vault service to protect data at rest for its databases. Users now have the control to create and manage TDE master keys within the OCI Vault that protect your Exadata databases.
- [Using Non-default Encryption Algorithms for TDE Tablespace Encryption](#)

KMS Integration (HSM keys)

Oracle Exadata Database Service on Dedicated Infrastructure has integration with the OCI Vault service to protect data at rest for its databases. Users now have the control to create and manage TDE master keys within the OCI Vault that protect your Exadata databases.

With this feature, users have the option to start using the OCI Vault Service to store and manage the master encryption keys. The OCI Vault keys used for protecting databases are stored in a highly available, durable, and managed service. OCI vault integration for ExaDB-D is only available after Oracle Database 11g release 2 (11.2.0.4).

With OCI Vault Service integration with ExaDB-D, customers can now:

- Centrally control and manage your TDE master keys
- Have their TDE master keys stored in a highly available, durable and managed service wherein the keys are protected by hardware security modules (HSM) that meet Federal Information Processing Standards (FIPS) 140-2 Security Level 3 security certification.
- Rotate their encryption keys periodically to maintain security compliance and regulatory needs.
- Migrate from Oracle-managed keys to customer-managed keys for their existing databases.
- The Key version will only be assigned to the container database (CDB), and not to its pluggable database (PDB). PDB will be assigned an automatically generated new key version.

Related Topics

- [Announcing Customer-Managed Encryption Keys for Oracle Exadata Cloud Service](#)
- [Manage Databases on Exadata Cloud Infrastructure](#)

Using Non-default Encryption Algorithms for TDE Tablespace Encryption

In the published Oracle Advanced Security Guide ([section Encrypting Columns in Tables](#)), the methodology to create a table to encrypt columns using a non-default encryption algorithm.

EU Data Act Compliance

Oracle supports the goal of fair and transparent data sharing and access within the European Union.

To learn more about how to export data from Oracle Cloud Infrastructure (OCI) Database as a Service (DBaaS) to on-premises systems or other cloud environments, refer to our compliance page in the base database service documentation.

Related Topics

- [Export Data from OCI Database Services](#)

Troubleshooting Exadata Cloud Infrastructure Systems

These topics cover some common issues you might run into and how to address them.

- [Known Issues for Exadata Cloud Infrastructure](#)
General known issues.
- [Troubleshoot Network Connectivity](#)
To determine if a VM Cluster is properly configured to access the Oracle Cloud Infrastructure (OCI) Services Network, you need to perform the following steps on each virtual machine in the VM Cluster.
- [Backup Failures in Exadata Database Service on Dedicated Infrastructure](#)
If your Exadata managed backup does not successfully complete, you can use the procedures in this topic to troubleshoot and fix the issue.
- [Troubleshooting Oracle Data Guard](#)
Learn to identify and resolve Oracle Data Guard issues.
- [Patching Failures on Exadata Cloud Infrastructure Systems](#)
- [Obtaining Further Assistance](#)
- [Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup](#)
- [Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel](#)

Known Issues for Exadata Cloud Infrastructure

General known issues.

- [CPU Offline Scaling Fails](#)
- [Adding a VM to a VM Cluster Fails](#)

CPU Offline Scaling Fails

Description: CPU offline scaling fails with the following error:

```
** CPU Scale Update **An error occurred during module execution. Please refer to the log file for more information
```

Cause: After provisioning a VM cluster, the `/var/opt/oracle/cprops/cprops.ini` file, which is automatically generated by the database as a service (DBaaS) is not updated with the `common_dcs_agent_bindHost` and `common_dcs_agent_port` parameters and this causes CPU offline scaling to fail.

Action: As the `root` user, manually add the following entries in the `/var/opt/oracle/cprops/cprops.ini` file.

```
common_dcs_agent_bindHost=<IP_Address>  
common_dcs_agent_port=7070
```

Note

The `common_dcs_agent_port` value is **7070** always.

Run the following command to get the IP address:

```
netstat -tunlp | grep 7070
```

For example:

```
netstat -tunlp | grep 7070
tcp 0 0 <IP address 1>:7070 0.0.0.0:* LISTEN 42092/java
tcp 0 0 <IP address 2>:7070 0.0.0.0:* LISTEN 42092/java
```

You can specify either of the two IP addresses, `<IP address 1>` or `<IP address 2>` for the `common_dcs_agent_bindHost` parameter.

Adding a VM to a VM Cluster Fails

Description: When adding a VM to a VM cluster, you might encounter the following issue:

```
[FATAL] [INS-32156] Installer has detected that there are non-readable files in oracle
home.
CAUSE: Following files are non-readable, due to insufficient permission oracle.ahf/data/
scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc
ACTION: Ensure the above files are readable by grid.
```

Cause: Installer has detected a non-readable trace file, `oracle.ahf/data/scaqak03dv0104/diag/tfa/tfactl/user_root/tfa_client.trc` created by Autonomous Health Framework (AHF) in Oracle home that causes adding a cluster VM to fail.

AHF ran as `root` created a `trc` file with `root` ownership, which the `grid` user is not able to read.

Action: Ensure that the AHF trace files are readable by the `grid` user before you add VMs to a VM cluster. To fix the permission issue, run the following commands as `root` on all the existing VM cluster VMs:

```
chown grid:oinstall /u01/app/19.0.0.0/grid/srvm/admin/logging.properties
```

```
chown -R grid:oinstall /u01/app/19.0.0.0/grid/oracle.ahf*
```

```
chown -R grid:oinstall /u01/app/grid/oracle.ahf*
```

Troubleshoot Network Connectivity

To determine if a VM Cluster is properly configured to access the Oracle Cloud Infrastructure (OCI) Services Network, you need to perform the following steps on each virtual machine in the VM Cluster.

Validation check for Identity and Access management connectivity:

- ssh to a virtual machine on your ExaDB-D VM Cluster as `opc` user.
- Execute the command: `curl https://identity.<region>.oci.oraclecloud.com` here `<region>` corresponds to the OCI region where your VM Cluster is deployed. If your VM Cluster is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`. The CURL command will now look like `curl https://identity.us-ashburn-1.oci.oraclecloud.com`.
- If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response as follows:

```
{ "code" : "NotAuthorizedOrNotFound", "message" : "Authorization failed or requested resource not found." }
```
- The SSH session will hang and will eventually timeout if your network is not configured for accessing the OCI Services
- Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Validation check for Object Storage Service (OSS) connectivity:

- ssh to a virtual machine on your ExaDB-D VM Cluster as `opc` user.
- Execute the command: `curl https://objectstorage.<region>.oraclecloud.com`, here `<region>` corresponds to the OCI region where your VM Cluster is deployed. If your VM Cluster is deployed in the Ashburn region you need to use `us-ashburn-1` for `<region>`. The curl command will now look like `curl https://objectstorage.us-ashburn-1.oraclecloud.com`.
- If your Virtual Cloud Network (VCN) is properly configured for accessing the OCI Services Network, you will get an immediate response as follows:

```
{ "code" : "NotAuthorizedOrNotFound", "message" : "Authorization failed or requested resource not found." }
```
- The SSH session will hang and will eventually timeout if your network is not configured for accessing the OCI Services
- Depending on your VCN setup, you will need to follow the steps outlined in the action section below to configure access to the OCI Services Network.

Action:

- **This action is applicable to customers who have deployed their VM Cluster on a private subnet.**
If you haven't already configured a Service Gateway to reach the OCI Services Network, use the instructions in the documentation to configure a Service Gateway for use by the VM Cluster to reach the OCI Services. For more information, see [Option 2: Private Subnets](#).
- **This action is applicable to customers who have deployed their VM Cluster on a public subnet.**
If you haven't already configured an Internet Gateway to reach the OCI Services Network, use the instructions in the documentation to configure the Internet Gateway for use by the VM Cluster to reach OCI Services. For more information, see [Option 1: Public Client Subnet with Internet Gateway](#).

Once you configure your VCN to reach the OCI Services network following the above instructions, execute the steps in both the **Validation check** sections to ensure that you have established connectivity to the OCI Services network from your VM Cluster.

Additional Information:

You can find instructions to update a service gateway [here](https://docs.oracle.com/en-us/iaas/Content/Network/Tasks/servicegateway.htm#switch_label) (https://docs.oracle.com/en-us/iaas/Content/Network/Tasks/servicegateway.htm#switch_label)

Backup Failures in Exadata Database Service on Dedicated Infrastructure

If your Exadata managed backup does not successfully complete, you can use the procedures in this topic to troubleshoot and fix the issue.

The most common causes of backup failure are the following:

- The host cannot access Object Storage
- The database configuration on the host is not correct

The information that follows is organized by the error condition. If you already know the cause, you can skip to the section with the suggested solution. Otherwise, use the procedure in [Determining the Problem](#) to get started.

- [Determining the Problem](#)
In the Console, a failed database backup either displays a status of **Failed** or hangs in the **Backup in Progress** or **Creating** state.
- [Database Service Agent Issues](#)
Your Oracle Cloud Infrastructure database makes use of an agent framework to allow you to manage your database through the cloud platform. Use the following to check and restart the `dbcsagent`.
- [Object Store Connectivity Issues](#)
Backing up your database to Oracle Cloud Infrastructure Object Storage requires that the host can connect to the applicable Swift endpoint.
- [Host Issues](#)
One or more of the following conditions on the database host can cause backups to fail:
- [Database Issues](#)
An improper database state or configuration can lead to failed backups.
- [TDE Wallet and Backup Failures](#)
Learn to identify the root cause of TDE wallet and backup failures.

Determining the Problem

In the Console, a failed database backup either displays a status of **Failed** or hangs in the **Backup in Progress** or **Creating** state.

If the error message does not contain enough information to point you to a solution, you can gather more information by using `dbaascli` and by viewing the log files. Then, refer to the applicable section in this topic for a solution.

NOT_SUPPORTED

Database backups can fail during the `RMAN` configuration stage or during a running `RMAN` backup job. `RMAN` configuration tasks include validating backup destination connectivity, backup module installation, and `RMAN` configuration changes. The log files you examine depend on which stage the failure occurs.

1. Log on to the host as the `oracle` user.
2. Check the applicable log file:

- To identify the job ID of an automated backup, use the `dbaascli database backup --dbname <dbname> --showHistory` command. This displays the history of all backup jobs, including their corresponding job IDs.
- Job logs are available at `/var/opt/oracle/log/dtrs/jobs/`, named using the format `<job_id>.log`. If a job fails, a corresponding debug log `<job_id>.debug` is also generated in the same location.
- You can find the corresponding `RMAN` command execution logs for backup, recovery, and configuration operations in the `/var/opt/oracle/log/<dbname>/dtrs/rman/bkup` directory.

Note

Make sure to review the log files on all compute nodes.

Database Service Agent Issues

Your Oracle Cloud Infrastructure database makes use of an agent framework to allow you to manage your database through the cloud platform. Use the following to check and restart the `dbcsagent`.

Occasionally you might need to restart the `dbcsagent` program if it has the status of **stop/waiting** to resolve a backup failure. View the `/opt/oracle/dcs/log/dcs-agent.log` file to identify issues with the agent.

NOT_SUPPORTED

1. From a command prompt, check the status of the agent:

```
systemctl status dbcsagent.service
```

2. If the agent is in the **stop/waiting** state, try to restart the agent:

```
systemctl start dbcsagent.service
```

3. Check the status of the agent again to confirm that it has the **stop/running** status:

```
systemctl status dbcsagent.service
```

Object Store Connectivity Issues

Backing up your database to Oracle Cloud Infrastructure Object Storage requires that the host can connect to the applicable Swift endpoint.

Though Oracle controls the actual Swift user credentials for the storage bucket for managed backups, verifying general connectivity to Object Storage in your region is a good indicator that object store connectivity is not the issue. You can test this connectivity by using another Swift user.

NOT_SUPPORTED

1. Create a Swift user in your tenancy. See [Working with Auth Tokens](#).

2. With the user you created in the previous step, use the following command to verify the host can access the object store.

```
curl -v -X HEAD -u <user_ID>:'<auth_token>' https://  
swiftobjectstorage.<region_name>.oraclecloud.com/v1/  
<object_storage_namespace>
```

See [Object Storage FAQ](#) for the correct region to use. See [Understanding Object Storage Namespaces](#) for information about your Object Storage namespace.

3. If you cannot connect to the object store, refer to [Prerequisites for Backups on Exadata Cloud Service](#) topic for information on configuring object store connectivity.

Host Issues

One or more of the following conditions on the database host can cause backups to fail:

NOT_SUPPORTED

If an interactive command such as `oraenv`, or any command that might return an error or warning message, was added to the `.bash_profile` file for the `grid` or `oracle` user, Database service operations like automatic backups can be interrupted and fail to complete. Check the `.bash_profile` file for these commands, and remove them.

NOT_SUPPORTED

Backup operations require space in the `/u01` directory on the host file system. Use the `df -h` command on the host to check the space available for backups. If the file system has insufficient space, you can remove old log or trace files to free up space.

NOT_SUPPORTED

Your system might not have the required version of the backup module (`opc_installer.jar`). See [Unable to use Managed Backups in your DB System](#) for details about this known issue. To fix the problem, you can follow the procedure in that section or simply update your DB system and database with the latest bundle patch.

NOT_SUPPORTED

[Customizing](#) the site profile file (`$ORACLE_HOME/sqlplus/admin/glogin.sql`) can cause managed backups to fail in Oracle Cloud Infrastructure. In particular, interactive commands can lead to backup failures. Oracle recommends that you not modify this file for databases hosted in Oracle Cloud Infrastructure.

Database Issues

An improper database state or configuration can lead to failed backups.

NOT_SUPPORTED

The database must be active and running (ideally on all nodes) while the backup is in progress.

NOT_SUPPORTED

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d <db_unique_name> -verbose
```

The system returns a message including the database's instance status. The instance status must be `Open` for the backup to succeed. If the database is not running, use the following command to start it:

```
srvctl start database -d <db_unique_name> -o open
```

If the database is mounted but does not have the `Open` status, use the following commands to access the SQL*Plus command prompt and set the status to `Open`:

```
sqlplus / as sysdba  
alter database open;
```

NOT_SUPPORTED

When you provision a new database, the archiving mode is set to `ARCHIVELOG` by default. This is the required archiving mode for backup operations. Check the archiving mode setting for the database and change it to `ARCHIVELOG`, if applicable.

NOT_SUPPORTED

Open an SQL*Plus command prompt and enter the following command:

```
select log_mode from v$database;
```

If you need to set the archiving mode to `ARCHIVELOG`, start the database in `MOUNT` status (and not `OPEN` status), and use the following command at the SQL*Plus command prompt:

```
alter database archivelog;
```

Confirm that the `db_recovery_file_dest` parameter points to `+RECO`, and that the `log_archive_dest_1` parameter is set to `USE_DB_RECOVERY_FILE_DEST`.

For RAC databases, one instance must have the `MOUNT` status when enabling archivelog mode. To enable archivelog mode for a RAC database, perform the following steps:

1. Shut down all database instances:

```
srvctl stop database -d
```

2. Start one of the database instances in mount state:

```
srvctl start instance -d <db_unique_name> -i <instance_name> -o mount
```

3. Access the SQL*Plus command prompt:

```
sqlplus / as sysdba
```

4. Enable archive log mode:

```
alter database archivelog;  
exit;
```

5. Stop the database:

```
srvctl stop instance -d <db_unique_name> -i <instance_name>
```

6. Restart all database instances:

```
srvctl start database -d <db_unique_name>
```

7. At the SQL*Plus command prompt, confirm the archiving mode is set to: ARCHIVELOG:

```
select log_mode from v$database;
```

NOT_SUPPORTED

Backups can fail when the database instance has a stuck archiver process. For example, this can happen when the flash recovery area (FRA) is full. You can check for this condition using the `srvctl status database -db <db_unique_name> -v` command. If the command returns the following output, you must resolve the stuck archiver process issue before backups can succeed:

```
Instance <instance_identifier> is running on node *<node_identifier>.  
Instance status: Stuck Archiver
```

Refer to [ORA-00257:Archiver Error \(Doc ID 2014425.1\)](#) for information on resolving a stuck archiver process.

After resolving the stuck process, the command should return the following output:

```
Instance <instance_identifier> is running on node *<node_identifier>.  
Instance status: Open
```

If the instance status does not change after you resolve the underlying issue with the device or resource being full or unavailable, try restarting the database using the `srvctl` command to update the status of the database in the clusterware.

NOT_SUPPORTED

Editing certain RMAN configuration parameters can lead to backup failures in Oracle Cloud Infrastructure. To check your RMAN configuration, use the `show all` command at the RMAN command line prompt.

See the following list of parameters for details about RMAN the configuration settings that should not be altered for databases in Oracle Cloud Infrastructure.

NOT_SUPPORTED

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 30 DAYS;
```

```
CONFIGURE CONTROLFILE AUTOBACKUP ON;
```

```
CONFIGURE DEVICE TYPE 'SBT_TAPE' PARALLELISM 5 BACKUP TYPE TO COMPRESSED  
BACKUPSET;
```

```
CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;
```

```
CONFIGURE CHANNEL DEVICE TYPE 'SBT_TAPE' PARMS 'SBT_LIBRARY=/var/opt/oracle/  
dbaas_acfs/<db_name>/opc/libopc.so, ENV=(OPC_PFILE=/var/opt/oracle/dbaas_acfs/  
<db_name>/opc/opc<db_name>.ora)';
```

```
CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO 'SBT_TAPE';
```

```
CONFIGURE CHANNEL DEVICE TYPE DISK MAXPIECESIZE 2 G;
```

```
CONFIGURE ENCRYPTION FOR DATABASE ON;
```

NOT_SUPPORTED

RMAN backups fail when an object store wallet file is lost. The wallet file is necessary to enable connectivity to the object store.

NOT_SUPPORTED

1. Get the name of the database with the backup failure using SQL*Plus:

```
show parameter db_name
```

2. Determine the file path of the backup config parameter file that contains the RMAN wallet information at the Linux command line:

```
locate opc_<database_name>.ora
```

For example:

```
find / -name "opctestdb30.ora" -print /var/opt/oracle/dbaas_acfs/  
testdb30/opctestdb30.ora
```

3. Find the file path to the wallet file in the backup config parameter file by inspecting the value stored in the `OPC_WALLET` parameter. To do this, navigate to the directory containing the backup config parameter file and use the following `cat` command:

```
cat opc<database_name>.ora
```

For example:

```
cd /var/opt/oracle/dbaas_acfs/testdb30/opc/

ls -altr *.ora
opctestdb30.ora

cat opctestdb30.ora
OPC_HOST=https://swiftobjectstorage.us-phoenix-1.oraclecloud.com/v1/
dbbackupphx
OPC_WALLET='LOCATION=file:/var/opt/oracle/dbaas_acfs/testdb30/opc/
opc_wallet CREDENTIAL_ALIAS=alias_opc'
OPC_CONTAINER=BUG3TFsSi8QzjWfuTxqqExample
 OPC_DEFERRED_DELETE=false
```

4. Confirm that the `cwallet.sso` file exists in the directory specified in the `OPC_WALLET` parameter, and confirm that the file has the correct permissions. The file permissions should have the octal value of "600" (`-rw-----`). Use the following command:

```
ls -ltr /var/opt/oracle/dbaas_acfs/<database_name>/opc/opc_wallet
```

For example:

```
ls -altr /var/opt/oracle/dbaas_acfs/testdb30/opc/opc_wallet
-rw----- 1 oracle oinstall 0 Oct 29 01:59 cwallet.sso.lck
-rw----- 1 oracle oinstall 111231 Oct 29 01:59 cwallet.sso
```

TDE Wallet and Backup Failures

Learn to identify the root cause of TDE wallet and backup failures.

NOT_SUPPORTED

For backup operations to work, the `$ORACLE_HOME/network/admin/sqlnet.ora` file must contain the `ENCRYPTION_WALLET_LOCATION` parameter formatted exactly as follows:

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/<database_name>/
tde_wallet)))
```

NOT_SUPPORTED

Use the `cat` command to check the TDE wallet location specification. For example:

```
$ cat $ORACLE_HOME/network/admin/sqlnet.ora
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE)
(METHOD_DATA=(DIRECTORY=/var/opt/oracle/dbaas_acfs/<database_name>/
tde_wallet)))
```

NOT_SUPPORTED

Database backups fail if the TDE wallet is not in the proper state. The following scenarios can cause this problem:

NOT_SUPPORTED

If the database was started using SQL*Plus, and the `ORACLE_UNQNAME` environment variable was not set, the wallet is not opened correctly.

To fix the problem, start the database using the `srvctl` utility:

```
srvctl start database -d <db_unique_name>
```

NOT_SUPPORTED

In a multitenant environment for Oracle Database versions that support PDB-level keystore, each PDB has its own master encryption key. For Oracle 18c databases, this encryption key is stored in a single keystore used by all containers. (Oracle Database 19c does not support a keystore at the PDB level.) After you create or plug in a new PDB, you must create and activate a master encryption key for it. If you do not do so, the `STATUS` column in the `v$encryption_wallet` view shows the value `OPEN_NO_MASTER_KEY`.

To check the master encryption key status and create a master key, do the following:

1. Review the the `STATUS` column in the `v$encryption_wallet` view, as shown in the following example:

```
SQL> alter session set container=pdb2;
Session altered.

SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from
v$encryption_wallet;

WRL_TYPE      WRL_PARAMETER
STATUS                WALLET_TYPE
-----
FILE            /var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/
OPEN_NO_MASTER_KEY AUTOLOGIN
```

2. Confirm that the PDB is in READ WRITE open mode and is not restricted, as shown in the following example:

```
SQL> show pdbs

CON_ID CON_NAME      OPEN MODE          RESTRICTED
-----
2      PDB$SEED      READ ONLY         NO
3      PDB1          READ WRITE        NO
4      PDB2          READ WRITE        NO
```

The PDB cannot be open in restricted mode (the `RESTRICTED` column must show `NO`). If the PDB is currently in restricted mode, review the information in the `PDB_PLUG_IN_VIOLATIONS` view and resolve the issue before continuing. For more information on the `PDB_PLUG_IN_VIOLATIONS` view and the restricted status, review the [Oracle Multitenant Administrator's Guide](#) on pluggable database for your Oracle Database version.

3. Create and activate a master encryption key for the PDB:

- Set the container to the PDB:

```
ALTER SESSION SET CONTAINER = <pdb>;
```

- Create and activate a master encryption key in the PDB by executing the following command:

```
ADMINISTER KEY MANAGEMENT SET KEY USING TAG '<tag>'
FORCE KEYSTORE IDENTIFIED BY <keystore-password> WITH BACKUP USING
'<backup_identifier>';
```

Note the following:

- The `USING TAG` clause is optional and can be used to associate a tag with the new master encryption key.
- The `WITH BACKUP` clause is optional and can be used to create a backup of the keystore before the new master encryption key is created.

You can also use the `dbaascli` commands `dbaascli tde status` and `dbaascli tde rotate masterkey` to investigate and manage your keys.

4. Confirm that the status of the wallet has changed from `OPEN_NO_MASTER_KEY` to `OPEN` by querying the `v$encryption_wallet` view as shown in step 1.

NOT_SUPPORTED

Configuration parameters related to the TDE wallet can cause backups to fail.

NOT_SUPPORTED

Confirm that the wallet status is `open` and the wallet type is `auto login` by checking the `v$encryption_wallet` view. For example:

```
SQL> select status, wrl_parameter, wallet_type from v$encryption_wallet;
```

```
STATUS  WRL_PARAMETER                                WALLET_TYPE
-----
OPEN    /var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/ AUTOLOGIN
```

For pluggable databases (PDBs), ensure that you switch to the appropriate container before querying `v$encryption_wallet` view. For example:

```
$ sqlplus / as sysdba

SQL> alter session set container=pdb1;

Session altered.
```

```
SQL> select WRL_TYPE,WRL_PARAMETER,STATUS,WALLET_TYPE from
v$encryption_wallet;
```

WRL_TYPE	WRL_PARAMETER	STATUS	WALLET_TYPE
FILE	/var/opt/oracle/dbaas_acfs/testdb30/tde_wallet/	OPEN	AUTOLOGIN

NOT_SUPPORTED

The TDE wallet file (`ewallet.p12`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example as the `root` user:

```
# ls -altr /var/opt/oracle/dbaas_acfs/<database_name>/tde_wallet/ewallet.p12
total 76
-rw----- 1 oracle oinstall 5467 Oct 1 20:17 ewallet.p12
```

The TDE wallet file should have file permissions with the octal value "600" (`-rw-----`), and the owner of this file should be a part of the `oinstall` operating system group.

NOT_SUPPORTED

The auto login wallet file (`cwallet.sso`) can cause backups to fail if it is missing, or if it has incompatible file system permissions or ownership. Check the file as shown in the following example as the `root` user:

```
# ls -altr /var/opt/oracle/dbaas_acfs/<database_name>/tde_wallet/cwallet.sso
total 76
-rw----- 1 oracle oinstall 5512 Oct 1 20:18 cwallet.sso
```

The auto login wallet file should have file permissions with the octal value "600" (`-rw-----`), and the owner of this file should be a part of the `oinstall` operating system group.

Troubleshooting Oracle Data Guard

Learn to identify and resolve Oracle Data Guard issues.

When troubleshooting Oracle Data Guard, you must first determine whether the problem occurs during the Data Guard setup and initialization or during Data Guard operation, when lifecycle commands are entered. The steps to identify and resolve the issues are different, depending on the scenario in which they are used.

There are three lifecycle operations: `switchover`, `failover`, and `reinstat`. The Data Guard broker is used for all of these commands. The broker command line interface (`dgmgrl`) is the main tool used to identify and troubleshoot the issues. Although you can use logfiles to identify root causes, `dgmgrl` is faster and easier to use to check and identify an issue.

Setting up and enabling Data Guard involves multiple steps. Log files are created for each step. If any of the steps fail, review the relevant log file to identify and fix the problem.

- Validation of the primary cloud VM Cluster and database
- Validation of the standby cloud VM Cluster

- Recreating and copying files to the standby database (passwordfile and wallets)
- Creating Data Guard through Network (RMAN Duplicate command)
- Configuring Data Guard broker
- Finalizing the setup
- [Troubleshooting Data Guard using logfiles](#)
The tools used to identify the issue and the locations of relevant logfiles are different, depending on the scenario in which they are used.
- [Troubleshooting the Data Guard Setup Process](#)
The following errors might occur in the different steps of the Data Guard setup process. While some errors are displayed within the Console, most of the root causes can be found in the logfiles

Troubleshooting Data Guard using logfiles

The tools used to identify the issue and the locations of relevant logfiles are different, depending on the scenario in which they are used.

Use the following procedures to collect relevant log files to investigate issues. If you are unable to resolve the problem after investigating the log files, contact My Oracle Support.

Note

When preparing collected files for Oracle Support, bundle them into a compressed archive, such as a ZIP file.

NOT_SUPPORTED

On each compute node associated with the Data Guard configuration, gather log files pertaining to the problem you experienced.

- Enablement stage log files (such as those documenting the Create Standby Database operation) and the logs for the corresponding primary or standby system.
- Enablement job ID logfiles. For example: 23.
- Locations of enablement log files by enablement stage and Exadata system (primary or standby).
- Database name logfiles (`db_name` or `db_unique_name`, depending on the file path).

Note

Check all nodes of the corresponding primary and standby Exadata systems. Commands executed on a system may have been run on any of its nodes.

NOT_SUPPORTED

Data Guard Deployer (`DGDeployer`) is the process that performs the configuration. When configuring the primary database, it creates the `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file.

This log should contain the root cause of a failure to configure the primary database.

NOT_SUPPORTED

- The primary log from the dbaasapi command-line utility is: `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log`. Look for entries that contain `dg_api`.
- One standby log from the dbaasapi command-line utility is: `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log`. In this log, look for entries that contain `dg_api`.
- The other standby log is: `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`. This log is the Data Guard configuration log.

NOT_SUPPORTED

- The Oracle Cloud Deployment Engine (ODCE) creates the `/var/opt/oracle/log/<dbname>/ocde/ocde.log` file. This log should contain the cause of a failure to create the standby database.
- The dbaasapi command line utility creates the `var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log` file. Look for entries that contain `dg_api`.
- The Data Guard configuration log file is `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`.

NOT_SUPPORTED

- `DGdeployer` is the process that performs the configuration. It creates the following `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file. This log should contain the root cause of a failure to configure the standby database.
- The dbaasapi command-line utility creates the `/var/opt/oracle/log/dbaasapi/db/dg/<job_ID>.log` file. Look for entries that contain `dg_api`.
- The Data Guard configuration log is `/var/opt/oracle/log/<dbname>/dgcc/dgcc.log`.

NOT_SUPPORTED

`DGdeployer` is the process that performs the configuration. While configuring Data Guard, it creates the `/var/opt/oracle/log/<dbname>/dgdeployer/dgdeployer.log` file. This log should contain the root cause of a failure to configure the primary database.

NOT_SUPPORTED

On each node of the primary and standby sites, gather log files for the related database name (`db_name`).

Note

Check all nodes on both primary and standby Exadata systems. A lifecycle management operation may impact both primary and standby systems.

NOT_SUPPORTED

- **Database alert log:** `/u02/app/oracle/diag/rdbms/<dbname>/<dbinstance>/trace/alert_<dbinstance>.log`
- **Data Guard Broker log:** `/u02/app/oracle/diag/rdbms/<dbname>/<dbinstance>/trace/drc<dbinstance>.log`

- **Cloud tooling log file for Data Guard:** `/var/opt/oracle/log/<dbname>/odg/odg.log`

Troubleshooting the Data Guard Setup Process

The following errors might occur in the different steps of the Data Guard setup process. While some errors are displayed within the Console, most of the root causes can be found in the logfiles

NOT_SUPPORTED

The password entered for enabling Data Guard didn't match the primary admin password for the SYS user. This error occurs during the Validate Primary stage of enablement.

NOT_SUPPORTED

The database may not be running. This error occurs during the Validate Primary stage of enablement. Check with `srvctl` and `sql` on the host to verify that the database is up and running on all nodes.

NOT_SUPPORTED

The primary database could not be configured. Invalid Data Guard commands or failed listener reconfiguration can cause this error.

NOT_SUPPORTED

The TDE wallet could not be created. The Oracle Transparent Database Encryption (TDE) keystore (wallet) files could not be prepared for transportation to the standby site. This error occurs during the create TDE Wallet stage of enablement. Either of the following items can cause failure at this stage:

- The TDE wallet files could not be accessed
- The enablement commands could not create an archive containing the wallet files

Troubleshooting procedure:

1. Ensure that the cluster is accessible. To check the status of a cluster, run the following command:

```
crsctl check cluster -all
```

2. If the cluster is down, run the following command to restart it:

```
crsctl start crs -wait
```

3. If this error occurs when the cluster is accessible, check the logs for create TDE Wallet (enablement stage) to determine cause and resolution for the error.

NOT_SUPPORTED

The archive containing the TDE wallet was likely not transmitted to the standby site. Retrying usually solves the problem.

NOT_SUPPORTED

- The primary and standby sites may not be able to communicate with each other to configure the standby database. These errors occur during the configure standby database

3. Update the applicable passwords to match. Log on to the system host as the `opc` user and run the following commands:

- a. To change the SYS password:

```
sudo dbaascli database changepassword --dbname <database_name>
```

- b. To change the TDE wallet password:

```
sudo dbaascli tde changepassword --dbname <database_name>
```

NOT_SUPPORTED

For possible causes and resolutions to TDE wallet issues, see [TDE Wallet and Backup Failures](#).

NOT_SUPPORTED

When the switchover, failover, and reinstate commands are run, multiple error messages may occur. Refer to the Oracle Database documentation for these error messages.

Note

Oracle recommends using the Data Guard broker command line interface (`dgmgrl`) to validate the configurations.

1. As the `oracle` user, connect to the primary or standby database with `dgmgrl` and verify the configuration and the database:

```
dgmgrl sys/<pwd>@<database>  
DGMGRL> VALIDATE CONFIGURATION VERBOSE  
DGMGRL> VALIDATE DATABASE VERBOSE <PRIMARY>  
DGMGRL> VALIDATE DATABASE VERBOSE <STANDBY>
```

2. Consult the Oracle Database documentation to check for the respective error message. For example:

- **ORA-16766:** Redo apply is stopped.
- **ORA-16853:** Apply lag has exceeded specified threshold.
- **ORA-16664:** Unable to receive the result from a member (under the standby database).
- **ORA-12541:** TNS: no listener (under the primary database)

For cause and resolution, review the errors in [Database Error Messages](#).

Patching Failures on Exadata Cloud Infrastructure Systems

Patching operations can fail for various reasons. Typically, an operation fails because a database node is down, there is insufficient space on the file system, or the virtual machine cannot access the object store.

- [Determining the Problem](#)
In the Console, you can identify a failed patching operation by viewing the patch history of an Exadata Cloud Infrastructure system or an individual database.
- [Troubleshooting and Diagnosis](#)
Diagnose the most common issues that can occur during the patching process of any of the Exadata Cloud Infrastructure components.

Determining the Problem

In the Console, you can identify a failed patching operation by viewing the patch history of an Exadata Cloud Infrastructure system or an individual database.

A patch that was not successfully applied displays a status of `Failed` and includes a brief description of the error that caused the failure. If the error message does not contain enough information to point you to a solution, you can use the database CLI and log files to gather more data. Then, refer to the applicable section in this topic for a solution.

Troubleshooting and Diagnosis

Diagnose the most common issues that can occur during the patching process of any of the Exadata Cloud Infrastructure components.

- [Database Server VM Issues](#)
One or more of the following conditions on the database server VM can cause patching operations to fail.
- [Oracle Grid Infrastructure Issues](#)
One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.
- [Oracle Databases Issues](#)
An improper database state can lead to patching failures.

Database Server VM Issues

One or more of the following conditions on the database server VM can cause patching operations to fail.

Database Server VM Connectivity Problems

Cloud tooling relies on the proper networking and connectivity configuration between virtual machines of a given VM cluster. If the configuration is not set properly, this may incur in failures on all the operations that require cross-node processing. One example can be not being able to download the required files to apply a given patch.

Given the case, you can perform the following actions:

- Verify that your DNS configuration is correct so that the relevant virtual machine addresses are resolvable within the VM cluster.
- Refer to the relevant Cloud Tooling logs as instructed in the *Obtaining Further Assistance* section and contact Oracle Support for further assistance.

Oracle Grid Infrastructure Issues

One or more of the following conditions on Oracle Grid Infrastructure can cause patching operations to fail.

Oracle Grid Infrastructure is Down

Oracle Clusterware enables servers to communicate with each other so that they can function as a collective unit. The cluster software program must be up and running on the VM Cluster for patching operations to complete. Occasionally you might need to restart the Oracle Clusterware to resolve a patching failure.

In such cases, verify the status of the Oracle Grid Infrastructure as follows:

```
./crsctl check cluster
CRS-4537: Cluster Ready Services is online
CRS-4529: Cluster Synchronization Services is online
CRS-4533: Event Manager is online
```

If Oracle Grid Infrastructure is down, then restart by running the following commands:

```
crsctl start cluster -all
```

```
crsctl check cluster
```

Oracle Databases Issues

An improper database state can lead to patching failures.

Oracle Database is Down

The database must be active and running on all the active nodes so the patching operations can be completed successfully across the cluster.

Use the following command to check the state of your database, and ensure that any problems that might have put the database in an improper state are resolved:

```
srvctl status database -d db_unique_name -verbose
```

The system returns a message including the database instance status. The instance status must be **Open** for the patching operation to succeed.

If the database is not running, use the following command to start it:

```
srvctl start database -d db_unique_name -o open
```

Obtaining Further Assistance

If you were unable to resolve the problem using the information in this topic, follow the procedures below to collect relevant database and diagnostic information. After you have collected this information, contact Oracle Support.

- [Collecting Cloud Tooling Logs](#)
Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.
- [Collecting Oracle Diagnostics](#)

Related Topics

- [Oracle Support](#)

Collecting Cloud Tooling Logs

Use the relevant log files that could assist Oracle Support for further investigation and resolution of a given issue.

DBAASCLI Logs

```
/var/opt/oracle/log/dbaascli
```

- `dbaascli.log`

Collecting Oracle Diagnostics

To collect the relevant Oracle diagnostic information and logs, run the `dbaascli diag collect` command.

For more information about the usage of this utility, see *DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check*.

Related Topics

- [DBAAS Tooling: Using dbaascli to Collect Cloud Tooling Logs and Perform a Cloud Tooling Health Check](#)

Standby Database Fails to Restart After Switchover in Oracle Database 11g Oracle Data Guard Setup

Description: After performing the switchover, the new standby (old primary) database remains shut down and fails to restart.

Action: After performing switchover, do the following:

1. Restart the standby database using the `srvctl start database -db <standby dbname>` command.
2. Reload the listener as `grid` user on all primary and standby cluster nodes.
 - To reload the listener using high availability, download and apply patch **25075940** to the Grid home, and then run `lsnrctl reload -with_ha`.
 - To reload the listener, run `lsnrctl reload`.

After reloading the listener, verify that the `<dbname>_DGMRG` services are loaded into the listener using the `lsnrctl status` command.

To download patch 25075940

1. Log in to My Oracle Support.
2. Click **Patches & Updates**.

3. Select **Bug Number** from the **Number/Name or Bug Number (Simple)** drop-down list.
4. Enter the bug number **34741066**, and then click **Search**.
5. From the search results, click the name of the latest patch.
You will be redirected to the **Patch 34741066: LSNRCTL RELOAD -WITH_HA FAILED TO READ THE STATIC ENTRY IN LISTENER.ORA** page.
6. Click **Download**.

Intermittent Failure in PDB Creation When Multiple PDBs are Getting Created in Parallel

Description: PDB creation might fail for subset of PDBs when multiple PDBs are getting created in parallel.

Cause: PDB creation might notice following error intermittently.

```
ORA-03113: end-of-file on communication channel
```

Action: Retry the failed PDB creation.